

Splitting Concepts*

Gualtiero Piccinini and Sam Scott†‡

A common presupposition in the concepts literature is that concepts constitute a singular natural kind. If, on the contrary, concepts split into more than one kind, this literature needs to be recast in terms of other kinds of mental representation. We offer two new arguments that concepts, in fact, divide into different kinds: (a) concepts split because different kinds of mental representation, processed independently, must be posited to explain different sets of relevant phenomena; (b) concepts split because different kinds of mental representation, processed independently, must be posited to explain responses to different kinds of category. Whether these arguments are sound remains an open empirical question, to be resolved by future empirical and theoretical work.

1. Introduction. In the past 35 years, psychologists, philosophers, and linguists have generated a vast interdisciplinary literature on the nature of concepts. This literature has produced three main families of psychological theories: the prototype, exemplar, and theory theories of concepts (e.g., Hampton [1993] for prototypes, Nosofsky [1988] for exemplars, and Gopnik and Meltzoff [1997] for theories). Although advocates of these three accounts disagree about the nature of concepts, the debates between them usually presuppose that concepts are a singular natural kind of mental representation, as opposed to a set of different natural kinds. In other words, theorists disagree on which kind of internal structure constitutes concepts, but they agree that there is only one such kind of internal

*Received April 2005; revised May 2006.

†To contact the authors, please write to: Gualtiero Piccinini, Department of Philosophy, Washington University in St. Louis, One Brookings Drive, St. Louis, MO 63130-4899; e-mail: gpiccini@artsci.wustl.edu. Sam Scott, 11-1317 King Street West, Toronto, ON, M6K 1H2, Canada; e-mail: SamScott@Canada.com.

‡The names of the authors are in alphabetical order. A version of this article was presented at the meeting of the Society for Philosophy and Psychology, in Winston-Salem, North Carolina, June 2005. Thanks to our commentator, Dan Ryder, and our audience, especially Ken Aizawa and Carl Gillett. Thanks also to Carl Craver, John Heil, Edouard Machery, Dan Weiskopf, and the referees for very helpful comments.

Philosophy of Science, 73 (October 2006) 390–409. 0031-8248/2006/7304-0002\$10.00
Copyright 2006 by the Philosophy of Science Association. All rights reserved.

structure with different instances for each category.¹ In this article, we discuss three arguments that this presupposition is false. If any of these arguments turns out to be correct, then the *concept* kind splits into several natural kinds, and as a consequence many debates over the nature of concepts are beside the point and perhaps even vacuous. At the very least, if an argument for concept splitting could be shown to be correct, the literature on concepts would need to be recast in terms of other kinds of mental representation.

The first argument for splitting concepts comes from a stimulating paper by Edouard Machery (2005). Machery argues that the *concept* kind splits on the grounds that several different kinds of mental representation are involved in many of the same higher cognitive tasks traditionally thought to involve concepts. But as we will show, this argument falls short, because the evidence Machery presents is that several different kinds of mental representation are *jointly* involved in the same cognitive tasks. Such evidence is consistent with the view that the mental representations postulated by Machery are *components* of the same natural kind. If this interpretation of Machery's evidence is correct, then his argument fails to establish that the concept kind needs to be split.

Nevertheless, there may still be good reasons to endorse concept splitting. After discussing Machery's argument, we will outline two alternative arguments of our own: (a) concepts split because different kinds of mental representation, processed independently, must be posited to explain different sets of relevant phenomena; (b) concepts split because different kinds of mental representation, processed independently, must be posited to explain responses to different kinds of category. But first, we need to say more about what kinds are, what concepts are, and what might justify splitting them.

2. Kind Splitting. Following Machery, we adopt Richard Boyd's notion of natural kinds as *homeostatic property clusters* (HPC), which we believe captures what is relevant to special sciences like psychology. The HPC kinds are defined by (i) the possession of a cluster of often co-occurring properties that play a role in induction and explanation and (ii) the existence of causal mechanisms (developmental, evolutionary, etc.) that bring about the co-occurrence of the property cluster (Boyd 1989, 1991). Candidate HPC natural kinds can be given *programmatic* definitions, which functionally specify the anticipated causal or explanatory role to be played by members of the kind in question (Boyd 1999, 70). With this

1. A category is the group of things in the world that a concept represents.

notion of natural kinds in place, we can consider circumstances under which scientists have split natural kinds in the past.

As a first example, we borrow from Carl Craver's (2004) discussion of memory. Human beings have a considerable ability to retain information. To explain this phenomenon, psychologists postulated the existence of something called *memory*, programmatically defined to be whatever internal mechanism explains the phenomenon of information retention. It was then discovered that there are distinct aspects of information retention—for example, the retention of factual knowledge versus the retention of skills—that can be independently disrupted. It was further discovered that different neural processes and mechanisms are independently involved in the different aspects of information retention. This kind of double dissociation between two sets of phenomena and associated mechanisms is among the evidence that led psychologists and neuroscientists to split the putative natural kind, *memory*, into several distinct kinds, such as procedural and declarative memory.

Sometimes, kind splitting leads to the elimination of a putative natural kind and the abandonment, within science, of the corresponding natural kind term. For example, consider jade. The explanandum phenomenon is the characteristic hardness, color, and other surface properties of certain stones. The putative explanans is a microphysical composition and structure, with an associated geological formation process, that gives rise to the phenomenon. As it turns out, there are two different microphysical structures—each with its own cluster of microphysical properties and formation process—that give rise to different instances of the explanandum. They are jadeite ($\text{NaAlSi}_2\text{O}_6$) and nephrite, which is a form of the mineral actinolite ($\text{Ca}_2(\text{Mg, Fe})_5\text{Si}_8\text{O}_{22}(\text{OH})_2$). Although jadeite and nephrite share some surface properties, their microphysical properties and formation mechanisms are sufficiently different to warrant different places in the mineralogical taxonomy. Jade, by contrast, no longer holds a place in mineralogy—it is not a natural kind.

Finally, kind splitting may lead to the retention of an old natural kind term, but with a significant change of reference. Consider the folk-biological kind “fish,” which includes whales, dolphins, and other mammals.² The explanandum is a set of observable properties of water-dwelling organisms with a characteristic streamlined shape. The putative explanans is a set of internal mechanisms and developmental and evolutionary processes that explain the phenomena. On closer examination, such organisms divide into two groups, which possess at least two importantly different

2. About one-quarter of college students persist in counting whales as fish despite the widely accepted and understood scientific taxonomy according to which they are mammals (McCloskey and Glucksberg 1978; Scott 2003a).

clusters of properties and mechanisms (e.g., lungs vs. gills, hair vs. scales, live birth vs. egg laying, etc.). Furthermore, the co-occurrence of each cluster of properties is explained by separate evolutionary and developmental mechanisms. These differences have enough scientific relevance that the old putative natural kind term “fish” has been eliminated in favor of two new ones. At least among the scientifically educated, the term “fish” is now restricted to cover only animals with one cluster of properties. Animals with the other cluster of properties are classified as cetaceans, which are an order of mammals.

These examples suggest that kinds are not split merely because they are constituted by inhomogeneous components. Kind splitting occurs when different, largely independent aspects of what, pretheoretically, appeared to be one phenomenon (retention of facts vs. skills, hardness and color of one vs. another set of stones, survival and reproduction of one vs. another set of animals) are explained by at least two clusters of internal properties (i) that differ for the purposes of induction or explanation and/or (ii) whose co-occurrence is explained by different causal mechanisms.³

In light of the above, to properly support concept splitting based on a notion of natural kinds as homeostatic property clusters, it must be shown that the phenomena antecedently associated with concepts divide into different, largely independent groups; that these groups of phenomena are explained by different kinds of internal structures possessing different clusters of scientifically relevant properties; and possibly that the co-occurrence of those clusters is explained by different causal mechanisms. Evidence that would support concept splitting includes evidence that the different kinds of internal structure corresponding to the notion of *concept* are stored in separate parts of the brain, are acquired through different mechanisms, can be independently disrupted, have different evolutionary histories, and so on.

Showing some of this would be akin to the strategy followed by Paul Griffiths regarding emotions (1997, 2004). Griffiths argues that emotions are not a natural kind because there is no one mechanism that will explain all of the phenomena that we pretheoretically explain by invoking the term “emotion,” and this is true even for specific emotions—for example, different episodes of what we call “anger” require different mechanisms, operating largely independently of one another, for their explanation. Of course, some of these mechanisms may be jointly involved in some (though not all) relevant phenomena, such as certain social emotions. Even so, they remain separate mechanisms operating together rather than as parts of the same mechanism. This is because, if Griffiths is right, they are

3. For a similar view of when natural kinds are and are not split, see Craver (2004), from whom we took the term “kind splitting.”

realized by different neural systems, are acquired through different developmental mechanisms, and have different evolutionary histories.

If concepts split into different kinds of mental representation, what should be the fate of the natural kind term “concept”? Should it be retained to name a superordinate kind, as in the case of “memory”? Should it be expelled from science, as it happened to “jade”? Or should its reference change, like the term “fish”? The answer to this question depends in part on empirical matters. If concepts split into kinds that have enough scientifically relevant properties in common, all those kinds will still be classified together under a superordinate kind. Otherwise, scientists may conclude that concepts are not a natural kind, and they may or may not retain the term “concept” with a changed reference. At this stage, there is no consensus on which natural kinds must be postulated to explain the relevant phenomena, so it is premature to judge what ought to happen to the term. Our goal here is mainly to map out the kinds of arguments and evidence that may lead to the eventual splitting of concepts.

3. Concepts. A good argument for concept splitting requires not only an adequate argumentative strategy but also an adequate programmatic definition of “concept.” Machery defines concepts as “bodies of knowledge used by default in our higher cognitive processes” (2005, 14). For present purposes, this definition is too inclusive. The cognitive science literature is replete with evidence for many different kinds of mental representation (“body of knowledge”) used by default in higher cognitive processes, including phonemes, production rules, spatial maps, and so on.⁴ These all fit Machery’s definition of concept, and they are at least as dissimilar from one another as the various representations postulated by theories of concepts.⁵ Furthermore, they are each postulated to explain *different* cognitive phenomena. Therefore, they do not constitute a natural kind, not even one constituted by different component structures. These ob-

4. Machery (2005) appears to use the term “body of knowledge” interchangeably with “mental representation.” We do not endorse all his terminological choices, but we follow them in order to engage with his argument.

5. If you doubt that our examples belong in “higher cognitive processes” (whatever that means), please replace them with examples you prefer. It may still be objected that although phonemes, production rules, spatial maps, etc., are used by default in *some* higher cognitive processes, they are not used by default in all or even *most* higher cognitive processes. But first, notice that Machery’s definition does not specify in how many higher cognitive processes in which a structure must be involved to qualify as a concept; second, it is not clear why we should require that for a structure to qualify as a concept, it must be involved in all or most higher cognitive processes. It may well be that no structure has this property, but this does not seem an especially compelling reason to deny that there are concepts.

servations, coupled with Machery's definition of "concept," can easily be turned into an argument for splitting concepts.

- P1.** Different cognitive theories posit different kinds of mental representation used by default in our higher cognitive processes (i.e., phonemes, production rules, spatial maps, etc.) to explain different phenomena.
- P2.** These kinds of mental representation possess different clusters of scientifically relevant properties.
- C1.** Therefore, the mental representations used by default in our higher cognitive processes split into different natural kinds.
- C2.** Therefore, concepts split into different natural kinds.

But nobody takes phonemes, production rules, and images to be concepts, precisely *because* they are deployed in very different types of cognitive processes. The argument succeeds, but only based on a faulty definition of "concept." If we are going to determine whether concepts split into different natural kinds in an interesting sense, we need a more restrictive notion. We must find a programmatic definition that fits, more or less, the current explanatory role played by concepts in the psychological literature.

Concepts are traditionally defined as the simplest constituents of *thought* (Locke's "materials of reason and knowledge"). Identifying concepts in terms of their role within thoughts is still common in much of the philosophical literature. But the notion of thought is about as inclusive as that of higher cognitive process, so this way of identifying concepts suffers from the same problem. As shown by the following argument, thoughts are not a singular natural kind.

Definition: *thought* is the folk name for higher cognitive processes.

- P1.** Scientific psychological theories posit different kinds of higher cognitive process (e.g., image manipulation, rule space search, reasoning by analogy, etc.) to explain different phenomena.
- P2.** These kinds of higher cognitive process possess different clusters of scientifically relevant properties.
- C1.** Therefore, higher cognitive processes split into different natural kinds.
- C2.** Therefore, thoughts split into different natural kinds.

Since thoughts are not a singular natural kind, we cannot define concepts by employing the notion of thought without begging the question under discussion. If we are going to identify the subject matter of concepts in a useful way, we need to work harder.

To identify concepts in the relevant sense, we need to look at the role they play within contemporary psychological theories. Within psychology,

concepts are mental particulars posited as a specific part of the explanation for certain cognitive phenomena. For present purposes, a concept is anything that plays the relevant role within the explanation of those phenomena. Thus understood, concepts are a singular natural kind if and only if all such entities turn out to share a large set of scientifically relevant properties as a result of the same causal mechanism.

The role that concepts play within psychological explanation should be specified in as theory neutral a way as possible. The following is relatively uncontroversial: for each kind of psychological process, there is a corresponding kind of psychological structure—perception and imagistic reasoning operate on mental images, rule-based processes operate on declarative representations, and so forth. Concepts, then, are the constituents of structures posited in the explanation of certain psychological phenomena. It remains to determine what those phenomena are. In the psychological literature, concepts have been invoked to explain at least the following six sets of phenomena.

1. *Discrimination*: The ability to respond differentially to objects, properties, and events that have something salient in common. For example, identifying and appropriately responding to conspecifics based on shape, color, smell, and so on.
2. *Nonlinguistic inference*: The general ability to draw inferences about classes of objects in the world. For example, inferring from one or more nasty experiences with red mushrooms that red mushrooms should be avoided.
3. *Categorization*: The arbitrary association of stimuli and responses, including the ability to produce an appropriate label (i.e., a word) in response to a class of stimuli as well as to choose an appropriate object or behavior in response to a label given as stimulus.⁶
4. *Word and sentence understanding*: The processing of words and sentences in accordance with their use by a linguistic community, so as to retrieve information that is often called the semantic or cognitive content of a word or sentence.
5. *Linguistic inference*: The performance of inferences between words (such as the inference from “red” to “colored”) and between sentences (such as *modus ponens*).
6. *Lexical combination*: The ability to respond to combinations of lexical items, such as nouns and adjectives, in a way that is appropriate to the combined cognitive contents of the constituents. For instance,

6. What is studied under the label *categorization* in the psychological literature often requires linguistic understanding and inference. We are here using *categorization* in a more restrictive sense.

the ability to describe the features of a pet fish given the understanding of “pet” and “fish.” Lexical combination may be seen as part of the ability to understand language and draw linguistic inferences, but it is sufficiently salient in the concepts literature—where it is assumed to be explained by *concept* combination—that it deserves to be listed separately.

In summary, we take concepts to be the main constituents of those psychological structures whose processing explains the six phenomena listed above—a list that, on our reading of the psychological literature, is fairly uncontroversial.⁷ We will refer to the explanation of phenomena 1–6 as our *desiderata* 1 through 6 on concepts.⁸

Some authors (e.g., Fodor 1998; Prinz 2002) include two further desiderata for a theory of concepts: that it accounts for reference and that it accounts for publicity. To a first approximation, reference is the relation between words and things in the world such that words apply correctly to all and only members of some category of things; publicity is the relation between speakers of a language such that they share the semantic contents of words. Important as these desiderata might be, they have two shortcomings. First, they are difficult to specify in terms of psychological phenomena. Second, the extent to which they are based on genuine psychological phenomena is controversial, especially among psychologists and linguists (e.g., Lakoff 1987; Chomsky 2000; Murphy 2002). Because of this, it is unclear whether and how they should constrain psychological theories. Like Machery, we are concerned with purely psychological theories of concepts. Hence, we leave reference and publicity aside.

Our programmatic definition of concepts allows us to specify more precisely several ways that concepts might split. We should recall that in the psychological literature, the entities or properties that concepts represent are usually called “categories.” First, concepts would split if different aspects of 1–6 were explained by different kinds of mental representation within each category. Second, concepts would split if different subsets of 1–6 were explained by different kinds of mental representation for all categories. Finally, concepts would split if all of 1–6 were explained by different kinds of mental representation for different kinds of categories. Machery alludes to all these argumentative lines, but he pursues and defends only the first. In the next section, we will discuss the prospects

7. If concepts are the *main* constituents, what are the other constituents? We are leaving open the possibility that the relevant cognitive processes operate on concepts as well as nonconceptual components. For instance, there may be nonconceptual components whose function is to link two concepts together within a larger cognitive structure.

8. Authors whose desiderata for a theory of concepts are approximately in line with ours include Smith and Medin (1981), Komatsu (1992), and Murphy (2002).

of his version of the first line. In the subsequent and last sections, we will offer our versions of the second and third lines.

4. The Argument from Explanatory Necessity. Machery's *Argument from Explanatory Necessity* may be paraphrased as follows:

- P1.** If concepts are a singular natural kind, then there is one natural kind of mental representation that satisfies desiderata 1–6.
- P2.** Different aspects of desiderata 1–6 require the postulation of different kinds of mental representation (i.e., at least prototypes, exemplars, and folk theories).
- P3.** These kinds of mental representation possess different clusters of scientifically relevant properties.
- C1.** There is no single kind of mental representation that satisfies all aspects of desiderata 1–6.
- C2.** Concepts split into different natural kinds.

Machery's actual conclusion is stronger than C2. He argues that "concepts are not a natural kind" (2005, 14). As we indicated in Section 2, however, we think that this conclusion is premature. Even if concepts split into several natural kinds, it is too early to divine whether those natural kinds will deserve to be classified together as concepts (analogously to memory systems). For this reason, we have weakened Machery's conclusion. Even so, we think his argument needs work.

P1 expresses the basic assumption of the concepts literature that is the target of this article. P3 could be challenged, but we will waive any concerns because supporters of the three main theories of concepts would likely assent to it. If we could accept P2, this would get us as far as C1, which follows from P2 and P3 by the kind of considerations on kind splitting that we offered in Section 2. From there, C2 follows by our programmatic definition of concepts.

As we formulated it, P2 is strong enough to make the argument go through. But the evidence Machery provides only supports something weaker than P2, according to which prototypes, exemplars, and folk theories are all needed to explain the *same* phenomenon. Machery appeals to the literature on concept combination. Concept combination occurs when two or more simple concepts are combined to yield a new, complex concept. In Section 2, we individuated concept combination as a putative explanans of the phenomenon of *lexical* combination, so as to avoid commitment to the existence of concepts as a singular natural kind. For if concepts are not a singular natural kind, there may be no natural phenomenon of concept combination that is amenable to scientific induction and explanation. Compare the following: since jade is not a natural kind, there are no natural phenomena individuated in terms of jade.

An adequate account of concept combination is generally considered to be a litmus test for theories of concepts, as the combination of concepts is believed to explain the productivity and compositionality of declarative thought and language (e.g., Fodor 1998; Laurence and Margolis 1999; Murphy 2002; Prinz 2002). If a theory of concepts is incapable of handling concept combination, then that theory fails.⁹

Drawing largely on the empirical evidence cited by Murphy (2002, Chapter 12) and leaning heavily on the evidence from conjunctive combinations gathered by Hampton (1987), Machery argues that prototypes, exemplars, and theories all jointly play a role during concept combination. Postulating prototypes is required to explain the empirical facts about property inheritance in concept combinations (e.g., the most typical properties of the constituent concepts tend to be those that are most typical in a combined concept). Postulating mental theories is required to explain why in some concept combinations, certain features are never inherited (the theory states that they are impossible for the combined concept) and certain features are always inherited (the theory states that they are necessary for the combined concept). And postulating a role for stored exemplars is required to explain why certain features emerge in certain combinations. This interpretation of the evidence is likely to encounter resistance from authors committed to one or another theory of concepts. Murphy (2002) seems to endorse it, but other theorists are likely to reinterpret Machery's evidence for P2 in line with their preferred theories. We accept Machery's interpretation of the evidence for the sake of the argument.

So, supposing that all of this is true, we may conclude that prototypes, exemplars, and folk theories are all needed to account for desideratum 6. That is, one phenomenon appears to involve at least three different kinds of mental representation that are sufficiently different from one another as to constitute different natural kinds. Still, this falls short of supporting concept splitting. Rather, it suggests what a number of psychologists are inclined to conclude, namely, that all these different kinds of mental representation are parts of a larger whole, which is what deserves the label *concept*. Several hybrid proposals of this kind are already present in the psychological literature, as Machery himself points out.¹⁰ So rather than unequivocally supporting concept splitting, Machery's argument is also consistent with a hybrid theory of concepts.

Machery responds that the hybrid view is a "terminological variant"

9. There is some disagreement over what counts as handling concept combination in the right way, but it makes no difference to the present discussion.

10. Machery cites Gelman (2004); see also Michalski (1993) and Anderson and Betz (2001).

of his view (2005, 30). This reply is too quick. There is a big difference between a hybrid view of concepts and the view that concepts split into different natural kinds. The former says not only that there are concepts but also that they share a large set of scientifically relevant properties whose co-occurrence is explained by the same causal mechanisms. The latter says either that there are no concepts or that if there are any, either they do not have enough relevant properties in common to support psychological generalizations and explanations or their co-occurrence is not explained by the same causal mechanisms.

If the evidence against the assumption that concepts are a natural kind is just that several kinds of mental representation are employed together within the same processes, the most prudent theoretical move is to lump them into one larger whole and label that whole “concept.” This is not to say that Machery’s argument from explanatory necessity is unsound; it may or may not be. But anyone who wishes to bolster the argument should at least provide evidence that the different kinds of mental representation are needed separately to explain different phenomena. We will not attempt to do this. Instead, we shall outline two alternative arguments that concepts split into different natural kinds.

5. Two Ways in Which Concepts Might Split. Our first argument is that there are at least two subsets of desiderata 1–6 for which we are forced to posit different kinds of mental representation. The second is that there are at least two classes of categories for which we are forced to posit different kinds of mental representation. Both arguments propose that concepts decompose into at least two different kinds of mental representation, whose clusters of properties are sustained by different causal mechanisms and are implicated in different phenomena—they either satisfy different desiderata or satisfy the same desiderata but for different classes of categories. As a consequence, both arguments support genuine concept splitting (as opposed to a hybrid theory of concepts). Although we will briefly indicate why we find the premises of these arguments plausible, whether the arguments are sound remains an open empirical question.

5.1. The Argument from Language.

- P1.** If concepts are a singular natural kind, then there is one kind of mental representation (possibly a hybrid structure) that satisfies desiderata 1–6.
- P2.** Desiderata 1–3 require the postulation of one kind of mental representation.
- P3.** Desiderata 4–6 require the postulation of a different kind of mental representation.

- P4.** These two kinds of mental representation possess different clusters of scientifically relevant properties.
- C1.** There is no single kind of mental representation (hybrid or not) that satisfies desiderata 1–6.
- C2.** Concepts split into different natural kinds.

As we already mentioned, P1 is a basic assumption of the concepts literature. Should we believe P2 through P4?

The phenomena putatively explained by concepts can be divided into two groups: those that require the ability to understand and process a natural language and, thus, are only found in linguistically competent human beings, and those that do not require linguistic ability and are in fact found in most humans and nonhuman animals. The latter include phenomena 1–3: discrimination, nonlinguistic inference, and categorization (in our restricted sense). The former include phenomena 4–6: language understanding, linguistic inference, and lexical combination.

There are good reasons, rehearsed in many linguistics and psychology textbooks, to think that the processes and representations that underlie human language are different in kind from those that underlie other cognitive abilities. Human language is widely thought to be the result of a species-specific, genetically and developmentally canalized, psychological faculty (e.g., Pinker 1994). In addition, it is widely believed that the mechanisms responsible for language processing are localized in specific parts of the brain and can be disrupted independently of other cognitive mechanisms (e.g., Gazzaniga et al. 2002). Hence, cognitive abilities that are linked to the acquisition and processing of a natural language may well differ in kind from, and operate over different kinds of representation than, those that do not.

The claim that language requires proprietary cognitive processes is typically made with respect to syntactic competence, but what is at issue here may seem to be only (linguistic) semantic competence—concepts are typically invoked within semantics, not syntax. Hence, it may appear that by offering evidence that pertains to syntax, we are confusing syntax and semantics. But the fact that concepts are typically invoked within semantics and not syntax is a terminological accident with no bearing on the present issue. The issue is whether the representations that are needed to explain phenomena 1–6 split into different kinds. There is little doubt that phenomena 3–6 require linguistic competence, that linguistic competence requires syntactic competence, and that syntactic competence requires special kinds of representations, which are likely to differ in kind from those of other cognitive processes. We are merely pointing out that these statements, which are hardly controversial, have consequences for the theory of concepts—namely, the theory of the representations that

explain phenomena 1–6. The fact that psychologists and linguists do not *call* syntactic representations *conceptual* may help explain why they have not noticed these consequences, but it does nothing to diminish them.

Furthermore, our point is likely to apply to what are usually called (linguistic) semantic representations. Unfortunately, there is no consensus on the relationship between syntactic and semantic competence, and a detailed discussion of this issue would take us too far afield. But briefly, the mechanisms for syntactic and semantic competence might be related in one of three ways. First, there may be only one type of representation for both syntax and semantics. This would support our hypothesis on the grounds that syntactic competence is explained by processes and representations of a special kind. Second, there may be two kinds of representation, but the mechanism for semantic competence must rely on syntactic representations to build or operate on semantic representations. This would also support our hypothesis. Finally, the mechanism for semantic competence might be able to construct and use its representations independently of syntactic representations. This is still consistent with our hypothesis, provided that the mechanism for semantic competence operates on proprietary representations that differ in kind from other mental representations. Considering all the options, (linguistic) semantic representations are likely to be of a different kind than nonlinguistic representations.

In summary, P2 through P4 are plausible with respect to semantic representations and should be uncontroversial with respect to syntactic representations. Whether the latter are usually called *conceptual* in the literature makes no difference to our argument. Hence, from now on, we will ignore the distinction between semantic and syntactic representations.

Call the representations that are needed for acquiring and processing language *linguistic representations*. The idea is not that acquiring and processing language is independent of nonlinguistic representations, only that in addition to nonlinguistic representations, linguistic cognition requires a proprietary kind of representation.¹¹ If these linguistic representations satisfy desiderata 4–6 and are qualitatively different from those that satisfy desiderata 1–3, then the term “concept,” as currently used in the literature, encompasses at least two different kinds of representation: the kind that humans share with other animals and another kind—linguistic representation—that is unique to humans as language-using creatures. This is one way to argue that concepts split into different natural kinds.

11. More specifically, the present hypothesis is consistent with the possibility that linguistic and nonlinguistic cognition share important properties. For instance, as Dan Weiskopf reminded us, prototype effects occur in both nonlinguistic tasks, such as the categorization of meaningless dot patterns and pictures, and linguistic tasks, such as lexical combination and sentence verification tasks.

Note that the argument from language does not depend on any controversial thesis about the role played by language in cognition. Views range from the purely communicative conception of language, according to which language is not required for most cognitive processes, to the many variants of the cognitive conception of language, according to which language is the medium for some or all of human cognition.¹² The consequences of our argument vary depending on what the relationship between language and cognition turns out to be. If language is only a tool for communication, then linguistic representations may be needed only for explaining language processing itself, as opposed to other cognitive processes.¹³ Most cognitive scientists, however, endorse some form of the cognitive conception of language, according to which language is involved in other cognitive processes, such as reasoning or decision making. If so, linguistic representations play a larger explanatory role. Fortunately, we can remain neutral on the exact explanatory role of linguistic representations. All we need is that explaining phenomena 4–6, which involve language processing, requires the postulation of a special sort of mental representation.

The strategy followed in this argument is a general one. If there are more subtle ways of breaking down the phenomena and the processes responsible for them, and if the processes are different enough that they operate on different kinds of representation, then there are more ways of breaking down concepts into different natural kinds. For instance, phenomena 1–3 are exhibited by different species to different degrees. In more traditional terms, some nonlinguistic species are conceptually more sophisticated than others. If some of these differences are a matter of kind and not of degree, then there might be other ways to split concepts into different natural kinds of representation. Again, this is an empirical question—we will not pursue it any further here.

It might be tempting to respond to the above argument by restricting the term “concept” to the constituents of structures that underlie linguistic

12. For a useful review of positions, see Carruthers (2002), esp. Sections 2 and 3.

13. An anonymous referee has objected that under a purely communicative conception of language, linguistic representations need not be different in structure and use from those involved in other cognitive activities; hence, in this case, there would be no need to split concepts. But the hypothesis under discussion is precisely that in order to explain the acquisition and processing of language, with its peculiar syntactic and semantic structure, we need to posit representations that are different in kind from those involved in the cognitive processes of which nonlinguistic creatures are capable. This point is not affected by the degree to which language shapes cognition in linguistic creatures.

cognition, analogously to what biologists did with “fish.”¹⁴ Even so, there are reasons for skepticism that concepts are a singular natural kind.

5.2. *The Argument from Typicality Effects.*

- P1.** If concepts are a singular natural kind, then there is one kind of mental representation (possibly a hybrid structure) that satisfies desiderata 1–6 for all categories.
- P2.** In order to explain some or all of desiderata 1–6, we must postulate the existence of different kinds of mental representation for different categories.
- P3.** These kinds of mental representation possess different clusters of scientifically relevant properties.
- C1.** There is no single kind of mental representation that satisfies desiderata 1–6 for all categories.
- C2.** Concepts split into different natural kinds.

Again, P1 has already been argued for. Is there any reason to believe P2 and P3?

The most robustly established set of psychological results regarding referring concepts is the set of typicality effects first documented by Rosch and her colleagues (Rosch 1973, 1975; Rosch and Mervis 1975). Rosch found that people not only categorize the world but also willingly make judgments as to how typical an exemplar or instance is of a given category. For example, in the category *fruit*, apples are considered by most North Americans to be much better (more typical) exemplars than figs.¹⁵ Furthermore, typicality judgments are predictive of reaction time in speeded categorization and correlate with independent measures of within-category similarity. These basic results, which are relevant to our desideratum 3, have been reproduced many times (McCloskey and Glucksberg 1978; Hampton 1979, 1981; Hampton and Gardiner 1983; Malt and Smith 1984; Barsalou 1985; Chang 1986; Malt 1994). Regarding desideratum 4, typicality also affects word-order choice in sentence production (Kelly, Bock, and Keil 1986). And regarding desideratum 5, typicality has also been found to have an effect on both inductive (Kahneman and Tversky 1973; Rips 1995) and deductive (Cherniak 1984) judgments.

Typicality effects are robust, reproducible, and indicative of the structure of concepts, although exactly what typicality effects indicate *about*

14. We believe that some philosophers will be sympathetic to this suggestion, but most psychologists would probably reject it. We are going to ignore it in the rest of the article.

15. The extent to which participants agree in their ratings has been challenged by Barsalou (1987). For a detailed response, see Scott (2003a, 2003b).

conceptual structure is up for debate. The basic set of results is explainable by any theory that makes room for a similarity-based structural component of some kind (e.g., exemplars or prototypes). But whereas typicality effects have been convincingly demonstrated for naturally occurring, lexicalized, referring object concepts (*fruit, weapon, fish, furniture, etc.*), they do not automatically generalize to other kinds of concept.

Hampton (1981) was unable to demonstrate a correlation between rated typicality and within-category similarity for three of the eight abstract concepts he examined. Barsalou (1983, 1985) showed that ad hoc concepts (those that are not stored in long-term memory but are constructed as needed in short-term memory, such as *things to do for weekend entertainment*) do not structure around within-category similarity. Finally, Armstrong, Gleitman, and Gleitman (1983) famously pointed out that there is intuitively a class of well-defined logical and mathematical concepts (e.g., *odd number*) structured as definitions rather than prototypes or exemplars.¹⁶ For all three of these kinds of concept, participants will generate graded typicality judgments, but these judgments have not been shown to be linked to an underlying evaluation of similarity. This provides a reason to believe that there might be at least two kinds of concept: (a) similarity-based concepts, including concrete object concepts and some abstract concepts; and (b) nonsimilarity-based concepts, including ad hoc and goal directed concepts, some abstract concepts, and perhaps logical and mathematical concepts.

Like the strategy behind our first argument, the strategy followed in the argument from typicality effects is a general one. Instead of two kinds of concept, each corresponding to two classes of categories, there might be many kinds of concept, each corresponding to different classes of categories. For instance, among concrete object concepts, natural kind concepts and artifact concepts might turn out to possess sufficiently different clusters of scientifically relevant properties that they warrant splitting concepts one more time. As before, we will leave this to be decided by future empirical work.¹⁷

Our two arguments are fully compatible with each other. Either one of them, or both, might be sound. It is suggestive to note here that the kinds of concept for which it is difficult to demonstrate or intuitively odd to suppose a similarity-based structure are just the kinds of concept that we

16. It should be noted, however, that Armstrong et al. (1983) did not prove this result. They collected typicality ratings for their intuitively well-defined concepts but did not attempt experimentally to rule out similarity as a structural determinant for these concepts.

17. See Medin, Lynch, and Solomon (2000) for a review of older psychological work on kinds of concepts.

might expect nonlinguistic creatures to lack: abstract concepts, which require the ability to reason about the nonconcrete; ad hoc concepts, which are highly compositional in nature; and Armstrong et al.'s (1983) logical and mathematical concepts, which are almost certainly unavailable to nonlinguistic creatures.

6. Conclusions. We have argued that concepts, as understood by psychologists, are mental representations posited to explain six phenomena: discrimination, nonlinguistic inference, categorization, linguistic understanding and inference, and lexical combination. Most philosophers and psychologists have assumed that concepts are a singular natural kind. But there is evidence that different kinds of mental representation (which do not share a large set of properties relevant to induction and explanation) are needed to explain different phenomena. If so, then concepts split into several natural kinds.

We have argued that there are at least two fault lines along which this is likely to happen. One is the line between different processes thought to involve concepts, such as linguistic and nonlinguistic cognition: mental representations available only to linguistic creatures may differ in kind from those shared by both linguistic and nonlinguistic creatures. The other is the line between mental representations involved in the same processes, but for different categories: for instance, only concrete object concepts and some abstract concepts appear to possess a similarity-based structure. We offered some reasons to believe that each of these possibilities obtains, but we consider these to be open empirical questions.

A third line along which concepts may split is that between different representations of the same category involved in higher cognitive processes. Machery (2005) argues that concepts should be split into different kinds of representation because different kinds of representation are jointly involved in the same higher cognitive tasks (such as lexical combination). We argued that this is insufficient motivation for splitting concepts, because all the relevant representations, taken together, may be considered components of a hybrid natural kind. But if it could be shown that different representations of the same category are independently involved in explaining different higher cognitive phenomena, then concepts would have to be split along this line too.¹⁸

A lesson for the concepts literature is that apparently contradictory theories and results may be reconciled. Historically, theories of concepts have often been criticized for their supposed inability to account for some previously unconsidered cognitive process (e.g., Fodor [1998] on concept

18. Weiskopf's argument for "concept pluralism" appears to be headed in this direction (Weiskopf 2005).

combination) or type of category (e.g., Armstrong et al. [1983] for intuitively well-defined categories; Barsalou [1985] for ad hoc categories). But if concepts split into different natural kinds, then such criticisms may be inappropriate. The solution may lie in postulating different kinds of mental representation for different cognitive phenomena or different types of category. More generally, theories of concepts will need to be replaced by theories of whichever kinds of mental representation are needed to explain the relevant phenomena.

Beyond the concepts literature, there is an important lesson for theorists and experimentalists who have previously relied on concepts. They should not take it for granted that there is a single, unified natural kind corresponding to the notion of concept. If they cannot defend the assumption that concepts are a singular natural kind, they should make clear how many kinds of mental representation they need in order to explain the phenomena that interest them, and which kinds they are.

REFERENCES

- Anderson, John R., and Jonathan Betz (2001), "A Hybrid Model of Categorization," *Psychonomic Bulletin and Review* 8 (4): 629–647.
- Armstrong, Sharon Lee, Lila R. Gleitman, and Henry Gleitman (1983), "What Some Concepts Might Not Be," *Cognition* 13: 263–308.
- Barsalou, Lawrence W. (1983), "Ad Hoc Categories," *Memory and Cognition* 11 (3): 211–227.
- (1985), "Ideals, Central Tendency, and Frequency of Instantiation as Determinants of Graded Structure in Categories," *Journal of Experimental Psychology: Learning, Memory, and Cognition* 11 (4): 629–654.
- (1987), "The Instability of Graded Structure: Implications for the Nature of Concepts," in Ulric Neisser (ed.), *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*. Cambridge: Cambridge University Press, 101–140.
- Boyd, Richard (1989), "What Realism Implies and What It Does Not," *Dialectica* 43: 5–29.
- (1991), "Realism, Anti-foundationalism, and the Enthusiasm for Natural Kinds," *Philosophical Studies* 61 (1–2): 127–148.
- (1999), "Kinds, Complexity and Multiple Realization," *Philosophical Studies* 95: 67–98.
- Carruthers, Peter (2002), "The Cognitive Functions of Language," *Behavioral and Brain Sciences* 25: 657–726.
- Chang, Tien Ming (1986), "Semantic Memory: Facts and Models," *Psychological Bulletin* 99 (2): 199–220.
- Cherniak, Christopher (1984), "Prototypicality and Deductive Reasoning," *Journal of Verbal Learning and Verbal Behavior* 23: 625–642.
- Chomsky, Noam (2000), *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Craver, Carl (2004), "Dissociable Realization and Kind Splitting," *Philosophy of Science* 71 (4): 960–971.
- Fodor, Jerry A. (1998), *Concepts: Where Cognitive Science Went Wrong*. Oxford: Clarendon.
- Gazzaniga, Michael S., Richard B. Ivry, George R. Mangun, and Tamara Y. Swaab (2002), "Language and the Brain," in Michael S. Gazzaniga, Richard B. Ivry, and George R.

- Mangun (eds.), *Cognitive Neuroscience: The Biology of the Mind*. 2nd ed. New York: Norton, 351–399.
- Gelman, R. (2004), “Cognitive Development,” in Hal Pashler and Douglas L. Medin (eds.), *Stevens' Handbook of Experimental Psychology*, vol. 2, *Memory and Cognitive Processes*. 3rd ed. New York: Wiley, 599–621.
- Gopnik, Alison, and Andrew Meltzoff (1997), *Words, Thoughts, and Theories*. Cambridge, MA: MIT Press.
- Griffiths, Paul E. (1997), *What Emotions Really Are: The Problem of Psychological Categories*. Chicago: University of Chicago Press.
- (2004), “Is Emotion a Natural Kind?” in Robert C. Solomon (ed.), *Philosophers on Emotion*. Oxford: Oxford University Press, 233–249.
- Hampton, James A. (1979), “Polymorphous Concepts in Semantic Memory,” *Journal of Verbal Learning and Verbal Behavior* 18: 441–461.
- (1981), “An Investigation of the Nature of Abstract Concepts,” *Memory and Cognition* 9 (2): 149–156.
- (1987), “Inheritance of Attributes in Natural Concept Conjunctions,” *Memory and Cognition* 15: 55–71.
- (1993), “Prototype Models of Concept Representation,” in Ivan Van Mechelen, James Hampton, Ryszard S. Michalski, and Peter Theuns (eds.), *Categories and Concepts: Theoretical Views and Inductive Data Analysis*. London: Academic Press, 67–96.
- Hampton, James A., and Margaret M. Gardiner (1983), “Measures of Internal Category Structure: A Correlational Analysis of Normative Data,” *British Journal of Psychology* 74: 491–516.
- Kahneman, Daniel, and Amos Tversky (1973), “On the Psychology of Prediction,” *Psychological Review* 80: 237–251.
- Kelly, Michael H., J. Kathryn Bock, and Frank C. Keil (1986), “Prototypicality in a Linguistic Context: Effects on Sentence Structure,” *Journal of Memory and Language* 25: 59–74.
- Komatsu, Lloyd K. (1992), “Recent Views of Conceptual Structure,” *Psychological Bulletin* 112 (3): 500–526.
- Lakoff, George (1987), “Cognitive Models and Prototype Theory,” in Ulric Neisser (ed.), *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*. Cambridge: Cambridge University Press.
- Laurence, Stephen, and Eric Margolis (1999), “Concepts and Cognitive Science,” in Eric Margolis and Stephen Laurence (eds.), *Concepts: Core Readings*. Cambridge, MA: MIT Press, 3–81.
- Machery, Edouard (2005), “Concepts Are Not a Natural Kind,” *Philosophy of Science* 72: 444–467.
- Malt, Barbara C. (1994), “Water Is Not H₂O,” *Cognitive Psychology* 27: 41–70.
- Malt, Barbara C., and Edward E. Smith (1984), “Correlated Properties in Natural Categories,” *Journal of Verbal Learning and Verbal Behavior* 23: 250–269.
- McCloskey, Michael E., and Sam Glucksberg (1978), “Natural Categories: Well Defined or Fuzzy Sets?” *Memory and Cognition* 6 (4): 462–472.
- Medin, Douglas L., Elizabeth B. Lynch, and Karen O. Solomon (2000), “Are There Kinds of Concepts?” *Annual Review of Psychology* 51: 121–147.
- Michalski, Ryszard S. (1993), “Beyond Prototypes and Frames: The Two-Tiered Concept Representation,” in Iven Van Mechelen, James Hampton, Ryszard S. Michalski, and Peter Theuns (eds.), *Categories and Concepts: Theoretical Views and Inductive Data Analysis*. New York: Academic Press, 141–172.
- Murphy, Gregory L. (2002), *The Big Book of Concepts*. Cambridge, MA: MIT Press.
- Nosofsky, R. M. (1988), “Exemplar-Based Accounts of Relations between Classification, Recognition, and Typicality,” *Journal of Experimental Psychology: Learning, Memory and Cognition* 4: 700–708.
- Pinker, Steven (1994), *The Language Instinct*. New York: Morrow.
- Prinz, Jesse J. (2002), *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge, MA: MIT Press.

- Rips, Lance J. (1995), "The Current Status of Research on Concept Combination," *Mind and Language* 10: 72–104.
- Rosch, Eleanor H. (1973), "On the Internal Structure of Perceptual and Semantic Categories," in Timothy E. Moore (ed.), *Cognitive Development and the Acquisition of Language*. New York: Academic Press, 111–144.
- (1975), "Cognitive Representations of Semantic Categories," *Journal of Experimental Psychology: General* 104 (3): 192–233.
- Rosch, Eleanor, and Carolyn B. Mervis (1975), "Family Resemblances: Studies in the Internal Structure of Categories," *Cognitive Psychology* 7: 573–605.
- Scott, Sam (2003a), *Non-referring Concepts*. PhD Dissertation. Ottawa: Carleton University. Also available as Technical Report 2003-03. Ottawa: Carleton University, Cognitive Science Department. <http://www.carleton.ca/ics/TechReports>.
- (2003b), "Some New Evidence for Concept Stability," in Richard Alterman and David Kirsh (eds.), *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.
- Smith, Edward E., and Douglas L. Medin (1981), *Categories and Concepts*. Cambridge, MA: Harvard University Press.
- Weiskopf, Daniel A. (2005), "Concept Pluralism," unpublished paper.