

Concepts

Eric Margolis & Stephen Laurence

For the official and latest version of this article, please visit the Stanford Encyclopedia website:

<https://plato.stanford.edu/entries/concepts/>

Citation information:

Margolis, Eric and Stephen Laurence, “Concepts”, *The Stanford Encyclopedia of Philosophy* (Fall 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <<https://plato.stanford.edu/archives/fall2023/entries/concepts/>>.

Concepts are the building blocks of thoughts. Consequently, they are crucial to such psychological processes as categorization, inference, memory, learning, and decision-making. This much is relatively uncontroversial. But the nature of concepts—the kind of things concepts are—and the constraints that govern a theory of concepts have been the subject of much debate (Margolis & Laurence 1999, Margolis & Laurence 2015). This is due, at least in part, to the fact that disputes about concepts often reflect deeply opposing approaches to the study of the mind, to language, and even to philosophy itself. In this entry, we provide an overview of theories of concepts, and outline some of the disputes that have shaped debates surrounding the nature of concepts. The entry is organized around five significant issues that are focal points for many theories of concepts. Not every theory of concepts takes a stand on each of the five, but viewed collectively these issues show why the theory of concepts has been such a rich and lively topic in recent years. The five issues are: (1) the ontology of concepts, (2) the structure of concepts, (3) empiricism and nativism about concepts, (4) concepts and natural language, and (5) concepts and conceptual analysis.

1. The ontology of concepts

We begin with the issue of the ontological status of a concept. The three main options are to identify concepts with mental representations, with abilities, and with abstract objects such as Fregean senses.

1.1. Concepts as mental representations

The first of these views maintains that concepts are psychological entities, taking as its starting point the representational theory of the mind (RTM). According to RTM, thinking occurs in an internal system of representation. Beliefs and desires and other propositional attitudes enter into mental processes as internal symbols. For example, Sue might believe that Dave is taller than Cathy, and also believe that Cathy is taller than Ben, and together these may cause Sue to believe that Dave is taller than Ben. Her beliefs would be constituted by mental representations that are about Dave, Cathy and Ben and their relative heights. What makes these beliefs, as opposed to desires or other psychological states, is that the symbols have the characteristic causal-functional role of beliefs. (RTM is usually presented as taking beliefs and other propositional attitudes to be relations between an agent and a mental representation (e.g., Fodor 1987). But given that the

relation in question is a matter of having a representation with a particular type of functional role, it is simpler to say that occurrent beliefs just are mental representations with this functional role.)

Many advocates of RTM take the mental representations involved in beliefs and other propositional attitudes to have internal structure. Accordingly, the representations that figure in Sue's beliefs would be composed of more basic representations. For theorists who adopt the mental representation view of concepts, concepts are identified with these more basic representations.

Early advocates of RTM (e.g., Locke (1690/1975) and Hume (1739/1978)) called these more basic representations *ideas*, and often took them to be mental images. But modern versions of RTM assume that much thought is not grounded in mental images. The classic contemporary treatment maintains, instead, that the internal system of representation has a language-like syntax and a compositional semantics. According to this view, much of thought is grounded in word-like mental representations. This view is often referred to as the *language of thought hypothesis* (Fodor 1975). However, the analogy with language isn't perfect; obviously, the internal symbol system must lack many of the properties associated with a natural language. Nonetheless, like a natural language, the internal system's formulae are taken to have subject/predicate form and to include logical devices, such as quantifiers and variables. In addition, the content of a complex symbol is supposed to be a function of its syntactic structure and the contents of its constituents. Returning to Sue's beliefs, the supposition is that they are composed of such symbols as DAVE, CATHY and TALLER and that her beliefs represent what they do in virtue of the contents of these symbols and how they are arranged.

The mental representation view of concepts is the default position in cognitive science (Carey 2009, Pinker 2007) and enjoys widespread support in the philosophy of mind, particularly among philosophers who view their work as being aligned with research in cognitive science (e.g., Carruthers 2006, Millikan 2000, Fodor 2003, Margolis & Laurence 2007).¹ Supporters of this view argue for it on explanatory grounds. They maintain that concepts and structured mental representations play a crucial role in accounting for the productivity of thought (i.e., the fact that human beings can entertain an unbounded number of thoughts), in explaining how mental processes can be both rational and implemented in the brain, and in accommodating the need for structure-sensitive mental processes (Fodor 1987; see also the entry language of thought hypothesis).

Critics of this view argue that it is possible to have propositional attitudes without having explicitly entertained the relevant mental representations. Daniel Dennett (1977), for example, argues that most people believe *zebras don't wear overcoats in the wild*—and a million other similar facts—even though they have never stopped to consider such matters. Dennett also notes that computing systems can lack representations corresponding to the explanations we cite in characterizing and predicting their behavior. For example, it may make perfect sense to say of a chess-playing computer that it thinks that it is good to get one's queen out early, even though we

¹ In an earlier version of this entry, we cited Millikan (2000) as an example of the abilities view of concepts. But since Millikan thinks of concepts as abilities to produce and use representations, it is perhaps best to consider her a proponent of the mental representation view regarding ontology.

know from how the computer is programmed that it has no representation with that very content (see Dennett 1978, 1987 for these and related criticisms and Fodor 1987 and Manfredi 1993 for responses).

Other critics claim that RTM is too closely associated with commonsense psychology, which they argue should be abandoned as a stagnant and degenerate research program (Churchland 1981; see entry eliminative materialism), or that developments in computational modeling (esp. connectionism and dynamic systems theory) offer alternatives, particularly to the language of thought version of RTM (e.g., see Clark 1993, Van Gelder 1995, Elman et al. 1996, McClelland et al. 2010; see Fodor & Pylyshyn 1988, Marcus 2001, and Gallistel & King 2009 for critical discussion of theories that don't employ combinatorial structure.)

1.2 Concepts as abilities

According to the abilities view, it's wrong to maintain that concepts are mental particulars—concepts are neither mental images nor word-like entities in a language of thought. Rather, concepts are abilities that are peculiar to cognitive agents (e.g., Dummett 1993, Bennett & Hacker 2008, Kenny 2010). The concept CAT, for example, might amount to the ability to discriminate cats from non-cats and to draw certain inferences about cats.

While the abilities view is maintained by a diverse group of philosophers, the most prominent reason for adopting the view is a deep skepticism about the existence and utility of mental representations, skepticism that traces back Ludwig Wittgenstein (1953/1958). One of the most influential arguments along these lines claims that mental representations are explanatorily idle because they reintroduce the very sorts of problems they are supposed to explain. For example, Michael Dummett cautions against trying to explain knowledge of a first language on the model of knowledge of a second language. In the case of a second language, it is reasonable to suppose that understanding the language involves translating its words and sentences into words and sentences of one's first language. But according to Dummett, one can't go on to translate words and sentences of one's first language into a prior mental language. "[T]here is really no sense to speaking of a concept's coming into someone's mind. All we can think of is some image coming to mind which we take as in some way representing the concept, and this gets us no further forward, since we still have to ask in what his associating that concept with that image consists" (Dummett 1993, p. 98). In other words, the mental representation itself is just another item whose significance bears explaining. Either we are involved in a vicious regress, having to invoke yet another layer of representation (and so on indefinitely) or we might as well stop with the external language and explain its significance directly. (For critical discussion of this type of regress argument, see Fodor 1975, Crane 2015, Laurence & Margolis 1997).

Not surprisingly, critics of the abilities view argue in the other direction. They note difficulties that the abilities view inherits by its rejection of mental representations. One is that the view is ill-equipped to explain the productivity of thought; another is that it can say little about mental processes. And if proponents of the abilities view remain neutral about the existence of mental representations, they open themselves to the criticism that explication of these abilities is best given in terms of underlying mental representations and processes (see Fodor 1968 and Chomsky 1980 for general discussion of the anti-intellectualist tradition in the philosophy of mind).

1.3 Concepts as abstract objects

An alternative conception of concepts takes concepts to be abstract objects of one type or another. The idea behind this view is that concepts are the meanings (or “contents”) of words and phrases as opposed to mental objects or mental states. This type of view has most prominently been associated with the view that concepts are Fregean senses (e.g., Peacocke 1992, Zalta 2001), so it is this version of the view that concepts are abstract objects that we will focus on here. For proponents of this view, concepts, as meanings, mediate between thought and language, on the one hand, and referents, on the other. An expression without a referent (“Pegasus”) needn’t lack a meaning, since it still has a sense. Similarly, the same referent can be associated with different expressions (e.g., “Eric Blair” and “George Orwell”) because they convey different senses. Senses are more discriminating than referents. Each sense has a unique perspective on its referent—a mode of presentation which represents the referent in a particular way. Differences in cognitive content trace back to differences in modes of presentation. It’s for this reason that the thought that *George Orwell is Eric Blair* lacks the triviality of the thought that *George Orwell is George Orwell* even though George Orwell and Eric Blair are the same person. Philosophers who take concepts to be senses particularly emphasize this feature of senses. Christopher Peacocke, for example, locates the subject matter of a theory of concepts as follows: “Concepts *C* and *D* are distinct if and only if there are two complete propositional contents that differ at most in that one contains *C* substituted in one or more places for *D*, and one of which is potentially informative while the other is not” (Peacocke 1992, p. 2). In other words, *C* and *D* embody differing modes of presentation. (See the entry Frege for discussion of the sense/reference distinction and for more on the explanatory functions associated with senses. To avoid terminological confusion, we should note that Frege himself did not use the term “concept” for senses, but rather for the referents of predicates. Similarly, it is worth noting that Frege uses the term “thought” to stand for propositions, so for Frege thoughts are not psychological states at all but rather the meanings of psychological states.)

The view that concepts are Fregean senses, like the abilities view, is generally held by philosophers who are opposed to identifying concepts with mental representations. Peacocke himself doesn’t go so far as to argue that mental representations are explanatorily idle, but he does think that mental representations are too fine-grained for philosophical purposes. “It is possible for one and the same concept to receive different mental representations in different individuals” (Peacocke 1992, p. 3). He is also concerned that identifying concepts with mental representations rules out the possibility of there being concepts that human beings have never entertained, or couldn’t ever entertain:

If we accept that a thinker’s possession of a concept must be realized by some subpersonal state involving a mental representation, why not say simply that the concept is the mental representation? Just this proposal is made by Margolis and Laurence (1999, 77). Mental representations that are concepts could even be typed by the corresponding possession condition of the sort I favour. This seems to me an entirely legitimate notion of a kind of mental representation; but it is not quite the notion of a concept. It can, for instance, be true that there are concepts human beings may never acquire, because of their intellectual limitations, or because the sun will expand to eradicate human life before humans reach a stage at which they can acquire these concepts. ‘There are concepts that will never be acquired’ cannot mean or imply ‘There are mental representations which are not mental representations

in anyone's mind'. If concepts are individuated by their possession conditions, on the other hand, there is no problem about the existence of concepts that will never be acquired. They are simply concepts whose possession conditions will never be satisfied by any thinkers. (Peacocke, 2005, p. 169).

Advocates of the mental representation view would respond to these arguments by invoking the type/token distinction with respect to mental representations. Just as two speakers can each produce distinct instances (that is, tokens) of the same type of word (e.g., the word "chair"), different thinkers' minds can produce distinct instances (that is, tokens) of the same type of mental representation (e.g., the mental representation CHAIR). How does this help with the present objection? According to the mental representation view, to say that there are concepts that haven't been acquired by anyone is just to say that, for some types of mental representations, there has not yet been a token in anyone's mind; to say that some concepts will never be acquired by anyone is just to say that, for some types of mental representations, there never will be a token in anyone's mind (Margolis & Laurence 2007).

Critics of the sense-based view have questioned the utility of appealing to such abstract objects (Quine 1960). One difficulty stems from the fact that senses, as abstract entities, stand outside of the causal realm. The question then is how we can access these objects. Advocates of the Fregean sense view describe our access to senses by means of the metaphor of "grasping"—we are said to grasp the sense of an expression. But grasping here is just a metaphor for a cognitive relation that needs to be explicated. Moreover, though senses are hypothesized as providing different modes of presentation for referents, it is not clear why senses themselves do not generate the mode of presentation problem (Fodor 1998). Since they are external to our minds, just as referents typically are, it isn't clear why we can't stand in different epistemic relations towards them just as we can to referents. In the same way that we can have different modes of presentation for a number (the only even prime, the sum of one and one, Tim's favorite number, etc.), we ought to be able to have different modes of presentation for a given sense.

1.4 Is the issue merely terminological?

Stepping back from the details of these three views, there is no reason, in principle, why the different views of concepts couldn't be combined in various ways. For instance, one might maintain that concepts are mental representations that are typed in terms of the Fregean senses they express (see Margolis & Laurence 2007 for discussion).

One might also question whether the dispute about ontology is a substantive dispute. Perhaps there is only a terminological issue about which things ought to be granted the label "concepts". If so, why not just call mental representations "concepts₁", the relevant abilities "concepts₂", senses "concepts₃", and leave it at that?

However, the participants in the dispute don't generally view it as a terminological one. Perhaps this is because they associate their own theories of concepts with large-scale commitments about the way that philosophers should approach the study of mind and language. Undoubtedly, from Dummett's perspective, philosophers who embrace the mental representation view also embrace RTM, and RTM, as he sees it, is fundamentally misguided. Likewise, from Fodor's perspective, RTM is critical to the study of the mind, so an approach like Dummett's, which disallows RTM,

places inappropriate a priori constraints on the study of the mind. Given that the disagreement about concepts is so closely tied to such clearly substantive explanatory disagreements, the debate about what concepts are would itself seem to be a substantive disagreement about what sorts of entities are best suited to playing the central explanatory roles associated with concepts.

Of course, it is possible to introduce new theoretical terms (“concepts1”, “concepts1”, and “concepts3”) for the different theoretical posits made by the different approaches we have been considering (the ability view, the mental representation view, and the abstract object view), and then reconstruct the debate regarding these different approaches as a merely terminological disagreement about which of these terms we should use. But notice that it is possible to do the same thing even for clear-cut substantive debates. For example, we could introduce new terms for moral goodness as understood by deontologists (“good1”) and as understood by consequentialists (“good2”), and then reconstruct the debate about what moral goodness consists in as a merely terminological disagreement about which term we should use, “good1” or “good2”. Or we could introduce “human1” for humans as understood by standard evolutionary accounts (according to which humans have nonhuman primate ancestors) and “human2” for humans understood by creationists (according to which they don’t), and reconstruct the debate between these different approaches as a merely terminological issue about which term we should use, “human1” or “human2”. The danger, of course, is that this would drain the very idea of a *merely terminological debate* of its utility and obscure very real differences between different sorts of disagreements (e.g., stipulative vs. substantive) that would all end up counting as “merely terminological”.

In the remainder of this entry we will frame questions from the perspective of the mental representation view, but readers who prefer a different view can reframe the issues in their preferred terms. For example, the questions “Can there be concepts without language?” is typically understood on the mental representation view as asking whether a pre-linguistic or non-linguistic agent can entertain mental representations of a particular kind. If one adopts the view that concepts are abstract objects, the corresponding question might be whether a pre-linguistic or non-linguistic agent can stand in the concept possession relation to concepts understood as abstract objects of a certain type. And if one adopts the view that concepts are abilities, the corresponding question might be whether a pre-linguistic or non-linguistic agent can have the abilities that are constitutive of concept possession.

2. The structure of concepts

Just as thoughts are composed of concepts, many concepts are themselves complex entities that are composed of other concepts or more basic representational components. In this section, we look at different proposals about the structure of what are often called *lexical concepts*. Roughly speaking, these are concepts that tend to be associated with individual words in natural language—concepts like BIRD and WALK as opposed to manifestly complex concepts like NOCTURNAL ANIMALS THAT LIVE IN TREES or TO WALK AIMLESSLY ON A NEWLY PAVED ROAD. With BIRD and WALK and similar concepts, it isn’t obvious what type of internal structure (if any) they have, and this has led to a great deal of controversy (Margolis & Laurence 1999, Murphy 2002). Accounts of the structure of lexical concepts have

focused on concepts for categories of objects, actions, and events rather than concepts for individuals. Our discussion will share this focus (see, e.g., Blok et al. 2005 and Levine 2010 for discussion of concepts of individuals).

2.1 The classical theory

In one way or another, all theories regarding the structure of concepts are developments of, or reactions to, the *classical theory of concepts*. According to the classical theory, a lexical concept *C* has definitional structure in that it is composed of simpler concepts that express necessary and sufficient conditions for falling under *C*. The stock example is the concept BACHELOR, which is traditionally said to have the constituents UNMARRIED and MAN. If the example is taken at face value, the idea is that something falls under BACHELOR if it is an unmarried man and only if it is an unmarried man. According to the classical theory, lexical concepts generally will exhibit this same sort of definitional structure. This includes such philosophically interesting concepts as TRUTH, GOODNESS, FREEDOM, and JUSTICE.

Before turning to other theories of conceptual structure, it's worth pausing to see what's so appealing about classical or definitional structure. Much of its appeal comes from the way it offers unified treatments of concept acquisition, categorization, and reference determination. In each case, the crucial work is being done by the very same components. Concept acquisition can be understood as a process in which new complex concepts are created by assembling their definitional constituents. Categorization can be understood as a psychological process in which a complex concept is matched to a target item by checking to see if each and every one of its definitional constituents applies to the target. And reference determination, we've already seen, is a matter of whether the definitional constituents do apply to the target.

These considerations alone would be enough to show why the classical theory has been held in such high regard. But the classical theory receives further motivation through its connection with a philosophical method that goes back to antiquity and that continues to exert its influence over contemporary thought. This is the method of *conceptual analysis*. Paradigmatic conceptual analyses offer definitions of concepts that are to be tested against potential counterexamples that are identified via thought experiments. Conceptual analysis is supposed to be a distinctively a priori activity that many take to be the essence of philosophy. To the extent that paradigmatic conceptual analyses are available and successful, this will convey support for the classical theory. Conversely, if the definitions aren't there to be discovered, this would seem to put in jeopardy a venerable view of what philosophy is and how philosophical investigations ought to proceed (see section 5 below).

The classical theory has come under considerable pressure in the last forty years or so, not just in philosophy but in psychology and other fields as well. For psychologists, the main problem has been that the classical theory has difficulty explaining a robust set of empirical findings. At the center of this work is the discovery that certain categories are taken to be more representative or typical and that typicality scores correlate with a wide variety of psychological data (for reviews, see Smith & Medin 1981, Murphy 2002). For instance, apples are judged to be more typical than plums with respect to the category of fruit, and correspondingly apples are judged to have more features in common with fruit. There are many other findings of this kind. One other is that more

typical items are categorized more efficiently. For example, subjects are quicker to judge that apples are a kind of fruit than to judge that plums are. The problem isn't that the classical theory is inconsistent with results like these but that it does nothing to explain them.

In philosophy, the classical theory has been subjected to a number of criticisms but perhaps the most fundamental is that attempts to specify definitions for concepts have a poor track record. Quite simply, there are too few examples of successful definitional analyses, and certainly none that are uncontroversial (Wittgenstein 1953/1958, Fodor 1981). The huge literature on the analysis of knowledge is representative of the state of things. Since Edmund Gettier (1963) first challenged the traditional definition of KNOWLEDGE (as JUSTIFIED TRUE BELIEF), there has been widespread agreement among philosophers that the traditional definition is incorrect or at least incomplete (see the entry the analysis of knowledge)). But no one can seem to agree on what the correct definition is. Despite the enormous amount of effort that has gone into the matter, and the dozens of papers written on the issue, we are still lacking a satisfactory and complete definition. It could be that the problem is that definitions are hard to come by. But another possibility—one that many philosophers are now taking seriously—is that our concepts lack definitional structure.

2.2 The prototype theory

What other type of structure could they have? A non-classical alternative that emerged in the second half of the twentieth century is the *prototype theory* (e.g., Hampton 2006). According to this theory, a lexical concept *C* doesn't have definitional structure but has probabilistic structure in that something falls under *C* just in case it satisfies a sufficient number of properties encoded by *C*'s constituents. The prototype theory has its philosophical roots in Wittgenstein's (1953/1958) famous remark that the things covered by a term often share a family resemblance, and it has its psychological roots in Eleanor Rosch's groundbreaking experimental treatment of much the same idea (Rosch & Mervis 1975, Rosch 1978). The prototype theory is especially at home in dealing with the typicality effects that were left unexplained by the classical theory. One standard strategy is to maintain that, on the prototype theory, categorization is to be understood as a similarity comparison process, where similarity is computed as a function of the number of constituents that two concepts hold in common. On this model, the reason apples are judged to be more typical than plums is that the concept APPLE shares more of its constituents with FRUIT. Likewise, this is why apples are judged to be a kind of fruit faster than plums are.

The prototype theory does well in accounting for a variety of psychological phenomena and it helps to explain why definitions may be so hard to produce. But the prototype theory has its own problems and limitations. One is that its treatment of categorization works best for quick and unreflective judgments. Yet when it comes to more reflective judgments, people go beyond the outcome of a similarity comparison. If asked whether a dog that is surgically altered to look like a raccoon is a dog or a raccoon, the answer for most of us, and even for children, is that it is remains a dog (see Keil 1989, Gelman 2003 for discussion). Another criticism that has been raised against taking concepts to have prototype structure concerns compositionality. When a patently complex concept has a prototype structure, it often has emergent properties, ones that don't derive from the prototypes of its constituents (e.g., PET FISH encodes properties such as brightly colored, which have no basis in the prototype structure for either PET or FISH). Further, many patently complex concepts don't even have a prototype structure (e.g., CHAIRS THAT

WERE PURCHASED ON A WEDNESDAY) (Fodor & Lepore 1996, Fodor 1998; for responses to the arguments from compositionality, see Prinz 2002, Robbins 2002, Hampton & Jönsson 2012, and Del Pinal 2016).

One general solution that addresses all of these problems is to hold that a prototype constitutes just part of the structure of a concept. In addition, concepts have *conceptual cores*, which specify the information relevant to more considered judgments and which underwrite compositional processes. Of course, this just raises the question of what sort of structure conceptual cores have. One suggestion is that conceptual cores have classical structure (Osherson & Smith 1981, Landau 1982). This won't do, however, since it just raises once again most of the problems associated with the classical theory (Laurence & Margolis 1999).

2.3 The theory theory

Another and currently more popular suggestion is that cores are best understood in terms of the *theory theory of concepts*. This is the view that concepts stand in relation to one another in the same way as the terms of a scientific theory and that categorization is a process that strongly resembles scientific theorizing (see, e.g., Carey 1985, 2009, Gopnik & Meltzoff 1997, Keil 1989). It's generally assumed, as well, that the terms of a scientific theory are interdefined so that a theoretical term's content is determined by its unique role in the theory in which it occurs.

The theory theory is especially well-suited to explaining the sorts of reflective categorization judgments that proved to be difficult for the prototype theory. For example, theory theorists maintain that children override perceptual similarity in assessing the situation where the dog is made to look like a raccoon, claiming that even children are in possession of a rudimentary biological theory. This theory tells them that being a dog isn't just a matter of looking like a dog. More important is having the appropriate hidden properties of dogs—the dog essence (see Atran & Medin 2008 on folkbiology). Another advantage of the theory theory is that it is supposed to help to explain important aspects of conceptual development. Conceptual change in childhood is said to follow the same pattern as theory change in science.

One problem that has been raised against the theory theory is that it has difficulty in allowing for different people to possess the same concepts (or even for the same person to have the same concept over time). The reason is that the theory theory is *holistic*. A concept's content is determined by its role in a theory, not by its being composed of just a handful of constituents. Since beliefs that enter people's mental theories are likely to be different from one another (and are likely to change), there may be no principled basis for comparison (Fodor & Lepore 1992). Another problem with the theory theory concerns the analogy to theory change in science. The analogy suggests that children undergo radical conceptual reorganization in the course of cognitive development, but many of the central case studies have proved to be controversial on empirical grounds, with evidence that the relevant concepts are implicated in core knowledge systems that are enriched in development but not fundamentally altered (see Spelke & Kinzler 2007 on core knowledge). However, there are certain specific examples where radical conceptual reorganization is plausible, for instance, when children eventually develop a theory of matter that allows them to differentiate weight from density, and air from nothing (Carey 2009).

2.4 Conceptual atomism

A radical alternative to all of the theories we've mentioned so far is *conceptual atomism*, the view that lexical concepts have no semantic structure (Fodor 1998, Millikan 2000). According to conceptual atomism, the content of a concept isn't determined by its relation to other concepts but by its relation to the world.

Conceptual atomism follows in the anti-descriptivist tradition that traces back to Saul Kripke, Hilary Putnam, and others working in the philosophy of language (see Kripke 1972/80, Putnam 1975, Devitt 1981). Kripke, for example, argues that proper names function like mere tags in that they have no descriptive content (Kripke 1972/80). On a description theory one might suppose that "Gödel" means something like *the discoverer of the incompleteness of arithmetic*. But Kripke points out we could discover that Schmitt really discovered the incompleteness of arithmetic and that Gödel could have killed Schmitt and passed the work off as his own. The point is that if the description theory were correct, we would be referring to Schmitt when we say "Gödel". But intuitively that's not the case at all. In the imagined scenario, the sentence "Gödel discovered the incompleteness of arithmetic" is saying something false about Gödel, not something trivially true about the discoverer of the incompleteness of arithmetic, whoever that might be (though see Machery et al. 2004 on whether this intuition is universal and Genone 2012 for discussion of the role of intuitions in theories of reference). Kripke's alternative account of names is that they achieve their reference by standing in a causal relation to their referents. Conceptual atomism employs a similar strategy while extending the model to all sorts of concepts, not just ones for proper names.

2.5 Pluralism and eliminativism

At present, the nature of conceptual structure remains unsettled. Perhaps part of the problem is that more attention needs to be given to the question of what explanatory work conceptual structure is supposed to do and the possibility that there are different types of structure associated with different explanatory functions. We've seen that conceptual structure is invoked to explain, among other things, typicality effects, reflective categorization, cognitive development, reference determination, and compositionality. But there is no reason to assume that a single type of structure can explain all of these things. As a result, there is no reason why philosophers shouldn't maintain that concepts have different types of structure. For example, notice that atomism is largely motivated by anti-descriptivism. In effect, the atomist maintains that considerable psychological variability is consistent with concepts entering into the same mind-world causal relations, and that it's the latter that determines a concept's reference. But just because the mechanisms of reference determination permit considerable psychological variability doesn't mean that there aren't, in fact, significant patterns for psychologists to uncover. On the contrary, the evidence for typicality effects is impressive by any measure. For this reason, it isn't unreasonable to claim that concepts do have prototype structure even if that structure has nothing to do with the determination of a concept's referent. Similar considerations suggest that concepts may have theory-structure and perhaps other types of structure as well (see Laurence & Margolis 1999 on different types of conceptual structure).

One way of responding to the plurality of conceptual structures is to suppose that concepts have multiple types of structure. This is the central idea behind *conceptual pluralism*. According to one version of conceptual pluralism, suggested by Laurence & Margolis (1999), a given concept will have a variety of different types of structure associated with it as components of the concept in question. For example, concepts may have atomic cores that are linked to prototypes, internalized theories, and so on. On this approach, the different types of structure that are components of a given concept play different explanatory roles. Reference determination and compositionality have more to do with the atomic cores themselves and how they are causally related to things outside of the mind, while rapid categorization and certain inferences depend on prototype structure, and more considered inferences and reasoning depend upon theory structure. Many variants on this general proposal are possible, but the basic idea is that, while concepts have a plurality of different types of structure with different explanatory roles, this differing structure remains unified through the links to an atomic representation that provides a concept's reference. One challenge for this type of account is to delineate which of the cognitive resources that are associated with a concept should be counted as part of its structure and which should not. As a general framework, the account is neutral regarding this question, but as the framework is filled in, clarification will be needed regarding the status of potential types of structure.

A different form of pluralism about conceptual structure doesn't employ atomic cores but simply says that the prototype, theory, etc. are all themselves concepts (Weiskopf 2009). Rather than holding that a single concept (e.g., the concept CAT) has multiple types of structure as components, as in the first form of pluralism, this form takes each type of structure to *be* a concept on its own, resulting in a plurality of concepts (CAT1, CAT2, CAT3, etc). On this view, it is wrong to suppose that there is such a thing as *the* concept CAT. Instead, there are many cat-concepts, each with a different type of structure, where each is involved in just a subset of the high-level psychological processes associated with cats. CAT1, for example, might explain some instances of categorization and some inferences, while CAT2, CAT3, etc. explain others. What's more, on this form of pluralism, people might also differ with respect to which kinds of cat-concepts they possess. And even if two people have a cat-concept with the same general type of structure (e.g., prototype structure), the concepts might still be rather different (treating prototypical cats as having rather different sorts of properties). In response to this second form of pluralism, some philosophers have argued that CAT1, CAT2, etc. may be better understood as different senses of a single CAT concept, on analogy with the different senses of a polysemous linguistic expression (e.g. the *physical object* sense and the *content* sense of "book") (Vicente & Manrique 2016). Regardless of whether one accepts this point about polysemy, an important challenge facing this second ("multiple concepts") version of pluralism is to explain why all of the different cat-concepts count as cat-concepts—that is, to explain what unifies the plurality of cat-concepts. A natural answer to this challenge is that what unifies them is that they all refer to the same category, the category of cats. But it is not so clear that they can all refer to the same category given the differences between the different cat-concepts and the way that they function in cognition. For example, a standard prototype structure would capture prototypical cats and exclude the highly unusual, atypical cats that a theory structure would cover, and consequently the two concepts would refer to distinct (though related) categories.

In all of its forms, pluralism about conceptual structure recognizes that concepts have diverse functions and that a corresponding variety of types of representations are needed to fulfill these

functions. These same considerations have led some theorists to advocate *concept eliminativism*—the view that there are no concepts (Machery 2009). The reasoning behind concept eliminativism is that *concept* should be understood to be a natural kind if concepts exist at all, and that natural kinds ought to have significant commonalities that can be discovered using empirical methods, including commonalities that go well beyond the criteria that are initially used to characterize them. But according to concept eliminativists, there are no such commonalities that hold among the types of representations that pluralists embrace. Perhaps we need prototypes and theories and other types of representations for distinct higher-level cognitive processes, but they are too diverse to warrant the claim that they constitute a single kind. On this view, then, we should simply abandon the theoretical construct of a concept and refer only to more fine-grained types of representations, such as prototypes and theories. Opponents of concept eliminativism have responded to the eliminativist's challenge in a number of ways. Some have argued that Machery's criteria for elimination are simply too strong and that *concept*, understood as a higher-level kind or perhaps a functional kind, has great utility in psychological models of cognitive processes (e.g., Hampton 2010, Lalumera 2010, Strohminger & Moore 2010). Others have argued that Machery's criteria for something's being a natural kind are too restrictive and that his view would have the consequence of ruling out clear cases of legitimate higher-level kinds in science generally (e.g., Gonnerman & Weinberg 2010, Margolis & Laurence 2010). And others have argued that even if we grant Machery's stringent criteria for being a natural kind, elimination wouldn't follow, as concepts *are* natural kinds according to his criteria (Samuels & Ferreira 2010, Weiskopf 2010). (For further critical discussion of eliminativism, see the peer commentary that appears with Machery 2010 and the author's response.)

3. Empiricism and nativism about concepts

One of the oldest questions about concepts concerns whether there are any innate concepts and, if so, how much of the conceptual system is innate. Empiricists maintain that there are few if any innate concepts and that most cognitive capacities are acquired on the basis of a few relatively simple general-purpose cognitive mechanisms. Nativists, on the other hand, maintain that there may be many innate concepts and that the mind has a great deal of innate differentiation into complex domain-specific subsystems.

3.1 Renewed interest in the empiricism/nativism dispute

In recent years, the debate over innate concepts has been reinvigorated as advances in cognitive science have provided philosophers with new tools for revisiting and refining the traditional dispute (see, e.g., Elman et al. 1996, Carruthers, Laurence, & Stich 2005, 2006, 2007, Johnson 2010). Philosophers have greatly benefited from empirical studies in such diverse fields as developmental psychology, evolutionary psychology, cognitive anthropology, neuroscience, linguistics, and ethology. Part of the philosophical interest of this work is that, while the scientists themselves take sides on the empiricist-nativist dispute, their theories and data are often open to interpretation.

As an example, one of the earliest lines of investigation that appeared to support traditional nativist conceptions of the mind was the study of language (Pinker 1994). Noam Chomsky and his followers argued that language acquisition succeeds even though children are only exposed to severely limited evidence about the structure of their language (Chomsky 1967, 1975, 1988; see also Laurence & Margolis 2001). Given the way that the final state (e.g., knowledge of English) outstrips the data that are available to children, we can only postulate that the human mind brings to language acquisition a complex set of language-specific dispositions. For Chomsky, these dispositions are grounded in a set of innate principles that constrain all possible human natural languages, viz., universal grammar (see Baker 2001 and Roberts 2017 on universal grammar).

Not surprisingly, many philosophers have questioned Chomsky's position. The ensuing debate has helped to sharpen the crucial arguments and the extent to which nativist models should continue to command their central place in linguistic theory. (On the empiricist side, see Cowie 1999, Prinz 2002, and Sampson 2005; on the nativist side, see Laurence & Margolis 2001 and Pietroski & Crain 2012; see also the entry innateness and language and the entry philosophy of linguistics). For instance, one of Fiona Cowie's criticisms of Chomsky's poverty of the stimulus argument is that any induction establishes a conclusion that outstrips the available data; hence, going beyond the data in the case of language acquisition doesn't argue for innate *language-specific* dispositions—or else there would have to be a specific innate disposition for every induction we make (for an earlier version of this argument, see Putnam 1967, Goodman 1969). Both Laurence & Margolis (2001) and Crain & Pietroski (2001) respond by teasing out the various ways in which the problem of language acquisition goes beyond general problems about induction.

3.2 Empiricism about concepts

Traditionally, empiricists have argued that all concepts derive from sensations. Concepts were understood to be formed from copies of sensory representations and assembled in accordance with a set of general-purpose learning rules, e.g., Hume's principles of association (Hume 1739/1978). On this view, the content of any concept must be analyzable in terms of its perceptual basis. Any purported concept that fails this test embodies a confusion. Thus David Hume ends his *Enquiry* with the famous remark:

When we run over libraries, persuaded of these principles, what havoc must we make? If we take in our hand any volume; of divinity or school metaphysics, for instance; let us ask, *Does it contain any abstract reasoning concerning quantity or number?* No. *Does it contain any experimental reasoning concerning matter of fact and existence?* No. Commit it then to the flames: For it can contain nothing but sophistry and illusion. (1748/1975, p. 165)

A similar doctrine was maintained by the logical positivists in the early Twentieth Century, though the positivists couched the view in linguistic terms (Ayer 1959). Their principle of verification required for a sentence or statement to be meaningful that it have empirical consequences, and, on some formulations of the principle, that the meaning of a sentence is the empirical procedure for confirming it (see the entry Vienna Circle). Sentences that have no empirical consequences were deemed to be meaningless. Since a good deal of philosophy purports to express propositions that transcend all possible experience, the positivists were happy

to say that these philosophical doctrines are entirely devoid of content and are composed of sentences that aren't merely false but are literally gibberish.

Despite the unpopularity of verificationism (though see Dummett 1993, Wright 1989, and Dennett 1991), a growing number of philosophers are attracted to modified forms of empiricism, forms that primarily emphasize psychological relations between the conceptual system and perceptual and motor states, not semantic relations. An example is Lawrence Shapiro's defense of the claim that the type of body that an organism has profoundly affects its cognitive operations as well as the way that the organism is likely to conceptualize the world (Shapiro 2004). Shapiro's claim is directed against philosophical theories that willfully ignore contingent facts about human bodies as if a human mind could inhere in wildly different body types. Drawing on a number of empirical research programs, Shapiro cites examples that appear to support what he calls *the embodied mind thesis*, viz., that "minds profoundly reflect the bodies in which they are contained" (Shapiro 2004, p. 167).

Jesse Prinz (2002) also defends a modified form of empiricism. Prinz claims that "all (human) concepts are copies or combinations of copies of perceptual representations" (Prinz 2002, p. 108). Though the reference to copies is a nod to Hume, Prinz certainly doesn't buy into Hume's verificationism. In fact, Prinz adopts a causal theory of content of the kind that is usually associated with atomistic theories of concepts (e.g., Fodor 1998); thus Prinz's theory of intentional content doesn't require a concept to inherit the specifically perceptual content of its constituents. Nonetheless, Prinz thinks that every concept derives from perceptual representations. Perhaps the best way to understand the claim is that the mental representations that are activated when someone thinks about something—no matter what the thought—are representations that originate in neural circuits with perceptual or motor functions and that the mental process is affected by that origin. Suppose, for example, that one is thinking about a hammer. Then she is either activating representations that inhere in visual circuits, or representations involved in circuits that control hand shape, etc., and her thought is affected in some way by the primary function of these circuits. Following Lawrence Barsalou (1999; see also Barsalou et al. 2003), Prinz characterizes concept possession as a kind of simulation "tantamount to entering a perceptual state of the kind one would be in if one were to experience the thing it represents" (Prinz 2002, p. 150).

One challenge to this view of cognition is its implication for abstract concepts. It's one thing to say that the concept HAMMER involves the activation of circuits related to hand shape; it's quite another to identify representations tied to a particular sensory modality underlying such concepts as TRUTH, DEMOCRACY, ENTROPY, and NINETEEN (Adams & Campbell 1999, Brewer 1999). Logical concepts are also a challenge. Prinz suggests as a perceptual basis for the concept of disjunction that it is based on feelings of hesitation. However, his more considered view seems to be that logical concepts are best understood as operations, not representations. The resulting theory is one in which thoughts lack logical form. The trouble is that this makes it difficult to see how to distinguish logically equivalent thoughts. A related problem is that, since composition for Prinz does not yield structurally complex representations, there seems to be nothing to distinguish the type of contents associated with judgements (propositional contents) from those associated lists or even single concepts (for related discussion Fodor 2003). Finally, there are difficulties regarding how to interpret behavioral and neurological evidence that is

supposed to support Prinz and Barsalou's case against amodal representations. For example, Machery (2007) points out that proponents of amodal representations typically suppose that imagery is useful in solving certain types of problems. So to argue against amodal representations, it is not enough to show that sensory representations show up in a task in which experimental subjects are not explicitly told to visualize a solution, since the use of sensory-based imagery in a task is perfectly compatible with there also being amodal representations in play. (For further critical discussion of the form of empiricism that is opposed to amodal representations, see Weiskopf 2007, Dove 2009, and Mahon 2015).

3.3 Nativism about concepts

Perhaps the most influential discussion of concepts in relation to the nativism/empiricism debate is Jerry Fodor's (1975, 1981) argument for the claim that virtually all lexical concepts are innate. Fodor (1975) argued that there are theoretical problems with all models of concept learning in that all such models treat concept learning as hypothesis testing. The problem is that the correct hypothesis invariably employs the very concept to be learned and hence the concept has to be available to a learner prior to the learning taking place. In his (1981), Fodor developed this argument by allowing that complex concepts (and only complex concepts) can be learned in that they can be assembled from their constituents during the learning process. He went on to argue that lexical concepts lack semantic structure and consequently that virtually all lexical concepts must be innate—a position known as *radical concept nativism*. Fodor's arguments have had a great deal of influence on debates about nativism and concept learning, especially amongst cognitive scientists. Few if any have endorsed Fodor's radical conclusions, but many have shaped their views of cognitive development at least in part in response to Fodor's arguments (Jackendoff 1989, Levin & Pinker 1991, Spelke & Tsivkin 2001, Carey 2009). And Fodor has convinced many that primitive concepts are in principle unlearnable (see, e.g., Pinker 2007). Fodor's arguments for this conclusion, however, can be challenged in a number of ways. The most direct way to challenge it is to construct an account of what it is to learn a primitive concept and to show that it is immune to Fodor's challenges (Margolis 1998, Laurence & Margolis 2002, Carey 2009).

Fodor himself changed his view about these issues when he revisited them in Fodor (2008). He now maintains that while considerations about the need for hypothesis testing show that no concepts can be learned, not even complex concepts, this does not require concepts to be innate (Fodor 2008). Instead, Fodor suggests that they are largely acquired via processes that are merely biological in that they don't admit of a psychological-level description. Though such a "biological" account of concept acquisition does offer an alternative to the innate/learned dichotomy, there are reasons for supposing that many concepts are learned all the same (Margolis & Laurence 2011). These include the fact that a person's conceptual system is highly sensitive to the surrounding culture. For example, the concept PURGATORY comes from cultural products such as books, stories, and sermons. But clearly these can only succeed in conveying the concept when mediated by the right sort of psychological processes. Acquiring such concepts is a cognitive-level achievement, not a merely biological one.

Although much of the philosophical discussion of concept nativism has focused on Fodor's radical concept nativism, it is important to note that other less radical approaches to the origins

of concepts exist and that these other approaches embrace a fair amount of learning—they just hold that much of this learning is structured by innate special-purpose acquisition systems (for example, Gallistel 1990, Keil 1989, Spelke & Newport 1998, Baillargeon 2008, Carey 2009, Margolis & Laurence 2013, Tooby & Cosmides 2016).

One further issue concerning innate concepts that is in dispute is whether the very idea of innateness makes sense. A common point among those who are skeptical of the notion is the observation that all traits are dependent upon interactions between genes and the environment and that there is no way to fully untangle the two (Elman et al. 1996, Griffiths 2002; see also Clark 1998 and Marcus 2004, and the entry on the distinction between innate and acquired characteristics). Nonetheless, there are clear differences between models of the mind with empiricist leanings and models of the mind with nativist leanings, and the notion of innateness may be thought to earn its usefulness by marking these differences (Margolis & Laurence (2013). For discussion of different proposals of what innateness is see Ariew (1999), Cowie (1999), Samuels (2002), Mallon & Weinberg (2006), Mameli (2008) and Khalidi (2015).

4. Concepts and natural language

We turn now to the issue of how concepts and thoughts relate to language.

4.1 Can there be concepts without language?

Some philosophers maintain that possession of natural language is necessary for having any concepts (Brandom 1994, Davidson 1975, Dummett 1993) and that the tight connection between the two can be established on a priori grounds. In a well known passage, Donald Davidson summarizes his position as follows:

We have the idea of belief only from the role of belief in the interpretation of language, for as a private attitude it is not intelligible except as an adjustment to the public norm provided by language. It follows that a creature must be a member of a speech community if it is to have the concept of belief. And given the dependence of other attitudes on belief, we can say more generally that only a creature that can interpret speech can have the concept of a thought.

Can a creature have a belief if it does not have the concept of belief? It seems to me it cannot, and for this reason. Someone cannot have a belief unless he understands the possibility of being mistaken, and this requires grasping the contrast between truth and error—true belief and false belief. But this contrast, I have argued, can emerge only in the context of interpretation, which alone forces us to the idea of an objective, public truth. (Davidson 1975, p. 170).

The argument links having beliefs and concepts with having the concept of belief. Since Davidson thinks that non-linguistic creatures can't have the concept of belief, they can't have other concepts as well. Why the concept of belief is needed to have other concepts is somewhat obscure in Davidson's writings (Carruthers 1992). And whether language is necessary for this particular concept is not obvious. In fact, there is an ongoing research program in cognitive science that addresses this very issue. A variety of non-linguistic tasks have been given to animals and infants to determine the extent to which they are able to attribute mental states to

others (see Andrews & Beck 2018 for work on animals and Baillargeon et al. 2015 for work on infants). These and related studies provide strong evidence that at least some aspects of theory of mind are nonlinguistic.

Davidson offers a pair of supplementary arguments that may elucidate why he is hesitant to turn the issue over to the cognitive scientists. He gives the example of a man engaging in a non-linguistic task where the man indicates his answer by making a choice, for example, selecting an apple over a pear. Davidson comments that until the man actually says what he has in mind, there will always be a question about the conceptualization guiding his choice. “Repeated tests may make some readings of his actions more plausible than others, but the problem will remain how to determine when he judges two objects of choice to be identical” (1975, p. 163). The second argument points to the difficulties of settling upon a specification of what a non-linguistic creature is thinking. “The dog, we say, knows that its master is home. But does it know that Mr. Smith (who is the master) is home? We have no real idea how to settle, or make sense of, these questions” (1975, p. 163). It’s not clear how seriously Davidson himself takes these arguments. Many philosophers have been unconvinced. Notice that both arguments turn on an underdetermination claim—e.g., that the interpretation of the man’s action is underdetermined by the non-linguistic evidence. But much the same thing is true even if we add what the man says (or to be more precise, if we add what the man utters). The linguistic evidence doesn’t guarantee a correct interpretation any more than the non-linguistic evidence does.

Davidson appears to be employing a very high standard for attributing concepts to animals. In effect, he is asking for proof that our attributions are correct. In contrast, most philosophers who are happy to attribute concepts to animals do so because of a wealth of data that are best explained by appealing to an internal system of representation (e.g., Bermudez 2003, Carruthers 2006, Camp 2009; for overviews within cognitive science, see Gallistel 1990, Bekoff, Allen, & Burghardt 2002, and Shettleworth 2010). For example, many species of birds cache food for later retrieval. Their very survival depends upon their ability to successfully recover, in some cases, more than 10,000 different caches in a single season. Researchers studying one species of caching birds have shown that not only do the birds represent the location of the food, but they integrate this information with information about the quality of the food, its perishability, and whether their caching was observed by other birds. Evidence here comes from demonstrations of selective retrieval and recaching of food items under experimentally controlled conditions. Birds will retrieve more perishable items first. When highly valued food items become highly perishable, they shift strategies to retrieve a higher percentage of less perishable food items. And birds that have themselves stolen food from other birds will selectively recache stored food when they are observed caching it (see Clayton, Bussey, & Dickinson 2003, Emery, Dally, & Clayton 2004). Experimental data of this kind provide evidence for particular concepts in birds (of food types, locations, and so on) as well as surprisingly sophisticated cognitive operations that make use of them.

There is a great deal of controversy among philosophers about the implications of this type of research. Proponents of RTM are, of course, entirely happy with the idea that the scientific theories of what birds are doing can be taken at face value. Other philosophers maintain that if the scientific theories say that birds are computing an algorithm for determining a caching strategy, then this can only be read as a *façon de parler*. Still others will grant that animals have

representations but go on to claim that these representation are of a lesser status, not to be confused with concepts (Brandom 1994, 2000, McDowell 1994).

This raises an interesting question about whether there is a motivated and principled difference between concepts in humans and mere representations in animals (Laurence & Margolis 2012). Philosophers who maintain that there is such a difference often cite the role of concepts in reasoning. For example, Robert Brandom claims that representations in animals do little more than act as reliable mechanisms of discrimination. These representations are supposed to be like thermometers, responding to specific environmental features yet without entering into appropriate inferential processes. However, it's not clear what counts as an appropriate inferential process, and certainly there is room for differing opinions on this point. Moreover, whatever reasoning amounts to, comparative psychology is replete with examples that suggest that animals are capable of far more than reliable detection. Animals may not be as smart as humans, but that doesn't mean they are as dumb as thermometers (see Hurley & Nudds 2006 and Carruthers 2006 on reasoning in animals).

4.2 The priority between language and concepts

Even if it's agreed that it is possible to have concepts in the absence of language, there is a dispute about how the two are related. Some maintain that concepts are prior to and independent of natural language, and that natural language is just a means for conveying thought (Fodor 1975, Pinker 1994). Others maintain that at least some types of thinking (and hence some concepts) occur in the internal system of representation constituting our natural language competence (Carruthers 1996, 2002, Spelke 2003) or that natural language augments our concepts in significant ways (Lupyan 2012).

The arguments for deciding between these two positions involve a mixture of theoretical and empirical considerations. Proponents of the first view have claimed that language is ambiguous in ways that thought presumably is not. For example, the natural language sentence *everyone loves someone* could be interpreted to mean that for each person, there is some person that they love, or to mean that everyone loves one and the same person (Pinker 1994). Proponents of the first view have also argued that since language itself has to be learned, thought is prior to language (Fodor 1975; Pinker 1994). A third and similar consideration is that people seem to be able formulate novel concepts which are given a linguistic label later; the concept comes first, the linguistic label second (Pinker 1994).

Proponents of the alternative view—that some thinking occurs in language—have pointed to the phenomenology of thought. It certainly seems as if we are thinking in language when we “hear” ourselves silently talking to ourselves (Carruthers 1996; see also Langland-Hassan & Vicente 2018 on the nature and significance of inner speech). Another type of consideration that proponents of this view highlight is the finding that success on certain tasks (e.g., spatial reorientation that relies on combining landmark information with geometrical information) is selectively impaired when the linguistic system is engaged but not when comparable attention is given to non-linguistic distractors. The suggestion is that solving these tasks requires thinking in one's natural language and that some of the crucial concepts must be couched linguistically (Hermer-Vazquez, Spelke, & Katsnelson 2001; Shusterman & Spelke 2005; Carruthers 2002; see

also Gleitman & Papafragou 2012 for a critical overview of a range of tasks where online processing has been thought to rely on linguistic representation).

4.3 Linguistic determinism and linguistic relativity

Finally, one further issue that bears mentioning is the status of various claims regarding linguistic determinism and linguistic relativity. *Linguistic determinism* is the doctrine that the language a person speaks both causes her to conceptualize the world in certain ways and limits what she can think about by imposing boundaries on her conceptual system; as a result, people who speak very different languages are likely to conceptualize the world in correspondingly different ways. *Linguistic relativity* is the weaker doctrine that the language one speaks influences how one thinks.

Linguistic determinism is historically associated with the writings of Benjamin Lee Whorf (Whorf 1956). Whorf was especially interested in the languages of the indigenous people of America. He famously argued that the Hopi both speak and think about time in ways that are incongruent with European languages and thought. Rather than viewing time as a continuum that flows evenly throughout the universe and that can be broken up into countable events occurring in the past, present, and future, the Hopi are supposed to focus on change as a process. Their conceptual system is also supposed to differ from ours in that it embodies a distinction between things that are or have been accessible to perception versus things that are not, where the latter category includes things in the future as well as mythical and mental constructs.

The claim that the Hopi lack our concept of time has not stood up to scrutiny. Whorf used clumsy translations of Hopi speech that concealed the extent to which they talk about time (references to yesterday, tomorrow, days of the week, lunar phases, etc.). More interestingly, Whorf provided no direct evidence of how the Hopi think. Instead, he used the circular reasoning that they don't *think* about time as we do because they don't *talk* about time as we do. In fact, the Hopi use numerous familiar devices for time keeping, such as calendar strings and sun dials, and their sensitivity to time is evident a wide variety of cultural practices (Pinker 1994).

Linguistic determinism isn't an especially promising doctrine and has few adherents these days, but linguistic relativity is the subject of a spirited debate (see Gumperz & Levinson 1996, Bowerman & Levinson 2001, and Gentner & Goldin-Meadow 2003). Some recent examples of particular interest include whether language influences how we conceptualize spatial frames of reference (in non-linguistic spatial reasoning) (e.g., Li & Gleitman 2002, Levinson et al. 2002, Levinson 2003, Li et al. 2011), spatial relations (in non-linguistic reasoning) (e.g., Choi & Bowerman 1991, Hespos & Spelke 2004), sex (the impact of grammatical gender) (e.g., Boroditsky, Schmidt, & Phillips 2003) and number (e.g., Gordon 2004, Pica et al. 2004, Laurence & Margolis 2007).

5. Concepts and conceptual analysis

Some of the deepest divides in contemporary philosophy concern the limits of empirical inquiry, the status of conceptual analysis, and the nature of philosophy itself (see, e.g., Chalmers 1996,

Jackson 1998, DePaul & Ramsey 1998, Block & Stalnaker 1999, Williamson 2007, Stich 2012, Machery 2017, Strevens 2018). And concepts are right at the center of these disputes. For many, philosophy is essentially the a priori analysis of concepts, which can and should be done without leaving the proverbial armchair. We've already seen that in the paradigm case, an analysis embodies a definition; it specifies a set of conditions that are individually necessary and jointly sufficient for the application of the concept. For proponents of traditional conceptual analysis, the analysis of a concept is successful to the extent that the proposed definition matches people's intuitions about particular cases, including hypothetical cases that figure in crucial thought experiments.

5.1 Attractions of conceptual analysis

Conceptual analysis has been attractive to philosophers for a number of reasons. One is that it makes sense of a good deal of philosophical practice—what George Bealer (1998) calls *the standard justificatory procedure*. Thought experiments and the elicitation of intuitions are characteristic features of philosophical practice. If this practice is justified, then there has to be an understanding of what philosophy is that explains why. Conceptual analysis is supposed to provide just what's needed here. Intuitions can be said to be of value to philosophy precisely because they help us to get clearer about our concepts, especially concepts of intrinsic philosophical interest (JUSTICE, KNOWLEDGE, etc.).

A related attraction is that conceptual analysis explains how philosophy could be an a priori discipline, as many suppose it is. If philosophy is primarily about concepts and concepts can be investigated from the armchair, then the a priori character of philosophy is secured (Jackson 1998).

A third attraction of conceptual analysis is that conceptual analysis has been argued to be a necessary precursor for answering questions about ontological reduction, that is, the sort of reduction that takes place when it's argued that genes are DNA segments, that sensations are brain states, and so on (Chalmers 1996, Jackson 1998). According to one way of filling this view out, one has to begin with an a priori analysis of the higher-level concept, particularly an analysis that makes explicit its causal relations. One can then appeal to empirical findings regarding the things that actually have those causal relations. For example, neuroscience may reveal that such-and-such brain state has the causal relations that analysis reveals to be constitutive of our concept of pain. In the course of doing this, neuroscience is supposed to be showing us what pain is (Lewis 1966, Armstrong 1968). But neuroscience is only in a position to do this against the background of the philosophical work that goes into articulating the concept. (For detailed treatments of this view of reduction, see Chalmers 1996, 2010 and Jackson 1998—though it should be noted that Chalmers argues that PAIN, and other concepts of conscious mental states, cannot be analyzed solely in terms of their causal relations and concludes from this that consciousness itself is irreducible.) This work has generated a great deal of debate (e.g., Block & Stalnaker 1999, Levine 2001, Laurence & Margolis 2003, Yablo 2008, Papineau 2002).

A fourth attraction is that conceptual analysis may offer normative guidance (Goldman 1986). For instance, epistemologists face the question of whether our inferential practices are justified and, if so, what justifies them. One standard answer is that they can be justified if they conform

to our intuitions about what counts as a justified inference (Goldman 1986). In other words, an analysis of our concept of justification is supposed to be all that is needed in order to establish that a set of inference rules is justified. So if it ever turned out that different groups of people employed qualitatively different sets of inferential principles, we could establish the epistemically preferable one by showing that it does a better job of conforming to our concept of justification.

5.2 Objections to conceptual analysis

Many philosophers who are opposed to conceptual analysis identify their approach as being *naturalistic* (e.g., Devitt 1996, Kornblith 2002, Papineau 2013; see also the entry *naturalism*). A common theme of this work is that philosophy is supposed to be continuous with science and that philosophical theories are to be defended on largely explanatory grounds, not on the basis of a priori arguments that appeal to intuition. Accordingly, perceived difficulties with conceptual analysis provide arguments for naturalism.

One such argument centers around the failures of the classical theory of concepts. Earlier, in Section 2, we noted that paradigmatic conceptual analyses require concepts to have classical structure, an assumption that is increasingly difficult to maintain. For this reason, a number of philosophers have expressed skepticism about the viability of conceptual analysis as a philosophical method (e.g., Ramsey 1998, Stich 1992). Others, however, have called into question the connection between conceptual analysis and definitions (Chalmers & Jackson 2001).

Another objection to conceptual analysis is that the intuitions that philosophers routinely rely upon may not be shared. Anyone who teaches philosophy certainly knows that half the time students have the “wrong intuitions”. But who are we to say that they are wrong? And given that people disagree about their intuitions, these can hardly be treated as objective data (Cummins 1998).

Things become even more interesting if we branch out to other cultures. In a preliminary study of East Asian vs. Western intuitions, Jonathan Weinberg, Shaun Nichols, & Stephen Stich (2001) found that East Asians often have the “wrong intuitions” regarding variations on classic philosophical thought experiments, including Gettier-type thought experiments (though see Machery et al 2017 for evidence that some of these intuitions may in fact be universal). At the very least, this work suggests that philosophers should be cautious about moving from their own intuitions to claims about the proper analysis of a concept.

What’s more, the possibility of cultural diversity raises a troubling question for philosophers who want to establish normative claims on the basis of analyses of concepts, such as the concept of justification. Suppose, for example, that East Asian culture offers a different concept of justification than the one that is embedded in Western commonsense thought (assuming for sake of argument that there is a single concept of justification in each culture). In addition, suppose that East Asians employ different inferential practices than our own and that their practices do a fair job of conforming to their concept of justification and that ours do a fair job of conforming to our own. On what basis, then, are we to compare and evaluate these differing practices? Does

it really make sense to say that ours are superior on the ground that they conform better to *our* concept of justification? Wouldn't this just be a form of epistemic prejudice? After all, the question arises whether, given the two concepts of justification, ours is the one that ought to be used for performing normative epistemic evaluations (Stich 1990; for further discussion see Williamson 2005, Sosa 2009, Stich 2009, Weinberg, Nichols, & Stich 2001, Weinberg et al. 2010).

The critique of traditional philosophical analysis has also generated proposals for various types of normative, revisionary projects. Instead of asking which cases are or are not covered by a concept, they ask how a concept should be modified or what new concept should be adopted in its place, given the practical context in which the concept is used—for example, given the goal of promoting social equality (Haslanger 2012, Cappelen 2018).

Much is at stake in the debate between conceptual analysts and naturalists, and it is likely to be a central topic in the theory of concepts for the foreseeable future.

Bibliography

Adams, F. & Campbell, K. (1999). Modality and Abstract Concepts. *Behavioral and Brain Sciences*, 22: 610.

Andrews, K. & Beck, J. (2017). *The Routledge Handbook of Philosophy of Animal Minds*, London: Routledge.

Ariew, A. (1999). Innateness is Canalization: In Defense of a Developmental Account of Innateness, in V. Hardcastle (ed.), *Where Biology Meets Psychology: Philosophical Essays*, pp. 117–38. Cambridge, MA: MIT Press.

Armstrong, D. (1968). *A Materialist Theory of Mind*, London: Routledge.

Atran, S. & Medin, D. (2008). *The Native Mind and the Cultural Construction of Nature*, Cambridge, MA: MIT Press.

Ayer, A.J. (1959). *Logical Positivism*, New York: The Free Press.

Baker, M. (2001). *The Atoms of Language: The Mind's Hidden Rules of Grammar*, New York: Basic Books.

Baillargeon, R. (2008). Innate Ideas Revisited: For a Principle of Persistence in Infants' Physical Reasoning. *Perspectives on Psychological Science*, 3 (1): 2–13.

Baillargeon, R., Scott, R.M., He, Z., Sloane, S., Setoh, P., Jin, K.-S., Wu, D., & Bian, L. (2015). Psychological and Sociomoral Reasoning in Infancy, in M. Mikulincer & P.R. Shaver (eds.), *APA Handbook of Personality and Social Psychology: Volume 1. Attitudes and Social Cognition*, Washington, DC: American Psychological Association, pp. 79–150.

Barsalou, L. (1999). Perceptual Symbol Systems. *Behavioral and Brain Sciences*, 22: 577–609.

Barsalou, L., Simmons, W., Barbey, A., & Wilson, C. (2003). Grounding Conceptual Knowledge in Modality-Specific Systems. *Trends in Cognitive Science*, 7 (2): 84–89.

Bealer, G. (1998). Intuition and the Autonomy of Philosophy, in M. DePaul & W. Ramsey (eds.) *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Lanham, MD: Rowman & Littlefield, pp. 201–239.

Bekoff, M., Allen, C., & Burghardt, G. (2002). *The Cognitive Animal: Empirical and Theoretical Perspectives on Animal Cognition*, Cambridge, MA: MIT Press.

Bennett, M. & Hacker, P. (2008). *History of Cognitive Neuroscience*, Oxford: Wiley-Blackwell.

Bermudez, J. (2003). *Thinking Without Words*, Oxford: Oxford University Press.

Block, N. & Stalnaker, R. (1999). Conceptual Analysis and the Explanatory Gap, *Philosophical Review*, 108 (1): 1–46.

Blok, S., Newman, G. & Rips, L.J. (2005). Individuals and Their Concepts, in W.-K. Ahn, R.L. Goldstone, B.C. Love, A.B. Markman, & P. Wolff (eds.), *Categorization Inside and Outside the Laboratory: Essays in Honor of Douglas L. Medin*, Washington, DC: American Psychological Association, pp. 127–149.

Boroditsky, L., Schmidt, L., & Phillips, W. (2003). Sex, Syntax, and Semantics, in D. Gentner & S. Goldin-Meadow (eds.) *Language in Mind*, Cambridge, MA: MIT Press, pp. 61–79.

Bowerman, M. & Levinson, S. (2001). *Language Acquisition and Conceptual Development*, Cambridge: Cambridge University Press.

Brandom, R. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*, Cambridge, MA: Harvard University Press.

— (2000). *Articulating Reasons: An Introduction to Inferentialism*, Cambridge, MA: Harvard University Press.

Brewer, W. (1999). Perceptual Symbols: The Power and Limitations of a Theory of Dynamic Imagery and Structured Frames, *Behavioral and Brain Sciences*, 22: 611–12.

Camp, E. (2009). A Language of Baboon Thought?, in R.W. Lurz (ed.) *The Philosophy of Animal Minds*, Cambridge: Cambridge University Press, pp. 108–127.

Cappelen, H. (2018). *Fixing Language: An Essay on Conceptual Engineering*, Oxford: Oxford University Press.

Carey, S. (1985). *Conceptual Change in Childhood*, Cambridge, MA: MIT Press.

— (2009). *The Origin of Concepts*, Oxford: Oxford University Press.

Carruthers, P. (1992). *The Animals Issue: Moral Theory in Practice*, Cambridge: Cambridge University Press.

— (1996). *Language, Thought, and Consciousness*, Cambridge: Cambridge University Press.

- (2002). The Cognitive Functions of Language, *Behavioral and Brain Sciences*, 25: 657–74.
- (2006). *The Architecture of the Mind: Massive Modularity and the Flexibility of Thought*, New York: Oxford University Press.
- Carruthers, P., Laurence, S., & Stich, S. (2005). *The Innate Mind: Structure and Contents*, Oxford: Oxford University Press.
- (2006). *The Innate Mind: Culture and Cognition*, Oxford: Oxford University Press.
- (2007). *The Innate Mind: Foundations and the Future*, Oxford: Oxford University Press.
- Chalmers, D. (1996). *The Conscious Mind: In Search of a Fundamental Theory*, Oxford: Oxford University Press.
-
- (2010). *The Character of Consciousness*, Oxford: Oxford University Press.
- Chalmers, D. & Jackson, F. (2001). Conceptual Analysis and Reductive Explanation, *Philosophical Review*, 110: 315–61.
- Choi, S. & Bowerman, M. (1991). Learning to Express Motion Events in English and Korean: The Influence of Language-Specific Lexicalization Patterns, *Cognition*, 41: 83–121.
- Chomsky, N. (1967). Recent Contributions to the Theory of Innate Ideas, *Synthese*, 17: 2–11. Reprinted in S. Stich (ed.) *Innate Ideas*, Berkeley: University of California Press, pp. 121–31.
- (1975). *Reflections on Language*, New York: Pantheon.
- (1980). *Rules and Representations*, New York: Columbia University Press.
- (1988). *Language and the Problems of Knowledge: The Managua Lectures*, Cambridge, MA: MIT Press.
- Churchland, P. (1981). Eliminative Materialism and Propositional Attitudes, *Journal of Philosophy*, 78 (2): 67–90.
- Clark, A. (1993). *Associative Engines: Connectionism, Concepts, and Representational Change*, Cambridge, MA: MIT Press.
- (1998). Twisted Tales: Causal Complexity and Cognitive Scientific Explanation, *Minds and Machines*, 8: 76–99.
- Clayton, N., Bussey, T., & Dickinson, A. (2003). Can Animals Recall the Past and Plan for the Future? *Nature Reviews: Neuroscience*, 4: 685–91.
- Cowie, F. (1999). *What's Within: Nativism Reconsidered*, Oxford: Oxford University Press.
- Crain, S. & Pietroski, P. (2001). Nature, Nurture, and Universal Grammar, *Linguistics and Philosophy*, 24: 139–186.
- Crane, T. (2015). *The Mechanical Mind, Third Edition*, London: Routledge.

- Cummins, R. (1998). Reflection on Reflective Equilibrium, in M. DePaul & W. Ramsey (eds.), *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Lanham, MD: Rowman & Littlefield, pp. 113–27.
- Davidson, D. (1975). Thought and Talk. In his *Inquiries into Truth and Interpretation*, Oxford: Oxford University Press.
- Del Pinal, G. (2016). Prototypes as Compositional Components of Concepts, *Synthese*, 193 (9): 2899–2927.
- Dennett, D. (1977). A Cure for the Common Code. Reprinted in D. Dennett (1978) *Brainstorms*, Cambridge, MA: MIT Press.
- (1978). *Brainstorms: Philosophical Essays on Mind and Psychology*, Cambridge, MA: MIT Press.
- (1987). *The Intentional Stance*, Cambridge, MA: MIT Press.
- (1991). *Consciousness Explained*, Boston: Little, Brown and Company.
- DePaul, M. & Ramsey, W. (1998). *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Lanham, MD: Rowman & Littlefield.
- Devitt, M. (1981). *Designation*, New York: Columbia University Press.
- (1996). *Coming to Our Senses: A Naturalistic Program for Semantic Localism*, Cambridge: Cambridge University Press.
- Dove, G. (2009). Beyond Perceptual Symbols: A Call for Representational Pluralism, *Cognition*, 110: 412–431.
- Dummett, M. (1993). *Seas of Language*, Oxford: Oxford University Press.
- Elman, J., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D., and Plunkett, K. (1996). *Rethinking Innateness: A Connectionist Perspective on Development*, Cambridge, MA: MIT Press.
- Emery, N., Dally, J., & Clayton, N. (2004). Western Scrub-Jays (*Aphelocoma Californica*) Use Cognitive Strategies to Protect Their Caches from Thieving Conspecifics, *Animal Cognition*, 7: 37–43.
- Fodor, J. (1968). The Appeal to Tacit Knowledge, *Journal of Philosophy*, 65 (20): 627–640.
- (1975). *The Language of Thought*, Cambridge, MA: Harvard University Press.
- (1981). The Present Status of the Innateness Controversy. In his *Representations: Philosophical Essays on the Foundations of Cognitive Science*, Cambridge, MA: MIT Press.
- (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, Cambridge, MA: MIT Press.

- (1998). *Concepts: Where Cognitive Science Went Wrong*, New York: Oxford University Press.
- (2003). *Hume Variations*, Oxford: Oxford University Press.
- (2008). *LOT 2: The Language of Thought Revisited*, New York: Oxford University Press.
- Fodor, J. & Lepore, E. (1992). *Holism: A Shopper's Guide*, Oxford: Blackwell.
- (1996). The Red Herring and the Pet Fish: Why Concepts Still Can't Be Prototypes, *Cognition*, 58: 253–270.
- Fodor, J. & Pylyshyn, Z. (1988). Connectionism and Cognitive Architecture: A Critical Analysis, *Cognition*, 28: 3–71.
- Gallistel, C.R. & King, A. (2009). *Memory and the Computational Brain: Why Cognitive Science Will Transform Neuroscience*, Oxford: Wiley-Blackwell.
- (1990). *The Organization of Learning*, Cambridge, MA: MIT Press.
- Gelman, S. (2003). *The Essential Child: Origins of Essentialism in Everyday Thought*, Oxford: Oxford University Press.
- Genone, E. (1963). Theories of Reference and Experimental Philosophy. *Philosophy Compass*, 7 (2): 152–163.
- Gentner, D. & Goldin-Meadow, S. (2003). *Language in Mind*, Cambridge, MA: MIT Press.
- Gettier, E. (1963). Is Justified True Belief Knowledge? *Analysis*, 23: 121–123.
- Gleitman, L. & Papafragou, A. (2012). New Perspectives on Language and Thought, in K. Holyoak & R. Morrison (eds.), *Oxford Handbook of Thinking and Reasoning*, Oxford: Oxford University Press, pp. 543–568.
- Goldman, A. (1986). *Epistemology and Cognition*, Cambridge, MA: MIT Press.
- Gonnerman, C. & Weinberg, J. (2010). Two Uneliminated Uses for 'Concepts': Hybrids and Guides for Inquiry, *Behavioral and Brain Sciences*, 33: 211–212.
- Goodman, N. (1969). The Emporer's New Ideas, in S. Hook (ed.), *Language and Philosophy*, New York: NYU Press, pp. 138–42.
- Gopnik, A., & Meltzoff, A. (1997). *Words, Thoughts, and Theories*, Cambridge, MA: MIT Press.
- Gordon, P. (2004). Numerical Cognition Without Words: Evidence from Amazonia, *Science*, 306: 496–499.
- Griffiths, P. (2002). What Is Innateness? *Monist*, 85 (1): 70–85.
- Gumperz, J. & Levinson, S. (1996). *Rethinking Linguistic Relativity*, Cambridge: Cambridge University Press.

- Hampton, J. (2006). Concepts as Prototypes, in B.H. Ross (ed.), *Psychology of Learning and Motivation* (Volume 46), New York: Academic Press, pp. 79–113.
- (2010). Concept Talk Cannot Be Avoided, *Behavioral and Brain Sciences*, 33: 212–213.
- Hampton, J. & Jönsson, M. (2012). Typicality and Compositionality: The Logic of Combining Vague Concepts, in M. Werning, W. Hintzen, & E. Machery (eds.), *Oxford Handbook of Compositionality*. Oxford: Oxford University Press.
- Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique*, Oxford: Oxford University Press.
- Hermer-Vazquez, L., Spelke, E., & Katsnelson, A. (2001). Sources of Flexibility in Human Cognition: Dual-Task Studies of Space and Language, *Cognitive Psychology*, 39: 3–36.
- Hespos, S. & Spelke, E. (2004). Conceptual Precursors to Language, *Nature*, 430: 453–456.
- Hume, D. (1739/1978). *A Treatise of Human Nature*, Oxford: Oxford University Press.
- (1748/1975). *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, Oxford: Oxford University Press.
- Hurley, S. & Nudds, M. (2006) *Rational Animals?*, Oxford: Oxford University Press.
- Jackendoff, R. (1989). What is a Concept, that a Person May Grasp It? *Mind & Language*, 4: 68–102.
- Jackson, F. (1998). *From Metaphysics to Ethics: A Defense of Conceptual Analysis*, Oxford: Oxford University Press.
- Johnson, S.P. (2010). *Neuroconstructivism: The New Science of Cognitive Development*, Oxford: Oxford University Press.
- Keil, F. (1989). *Concepts, Kinds, and Cognitive Development*, Cambridge, MA: MIT Press.
- Kenny, A. (2010). Concepts, Brains, and Behaviour, *Grazer Philosophische Studien*, 81 (1): 105– 113.
- Khalidi, M.A. (2015). Innateness as a Natural Cognitive Kind, *Philosophical Psychology*, 29 (3): 319–333.
- Kornblith, H. (2002). *Knowledge and Its Place in Nature*, Oxford: Oxford University Press.
- Kripke, S. (1972/1980). *Naming and Necessity*, Cambridge, MA: Harvard University Press.
- Lalumera, E. (2010). Concepts are a Functional Kind, *Behavioral and Brain Sciences*, 33: 217–218.
- Landau, B. (1982). Will the Real Grandmother Please Stand Up? The Psychological Reality of Dual Meaning Representations, *Journal of Psycholinguistic Research*, 11 (1): 47–62.
- Langland-Hassan, P. & Vicente, A. (2018). *Inner Speech: New Voices*, Oxford: Oxford University Press.
- Laurence, S. & Margolis, E. (1997). Regress Arguments Against the Language of Thought, *Analysis*, 57: 60–66.

- (1999). Concepts and Cognitive Science, in E. Margolis & S. Laurence (eds.), *Concepts: Core Readings*, Cambridge, MA: MIT Press, pp. 3–81.
- (2001). The Poverty of the Stimulus Argument, *British Journal for the Philosophy of Science*, 52: 217–76.
- (2002). Radical Concept Nativism, *Cognition*, 86: 25–55.
- (2007). Linguistic Determinism and the Innate Basis of Number, in P. Carruthers, S. Laurence, & S. Stich (eds.), *The Innate Mind: Foundations and the Future*, New York: Oxford University Press, pp. 139–169.
- (2012). The Scope of the Conceptual, in E. Margolis, R. Samuels, & S. Stich (eds.), *The Oxford Handbook of Philosophy of Cognitive Science*, New York: Oxford University Press, pp. 291–317.
- Levin, B. & Pinker, S. (1991). Introduction, in B. Levin & S. Pinker (eds.), *Lexical and Conceptual Semantics*, Oxford: Blackwell.
- Levine, J. (2001). *Purple Haze: The Puzzle of Consciousness*, Oxford: Oxford University Press.
- (2010). Demonstrative Thought. *Mind & Language*, 25 (2): 169–195.
- Levinson, S. (2003). *Space in Language and Cognition: Explorations in Cognitive Diversity*, Cambridge: Cambridge University Press.
- Levinson, S., Kita, S., Haun, D., & Rasch, B. (2002). Returning the Tables: Language Affects Spatial Reasoning, *Cognition*, 84: 155–188.
- Lewis, D. (1966). An Argument for the Identity Theory, *Journal of Philosophy*, 63 (1): 17–25.
- Li, P. & Gleitman, L. (2002). Turning the Tables: Language and Spatial Reasoning, *Cognition*, 83: 265–294.
- Li, P., Abarbanell, L., Gleitman, L. & Papafragou, A. (2011). Spatial Reasoning in Tenejapan Mayans, *Cognition*, 120 (1): 33–53.
- Locke, J. (1690/1975). *An Essay Concerning Human Understanding*, New York: Oxford University Press.
- Lupyan, G. (2012). What Do Words Do? Toward a Theory of Language-Augmented Thought, in B.H. Ross (ed.), *Psychology of Learning and Motivation* (Volume 57), New York: Academic Press, pp. 155–297.
- Machery, E. (2007). Concept Empiricism: A Methodological Critique, *Cognition*, 104: 19–46.
- (2009). *Doing Without Concepts*, New York: Oxford University Press.
- (2010). Precis of *Doing Without Concepts*, *Behavioral and Brain Sciences*, 33: 195–244.
- (2017). *Philosophy within Its Proper Bounds*, Oxford: Oxford University Press.

- Machery, E., Mallon, R., Nichols, S. & Stich, S. (2004). Semantics, Cross-Cultural Style, *Cognition*, 92: B1–B12.
- Machery, E., Stich, S., Rose, D., Chatterjee, A., Karasawa, K., Struchiner, N., Sirker, S., Usui, N., & Hashimoto, T. (2017). Gettier Across Cultures. *Noûs*, 51 (3): 645–664.
- Mahon, B.Z. (2015). What is Embodied about Cognition? *Language, Cognition, and Neuroscience*, 30 (4): 420–429.
- Mallon, R. & Weinberg, J. (2006). Innateness as Closed Process Invariance, *Philosophy of Science*, 73 (3): 323–344.
- Mameli, M. (2008). On Innateness: The Clutter Hypothesis and the Cluster Hypothesis. *The Journal of Philosophy*, 105 (12): 719–736.
- Manfredi, P. (1993). Tacit Beliefs and Other Doxastic Attitudes. *Philosophia*, 22 (1): 95–117.
- Marcus, G. (2001). *The Algebraic Mind: Integrating Connectionism and Cognitive Science*, Cambridge, MA: MIT Press.
- (2004). *The Birth of the Mind: How a Tiny Number of Genes Creates the Complexities of Human Thought*, New York: Basic Books.
- Margolis, E. (1998). How to Acquire a Concept, *Mind & Language*, 13: 347–369.
- Margolis, E. & Laurence, S. (1999). *Concepts: Core Readings*, Cambridge, MA: MIT Press.
- (2007). The Ontology of Concepts — Abstract Objects or Mental Representations? *Noûs*, 41 (4): 561–593.
- (2010). Concepts and Theoretical Unification, *Behavioral and Brain Sciences*, 33: 219–220.
- (2011). Learning Matters: The Role of Learning in Concept Acquisition, *Mind & Language*, 26 (5): 507–539.
- (2013). In Defense of Nativism, *Philosophical Studies*, 165 (2): 693–718.
- (2015). *The Conceptual Mind: New Directions in the Study of Concepts*, Cambridge, MA: MIT Press.
- McClelland, J., Botvinick, M., Noelle, D., Plaut, D., Rogers, T., Seidenberg, M., & Smith, L. (2010). Letting Structure Emerge: Connectionist and Dynamic Systems Approaches to Cognition, *Trends in Cognitive Sciences*, 14: 348–56.
- McDowell, J. (1994). *Mind and World*, Cambridge, MA: Harvard University Press.
- Millikan, R. (2000). *On Clear and Confused Ideas*, Cambridge: Cambridge University Press.
- Murphy, G. (2002). *The Big Book of Concepts*, Cambridge, MA: MIT Press.
- Osherson, D., & Smith, E. (1981). On the Adequacy of Prototype Theory as a Theory of Concepts, *Cognition*, 9: 35–58.

- Papineau, D. (2002). *Thinking About Consciousness*, Oxford: Oxford University Press.
- (2013). The Poverty of Conceptual Analysis, in M. Haug (ed.), *The Armchair or the Laboratory?*, New York: Routledge, pp. 166–194.
- Peacocke, C. (1992). *A Study of Concepts*, Cambridge, MA: MIT Press.
- (2005). Rationale and Maxims in the Study of Concepts, *Noûs*, 39 (1): 167–78.
- Pica, P., Lemer, C., Izard, V., & Dehaene, S. (2004). Exact and Approximate Arithmetic in an Amazonian Indigene Group, *Science*, 306: 499–503.
- Pietroski, P. & Crain, S. (2012). The Language Faculty, in E. Margolis, R. Samuels, & S. Stich (eds.), *The Oxford Handbook of Philosophy of Cognitive Science*, New York: Oxford University Press, pp. 361–381.
- Pinker, S. (1994). *The Language Instinct: The New Science of Language and Mind*. London: Penguin.
- (2007). *The Stuff of Thought: Language as a Window into Human Nature*. London: Penguin.
- Prinz, J. (2002). *Furnishing the Mind: Concepts and Their Perceptual Basis*, Cambridge, MA.: MIT Press.
- Putnam, H. (1967). The ‘Innateness Hypothesis’ and Explanatory Models in Linguistics, *Synthese*, 17: 12–22; reprinted in H. Putnam, *Philosophical Papers* (Volume 2: Mind, Language and Reality), Cambridge: Cambridge University Press, pp. 107–116.
- (1975). The Meaning of ‘Meaning’, in H. Putnam, *Philosophical Papers. Volume 2: Mind, Language and Reality*, pp. 215–271, Cambridge: Cambridge University Press.
- Quine, W.V.O. (1960). *Word and Object*, Cambridge, MA: MIT Press.
- Ramsey, W. (1998). Prototypes and Conceptual Analysis, in M. DePaul & W. Ramsey (eds.), *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Lanham, MD: Rowman & Littlefield, pp. 161–77.
- Robbins, P. (2002). How to Blunt the Sword of Compositionality, *Noûs*, 36 (2): 313–34.
- Roberts, I. (2017). *The Oxford Handbook of Universal Grammar*, Oxford University Press.
- Rosch, E. (1978). Principles of Categorization, in E. Rosch & B. Lloyd (eds.), *Cognition and Categorization*, Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 27–48.
- Rosch, E., & Mervis, C. (1975). Family Resemblances: Studies in the Internal Structure of Categories, *Cognitive Psychology*, 7: 573–605.
- Sampson, G. (2005). *The “Language Instinct” Debate*, revised edition, London: Continuum.
- Samuels, R. (2002). Nativism in Cognitive Science, *Mind & Language*, 17 (3): 233–65.

Samuels, R. & Ferreira, M. (2010). Why *Don't* Concepts Constitute a Natural Kind, *Behavioral and Brain Sciences*, 33: 222–223.

Shapiro, L. (2004). *The Mind Incarnate*, Oxford: Oxford University Press.

Shettleworth, S. (2010). *Cognition, Evolution, and Behavior*, 2nd edition, New York: Oxford University Press.

Shusterman, A. & Spelke, E. (2005). Language and the Development of Spatial Reasoning, in P. Carruthers, S. Laurence, & S. Stich (eds.), *The Innate Mind: Structure and Contents*, New York: Oxford University Press, pp. 89–106.

Smith, E., & Medin, D. (1981). *Categories and Concepts*, Cambridge, MA: Harvard University Press.

Sosa, E. (2009). A Defense of the Use of Intuitions in Philosophy, in D. Murphy & M. Bishop (eds.), *Stich and His Critics*, Oxford: Wiley-Blackwell.

Spelke, E.S. (2003). What Makes Us Smart? Core Knowledge and Natural Language, in D. Gentner & S. Goldin-Meadow (eds.), *Language in Mind*, Cambridge, MA: MIT Press, pp 277–311.

Spelke, E.S. & Kinzler, K.D. (1992). Core Knowledge, *Developmental Science*, 10 (1): 89–96.

Spelke, E.S. & Newport, E. (1998). Nativism, Empiricism, and the Development of Knowledge, in W. Damon (series ed.) and R.M. Lerner (volume ed.), *Handbook of Child Psychology (Volume 1: Theoretical Models of Human Development)*, fifth edition, New York: Wiley, pp. 275–340.

Spelke, E.S. & Tsivkin, S. (2001). Initial Knowledge and Conceptual Change: Space and Number, in M. Bowerman & S. Levinson (eds.), *Language Acquisition and Conceptual Development*, Cambridge: Cambridge University Press.

Stich, S. (1990). *The Fragmentation of Reason*, Cambridge, MA: MIT Press.

— (1992). What Is a Theory of Mental Representation? *Mind*, 101 (402): 243–61.

— (2009). Reply to Sosa. In D. Murphy & M. Bishop (eds.). *Stich and His Critics*, Oxford: Wiley-Blackwell.

— (2012). *Collected Papers* (Volume 2: Knowledge, Rationality, and Morality, 1978–2010), Oxford: Oxford University Press.

Strevens, M. (2018). *Thinking Off Your Feet: How Empirical Psychology Vindicates Armchair Psychology*, Cambridge, MA: Harvard University Press.

Strohming, N. & Moore, B. (2010). Banishing the Thought, *Behavioral and Brain Sciences*, 33: 225–226.

Tooby, J. & Cosmides, L. (2016). The Theoretical Foundations of Evolutionary Psychology. In D. Buss (ed.) *The Handbook of Evolutionary Psychology, Second Edition*, Wiley, pp. 3–87.

- Van Gelder, T. (1995). What Might Cognition Be, If Not Computation? *Journal of Philosophy*, 91: 345–381.
- Vicente, A. & Manrique, F.M. (2016). The Big Concepts Paper: A Defense of Hybridism, *British Journal for the Philosophy of Science*, 67: 59–88.
- Weinberg, J., Gonnerman, C., Buckner, C., & Alexander, J. (2010). Are Philosophers Expert Intuiters? *Philosophical Psychology*, 23 (3): 331–355.
- Weinberg, J., Nichols, S., & Stich, S. (2001). Normativity and Epistemic Intuitions, *Philosophical Topics*, 29 (1&2): 429–460.
- Weiskopf, D. (2007). Concept Empiricism and the Vehicles of Thought, *Journal of Consciousness Studies*, 14: pp. 156–183.
- (2009). The Plurality of Concepts, *Synthese*, 169: 145–173.
- (2010). The Theoretical Indispensability of Concepts, *Behavioral and Brain Sciences*, 33: 228–229.
- Whorf, B. (1956). *Language, Thought, and Reality*, Cambridge, MA: MIT Press.
- Williamson, T. (2005). Armchair Philosophy, Metaphysical Modality and Counterfactual Thinking, *Proceedings of the Aristotelian Society*, 105: 1–23.
- (2007). *The Philosophy of Philosophy*, Oxford: Blackwell.
- Wittgenstein, L. (1953/1958). *Philosophical Investigations*, 3rd edition, G.E.M. Anscombe (trans.), Oxford: Blackwell.
- Wright, C. (1989). The Verification Principle: Another Puncture—Another Patch, *Mind*, 98: 611– 622.
- Yablo, S. (2008). *Thoughts: Papers on Mind, Meaning, and Modality*, Oxford University Press.
- Zalta, E. (2001). Fregean Senses, Modes of Presentation, and Concepts, *Philosophical Perspectives*, 15: 335–359.