To be published in *Behavioral and Brain Sciences* (in press)

Précis of: *Doing without Concepts*

Doing without Concepts published by Oxford University Press, 2009

Edouard Machery
Department of History and Philosophy of Science
University of Pittsburgh
1017CL
Pittsburgh, PA 15260
machery@pitt.edu

Abstract: The study of concepts is in an odd state of disarray. Cognitive scientists working on categorization, induction, and reasoning have discovered a dazzling amount of phenomena. New work on prototypes in the 1990s and early 2000s, innovative ideas on causal cognition in the first decade of the twenty-first century, the development of the neo-empiricist approach that assimilates the tokening of a concept to a multi-modal perceptual simulation, and the promising growth of the neuropsychology of concepts have rejuvenated the field. At the same time, this extraordinary amount of findings has yet to be organized in a coherent theoretical framework. The current theories of concepts—prototype theories, exemplar theories, theory theories, and neo-empiricist theories—fail to explain all the known phenomena, and there is very little agreement about what concepts are. *Doing without Concepts* attempts to provide such a theoretical framework. In this article, I review the main points and arguments developed at greater length in *Doing without Concepts*, and I conclude that abandoning the very notion of concept is probably required to remedy the state of disarray of the current psychology of concepts.

Keywords: concept, category, categorization, induction, concept combination, dissociation, empiricism, eliminative argument, natural kind

1. Regimenting the Use of "Concept" in Cognitive Science

Because cognitive scientists rarely spell out the notion of concept in detail, I begin by making explicit the notion of concept that is typically used within cognitive science. My goal in chapters 1 and 2 is threefold: To clarify this notion, to regiment the use of the term "concept," and to show that philosophers and cognitive scientists theorize about different things when developing theories of concepts.

The cognitive processes that underwrite cognitive competences are typically assumed to access some relevant information or knowledge. Some bodies of information are only accessed by particular processes: For instance, our implicit knowledge of the syntax of the natural languages we speak (e.g., English) is only accessed by the processes involved in parsing and in producing sentences. When this is the case, I will say that the relevant body of information is "proprietary to a particular cognitive process." By contrast, some information is "non-proprietary": It is accessed by the cognitive processes that underlie several distinct cognitive competences. Cognitive scientists often assume that the cognitive processes underlying our higher cognitive competences access the same bodies of knowledge. For instance, the processes underlying categorization, induction, and speech are hypothesized to access the same body of knowledge about dogs when people classify something as a dog, when they make an induction about dogs, and when they understand sentences containing the word "dog." This knowledge is assumed to be stored in long-term memory.

These preliminary points having been made, I propose to characterize the notion of concept as follows: Within cognitive science, a concept of *x* is a body of information about *x* that is stored in long-term memory and that is used by default in the processes underlying most, if not all, higher cognitive competences when they result in judgments about *x*. I call this characterization "C."

It is important to highlight some significant properties of concepts, so understood. First, concepts can be about classes of objects (dog), events (going to the dentist), substances (water), and individuals (Barack Obama). Second, concepts are non-proprietary: dog is used by the processes underlying categorization, induction, linguistic understanding, metaphor building, planning, and perhaps other competences. Third, the elements of information that are constitutive of a concept can vary over time and across individuals. Fourth, it might be unclear whether a given element of information about *x* belongs to a concept of *x*. Finally, concepts are used *by default* in the cognitive processes underlying higher cognitive competences (I call "Default" the hypothesis that some bodies of knowledge are retrieved by default when one is categorizing, reasoning, drawing analogies, making inductions, and so on). This entails that our concept(s) of, say,

dogs is (are) only a subset of our whole knowledge about dogs: The knowledge that is constitutive of dog is the knowledge about dogs that is retrieved by default from long-term memory when we reason about dogs, categorize things as dogs, etc. I call "background knowledge" about dogs the knowledge about dogs that is not part of the concept(s) of dogs. Our background knowledge about a category, a substance, some kind of events, etc., can be called upon occasionally when the default body of knowledge is insufficient to solve a cognitive task.

At this point, it is useful to spell out the notion of being used by default at greater length. A default body of knowledge about *x* is the body of knowledge that is presumptively taken to be relevant when one reasons about *x*, when one categorizes things as *x*, and so on. The knowledge that is stored in a concept of *x* is preferentially available when we think, reason, etc., about *x*. So to speak, it spontaneously comes to mind.

The proposed characterization of the notion of concept captures much of what is implicit in the use of the term "concept" in cognitive science. However, it is also clear that C is partly at odds with some characterizations of the notion of concept found in the literature, which are discussed at length in chapter 1 of *Doing without Concepts*. In any case, with the proposed characterization of the notion of concept, I do not merely aim at capturing the use of "concept" in cognitive science. I also want to regiment it: I contend that some bodies of knowledge are retrieved by default from long-term memory when one reason, categorize, etc. (see Section 3 for a defense of this claim), and I propose that "concept" should be used to refer to these bodies of knowledge.

2. Individuating Concepts

It is certainly possible for a given individual to have several concepts of the same category (e.g., several concepts of chairs) or of the same substance (e.g., several concepts of gold): For instance, one might think of chairs in several distinct ways, each of which corresponds to a different concept of chairs. This possibility raises the following question, which is examined in chapter 3 of *Doing without Concepts*: What does it mean for two bodies of knowledge about *x* (e.g., the knowledge that water is typically transparent and the knowledge that water is made of molecules of H_2O) to be part of the same concept of *x* (water) rather than of two distinct concepts ($water_1$ and $water_2$)?

I propose the following individuation criteria (respectively, Connection and Coordination). When two elements of information about *x*, A and B, fulfill either of these criteria, they belong to two distinct concepts:

(1) If retrieving A (e.g., water is typically transparent) from long-term memory and using it in a cognitive process (e.g., a categorization process) does

not facilitate the retrieval of B (water is made of molecules of H_2O) from long-term memory and to use it in some cognitive process, then A and B belong to two distinct concepts ($water_1$ and $water_2$).

(2) If A and B yield conflicting judgments (e.g., the judgment that some liquid is water and the judgment that this very liquid is not water) and if I do not view either judgment as defeasible in light of the other judgment (i.e., if I hold both judgments to be equally authoritative), then A and B belong to two distinct concepts ($water_1$ and $water_2$).

Let's clarify these two criteria a bit further. Connection unpacks the idea that the components of a given concept must be connected (or "linked" as I said in Machery, 2009): If the beliefs that water is typically transparent and that water is made of molecules of H_2O are part of the same concept, water, then using the first belief to classify some sample as water allows me to conclude that it is made of molecules of H_2O. Accessing one part of a concept makes the other parts accessible and poised to be used in cognitive processing (for related discussion, see Millikan, 2000, chapter 10).

Coordination expresses the idea that the parts of a single concept should not result in conflicting judgments that are both taken to be equally authoritative. Coordination is compatible with the parts of a given concept yielding conflicting judgments, provided that all but one are viewed as defeasible. To illustrate, although the belief that grandmothers have grey hair and the belief that grandmothers are mothers of parents would yield different judgments if they were used to categorize a young-looking grandmother (e.g., Sarah Palin), they could still both be part of the concept grandmother if one of the two judgments (Sarah Palin is not a grandmother) is defeated by the other one (Sarah Palin is a grandmother). In effect, the first belief would be treated as a heuristic that sometimes leads us astray.

These two criteria merely explain what it is for a given individual to have one or several concepts about some class of objects, some substance, etc. (within-person individuation). It does not explain what it is for two or more individuals to have the same concept of *x* (between-person individuation). Now, one might expect a theoretician of concepts to provide individuation criteria for this situation too. However, I doubt that these are needed for the psychology of concepts since they seem to play no explanatory role in psychology.

3. Defending the Proposed Notion of Concept

Some cognitive scientists and philosophers of psychology reject Default (viz. the assumption that some bodies of knowledge are retrieved by default when one is categorizing, reasoning, drawing analogies, and making inductions). In chapters 1 and 8

of the book, I rebut the criticisms mounted against this assumption.

Several findings suggest that typicality varies across contexts. Roth and Shoben (1983) have shown that depending on the linguistic context (e.g., when participants are presented with "Stacy volunteered to milk the *animal* whenever she visited the farm" or "Fran pleaded with her father to let her ride the *animal*"), participants judge that different animals are typical (cows and goats for the first linguistic context, horses and mules for the second). Similarly, Barsalou (1985) reports that judgments of typicality vary across contexts. In Study 2, participants' typicality judgments about the members of two groups differed when these groups were conceptualized differently (as physical education teachers and current event teachers on the one hand, and as two invented types of programmers—namely, Q programmers and Z programmers—on the other hand). This study also shows that when participants are familiarized with a given category in different contexts, their judgments of typicality vary. Barsalou (1987, 1993) also reports that the typicality of objects varies when participants are asked to take different points of view on these objects. For instance, people judge differently the typicality of birds when they take the point of view of someone from China and when they report their typicality evaluation from their own point of view. In addition, the correlation between typicality judgments across participants is low (circa .5) and lower than expected for a given subject on two different occasions (around .8).

Theorists have used such findings to challenge the idea that some bodies of knowledge are retrieved by default when one categorizes, draws inductions, reasons, etc. Barsalou concludes (1985, 646; see also Barsalou 1987; Smith and Samuelson 1997, 170)

> [P]eople may not retrieve the same concept from long-term memory everytime they deal with a particular category. Instead they may construct a diverse variety of concepts in working memory to represent a particular category across different situations such that the concept used to represent a category is rarely, if ever, the same. According to this view, long-term memory does not contain invariant concepts.

Theorists who reject Default have drawn two distinct conclusions. As discussed in chapter 8 of *Doing without Concepts*, some theorists, such as L. Smith (Smith & Samuelson, 1997), contend that Default is part of the very notion of concept, and they conclude that since there are no such things as bodies of knowledge retrieved by default from long-term memory, there are no such things as concepts. Smith and Samuelson (1997, 190) conclude that "a successful theory of categories (…) might require that we give up timeless abstractions such as concepts."

While agreeing that Default is part of the notion of concept typically used in cognitive science, others, such as Barsalou and Prinz, propose to redefine the notion of concept. Concepts should be thought of as the bodies of knowledge in working memory

that are used at a given time in a given task; they are constructed on the fly to deal with the peculiarities of the task at hand; and they typically vary from time to time (Barsalou, 1993, 29; Prinz, 2002; Malt & Sloman, 2007).

While Barsalou (1993), Prinz (2002), L. Smith (Smith & Samuelson, 1997, and perhaps Malt and Sloman (2007) take the body of evidence reviewed above to establish either conclusion, I demur on three grounds. First, the hypothesis that some bodies of knowledge are retrieved by default from long-term memory and used in the processes underlying the higher cognitive competences is consistent with some variation in the bodies of knowledge that are used at any given time. This variation can have two sources. When we reason about *x*, in addition to the default body of knowledge about *x*, we sometimes retrieve some specific elements from our background knowledge about *x*. In addition, once retrieved from memory the body of knowledge that is retrieved by default can be tailored to the peculiarities of the given situation. On this view, knowledge retrieval would be a two-step procedure: Retrieve the default body of knowledge from long-term memory; tailor it to the situation (Sperber and Wilson [1998] present similar ideas). Thus, the mere fact that performances in experimental tasks vary from time to time does not show that there are no bodies of knowledge retrieved by default from long-term memory. What would not be consistent with Default is a very large variability across contexts of the knowledge brought to bear on tasks. And indeed Barsalou claims that there is a "*tremendous variability* in performances…not only in category membership, but also in typicality, definitions, and probably most other categorization tasks" (1993, 34; my emphasis). However, as we shall now see, the relevant variability in performances is in fact moderate.

Second, the nature of the variation found by Barsalou, Malt, Sloman, and others is either irrelevant to evaluate Default or supportive of it. Let's consider first the pieces of evidence that turn out to be irrelevant to evaluate Default. Many findings about the context-sensitivity of typicality are misleading. Roth and Shoben's (1983) findings merely show that (unsurprisingly) people evaluate differently the typicality of target animals with respect to the category of animals that get milked and the category of animals one uses to ride. That typicality varies when evaluated from different points of view (Barsalou, 1993) does not show that people's concepts vary across occasions, since, in effect, one asks participants not to use their own concept of *x* to complete a task when one asks them to make judgments about *x* from someone else's point of view. Similarly, the variability of the typicality judgments made by different individuals says nothing about whether a given subject retrieves a default body of knowledge across occasions. Showing, as Barsalou (1985) did, that the typicality structure of a given class of objects can substantially vary when it is conceptualized differently (viz. as physical education teachers and current event teachers on the one hand, and as Q programmers and Z

programmers on the other hand) is interesting, but this finding does not to show that the concept of a given category varies across contexts and circumstances, since properly speaking current events teachers and Z programmers are two distinct categories, although they are composed of the same individuals. Showing that the typicality structure of a given category varies when people are familiarized with this category in different situations is also interesting, but again it says little about whether someone who is familiarized with a given category one way will rely on a default body of knowledge about this category. In addition, some results that allegedly undermine Default in fact support it. As noted, Barsalou (1987, 1993) reports that, on average, the test-retest reliability of typicality judgments is at least .8. It is also higher when participants are re-tested one hour and a day after the first test. Furthermore, Barsalou reports that the typicality of highly typical and atypical items does not change much. These results are evidence that, across occasions, a default concept is retrieved from long-term memory.

Finally, a large body of evidence supports Default. Consider linguistic understanding (Ziff, 1972, discussed in Murphy & Medin, 1985). Consider the sentence, "A cheetah can outrun a man." This sentence is meaningful, and most people would agree with it. However, as Murphy and Medin put it (1985, 303), it is true only if the represented cheetah is not "a 1-day old cheetah, or an aged cheetah with arthritis, or a healthy cheetah with a 100-pound weight on its back." But when we read "A cheetah can outrun a man," these representations of cheetahs do not come to mind. This phenomenon suggests that when a speaker utters "A cheetah can outrun a man" or when a hearer or a reader understands this sentence, she retrieves from memory a default body of knowledge about cheetahs. Perhaps one will object that when one reads the sentence, "A cheetah can outrun a man," one merely constructs a context-appropriate interpretation of "cheetah" rather than retrieving a default body of knowledge about cheetahs. If this were true, then people would also construct a context-appropriate representation of cheetahs if they had to decide whether the sentence, "A man can outrun a cheetah," is true. They would, e.g., imagine an old, three-legged cheetah, and the sentence, "A man can outrun a cheetah," would then be judged true too. However, I predict that, under time pressure, people would judge the sentence, "A man can outrun a cheetah," to be false. This would be evidence that in such conditions they retrieve the very default body of knowledge they retrieve when they read, "A cheetah can outrun a man." Naturally, with no time pressure, they could construct an interpretation of "cheetah" under which the sentence, "A man can outrun a cheetah," is true. But this is consistent with the existence of default bodies of knowledge, since, as I have proposed above, people can and do retrieve some additional information (viz. some information not contained in their concepts) from their background knowledge.

Behavioral studies also show that some information about a category, substance,

etc., is retrieved automatically in every context (Barsalou, 1982; Whithney, McKay, Kellas, & Emerson, 1985). Barsalou (1982) found that when people judge that a property (e.g., *stinks*) spontaneously "comes to mind" when they read a given noun (e.g., "skunk"), reaction times in a property-verification task are similar when the noun is presented in a relevant linguistic context ("the skunk stunk up the entire neighborhood") and when it is presented in an irrelevant linguistic context ("the skunk was under a large willow"). By contrast, reaction times are larger in the latter condition ("the roof had been renovated prior to the rainy season") than in the former condition ("the rook creaked under the weight of the repairman") when people judge that a property (e.g., *can be walked upon*) does not spontaneously come to mind when they read a given noun (e.g., "roof"). Barsalou calls the first kind of property "context-independent" and the second kind "context-dependent."

Neuropsychology provides further evidence in support of Default (although the relevant studies were not developed to test this hypothesis). After having trained participants with novel tools, Weisberg, van Turrennout, and Martin (2006) recorded the brain activation in a perceptual task (a visual matching task). To complete this task, one needs only appeal to some structural information about the shape of the novel tools; thus, one would expect the medial portion of the fusiform gyrus to be activated (for review, see Martin, 2007). Interestingly, activation was also recorded in the intraparietal sulcus, the premotor cortex, and the medial temporal gyrus, areas of the brain that are known to store information about the typical movements associated with tool use. It thus seems that the perceptual task resulted in the automatic retrieval of information that was not needed to solve the task, consistent with the idea that people have bodies of knowledge that they retrieve by default (for similar findings, see also James & Gauthier, 2003; Hoenig, Sim, Bochev, Herrnberger, & Kiefer, 2008).

4. Developing a psychological theory of Concepts

It is important to keep in mind that the notion of concept proposed above (viz. C) does not amount to a *theory* of concepts. Rather, C does two things: It spells out what is implicit in cognitive scientists' use of the term "concept," and it proposes to regiment this use. So, what does a psychological theory of concepts consist in?

As I explain in chapter 1 of *Doing without Concepts*, psychological theories of concepts typically attempt to identify the properties that are typical of concepts ("the general properties of concepts"). Five kinds of properties are of interest to cognitive scientists. First, cognitive scientists are interested in the nature of the information that is constitutive of concepts. For instance, cognitive scientists want to know whether concepts consist of some statistical information about the properties that are characteristic of a

class or of a substance, as prototype theorists have proposed (e.g., Hampton, 1979, 1981, 2006, 2007; Smith, 2002), or whether they consist of causal generalizations (e.g., Murphy & Medin, 1985; Gopnik & Meltzoff, 1997; Rehder, 2003; Griffiths, Steyvers, & Tenenbaum, 2007; Tenenbaum, Griffiths, & Niyogi, 2007). Second, cognitive scientists are interested in the nature of the processes that use concepts. For instance, some psychologists have argued that these processes are based on similarity (e.g., Hampton, 1993), while others disagree (e.g., Rips, 1989). Third, cognitive scientists develop hypotheses about the nature of the vehicles of concepts: Thus, neo-empiricists such as Barsalou and Prinz contend that the vehicle of concepts is similar to the vehicle of perceptual representations (Barsalou, 1999, 2008b, 2009; Prinz, 2002, 2005; Machery, 2006a). Fourth, for about a decade, cognitive scientists have attempted to identify the brain areas that are involved in possessing concepts (for review, see, e.g., Pulvermüller, 2005; Martin, 2007; Mahon & Caramazza, 2009). Finally, cognitive scientists have developed hypotheses about the processes of concept acquisition (e.g., Gopnik, 2003; Ashby & Maddox, 2004).

In addition to developing hypotheses about the general properties of concepts, cognitive scientists have shown some interest in distinguishing different types of concepts and in identifying the properties of these types of concepts. Medin, Lynch, and Solomon (2000) have rightly insisted on the importance of this task and on its relative neglect by cognitive scientists.

Why do cognitive scientists want a theory of concepts? Theories of concepts are meant to explain the properties of our cognitive competences. People categorize the way they do, they draw the inductions they do, and so on, because of the properties of the concepts they have. Thus, providing a good theory of concepts could go a long way toward explaining some important higher cognitive competences.

5. "Concept" in Cognitive Science and in Philosophy

The term "concept" is used in philosophy, particularly in the philosophy of mind, as well as in cognitive science. Chapter 2 of *Doing without Concepts* examines the relation between these two uses. It is common among philosophers to assume that "concept" is used in the same sense in philosophy and in cognitive science and that psychologists' theories of concepts aim at answering the issues philosophers are interested in (Rey, 1983, 1985, 2009; Margolis, 1994, 1995; Fodor, 1998, 2008; Laurence & Margolis, 1999; Margolis & Laurence, 2006; Edwards, 2009). In addition, it is common to hold that as answers to the issues of interest in philosophy, psychological theories of concepts are defective. Thus, Fodor concludes his review of Gregory Murphy's book, *The Big Book of Concepts*, as follows (2003, 4):

> It is part of our not knowing how the mind works that we don't know what concepts are or what it is to have one. Just about everything that current cognitive science says about either topic is wrong. (…) Gregory Murphy's book tells you most of what there is to the psychology of concepts. Read it, therefore, by all means; but don't even consider believing it.

It is also not uncommon to see some philosophical theories criticized for being unable to explain how we categorize, make inductions, and so on (Prinz, 2002; for discussion, see Edwards, 2009).

Philosophers' take on psychological theories is mistaken: Philosophical and psychological theories of concepts are not meant to answer the same questions and are thus not competing. Typically, by "concept," philosophers refer to that which allows people to have propositional attitudes (beliefs, desires, etc.) about the objects of their attitudes. The concept of a triangle is thus that which allows people to have propositional attitudes (beliefs, desires, etc.) about triangles. A theory of concepts in philosophy attempts to determine the conditions under which one can have propositional attitudes about the objects of their attitudes (Peacocke, 1992, 2008; Fodor, 1998, 2008), but not to explain the properties of our higher cognitive competences. By contrast, psychologists attempt to explain the properties of our categorizations, inductions, etc., but they do not attempt to determine the conditions under which people are able to have the propositional attitudes about the objects of their attitudes. Furthermore, psychologists do not need to hold, and typically do not hold, that we are able to have propositional attitudes about the objects of our attitudes by virtue of having specific bodies of knowledge about them. For instance, prototype theorists do not need to hold, and typically do not hold, that having a prototype is a condition for being able to have attitudes about the objects of our attitudes. In fact, prototype theorists are silent on this question.

The upshot of this argument should be clear. Although both philosophers and cognitive scientists use the term "concept," they are not talking about the same things. Cognitive scientists are talking about a certain kind of bodies of knowledge, while philosophers are talking about that which allows people to have propositional attitudes. Many controversies between philosophers and psychologists about the nature of concepts are thus vacuous.

6. The Heterogeneity Hypothesis versus the Received View

Cognitive scientists of concepts naturally acknowledge differences between concepts: The concept of dogs is clearly different from the concept of cats. More interesting, they also acknowledge differences between kinds of concepts: For instance, there has been much work in experimental and developmental psychology on the differences between

the concepts of animals and the concepts of artifacts (e.g., Gelman, 1988, 2003; Gelman & Markman, 1986; Bloom, 1996; Malt & Sloman, 2007). But above and beyond these differences, cognitive scientists often assume that concepts share many properties that are scientifically interesting. In chapter 3 of *Doing without Concepts*, I call this assumption "the received view." It is well expressed by Gregory Murphy:

> The psychology of concepts cannot by itself provide a full explanation of the concepts of all the different domains that psychologists are interested in. (…) The details of each of these must be discovered by the specific disciplines that study them (…). Nonetheless, the general processes of concept learning and representation may well be found in each of these domains. (Murphy, 2002, 2-3)

The received view has been instrumental in the debates that have marked the history of the psychology of concepts since the 1970s. Cognitive scientists who are committed to different theories of concepts (say, a particular prototype theory and a particular exemplar theory) have attempted to discover properties of our higher cognitive competences (e.g., the exemplar effect reported in Medin & Schaffer, 1978) that were easily explained by the theory they endorsed (e.g., the exemplar theory), but that were not easily explicable by the competing theory (prototype theories do not naturally explain the exemplar effect; for discussion, see Smith & Minda, 2000). This research strategy makes sense only if one supposes that a single theory of concepts should be able to account for all the relevant phenomena. If, contrary to the received view, the class of concepts divides into several kinds that have little in common, the distinct theories of concepts that characterize these kinds of concepts will account for different phenomena, and the fact that theory A, but not theory B, explains some phenomenon, such as the exemplar effect, will not necessarily constitute evidence against theory B.

As I explain in chapter 3, the received view stands in sharp contrast with a view about concepts developed in my own work (see also Machery, 2005): the heterogeneity hypothesis. According to this hypothesis, the class of concepts divides into several distinct kinds that have little in common—"the fundamental kinds of concepts." Because the class of concepts divides into distinct fundamental kinds, it is a mistake to assume that there are many general properties of concepts, and that a theory of concepts should attempt to describe these. Although the heterogeneity hypothesis can be developed in several ways (Machery, 2005, 2006b; Piccinini & Scott, 2006), I contend that a given category (e.g., dogs), a given substance (e.g., water), or a given kind of events (e.g., going to the dentist) is typically represented by several distinct concepts (e.g., dog_1 and dog_2). These coreferential concepts belong to the fundamental kinds of concepts. Each coreferential concept can be used to categorize, draw inductions, understand the relevant words, make analogies, etc. (Figure 1).

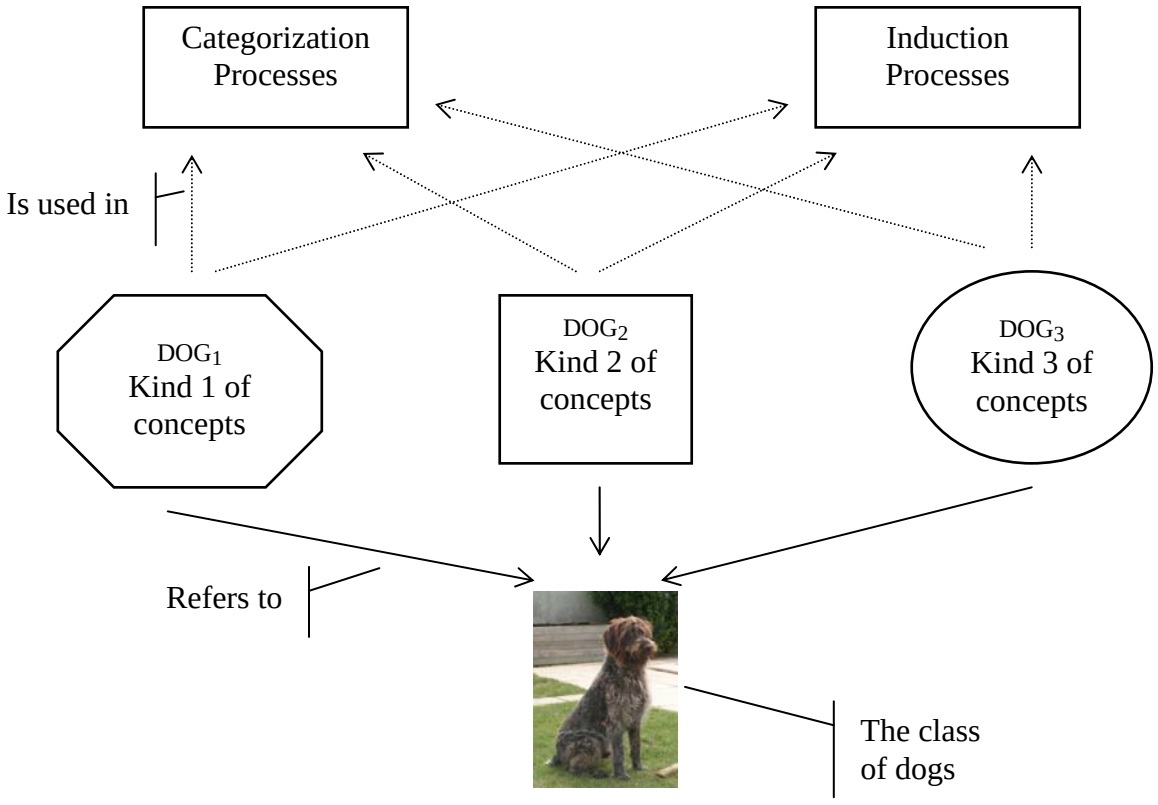

Figure 1: The Heterogeneity Hypothesis

In addition, the heterogeneity hypothesis contends that these concepts are often used in distinct processes. That is, we have several categorization processes, several induction processes, etc., each of which uses a distinct fundamental kind of concepts (Figure 2).

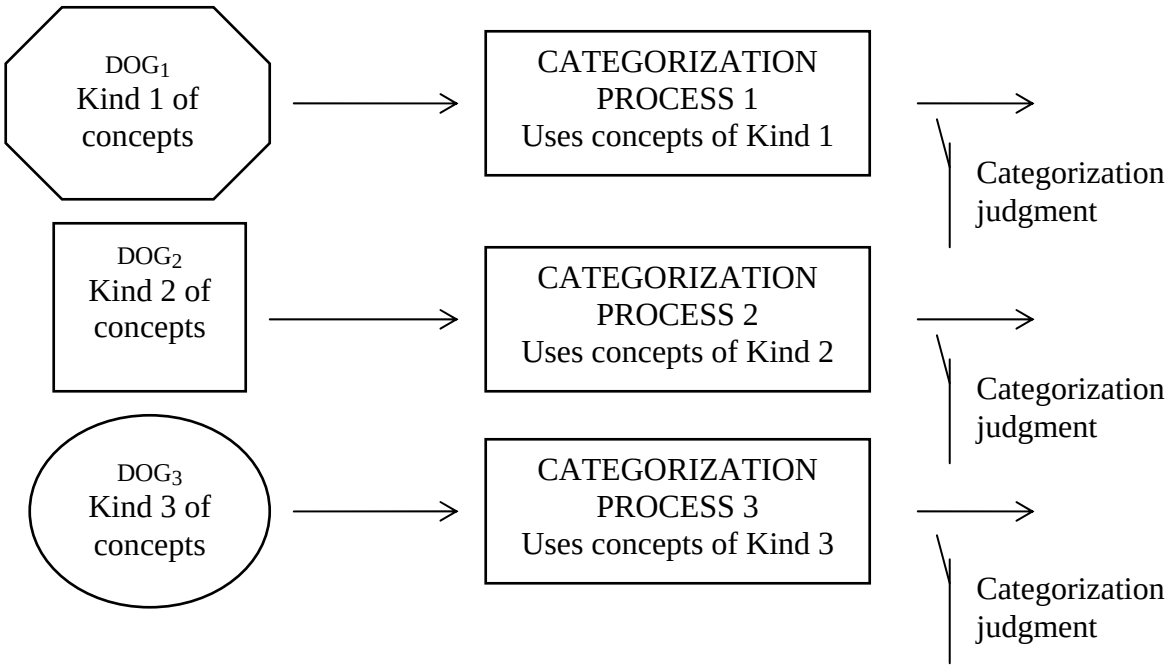

Figure 2: Several Processes Underlying a Given Cognitive Competence

If the heterogeneity hypothesis is correct, the class of concepts is not a natural kind. Natural kinds are classes whose members share many scientifically important properties in virtue of one or several causal mechanisms (Boyd, 1991, 1999; Griffiths, 1997; Machery, 2005). Water and dogs are natural kinds in this sense; for instance, samples of water have many properties in common in virtue of consisting of the same molecules of H_2O. In a given science, the scientific classificatory scheme is developed to identify the natural kinds in the relevant domain because identifying these kinds allows scientists to discover new generalizations. Scientific classificatory schemes are modified when they do not identify the relevant natural kinds (as happened during the chemical revolution in the eighteenth century), and scientific notions are often eliminated when it is found that they fail to pick out natural kinds (for discussion, see section 13). Because the hypothesized fundamental kinds of concepts have little in common, the class of concepts cannot be a natural kind if the heterogeneity hypothesis is correct.

7. What Kind of Evidence Could Support the Heterogeneity Hypothesis?

Chapter 5 of *Doing without Concepts* describe three kinds of evidence that can provide evidence for the heterogeneity hypothesis. I consider them in turn in this section.

Suppose that the class of concepts divides into several fundamental kinds, and suppose also that coreferential concepts are often used in distinct cognitive processes (i.e., distinct categorization processes, distinct induction processes). What properties would we then expect to observe in experimental tasks? First, if experimental conditions can be designed that trigger only one of the hypothesized categorization processes or only one of the hypothesized induction processes, we should expect some experimental findings to be best explained if the concepts used in the relevant experimental tasks are identical to a first fundamental kind of concepts, other experimental findings to be best explained if the concepts used in the relevant experimental tasks are identical to a second fundamental kind of concepts, and so on. For instance, if one hypothesizes that the fundamental kinds of concepts are exemplars and prototypes, then one might find categorization tasks where participants' categorization performances are best explained if the concepts used in these tasks are prototypes and other categorization tasks where participants' categorization performances are best explained if the concepts used in these tasks are exemplars.

Second, suppose that in some conditions, several of the hypothesized categorization (or induction) processes are triggered at the same time. Then, in some circumstances, these processes will produce congruent outputs (e.g., categorization judgments), while they will produce incongruent outputs in other circumstances. When the latter happens, participants would have to decide between conflicting judgments. Participants should thus be expected to be slower when the hypothesized processes are expected to yield conflicting outputs than when they are not. Test-retest reliability should also be expected to be lower in the experimental conditions where it is hypothesized that the hypothesized categorization (induction) processes will result in incongruent outputs than when it is hypothesized that they will result in congruent outputs. Noticeably, this kind of evidence (particularly, slower reaction times) has extensively been used in cognitive science to argue that a given task involves two independent cognitive processes (e.g., Greene, Sommerville, Nystrom, Darley, & Cohen, 2001).

Finally, experimental and neuropsychological dissociations can be used to determine whether a given task involves several processes. The epistemology of dissociations is intricate (Caramazza, 1986; Dunn & Kirsner, 1988, 2003; Shallice, 1988; Glymour, 1994; Plaut 1995; Van Orden, Pennington, & Stone, 2001; Ashby & Ell, 2002), but, I maintain, dissociations provide evidence about the number and nature of the processes underlying a given competence.

8. The Fundamental Kinds of concepts

Now that the nature of the evidence required to support the heterogeneity hypothesis has been clarified, it is time to lay my cards on the table: What are the fundamental kinds of concepts? And what is the evidence for their existence? In what follows, I will briefly describe the kind of evidence supporting the heterogeneity hypothesis, but due to limitations of space this will not amount to a comprehensive articulation of the evidence adduced in chapters 6 and 7 of *Doing without Concepts*.

In chapter 4, I propose that the class of concepts divides into at least three fundamental kinds of concepts—prototypes, exemplars, and theories. These three theoretical constructs are well known in the psychology of concepts since they correspond to the entities posited by the main theories of concepts that have been developed since the 1970s (for review, see Murphy, 2002). Although there are several distinct theories about what prototypes, exemplars, and theories are, these theories agree about the distinctive features of each type of concept. In substance, prototypes are bodies of statistical knowledge about a category, a substance, a type of event, etc. For instance, a prototype of dogs could store some statistical knowledge about the properties that are typical of dogs or/and the properties that are diagnostic of the class of dogs. According to prototype theories, when I categorize, draw an induction, make an analogy, etc., I spontaneously bring to mind the properties that are typical, diagnostic, etc., of the relevant category, substance, etc. Prototypes are typically assumed to be used in cognitive processes that compute the similarity between a prototype and other representations, such as the representations of the objects to be categorized, in a linear manner (Rosch & Mervis, 1975; Hampton, 1979, 1993, 2006, 2007; Smith, 2002). Exemplars are bodies of knowledge about individual members of a category (e.g., Fido, Rover), particular samples of a substance, particular instances of a kind of event (e.g., my last visit to the dentist). For instance, according to exemplar theories, a concept of dogs would consist of a set of bodies of knowledge about specific dogs (Rover, Fido). When I categorize, draw an induction, make an analogy, etc., I spontaneously bring to mind the properties of specific members of the relevant categories, of specific samples of the relevant substances, etc. Exemplars are typically assumed to be used in cognitive processes that compute the similarity between a set of exemplars and other representations, such as the representations of the objects to be categorized, in a non-linear manner (Medin & Schaffer, 1978; Nosofsky, 1986, 1993; Nosofsky & Stanton, 2005). Theories are bodies of causal, functional, generic, and nomological knowledge about categories, substances, types of events, etc. A theory of dogs would consist of some such knowledge about dogs. When I categorize, draw an induction, make an analogy, etc., I spontaneously bring to

mind this causal, functional, generic, and nomological knowledge. Recent work on causal knowledge suggests that theories might be used in cognitive processes that are similar to the algorithms involved in causal reasoning (Gopnik et al., 2004).

Thus, the heterogeneity hypothesis proposes that for many categories, substances, kinds of events, we typically have a prototype, a set of exemplars, and a theory about them. Thus, we might have a prototype of dogs, a set of exemplars of particular dogs, and a theory about dogs. Furthermore, prototypes, exemplars, and theories are often used in distinct processes. The heterogeneity hypothesis proposes that we have a prototype-based categorization process, an exemplar-based categorization process, and a theory-based categorization process. Note that the hypothesis is not merely that our knowledge about dogs includes some knowledge about their typical or diagnostic properties, some knowledge about some particular dogs, and some causal, functional, and generic knowledge (as Rey [2009] erroneously believes). This would be a fairly uncontroversial claim. Rather, the claim is that for most categories, substances, etc., we have several bodies of knowledge that are retrieved by default and that are often used in distinct cognitive processes (e.g., several distinct categorization processes).

The heterogeneity hypothesis also contends that the fundamental kinds of concepts have little in common. This is indeed the case if these fundamental kinds really consist of prototypes, exemplars, and theories. They consist of distinct kinds of knowledge, they are used in different kinds of processes, and they are probably acquired by distinct processes. Given what cognitive scientists working on concepts are interested in (Section 4), they count as very different kinds of entities.

One might perhaps object that prototypes, exemplars, and theories do have some properties in common. In particular, they are all bodies of knowledge, they are all stored in long-term memory, and they are all used in the processes underlying higher cognition. This, however, does not undermine the heterogeneity hypothesis, for the claim that prototypes, exemplars, and theories have little in common really states that the fundamental kinds of concepts have in common few properties that are *scientifically interesting* and *discovered empirically*. Prototypes, exemplars, and theories have in common numerous properties that are not of interest to cognitive scientists (e.g., they are all mental states). In addition, far from being discovered empirically, the mentioned commonalities between prototypes, exemplars, and theories (e.g., they are all bodies of knowledge, they are all stored in long-term memory, etc.) are in fact used to identify what concepts are.

So, what is the evidence for the claim that our long-term memory stores prototypes, exemplars, and theories? When one examines thirty years of research on categorization and induction, as I do in chapters 6 and 7, one finds out that in both areas of research, some phenomena are well explained if the concepts elicited by some

experimental tasks are prototypes, some phenomena are well explained if the concepts elicited by other experimental tasks are exemplars, and yet other phenomena are well explained if the concepts elicited by yet other experimental tasks are theories. As noted above, if one assumes that experimental conditions prime the reliance on one type of concepts (e.g., prototypes) instead of other types (e.g., exemplars and theories), this provides evidence for the heterogeneity hypothesis.

Let's illustrate this situation with the work on categorical induction—the capacity to conclude that the members of a category possess a property from the fact that the members of another category possess it and to evaluate the probability of this generalization (for review, see Heit, 2000; Murphy, 2002, chapter 8; Sloman & Lagnado, 2005; Feeney & Heit, 2007). A large number of phenomena suggest that prototypes or exemplars are sometimes involved in induction (Osherson, Smith, Wilkie, Lopez, & Shafir, 1990; Sloman, 1993). Similarity-based models of induction, which assume that processes underlying induction are defined over either prototypes or over exemplars, explain best two well-known findings about induction—the similarity effect and the typicality effect. Other phenomena are best explained if the concepts involved in the relevant experimental conditions are causal theories. Investigating the judgments made by tree experts (landscapers, taxonomists, and parks maintenance personnel) about the strength of inductive conclusions about trees, Profitt and colleagues (2000) found that rather than relying on typicality (as predicted, for instance, by Osherson and colleagues' similarity-coverage model), the pattern of answers and the justifications provided suggest that experts often base their judgments on theories about hypothetical causal mechanisms (see also López et al., 1997). As explained in Section 7, the fact that different properties of our inductive competence are best explained by theories positing different theoretical entities (viz. prototypes, exemplars, or theories) constitutes evidence for the existence of distinct kinds of concepts used in distinct processes. Strikingly, this conclusion is consistent with the emerging consensus among psychologists working on induction that people rely on several distinct induction processes (Profitt et al., 2000; Murphy, 2002; Sloman & Lagnado, 2005; Rehder, 2006).

A natural question raised by these findings concerns the conditions that prime the reliance on prototypes rather than exemplars and theories or on theories rather than prototypes and exemplars (and so on) in induction (see Section 12 below). Because cognitive scientists have rarely fully embraced the idea that there are several distinct kinds of concepts and several processes defined over them, there is little systematic work on this question (but see Rehder, 2006).

The research on categorization and concept learning, reviewed in chapter 6, tells an even clearer story, providing evidence for the existence of prototypes, exemplars, and theories that are used in distinct categorization processes. The research on concept

combination, reviewed in chapter 7, also shows that when people produce a complex concept, they appeal to exemplars, prototypes, and theories. However, in contrast to the research on induction and categorization, it appears that a single process uses prototypes, exemplars, and theories (instead of several distinct combination processes, each of which uses a distinct kind of concepts).

9. Neo-Empiricism

A number of cognitive scientists have recently developed a new approach to the nature of concepts (Glenberg, 1997; Barsalou, 1999, 2005, 2008a, 2009; Martin & Chao, 2001; Barsalou et al., 2003; Prinz, 2002, 2005; Gallese & Lakoff, 2005; Martin, 2007), which I have called "neo-empiricism" (Machery, 2006a, 2007). Although there are differences between neo-empiricist theories, they all endorse the two following theses:

(1) The knowledge that is stored in a concept is encoded in several perceptual and motor representational formats.

(2) Conceptual processing involves essentially reenacting some perceptual and motor states and manipulating these states.

Thesis 1 is about the format of concepts: Neo-empiricists claim that conceptual knowledge is encoded in perceptual and motor representational formats. By contrast, amodal theorists contend that our conceptual knowledge is encoded in a representational format that is distinct from the perceptual and motor representational formats (Barsalou et al., 2003, 85). This distinct representational format is usually thought of as being language-like, although, importantly, amodal representations need not form a language (see below). Thesis 2 concerns the nature of the cognitive processes underlying categorization, induction, deduction, analogy-making, planning, or linguistic comprehension. The central insight is that retrieving a concept from long-term memory during reasoning or categorization consists in tokening some perceptual representations, a process called "simulation" or "reenactment." Cognitive processing consists in manipulating these reenacted percepts (e.g., Barsalou, 1999, 578). Following Barsalou (1999), I will use the term "perceptual symbols" to refer to concepts understood in accordance with Theses 1 and 2.

Perceptual symbols might be one of the fundamental kinds of concepts, but I argue in chapter 4 of *Doing without Concepts* that the evidence provided so far falls short of establishing this. I have identified three main shortcomings of the research in support of neo-empiricism (see also Machery, 2007; for other arguments, see Machery, 2006a; Mahon & Caramazza, 2008; Dove, 2009).

First, what I have called "Anderson's problem" by reference to Anderson's (1978) work on the controversy between imagistic and propositional theories of thinking. Neo-

empiricists typically contrast the predictions made by amodal theories of concepts and the predictions made by neo-empiricist theories of concepts (e.g., Solomon & Barsalou, 2001, 2004; Pecher, Zeelenberg, & Barsalou, 2004; Yaxley & Zwaan, 2007), and they then attempt to show that the neo-empiricist predictions, but not the amodal predictions, are verified. The problem is that there is no such thing as the amodal prediction of concepts; rather, different amodal theories of concepts make different predictions, depending on what they assume about the processes that use amodal concepts. In numerous cases, some amodal theories of concepts make exactly the same predictions as the neo-empiricist theories of concepts developed by cognitive scientists such as Barsalou (for some examples, see Machery, 2007, 2009; Mahon & Caramazza, 2008). As a result, neo-empiricist findings do not distinguish between neo-empiricism and amodal theories of concepts in general. Rather, they provide evidence against specific amodal theories of concepts, while being naturally accommodated by other amodal theories of concepts.

Second, what I have called "the problem from imagery." Most proponents of amodal theories of concepts (e.g., Fodor, 1975; Simon, 1993) acknowledge that in some situations, people rely on imagery (e.g., visual imagery). For instance, we visualize our own home when we are asked how many windows it has. What proponents of amodal theories of concepts deny is that imagery is the only type of processes people have: People also have amodal concepts that are used in non-perceptual processes. The fact that proponents of amodal theories of concepts recognize the role and importance of imagery entails that when amodal theorists expect people to rely on imagery to solve a particular task, showing that people use imagery in this task fails to provide evidence for neo-empiricism and against amodal theories of concepts (for some examples, see Machery, 2007).

Third, what I have called "the generality problem." Neo-empiricists typically assume that all concepts are perceptual symbols. However, it could be that perceptual symbols constitute only a kind of concepts—a hypothesis that would naturally be consistent with the heterogeneity hypothesis. In fact, research suggests that at least some conceptual representations—viz. the representations of the magnitudes of classes of objects or sequences of sounds—are not perceptual, but amodal (Machery, 2007; Dove, 2009). Although these representations do not form a language and thus are different from the hypothesized representations of Fodor's (1975; 2008) language of thought, they are not perceptual either (Machery, 2006a; Dove, 2009). Dove (2009) has developed the generality problem in great detail, showing that the research in support of neo-empiricism has typically focused on a single kind of concepts—viz. "concrete or highly imageable concepts" (2009. 431)—and that neo-empiricist findings are unlikely to be found with concepts with low imageability such as abstract concepts.

Others have identified further difficulties. Reviewing a range of

neuropsychological work on concepts and various important behavioral studies, Mahon and Caramazza (2008) grant that the perceptual and motor systems are often activated during conceptual processing, but they insist that this activation falls short of supporting neo-empiricism, for it can be interpreted in two different ways. First, the interpretation preferred by neo-empiricists: The brain areas involved in perceptual and motor processing or the areas near those are activated because concepts are perceptual and motor representations and perceptual representations are realized in these areas. Second, the amodal interpretation of these findings: The activation of these brain areas results from the activation of other brain areas, not involved in perceptual processing, and from this activation spreading from the latter areas to the former (a well-known phenomenon). Both interpretations account equally well for the neo-empiricist findings.

Finally, let's say a few words about the neo-empiricist research in neuropsychology. A large number of fMRI studies show that tasks meant to tap into the processes underlying higher cognition (particularly, tasks involving the understanding of words) activate either the very brain areas involved in perceptual and motor processing or brain areas near those (see, e.g., Martin & Chao, 2001; Thompson-Schill, 2003; Pulvermüller, 2005; Martin, 2007; Kiefer et al., 2007; Simmons et al., 2007; Barsalou, 2008). However, in contrast to neo-empiricists' usual interpretation of these findings, I contend that much of the neuropsychological research on concepts challenges this approach. Because neo-empiricists insist that tokening a concept is tokening some perceptual representations, they are committed to the view that concept retrieval should activate our perceptual areas (Simmons et al., 2007, 2803). However, a typical finding in neuropsychology is that the brain areas activated are near, *and thus not identical to*, the brain areas involved in perceptual or motor processing (a point acknowledged by Simmons et al., 2007). Furthermore, in much neo-empiricist research on concepts in neuropsychology, the brain areas that are activated in the tasks meant to tap into the processes underlying higher cognition are anterior to the brains areas activated in perceptual processing (e.g., Martin, Haxby, Lalonde, Wiggs, & Ungerleider, 1995; Chao & Martin, 1999; Kable, Kan, Wilson, Thompson-Schill, & Chatterjee, 2005; Pulvermüller & Hauk, 2006). A plausible interpretation is that the brain areas activated in the tasks tapping into higher cognition are amodal representations, which are distinct from the perceptual representations activated in the tasks tapping into perceptual processes, but near them. To conclude, it might be that perceptual symbols are a fundamental kind of concepts, but research still fails to establish it beyond doubt.

10. Hybrid Theories of Concepts

Several hybrid theories of concepts have been developed since the 1970s, and there is a fair amount of differences between them, but they all agree on several crucial points. Hybrid theories of concepts grant the existence of several types of bodies of knowledge, but deny that these form distinct concepts; rather, these bodies of knowledge are the parts of concepts. Like the heterogeneity hypothesis, hybrid theories of concepts typically concur that these parts store different types of information. For instance, some hybrid theories (Gelman, 2004) have proposed that one part of a concept of *x* might store some statistical information about the *x's*, while another part stores some information about specific members of the class of *x's*, and a third part some causal, nomological, or functional information about the *x's*. Furthermore, they often contend that the distinct parts that compose a given concept are used in different processes (e.g., Osherson & Smith, 1981). For instance, the parts that compose a given hybrid concept might be used in distinct categorization processes, in distinct induction processes, and so on.

Although hybrid theories of concepts and the heterogeneity hypothesis agree on several points, they are far from being identical. In Section 2, I proposed two individuation criteria that specify when two bodies of knowledge about *x* form two distinct concepts rather than a single concept. Hybrid theories of concepts contend both that the different coreferential bodies of knowledge are connected and that they are coordinated. The heterogeneity hypothesis assumes that at least one of these two claims is false.

Evidence tentatively suggests that prototypes, set of exemplars, and theories are not coordinated. Malt's (1994) work on how people conceptualize water shows that people have at least two distinct concepts of water—a theoretical concept of water that identifies water with any substance composed of molecules of H_2O and a prototype that identifies water with any substance that shares some typical properties (origins, use, appearances). More recently, Machery and Seppälä (forthcoming) have shown that many participants (between 20% and 80% depending on the pair of sentences considered) are willing to endorse apparently contradictory sentences of the following form:

(1) In a sense, tomatoes are a fruit

(2) In a sense, tomatoes are not a fruit

(3) In a sense, whales are fish

(4) In a sense, whales are not fish.

That is, many participants say that both (1) and (2) and both (3) and (4) are true. Although there are several possible explanations of these findings, a plausible explanation is that people retrieve different concepts of tomato when they read (1) and (2). When they retrieve a theory of tomatoes, they answer that (1) is true, while they

answer that (2) is true when they retrieve a prototype of tomatoes. This suggests that their prototype and their theory form two distinct concepts rather than two parts of the same concept.

11. Multi-process theories

The heterogeneity hypothesis proposes that prototypes, exemplars, and theories are often used in distinct cognitive processes (e.g., distinct categorization processes). I call "multi-process theories" those theories that contend that a given cognitive competence (e.g., categorization, induction, or the capacity to make moral judgments) is underwritten by several distinct processes. Chapter 5 of *Doing without Concepts* is dedicated to examining this kind of cognitive theory. Dual-process theories, which have been embraced in social psychology, are a type of multi-process theory, characterized by a distinction between two types of processes (slow, analytic, intentional processes, and fast, automatic processes; for discussion, see also Gigerenzer, 1996). Gigerenzer and Todd's fast and frugal heuristics is another kind of multi-process theory (Gigerenzer, Todd, & the ABC research group, 1999).

The default hypothesis in cognitive science is that a cognitive competence is underwritten by a single cognitive process, and the burden of proof typically hangs on those who hold a multi-process theory for some cognitive competence. In light of the recent work on a range of cognitive competences, this state of affairs should be revised. Evidence suggests that cognitive competences are commonly underwritten by several distinct processes.

Postulating that a given cognitive competence is underwritten by distinct processes raises a host of questions that have rarely been explicitly confronted by proponents of multi-process theories (but should be). The two most important issues are the following ones:

A. In what conditions are the cognitive processes underlying a given cognitive competence triggered? Are they all always triggered? Are they rather triggered in distinct circumstances? Or, perhaps, in overlapping circumstances? If they are not all always triggered, what cues or processes determine their triggering? Is their triggering under intentional control?

B. If the cognitive processes that underlie a given cognitive competence are triggered in the same conditions, how does the mind choose between their outputs or integrate them?

It is fair to say that current multi-process theories, such as the dual-process theories, have typically failed to give clear answers to these questions. This limits their capacity to genuinely predict experimental outcomes.

What about the prototype-based, exemplar-based, and theory-based cognitive processes? In what conditions are they triggered? And if they are triggered simultaneously, how does the mind choose between their outputs? There is no systematic work on these issues; in fact, I hope that this book will invite cognitive scientists to systematically investigate them.

What is known can be presented briefly. It appears that the categorization processes can be triggered simultaneously (e.g., Allen & Brooks, 1991; Smith et al., 1998), but that some circumstances prime reliance on one of the categorization processes. Reasoning out loud seems to prime people to rely on a theory-based process of categorization (Smith & Sloman, 1994). Categorizing objects into a class with which one has little acquaintance seems to prime people to rely on exemplars (Smith & Minda, 1998). The same is true of these classes whose members appear to share few properties in common (Smith & Minda, 2000; Minda & Smith, 2001; Murphy, 2002). Very little is known about the induction processes except for the fact that expertise seems to prime people to rely on theoretical knowledge about the classes involved (López et al., 1997; Profitt et al., 2000).

12. Open Questions

One of the virtues of the heterogeneity hypothesis is to bring to the fore a range of questions that have not been systematically examined by cognitive scientists. I now summarize some of these issues.

First, psychologists should investigate the factors that determine whether an element of knowledge about *x* is part of the concept of *x* rather than being part of the background knowledge about *x*. Frequency of use is the only factor that has been systematically investigated (Barsalou, 1982). Other factors should be considered—including attention and explicit teaching.

Second, there are several prototype theories, several exemplar theories, and several theory theories. While evidence indicates that we have prototypes, exemplars, and theories, it remains however unclear which prototype theory, which exemplar theory, and which theory theory is correct. That is, the exact nature of prototypes, exemplars, and theories remains to be investigated. Cognitive scientists have typically attempted to show that concepts are prototypes, or that concepts are exemplars, or that concepts are theories, but they have paid little attention to investigating the nature of prototypes, exemplars, and theories in great detail. Similarly, it is unclear which prototype-based model of categorization (induction, etc.), which exemplar-based model of categorization (induction, etc.), and which theory-based model of categorization (induction, etc.) is correct (instead of comparing, say, a specific exemplar-based and a specific prototype-

based model of categorization, as has usually been done). Recently, some psychologists have taken up the important task of comparing the models of categorization and of induction developed by prototype theories (Sloman & Lagnado, 2005) as well as the models of categorization and of induction developed by theory theorists (Rehder & Kim, 2006). Such efforts should be systematically pursued.

Third, multi-process theories are also another important research area that requires systematic attention. I have sketched a framework for developing multi-process theories of the higher cognitive competences, identifying several key questions that need to be answered by proponents of these theories. Multi-process theories need to deal with some important issues that have not been fully solved. Of particular importance is the kind of evidence that can support multi-process theories. Among the three kinds of evidence I have distinguished, the legitimacy of dissociations remains controversial and should be investigated further. It is also plausible that other kinds of evidence can support multi-process theories. While contemporary psychologists often endorse dual-theories of cognition that distinguish System 1 and System 2-processes, there are numerous other types of multi-process theory. Furthermore, existing multi-process theories, such as dual-process theories, do not specify in which conditions the hypothesized processes are triggered and how their outputs are integrated, and, as a result, they are unable to yield clear predictions instead of mere post hoc accommodations.

Psychologists should also develop detailed multi-process theories of those cognitive competences that are the best candidates for being realized by several distinct processes—viz. categorization and induction (Section 11). So far, we know very little about how the distinct cognitive processes that realize competences such as categorization and induction are organized. We do not really know whether outside the lab the categorization (or induction) processes are triggered simultaneously or in distinct conditions. We do not really know what determine their triggering. And we do not know what happens to the outputs of the categorization (or induction) processes when these processes are simultaneously triggered.

13. Concept Eliminativism

Let’s take stock. We have seen that the class of concepts divides into several distinct kinds of concepts, namely, prototypes, exemplars, and theories, which have little in common. Categories, substances, events are often represented by several coreferential concepts (a prototype, a set of exemplars, and a theory). These are not parts of concepts, but are rather bona fide concepts. Prototypes, exemplars, and theories are also typically used in distinct cognitive processes, for instance, in distinct categorization processes, although little is known about the organization of these processes. It is rarely the case that

a given cognitive competence is underwritten by a single cognitive process; rather, the mind usually includes several distinct processes that do the same thing. Finally, the heterogeneity hypothesis focuses attention on a range of empirical questions, for which systematic empirical information is missing, partly because cognitive scientists have not considered the heterogeneity hypothesis seriously.

To conclude this article, I want to discuss a radical proposal inspired by the views about concepts presented in *Doing without Concepts* and in this article: Cognitive scientists might be better off renouncing the very notion of concept. Rather, they should use theoretical terms introduced to refer to the fundamental kinds of concepts—viz. "prototype," "exemplar," and "theory." The heterogeneity hypothesis contends that the class of concepts is not a natural kind: It does not support many causally grounded generalizations because the class of concepts divides into several fundamental kinds that have little in common. Furthermore, theoretical terms are often rejected when it is found that they fail to pick out natural kinds. To illustrate, some philosophers (Murphy & Stich, 1999, building on Griffiths, 1997) have proposed to eliminate the term "emotion" from the theoretical vocabulary of psychology on precisely these grounds. The proposal here is that "concept" should be eliminated from the vocabulary of cognitive science for the same reason.

Chapter 8 of *Doing without Concepts* examines the intricate and controversial logic of eliminativist arguments (see also Stich, 1996; Murphy & Stich, 1999; Mallon, Machery, Nichols, & Stich, 2009). Many eliminativist arguments attempt to conclude that there are no *x's* (for instance, no beliefs, e.g., Churchland, 1981 and Stich, 1983, or no races, for discussion, see, e.g., Mallon, 2006) from the fact that the definition of *x* is not satisfied: For instance, if nothing satisfies the definition of "belief" or "concept," it is concluded that there are no beliefs of concepts. However, such eliminativist arguments are bound to be unsuccessful because they are enmeshed with controversial issues concerning how words such as "belief" or "concept" refer (Mallon et al., 2009). What we need is another kind of eliminativist argument, which clarifies when it is legitimate to eliminate a scientific term from a scientific classificatory scheme.

In a nutshell, I propose that scientific terms should be eliminated on pragmatic grounds (this is what I have called "scientific eliminativism"): To determine whether *x* has a legitimate place in the vocabulary of a given science or whether it should be eliminated, one should examine whether using *x* helps to fulfill the goals of this science—particularly, whether it helps its classificatory purposes. Picking out natural kinds is the primary function of theoretical terms in many sciences (Quine, 1969; but perhaps not in all sciences: Russell, 1948). Thus, when it is found out that a scientific term fails to pick out a natural kind, there is a presumption that it should be eliminated from the relevant science. However, one still needs to consider and to weigh the costs and

benefits of eliminating this term. Perhaps, keeping this term might simplify communication between scientists. On the other hand, keeping this term might prevent the development of a more appropriate classificatory scheme (a common situation, I suspect). If the relevant term does not pick out a natural kind and if the benefits of keeping it do not clearly overweigh the costs, then one should eliminate it.

Because "concept" does not pick out a natural kind if the heterogeneity hypothesis is correct, there is a presumption that it should be eliminated from the theoretical vocabulary of psychology. Furthermore, the continued use of "concept" in cognitive science might invite cognitive scientists to look for commonalities shared by all concepts or to develop another theory that would encompass all the phenomena known about the processes underlying higher cognition. If the heterogeneity hypothesis is correct, these efforts would be wasted. By contrast, replacing "concept" with "prototype," "exemplar," and "theory" would bring to the fore the urgent open questions discussed in Section 12. For instance, speaking of a prototype-based categorization process, of an exemplar-based categorization process, and of a theory-based categorization process makes it clear that there are several categorization processes and brings to the fore the questions of the organization of these categorization processes.

Now, one might worry that eliminating the word "concept" would make communication among cognitive scientists cumbersome. To some extent, this is likely to be true, as is suggested by the frequent use of this term in this article. But I doubt that the elimination of "concept" would make communication too cumbersome; after all, when required, cognitive scientists can always appeal to the description "bodies of knowledge used in higher cognition." It seems likely that using such a description will not invite cognitive scientists (or at least not to the same extent) to look for commonalities among all bodies of knowledge used in higher cognition. If this is correct, then the costs resulting from the elimination of "concept" are limited, and cognitive scientists should eliminate the notion of concept from their theoretical apparatus.

References

Malt, B. C., & Sloman, S. A. (2007). Artifact categorization: The good, the bad, and the ugly. In E. Margolis & S. Laurence (Eds.), *Creations of the mind: Theories of artifacts and their representation* (pp. 85-123). Ocford: Oxford University Press.

Allen, S. W., & Brooks, L. R. (1991). Specializing the operation of an explicit rule. *Journal of Experimental Psychology: General*, 120, 3-19.

Ashby, F. G., & Ell, S. W. (2002). Single versus multiple systems of learning and memory. In J. Wixted & H. Pashler (Eds.), *Stevens' handbook of experimental psychology: Volume 4 Methodology in experimental psychology* (3rd ed., pp. 655-692). New York: Wiley.

Ashby, F. G., & Maddox, W. T. Human category learning. *Annual Review of Psychology*, 56, 149-178.

Anderson, J. R. (1978). Arguments concerning representations for mental imagery. *Psychological Review*, 85, 249-277.

Anderson, J. R., & Betz, J. (2001). A hybrid model of categorization. *Psychonomic Bulletin and Review*, 8, 629-647.

Barsalou, L. W. (1989). Intraconcept similarity and its implications for interconcept similarity. In S. Vosniadou and A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 76-121). Cambridge: Cambridge University Press.

Barsalou, L. W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A.C. Collins, S.E. Gathercole, and M.A. Conway (Eds.), *Theories of memory* (pp. 29-101). London: Lawrence Erlbaum Associates.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577-660.

Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition, 10*, 82-93.

Barsalou, L. W. (1985). Ideals, central tendency, and frequency of Instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11*, 629-654.

Barsalou, L. W. (1987). The instability of graded structures: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 101-138). Cambridge: Cambridge University Press.

Barsalou, L. W. (2008a). Cognitive and neural contributions to understanding the conceptual system. *Current Directions in Psychological Science, 17.*

Barsalou, L. W. (2008b). Grounded cognition. *Annual Review of Psychology, 59*, 617-

645.
Barsalou, L.W. (2009). Simulation, situated conceptualization, and prediction. *Philosophical Transactions of the Royal Society of London: Biological Sciences,* 364, 1281-1289.
Barsalou, L. W., Simmons, W. K., Barbey, A. K., & Wilson, C. D. (2003). Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences, 7,* 84-91.
Boulenger, V., Hauk, O., & Pulvermüller, F. (2009). Grasping ideas with the motor system: Semantic somatotopy in idiom comprehension. *Cerebral Cortex, 19,* 1905-1914.
Bloom, P. (1996). Intention, history, and artifact concepts. *Cognition,* 60, 1-29.
Boyd, R. (1991). Realism, anti-foundationalism and the enthusiasm for natural kinds. *Philosophical Studies,* 61, 127-148.
Boyd, R. (1999). Kinds, complexity and multiple realization. *Philosophical Studies,* 95, 67-98.
Caramazza, A. (1986). On drawing inferences about the structure of normal cognitive systems from the analysis of patterns of impaired performance: The case for single-patient studies. *Brain & Cognition,* 5, 41-66.
Chao, L. L., & Martin, A. (1999). Cortical representation of perception, naming, and knowledge of color. *Journal of Cognitive Neuroscience*, 11. 25-35.
Churchland, P. M. (1981). Eliminative materialism and the propositional attitudes. *Journal of Philosophy,* 78, 67-90.
Dove, G. (2009). Beyond perceptual symbols: A call for representational pluralism. *Cognition, 110,* 412-431.
Dretske, F. I. (1981). *Knowledge and the flow of information*. Oxford: Basil Blackwell.
Dunn, J. C., & Kirsner, K. (1988). Discovering functionally independent mental processes: The principle of reversed association. *Psychological Review,* 95, 91-101.
Dunn, J. C., & Kirsner, K. (2003). What can we infer from double dissociations? *Cortex*, 39, 1-7.
Edwards, K. (2009). What concepts do. *Synthese*, 170, 289-310.
Evans, J. (2007). *Hypothetical thinking: dual processes in reasoning and judgement.* New York: Routledge.
Evans, J., & Frankish, K. (Eds.) (2009). *In two minds: Dual processes and beyond*. Oxford: Oxford University Press.
Feeney, A., & Heit, E. (Eds.) (2007). *Inductive reasoning: Experimental, Developmental, and computational approaches*. Cambridge: Cambridge University Press.
Fodor, J. A. (1975). *The language of thought*. New York: Crowell.
Fodor, J. A. (2003) Is it a bird? Problems with old and new approaches to the theory of concepts. *TLS*, 17 Jan. (2003), 3-4.

Fodor, J. A. (1998). *Concepts, Where cognitive science went wrong*. New York: Oxford University Press.
Fodor, J. A. (2008). *LOT 2: The language of thought revisited*. Oxford : Oxford University Press.
Gallese, V., & Lakoff, G. (2005). The brain's concepts: the role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology, 21*, 455-479.
Gelman, R. (2004). Cognitive development. In H. Pashler and D. L. Medin (Eds.), *Stevens' Handbook of Experimental Psychology, Volume 3: Memory and Cognitive Processes* (533-560). New York: Wiley.
Gelman, S. A. (1988). The development of induction within natural kinds and artifacts categories. *Cognitive Psychology*, 20, 65-95.
Gelman, S. A. (2003). *The essential child. Origins of essentialism in everyday thought*. New York: Oxford University Press.
Gelman, S. A., & Markman, E. (1986) Categories and induction in young children. *Cognition*, 23, 183-209.
Gigerenzer, G., & Regier, T. P. (1996). How do we tell an association from a rule? *Psychological Bulletin*, 119, 23-26.
Gigerenzer, G., Todd, P. M., & the ABC Research Group (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
Glenberg, A. M. (1997). What memory is for? *Behavioral and Brain Sciences, 20*, 1-55.
Glymour, C. (1994). On the methods of cognitive neuropsychology. *British Journal for the Philosophy of Science*, 45, 815-835.
Goldman, A. I. (2006). *Simulating minds: The philosophy, psychology, and neuroscience of mindreading*. New York: Oxford University Press.
Goldstone, R. L., & Kersten, A.W. (2003). Concepts and categorization. In A. F. Healy & R.W. Proctor (Eds.), *Experimental psychology* (pp. 599-621). Volume 4 in I.B. Weiner (Ed.), *Handbook of psychology*. New York: Wiley.
Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts, and theories*. Cambridge, MA: MIT Press.
Gopnik, A., Glymour, C., Sobel, D., Schulz, L., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, 111, 1-31.
Gopnik, A. (2003). The theory theory as an alternative to the innateness hypothesis. In L. Antony & N. Hornstein (Eds.), *Chomsky and his critics*. Oxford: Blackwell.
Greene, J. D., & Haidt, J. (2002). How (and where) does moral judgment work? *Trends in Cognitive Sciences*, 6, 517-523.
Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral Judgment. *Science, 293*, 2105-

2108.
Greenwald, A. G., & Nosek, B. A. (2008). Attitudinal dissociation: What does it mean? In R. E. Petty, R. H. Fazio, & P. Brinol (Eds.), *Attitudes: Insights from the new implicit measures (pp. 65-82)*. Hillsdale, NJ: Erlbaum.
Griffiths, P. E. (1997). *What emotions really are*. Chicago: Chicago University Press.
Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, 114, 211-244.
Haidt, J. (2001) The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108, 814-834.
Hampton, J. A. (1979). Polymorphous concepts in semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 18, 441-461.
Hampton, J. A. (1981). An Investigation of the nature of abstract concepts. *Memory & Cognition*, 9, 149-156.
Hampton, J. A. (1993). Prototype models of concept representation. In I. Van Mechelen, J. A., Hampton, R. S. Michalski, & P. Theuns (Eds.), *Categories and concepts: Theoretical views and inductive data analysis* (pp. 67-95). London: Academic Press.
Hampton, J. A. (2006). Concepts as prototypes. In B. H. Ross (Ed), *The psychology of learning and motivation: Advances in research and theory* (Vol. 46, pp. 79-113). New York: Academic Press.
Hampton, J. A. (2007). Typicality, graded membership, and vagueness. *Cognitive Science*, 31, 355-384.
Hauk, O., Johnsrude, I., & Pulvermüller, F. (2004). Somatotopic representation of action words in human motor and premotor cortex. *Neuron*, *41*, 301-307.
Heit, E. (2000). Properties of inductive reasoning. *Psychonomic Bulletin and Review*, 7, 569-592.
Hoenig, K., Sim, E.-J., Bochev, V., Herrnberger, B., & Kiefer, M. (2008). Conceptual flexibility in the human nrain: Dynamic recruitment of semantic maps from visual, motor, and motion-related reas. *Journal of Cognitive Neuroscience*, *20*(10), 1799-1814.
Horn, B. K. P. (1975). Obtaining shape from shading information. In P. H. Winston (Ed.), *The psychology of computer vision* (pp. 115-155). New York: McGraw-Hill.
James, T. W., & Gauthier, I. (2003). Auditory and action semantic features activate sensory-specific perceptual brain regions. *Current Biology*, 13, 1792-1796.
Kable, J. W., Kan, I. P., Wilson, A., Thompson-Schill, S. L., & Chatterjee, A. (2005). Conceptual representations of action in the lateral temporal cortex. *Journal of Cognitive Neuroscience, 17*, 1855-1870.
Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds), *Heuristics & biases: The psychology of intuitive judgment* (pp. 49-81). New York.

Cambridge University Press.
Kahneman, D., & Frederick, S. (2005). A model of heuristic judgment. In K.J. Holyoak & R.G. Morrison (Eds.), *The Cambridge handbook of thinking and reasoning* (pp. 267-293). Cambridge: Cambridge University Press.
Keil, F. C. (1987). Category structure and patterns of conceptual change. In U. Neisser (Ed.), *Concepts and conceptual development: The ecological and intellectual factors in categorization* (pp. 175-200). Cambridge: Cambridge University Press.
Keil, F. C. (1989). *Concepts, kinds, and cognitive development*, Cambridge, MA: MIT Press.
Keil, F. C., Carter Smith, W., Simons, D. J., & Levin, D. T. (1998). Two dogmas of conceptual empiricism: implications for hybrid models of the structure of knowledge. *Cognition*, 65, 103-135.
Komatsu, L. K. (1992). Recent views of conceptual structure. *Psychological Bulletin*, 112, 500-526.
Kiefer, M., Sim, E-J, Liebich, S., Hauk, O., & Tanaka, J. (2007). Experience-dependent plasticity of conceptual representations in human sensory–motor areas. *Journal of Cognitive Neuroscience*, 19, 525-542.
Laurence, S., & Margolis, E. (1999). Concepts and cognitive science. In E. Margolis and S. Laurence (Eds.), *Concepts, core readings* (pp. 3-82). Cambridge, MA: MIT Press.
López, A., Atran, S., Coley, J. D., Medin, D. L., & Smith E. E. (1997). The tree of life: Universal and cultural features of folkbiological taxonomies and inductions. *Cognitive Psychology*, 32, 251-295.
Machery, E. (2005). Concepts are not a natural kind. *Philosophy of Science*, 72, 444-467.
Machery, E. (2006a). Two dogmas of neo-empiricism. *Philosophy Compass*, 1, 398-412.
Machery, E. (2006b). How to split concepts. Reply to Piccinini and Scott. *Philosophy of Science, 73*, 410-418.
Machery, E. (2007). Concept empiricism: A methodological critique. *Cognition*, 104, 19-46.
Machery, E. (2009). *Doing without concepts*. New York: Oxford University Press.
Machery, E., and Seppälä, S. (Forthcoming). Against hybrid theories of concepts, *Dialectica*.
Mahon, B. Z., & Caramazza, A. (2009). Concepts and categories: A cognitive neuropsychological perspective. *Annual Review of Psychology, 60*, 27-51.
Mallon, R. (2006). "Race": Normative, not metaphysical or semantic. *Ethics*, 116, 525-551.
Mallon, R., Machery, E., Nichols, S., and Stich, S. P. (2009). Against arguments from reference. *Philosophy and Phenomenological Research*.
Malt, B. C. (1994) Water is not H_2O. *Cognitive Psychology*, 27, 41-70.

Margolis, E. (1994). A reassessment of the shift from the classical theory of concepts to prototype theory. *Cognition*, 51, 73-89.
Margolis, E. (1995). The significance of the theory analogy in the psychological study of concepts. *Mind and Language*, 10, 45-71.
Margolis, E., & Laurence, S. (2006). Concepts. *Stanford Encyclopedia of Philosophy*. Available at http://plato.stanford.edu/entries/concepts/.
Martin, A. (2007). The representation of object concepts in the brain. *Annual Review of Psychology, 58*, 25-45.
Weiskopf, D. A. (2009). The plurality of concepts. *Synthese, 169*, 145-173.
Martin, A., & Chao, L. L. (2001). Semantic memory and the brain: Structure and processes. *Current Opinion in Neurobiology*, 11, 194-201.
Martin, A., Haxby, J. V., Lalonde, F. M., Wiggs, C. L., & Ungerleider, L. G. (1995). Discrete cortical regions associated with knowledge of color and knowledge of action. *Science*, 270, 102-105.
Medin, D. L. (1989). Concepts and conceptual structure. *American Psychologist*, 44, 1469-1481.
Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207-238.
Medin, D. L., Lynch, E. B., and Solomon, K. O. (2000). Are there kinds of concepts? *Annual Review of Psychology*, 51, 121-147.
Medin, D. L., Coley, J. D., Storms, G., & Hayes, B. (2003). A Relevance Theory of Induction. *Psychonomic Bulletin and Review*, 3 , 517-532.
Millikan, R. G. (2000). *On clear and confused ideas: An essay about substance concepts*. New York: Cambridge University Press.
Minda, J. P., & Smith, J. D. (2001). Prototypes in category learning: The effects of category size, category structure, and stimulus complexity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 775-799.
Murphy, D., & Stich, S. P. (1999). Griffiths, elimination and psychopathology. *Metascience*, 8, 13-25.
Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.
Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 104-114.
Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 14, 700-708.

Nosofsky, R. M. (1992). Exemplar-based approach to relating categorization, identification, and recognition. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 363-393). Hillsdale, NJ: Lawrence Erlbaum Associates.
Nosofsky, R. M., & Stanton, R. D. (2005). Speeded classification in a probabilistic category structure: Contrasting exemplar-retrieval, decision-boundary, and prototype models. *Journal of Experimental Psychology: Human Perception and Performance, 31,* 608-629.
Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 266-300.
Osherson, D. N., & Smith, E. E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*, 9, 35-58.
Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological review*, 97, 185-200.
Peacocke, C. (1992). *A study of concepts*. Cambridge, MA: MIT Press.
Peacocke, C. (2008). *Truly understood*. New York: Oxford University Press.
Pecher, D., Zeelenberg, R., & Barsalou, L.W. (2004). Sensorimotor simulations underlie conceptual representations: Modality-specific effects of prior activation. *Psychonomic Bulletin & Review*, 11, 164-167.
Piccinini, G., & Scott, S. (2006). Splitting concepts. *Philosophy of Science*, 73, 390-409.
Plaut, D.C. (1995). Double dissociation without modularity: Evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropsychology*, 17, 291-321.
Prinz, J. J. (2002). *Furnishing the Mind: Concepts and their Perceptual Basis.* Cambridge, MA: MIT Press.
Prinz, J. J. (2005). The return of concept empiricism. In H. Cohen & C. Lefebvre (Eds.), *Handbook of categorization in cognitive science* (pp. 679-695). Oxford: Elsevier.
Proffitt, J. B., Coley, J. D., & Medin, D. L. (2000). Expertise and category-based induction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 811-828.
Pulvermüller, F., & Hauk, O. (2006). Category-specific conceptual processing of color and form in left fronto-temporal cortex. *Cerebral Cortex, 16,* 1193-1201
Pulvermüller, F. (2005). Brain mechanisms linking language and action. *Nature Reviews Neuroscience*, 6, 576-582
Quine, W. V. O. (1960). *Word and object*. Cambridge, MA: MIT Press.
Quine, W. V. O. (1969). Natural kinds. In W. V. O. Quine, *Ontological relativity and other essays* (pp. 114-138). New York: Columbia University Press.
Rehder, B. (2003). A causal-model theory of conceptual representation and categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 1141-1159.

Rehder, B. (2006). When causality and similarity compete in category-based property induction. *Memory & Cognition*, 34, 3-16.
Rehder, B., & Kim, S. (2006). How causal knowledge affects classification: A generative theory of categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 659-683.
Rey, G. (1983). Concepts and stereotypes. *Cognition*, 15, 237-262.
Rey, G. (1985). Concepts and conceptions: A reply to Smith, Medin and Rips. *Cognition*, 19, 297-303.
Rey, G. (2009). Review of E. Machery, *Doing without Concepts*. *Notre Dame Philosophical Reviews*.
Rips, L. J. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behavior*, 1 4, 665-681.
Rips, L. J. (1989). Similarity, typicality, and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 21-59). Cambridge: Cambridge University Press.
Rips, L. J., Shoben, E. J., & Smith, E. E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, 12, 1-20.
Rosch, E., & Mervis, C. B. (1975). Family resemblance: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573-605.
Roth, E. M., & Shoben, E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology*, 15, 346-378.
Russell, B. (1948). *Human knowledge: Its scope and its limits*. London: Routledge.
Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge: Cambridge University Press.
Simmons, W. K., Ramjee, V., Beauchamp, M. S., McRae, K., Martin, A., & Barsalou, L. W. (2007). A common neural substrate for perceiving and knowing about color. *Neuropsychologia, 45*, 2802-2810.
Simon, H. A. (1995). Machine as mind. In K. M. Ford, C. Glymour, & P. J. Hayes (Eds.), *Android epistemology* (pp. 23-40). Cambridge, MA: MIT Press.
Sloman, S. A. (1993). Feature-based induction. *Cognitive Psychology*, 25, 231-280.
Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3-22.
Sloman, S. A. (1998). Categorical inference is not a tree: The myth of inheritance hierarchies. *Cognitive Psychology*, 35, 1-33.
Sloman, S. A., & Lagnado, D. (2005). The problem of induction. In K. Holyoak & R. Morrison (Eds.). *The Cambridge handbook of thinking & reasoning* (pp. 95-116). New York: Cambridge University Press.
Smith, E. E. (1989). Concepts and induction. In M. Posner (Ed.), *Foundations of*

cognitive science (pp. 510-526). Cambridge, MA: MIT Press.
Smith, E. E., & Medin, D. L. (1981) *Categories and concepts*. Cambridge, MA: Harvard University Press.
Smith, E. E., Osherson, D. N., Rips, L. J., & Keane, M. (1988). Combining concepts: A selective modification model. *Cognitive Science*, 12, 485-527.
Smith, E. E., & Sloman, S. A. (1994). Similarity- versus rule-based categorization. *Memory & Cognition*, 22, 377-386.
Smith, E. E., Patalano, A. L., & Jonides, J. (1998). Alternative strategies of categorization. *Cognition*, 65, 167-196.
Smith, J. D. (2002). Exemplar theory's predicted typicality gradient can be tested and disconfirmed. *Psychological Science*, 13, 437-442.
Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 1411-1436.
Smith, J. D., & Minda, J. P. (2000). Thirty categorization results in search of a model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 3-27.
Smith, L. B., & Samuelson, L. K. (1997). Perceiving and remembering: Category stability, variability and development. In K. Lamberts and D. Shanks (Eds.), *Knowledge, concepts, and categories* (161-196). Cambridge, MA: MIT Press.
Solomon, K. O., & Barsalou, L. W. (2001). Representing properties locally. *Cognitive Psychology*, 43, 129-169.
Solomon, K. O., & Barsalou, L. W. (2004). Perceptual simulation in property verification. *Memory & Cognition*, 32, 244-259.
Solomon, K. O., Medin, D. L., & Lynch, E. L. (1999). Concepts do more than categorize. *Trends in Cognitive Sciences*, 3, 99-105.
Stanovich, K. E. (1999). *Who is rational? Studies of individual differences in reasoning*. Mahweh, NJ: Erlbaum.
Stein, L. A. (1995). Imagination and situated cognition. In K. M. Ford, C. Glymour, & P. Hayes (Eds.), *Android epistemology* (pp. 167-182), Cambridge, MA: MIT Press.
Stich, S. P. (1983). *From folk psychology to cognitive science*. Cambridge, MA: MIT Press.
Stich, S. P. (1996). *Deconstructing the mind*. Oxford: Oxford University Press.
Tenenbaum, J.B., Griffiths, T. L., & Niyogi, S. (2007). Intuitive theories as grammars for causal inference. In A. Gopnik and L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 301-322). New York: Oxford University Press.
Thompson-Schill, S. L. (2003). Neuroimaging studies of semantic memory: Inferring "how" from "where". *Neuropsychologia*, 41, 280-292.
Van orden, G. C., Pennington, B. F., & Stone, G. O. (2001). What do double

dissociations prove? *Cognitive Science*, 25, 111-172.
Weisberg, J., van Turrennout, M., & Martin, A. (2006). A neural system for learning about object function. *Cerebral Cortex*, 17, 513-521.
Whitney, P., McKay, T., Kellas, G., & Emerson, W. A., Jr. (1985). Semantic activation of noun concepts in context. *Journal of Experimental Psychology: Learning Memory and Cognition*, 20, 804-823.
Yaxley, R.H., & Zwaan, R.A. (2007). Simulating visibility during language comprehension. *Cognition*, 105, 229-236.
Ziff, E. (1972). *Understanding understanding.* Ithaca, NY: Cornell University Press.