How to Split Concepts: A Reply to Piccinini and Scott*

Edouard Machery†‡

In "Concepts Are Not a Natural Kind" (2005), I argued that the notion of concept in psychology and in neuropsychology fails to pick out a natural kind. Piccinini and Scott (2006, in this issue) have criticized the argument I used to support this conclusion. They also proposed two alternative arguments for a similar conclusion. In this reply, I rebut Piccinini and Scott's main objection against the argument proposed in "Concepts Are Not a Natural Kind." Moreover, I show that the two alternative arguments developed by Piccinini and Scott are not promising for supporting the conclusion that concepts are not a natural kind.

1. Introduction. In psychology and in neuropsychology, concepts are assumed to be those bodies of knowledge that are stored in long-term memory and that are used by default when we categorize, induce, produce complex concepts, or draw analogies. In "Concepts Are Not a Natural Kind" (2005), I argued that the notion of concept in psychology and in neuropsychology fails to pick out a natural kind (see also Machery forthcoming). Rather, I proposed that each class of objects, such as the class of dogs, is represented by several coreferential concepts—namely, by a prototype, a set of exemplars, and a theory (mutatis mutandis, for events, substances, etc.). For instance, I proposed that we have a prototype of dog, a set of exemplars of dogs, and a theory about dogs. Prototypes, sets of exemplars, and theories have little in common with each other. Whence, concepts are not a natural kind. Or to use Piccinini and Scott's terminology, concepts split. I called this hypothesis "the Heterogeneity

*Received July 2006.

†To contact the author, please write to: Edouard Machery, Department of History and Philosophy of Science, University of Pittsburgh, 1017CL, Pittsburgh, PA 15260; e-mail: machery@pitt.edu.

‡I would like to thank Gualtiero Piccinini, Dan Weiskopf, and Wayne Wu for their comments on previous versions of this reply.

Philosophy of Science, 73 (October 2006) pp. 410–418. 0031-8248/2006/7304-0003$10.00
Copyright 2006 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/516812 Published online by Cambridge University Press

Hypothesis." Additionally, I briefly reviewed some evidence for this hypothesis, drawn from the empirical research on concept combination.

In their stimulating article "Splitting Concepts," Gualtiero Piccinini and Sam Scott have criticized the Heterogeneity Hypothesis. Before turning to their criticisms, it is worth emphasizing that we agree on several essential issues. Piccinini and Scott propose that most psychologists interested in concepts have erroneously looked for a uniform account of concepts. Psychologists have of course recognized differences between subclasses of concepts, for instance, between the concepts of animals and the concepts of artifacts. But they have typically proposed that in spite of these differences, all or most concepts are prototypes, sets of exemplars, or theories. Against this received view, Piccinini and Scott contend that the class of concepts is likely to include several subclasses that are very different from each other. Thus, no unified account of concepts will be found. In these respects, Piccinini and Scott's position is similar to the Heterogeneity Hypothesis.[1]

However, unlike Machery (2005), Piccinini and Scott fall short of arguing that concepts do not form a natural kind (2006, in this issue). They do not want to rule out that psychologists might find out numerous nontrivial properties that would be true of all or most concepts, even though the class of concepts includes very different subclasses. In what follows, for the sake of space, I will pass over this disagreement. Rather, I will successively focus on what I take to be the two other key points of contention. First, Piccinini and Scott disagree with me about how concepts should be split. They argue that the strategy I proposed in "Concepts Are Not a Natural Kind" does not justify splitting concepts. Second, Piccinini and Scott develop two alternative strategies for splitting concepts. They take these two strategies to be more promising than the Heterogeneity Hypothesis.

2. Rebutting Piccinini and Scott's Objection. In "Concepts Are Not a Natural Kind," I outlined three strategies for splitting concepts into kinds that have little in common (2005, 450–451; see also Machery forthcoming). According to the first strategy, concepts vary across cognitive tasks.[2] For example, the process(es) involved in induction might use a kind of concept that has little in common with the kind of concept used in the process(es) involved in categorization. Let us call this strategy "Task Pluralism." According to the second strategy, concepts vary depending on what they represent. Concepts of substances might have little in common with con-

1. For yet another perspective on this debate, see Weiskopf (2006), from whom the term "pluralism," used below, is drawn.

2. I am using Piccinini and Scott's terminology here. In Machery (forthcoming), I use "cognitive competence" to refer to what Piccinini and Scott call "task."

https://doi.org/10.1086/516812 Published online by Cambridge University Press

cepts of classes of three-dimensional, medium-sized objects; or concepts of animals might have little in common with concepts of artifacts. Let us call this second strategy "Scope Pluralism." Finally, each class (substance, etc.) might be represented by several coreferential concepts, each of which belongs to a very different kind of representation. This is the Heterogeneity Hypothesis. I took this third strategy to be the most promising.

Piccinini and Scott disagree. They contend that even if successful, this third strategy would not justify splitting concepts. Suppose, as I have suggested, that many classes (substances, etc.) are represented by a prototype, a set of exemplars, and a theory. According to Piccinini and Scott, this would not show that each class is represented by three distinct concepts that belong to three heterogeneous kinds, because this finding would be consistent with each class (substance, etc.) being represented by *a single concept with three components*.

At this point, the natural question is: What would it take to show that a prototype of dog, a set of exemplars of dogs, and a theory about dogs are three concepts rather than three components of a single concept? Piccinini and Scott propose that being involved in different tasks is strong evidence that prototypes, exemplars, and theories are distinct concepts. As evidence for the Heterogeneity Hypothesis, I described a model of concept combination in which the hypothesized coreferential concepts are jointly the inputs into a single cognitive process for producing complex concepts (Machery 2005, 457–462). Commenting on this model, Piccinini and Scott draw the following conclusion (2006, in this issue): "This argument falls short, because the evidence Machery presents is that several different kinds of mental representation are *jointly* involved in the same cognitive tasks. Such evidence is consistent with the view that the mental representations postulated by Machery are *components* of the same natural kind." Thus, because prototypes, exemplars, and theories are involved in the same tasks, Piccinini and Scott conclude that it is unclear whether these three bodies of knowledge constitute different concepts rather than different components of concepts.

There are two possible lines of reply to Piccinini and Scott's objection. As I suggested in Machery (2005, 463–464), I doubt that there is any substantial difference, rather than a merely terminological difference, between characterizing different bodies of knowledge, such as a prototype, a set of exemplars, and a theory, as being the components of a single concept and characterizing them as being different coreferential concepts. Suppose that one treats prototypes, exemplars, and theories as components of concepts. Instead of saying that a kind of concept, say prototypes, has such and such distinct properties, one would say that a component of concepts, that is, prototypes, has such and such distinct properties. Instead of saying that few properties are shared by the subclasses of

https://doi.org/10.1086/516812 Published online by Cambridge University Press

concepts—namely, prototypes, exemplars, and theories—one would say that few properties are shared by the components of concepts—namely, prototypes, exemplars, and theories. And so on.

In reply, Piccinini and Scott (2006, in this issue) claim that if prototypes, exemplars, and theories were components of concepts (rather than being concepts), one could formulate numerous generalizations about concepts. Unfortunately, Piccinini and Scott did not give any examples of these generalizations, and I do not see why calling prototypes, exemplars, and theories "components of concepts," instead of "concepts," would allow psychologists to formulate new scientific generalizations.

Rather than pushing this line further, I want to discuss Piccinini and Scott's principle that being involved in different tasks is strong evidence that prototypes, exemplars, and theories are distinct representations. I propose an alternative principle: different bodies of knowledge are distinct representations if these bodies of knowledge are involved in distinct processes.[3]

A brief explanation of the notions of task and process is in order. A task is characterized by its goal. Categorization—that is, deciding whether an object belongs to a given class—and categorical induction—deciding whether a property that is possessed by the members of a class is also possessed by the members of another class—are two tasks. A process is a specific way of completing a task. Importantly, there might be several distinct processes for a given task. For instance, there might be several distinct categorization processes. Thus, different bodies of knowledge, such as prototypes, exemplars, and theories, might be involved in the same tasks—for instance, categorization and induction—but be used in different categorization processes (e.g., a prototype-based process, an exempla-based process, and a theory-based process) and in different induction processes.[4]

The principle—different bodies of knowledge are distinct representations if they are involved in distinct processes—is to be preferred to Piccinini and Scott's principle because it is implicit in many psychologists' procedures for distinguishing representations from components of rep-

3. Additionally, these processes should not be organized in such a way that the outputs of one of them override the outputs of the other processes when these outputs conflict. I do not stress this point here (for further discussion, see Machery [forthcoming]).

4. An intricate issue is lurking here. The two competing principles for distinguishing components of representations from representations suppose that there is a principled way of individuating tasks and processes. However, it is unclear whether there is one. Lacking a principled way of individuating tasks and processes, I propose that one should rely on psychologists' intuitive understanding of tasks and processes. For instance, psychologists take it for granted that categorization is a task, that induction is a different task, and that an exemplar-based categorization process is a different process than a prototype-based categorization process.

https://doi.org/10.1086/516812 Published online by Cambridge University Press

resentations. For instance, according to the theory of implicit knowledge developed by the neuropsychologists Larry Squire and Barbara Knowlton, people form two distinct representations of classes of objects—an implicit representation and an explicit representation (e.g., Knowlton 1999). For Squire and Knowlton, whether implicit and explicit representations form two different kinds of representation, rather than two components of representations, does not hang on whether some tasks involve one of these representations, but not the other. Rather, implicit and explicit representations are assumed to be two different kinds of representation, because they are used in distinct processes. For instance, implicit and explicit representations are both assumed to be used when we categorize. But the categorization process that uses implicit representations is supposed to be distinct from the categorization process that uses explicit representations. The same is true of the hypothesized learning processes by which we are supposed to acquire these two kinds of representation.

The Heterogeneity Hypothesis proposes that the same is true of prototypes, exemplars, and theories. If the Heterogeneity Hypothesis is correct, prototypes, exemplars, and theories are used when we categorize, draw inductions, or produce complex concepts. Thus, prototypes, exemplars, and theories are not distinguished by reference to the tasks in which they are used. However, I propose that for many tasks, they are involved in distinct processes. Evidence tentatively suggests that prototypes, exemplars, and theories are used in distinct induction processes. More clearly, evidence shows that prototypes, exemplars, and theories are used in distinct categorization processes. That is, rather than having a single categorization process that would take several coreferential concepts as inputs, we have several categorization processes, each of which is defined over a distinct kind of concept (for some evidence see Section 3 and, at greater length, Machery [forthcoming]).

It is true that in the model of concept combination presented in Machery (2005), prototypes, exemplars, and theories were assumed to be the inputs into a single cognitive process. However, they are likely to be used by different subprocesses of the process sustaining concept combination, for prototypes, exemplars, and theories are assumed to fulfill different functions during concept combination. For instance, prototypes are supposed to be used to determine which properties are likely to be typical of the classes denoted by the complex concepts, while theories are supposed to be used to determine which properties cannot be possessed or are necessarily possessed by these classes.

To conclude, because the Heterogeneity Hypothesis proposes that prototypes, exemplars, and theories are used in distinct processes, this hypothesis is an appropriate strategy for splitting concepts. If evidence sup-

https://doi.org/10.1086/516812 Published online by Cambridge University Press

ports this hypothesis, prototypes, exemplars, and theories are distinct concepts rather than distinct components of concepts.

3. Two Alternative Strategies of Concept Splitting? Piccinini and Scott suggest that Task Pluralism and Scope Pluralism are the most promising strategies for splitting concepts.

I focus first on Piccinini and Scott's version of Task Pluralism. They distinguish between two types of cognitive task, namely, those that involve language, such as lexical combination (i.e., understanding linguistic compounds), and those that do not, such as perceptual discrimination.[5] They propose that two very different kinds of concept are probably used in these two types of task—"linguistic" concepts and "nonlinguistic" concepts (2006, in this issue).

I am not convinced. There is much evidence that linguistic concepts and nonlinguistic concepts have very similar properties. Consider, for instance, typicality (for a review, see Murphy [2002]; discussed in Piccinini and Scott [2006, in this issue]). It is well known that typical objects are categorized more quickly and more accurately than atypical objects. The class membership of typical objects is also learned more quickly than the class membership of atypical objects. The important point for present purposes is that these properties are common to linguistic concepts and to nonlinguistic concepts (e.g., Rosch and Mervis 1975). We decide more quickly that a robin is a bird than that a penguin is a bird. Similarly, when subjects learn to classify meaningless, abstract, and nonlexicalized figures into different categories and are then asked to classify new figures into these categories,[6] typical figures are classified more quickly and more accurately than atypical figures.

Discussing a similar objection, Piccinini and Scott reply that the fact that linguistic and nonlinguistic concepts have some properties in common does not show that they do not differ in other respects. They might also stress that by claiming that lexicalized concepts and nonlexicalized concepts form two different kinds of concept, they did not mean to deny that they share some properties.

This reply is problematic, however. There are obviously some differences between linguistic concepts and nonlinguistic concepts. In the first place, the former, but not the latter, are expressed by words and are retrieved from long-term memory when the words that express them are used. But these differences do not support Piccinini and Scott's proposal, because Piccinini and Scott do not merely contend that there are differences be-

5. One could question whether perceptual discrimination really involves concepts.

6. This is a typical design in the experimental psychology of concepts (Machery 2007).

https://doi.org/10.1086/516812 Published online by Cambridge University Press

tween the class of linguistic concepts and the class of nonlinguistic concepts—a moderate thesis that many psychologists have already put forward (e.g., Spelke 2003). More radically, they contend that the class of linguistic concepts and the class of nonlinguistic concepts form two very different subclasses of concepts. The fact that some phenomena as basic as the typicality phenomena are found both with linguistic concepts and with nonlinguistic concepts is evidence against this proposal. If Piccinini and Scott were correct, we would not expect typicality to be a significant variable for these two classes of concepts. Rather, the fact that the typicality phenomena are found both with linguistic concepts and with nonlinguistic concepts strongly suggests that we have some linguistic and some nonlinguistic prototypes (or some linguistic and nonlinguistic sets of exemplars). Thus, the distinction between linguistic concepts and nonlinguistic concepts does not distinguish subclasses of concepts that have little in common.

Let us turn now to Piccinini and Scott's version of Scope Pluralism. They propose that some entities (e.g., some abstract entities) are represented by "non-similarity-based concepts," while other entities (e.g., most classes of physical objects) are represented by "similarity-based concepts" (2006, in this issue). It might well be that some entities—such as some abstract entities—are represented by a distinct kind of concept. But I doubt that it is generally the case that some entities are represented by "non-similarity-based concepts," while other entities are represented by "similarity-based concepts."

The reason is that many entities are simultaneously represented by one or several "similarity-based concepts," for example, a prototype and a set of exemplars, and by a "non-similarity-based concept," either a theory or a definition. For instance, the word "uncle" expresses at least two concepts, a definition and a prototype. A male adult might literally be the uncle of a child, even though he is not the brother of one of the parents of this child. For, this male adult might satisfy the prototype of uncle (see Pinker and Prince [1999] for other examples).

Additional evidence comes from the experiments done by Allen and Brooks (1991; Experiment 1), who asked subjects to learn the concepts of two categories, A and B. Each category was defined by a disjunctive rule: a stimulus belonged to the category if it possessed a sufficient number of properties. Half of the subjects were told the rule (rule condition), and half were not (no-rule condition). During the learning phase, subjects were presented with members of A and members of B. They were asked to categorize these items, and they were given some feedback when their categorization was mistaken. Subjects were said to have learned the concepts of A and B when they were able to categorize the training stimuli. During the test phase, subjects were asked to categorize some stimuli seen

https://doi.org/10.1086/516812 Published online by Cambridge University Press

during training as well as some new stimuli. Some new stimuli (called "negative matches") were such that they satisfied the rule of a given category, say A, but were very similar to a member of B seen during training. Of interest was how subjects would categorize these negative matches.

In the rule condition, negative matches were more likely to be misclassified as members of B than the new stimuli that satisfied the rule for belonging to A and that were very similar to a member of A seen during training (called "positive matches"). This is evidence that exemplars of members of A and the rule for A were *both* used during categorization. Thus, in the rule condition, subjects formed two coreferential concepts—a set of exemplars (i.e., a "similarity-based" concept in Piccinini and Scott's [2006, in this issue] terminology) and a disjunctive rule (i.e., a "non-similarity-based" concept in Piccinini and Scott's terminology). These two concepts were used during categorization by those subjects in the rule condition. This is evidence against Piccinini and Scott's second argument for splitting concepts.

Moreover, in the rule condition, subjects were significantly slower when they had to classify the negative matches compared to the positive matches. A plausible interpretation of this finding is that when subjects had to categorize the negative matches, they had to resolve a conflict between two different categorization judgments (Allen and Brooks 1991, 7). This is evidence that, in Allen and Brooks's experiments, the exemplars and the rule were not merely the inputs into a single cognitive process but, rather, the inputs into two different processes yielding conflicting categorization judgments.

Noticeably, Allen and Brooks's behavioral findings were replicated by Smith, Patalano, and Jonides (1998), who also found that different brain areas are involved when subjects' categorization is driven by the similarity between a negative match and an exemplar and when it is driven by the negative match satisfying the rule.

4. Conclusion. Piccinini and Scott (2006, in this issue) have raised some important issues about how the class of concepts should be split. However, they have failed to mount a decisive challenge against the Heterogeneity Hypothesis. Piccinini and Scott have not shown that the strategy endorsed in "Concepts Are Not a Natural Kind" (2005) is inadequate for splitting concepts. For different bodies of knowledge are distinct representations if these bodies of knowledge are involved in distinct processes. And, although prototypes, exemplars, and theories are not distinguished by the tasks in which they are used, evidence suggests that they are the inputs into distinct categorization processes and distinct induction processes. Moreover, the two strategies developed by Piccinini and Scott are not

https://doi.org/10.1086/516812 Published online by Cambridge University Press

promising for splitting the class of concepts into classes that have little in common. Language does not draw a line between two kinds of concept that have little in common. Rather, evidence suggests that people have lexicalized and nonlexicalized prototypes (or sets of exemplars). It is also dubious whether some entities are represented by "non-similarity-based concepts," while other entities are represented by "similarity-based concepts." Rather, evidence suggests that many entities are represented by one or two similarity-based concepts, such as a prototype and a set of exemplars, and a non-similarity-based concept, such as a theory, as the Heterogeneity Hypothesis would have it.

REFERENCES

Allen, Scott W., and Lee R. Brooks (1991), "Specializing the Operation of an Explicit Rule," *Journal of Experimental Psychology: General* 120: 3–19.

Knowlton, Barbara J. (1999), "What Can Neuropsychology Tell Us about Category Learning," *Trends in Cognitive Sciences* 3 (4): 123–124.

Machery, Edouard (2005), "Concepts Are Not a Natural Kind," *Philosophy of Science* 72 (3): 444–467.

——— (2007), "100 Years of Psychology of Concepts: The Theoretical Notion of Concept and Its Operationalization," *Studies in History and Philosophy of Biological and Biomedical Science* 38: 63–84.

——— (forthcoming), *Doing without Concepts*. New York: Oxford University Press.

Murphy, Gregory L. (2002), *The Big Book of Concepts*. Cambridge, MA: MIT Press.

Piccinini, Gualtiero, and Sam Scott (2006), "Splitting Concepts," *Philosophy of Science* 73 (4), in this issue.

Pinker, Steven, and Alan Prince (1999), "The Nature of Human Concepts: Evidence from an Unusual Source," in Ray Jackendoff, Paul Bloom, and Karen Wynn (eds.), *Language, Logic, and Concepts*. Cambridge, MA: MIT Press, 221–261.

Rosch, Eleanor, and Carolyn B. Mervis (1975), "Family Resemblance: Studies in the Internal Structure of Categories," *Cognitive Psychology* 7: 573–605.

Smith, Edward E., Andrea L. Patalano, and John Jonides (1998), "Alternative Strategies of Categorization," *Cognition* 65: 167–196.

Spelke, Elizabeth S. (2003), "What Makes Us Smart? Core Knowledge and Natural Language," in Dedre Gentner and Susan Goldin-Meadow (eds.), *Language in Mind, Advances in the Study of Language and Thought*. Cambridge, MA: MIT Press, 277–311.

Weiskopf, Daniel A. (2006), "Concept Pluralism," unpublished paper.

https://doi.org/10.1086/516812 Published online by Cambridge University Press