

Two Kinds of Concept: Implicit and Explicit

GUALTIERO PICCININI *University of Missouri*

Preliminaries

In his refreshing and thought-provoking *Doing without Concepts* (2009), Edouard Machery argues that people possess different kinds of concept. This is probably true and important. Before I get to that, I will briefly disagree on two other points.

Machery argues that “philosophical theories of concepts” and “psychological theories of concepts” are about different things (*DwC*, p. 31).¹ To begin with, the expression “philosophical theory of concept” is somewhat obscure. Machery seems to use it as a synonym for “theory of concepts developed by a philosopher” (*DwC*, pp. 33, 34). Now, it may be true that *some* theories of concepts proposed by philosophers are about something different from what the theories proposed by psychologists are about. But *other* theories of concepts proposed by professional philosophers – including Machery! – are explicitly after the same thing (e.g., Fodor, which Machery cites). Roughly speaking, they are after the building blocks of thought. Slightly more precisely, concepts – the target of theories proposed by both psychologists and many philosophers – are representations posited to explain certain cognitive phenomena including recognition, naming, inference, and language understanding (cf. Piccinini and Scott, 396).² To be sure, philosophers tend to emphasize aspects of concepts, such as how concepts get to represent categories, that psychologists tend to de-emphasize, and vice versa. It doesn’t follow that *all*

Dialogue 50 (2011), 179–193.

© Canadian Philosophical Association/Association canadienne de philosophie 2011

doi:10.1017/S0012217311000187

members of the two classes of theories are about different things. So it is a mistake to conclude that in general, “philosophical” and “psychological” theories of concepts are about different things.

Here is another route to the same conclusion. In order to determine that “philosophical” theories of concepts and “psychological” theories of concepts are about different things – as Machery wishes to do – we need to classify theories of concepts as either philosophical or psychological. But Machery does not offer any principled way to demarcate philosophical from psychological theories.³ In fact, much work on concepts done by philosophers is explicitly aimed at the same empirical and theoretical study of cognition – cognitive science – in which psychologists are engaged. Unless Machery finds a principled and adequate way to draw the line between philosophical and psychological theories, the thesis that “philosophical” and “psychological” theories are about different things does not even have a clear meaning.

The second point I disagree with is that the term “concept” should be eliminated from our scientific vocabulary, because different kinds of concept do not constitute a natural kind (section 8.3). This prescription is premature. Machery argues that there are too few scientific generalizations ranging over all concepts for concepts to constitute a natural kind. Therefore, Machery concludes, the term “concept” should be replaced by terms such as “prototype,” “exemplar,” and “theory.” There are several problems with this inference. First, the number of scientific generalizations that range over all concepts is largely to be determined by future science, so the main premise is problematic. Second, whether the number of generalizations ranging over concepts is too low for concepts to constitute a natural kind is debatable. Perhaps one generalization – e.g., that all concepts of x represent x – is so important that it warrants grouping all concepts within the same natural kind. But it’s too early to decide whether this is the case. Third and most importantly, the notions of prototype, exemplar, and theory are too controversial and ambiguous for them (as a group) to be an adequate replacement for the notion of concept. Consider the many uses of the term “concept” – including Machery’s own discussion of *concept* learning (chap. 6), *concept* combination (section 7.2), and *conceptual* change (*DwC*, p. 247). What are we supposed to replace “concept” and “conceptual” with? Perhaps we could talk about the learning of / combination of / change in our *prototypes*, *exemplars*, and *theories*. That would be awkward to say the least, especially since Machery’s view that concepts split into prototypes, exemplars, and theories is far from being the consensus. Even Machery doesn’t go that far; he sticks with the standard terms. If Machery doesn’t follow his own advice, it’s hard to see why anyone else should.

The above points do not affect what I take to be the most fertile aspect of Machery’s book – the argument to the effect that there are different kinds of concept. This conclusion, which may be an important step toward the next generation of theories of cognition, is what I actually wish to discuss.

Splitting Concepts

Which kinds of concept are there? Machery argues that for each category, people possess three main kinds of concept: prototypes, exemplars, and theories. His argument is that different phenomena require postulating all these kinds of concept for each category. Dan Weiskopf runs a similar argument (Weiskopf, 155). I think there is much that is right in Machery and Weiskopf's argument – first and foremost, I agree that none of the structures that traditional psychological theories of concepts have proposed are likely to account for all relevant phenomena. But I also doubt that Machery and Weiskopf's account cuts concepts at their most important joints.

In this paper, I will sketch and briefly defend an alternative way to split concepts. Before giving my positive account, I will clear the terrain by briefly commenting on the three kinds of concept countenanced by Machery.

First, exemplars play at most a peripheral role in a theory of concepts. Exemplars properly so called are representations of particular objects and their properties.⁴ According to Machery, sets of exemplars are one kind of concept (of general categories; of course particular exemplars may well count as concepts of particular objects). On the contrary, sets of exemplars are largely out of place in a theory of concepts.

It's obvious that cognizers represent *some* particulars as such – for examples, their family members – so it's obvious that they possess something like exemplars. It's also plausible that in some cases, as Machery argues (*DwC*, pp. 173-4), one or a few exemplars participate in processes such as concept combination. However, it doesn't follow that sets of exemplars are (a kind of) concepts. In fact, that is quite implausible for two reasons. First, for most categories, cognizers possess few or no exemplars.⁵ I am pretty sure my mind stores no exemplar of lily, iguana, and most other categories, for the simple reason that no particular lily, iguana, etc., has ever struck my mind with enough salience and enough unique features for me to have a stable representation of it that distinguishes it from other members of the same category. Other cognizers seem to be like me in this respect. So for most categories, cognizers possess no set of exemplars that could possibly constitute a concept of it. Therefore, their concept of such categories must be something other than a set of exemplars. Second, when cognizers do possess exemplars, they usually possess too many for the exemplars to be used as a concept in a computationally feasible way. Consider the category of people (and its many subcategories). Surely each of us has at least hundreds and likely thousands of exemplars of people stored in memory. If our conceptual processes – recognition, naming, concept learning, induction, etc. – required us to manipulate all (or even a significant subset) of our exemplars, the processes would be computationally unfeasible.⁶ For these reasons, sets of exemplars are not concepts.

In summary, surely cognizers store exemplars and use them in their higher cognitive processes. Furthermore, exemplars are good candidates for the role

of concepts of individual objects. But it's unlikely that sets of exemplars constitute general concepts; that is, concepts of general categories such as lily and iguana.

Second, Machery's case that people represent both statistical and causal information about categories can be disentangled from his conclusion that people possess both prototypes and theories. Machery mounts a persuasive empirical argument that cognizers represent and use both statistical information and causal information about categories. Since prototypes and theories are postulated to store, respectively, statistical and causal information about categories, Machery concludes that cognizers possess both prototypes and theories. This is one possibility, but not the only one. Another possibility is that there is a kind of concept that can store both statistical and causal information, at least when both kinds of information are available.⁷ As the latter assumption is more parsimonious, I will make it.

The concepts I postulate here, which encode both statistical and causal information, have something in common with theories as most psychologists conceive of them but are not the same thing. For theories are generally assumed to encode *explicit* information about categories – information that includes not only causal information, which may be represented implicitly, but also modal information, nomological information, and definitional information, which probably requires explicit representation (more on the implicit-explicit distinction below). Such explicit information is acquired mostly through language, and thus it is unavailable to non-linguistic creatures. But non-linguistic creatures can surely categorize and make inferences about categories, so they must have concepts. Presumably, their concepts encode at least statistical information about categories, and possibly some causal information. I will call the resulting concepts *implicit* concepts.

Implicit concepts are similar to what psychologists call prototypes. I avoid the term “prototype,” however, both because my implicit concepts store some (implicit) causal information, which traditional prototype theories leave out, and because many traditional prototype theories are formulated in ways that suggest they contain explicit information. There is every reason to expect that human beings possess implicit concepts, although their implicit concepts may well encode more sophisticated causal information than most or all other cognizers can acquire. In the rest of this paper, I will argue that some human concepts are not implicit but *explicit*.⁸

Competence Pluralism

Machery does not discuss language processing even though it is an important area of scientific research in which concepts are posited to do explanatory work. Linguists, psycholinguists, and philosophers of language explain people's ability to understand and produce linguistic structures, in part, by their possessing appropriate internal representations – i.e., concepts. As Sam Scott and I argue, there are reasons to postulate a special kind of concept, *explicit*

concepts, to do this and other related explanatory jobs. Explicit concepts are postulated in addition to *implicit* ones. According to this view, which Machery dubs Competence Pluralism, implicit concepts explain one range of phenomena, whereas explicit concepts explain a separate range of phenomena. Before getting to the reasons in favour of Competence Pluralism, let me address Machery's objection.

Machery points out that subjects exhibit some of the same phenomena, such as typicality effects, for both non-lexicalized categories (i.e., categories for which subjects lack a word) and lexicalized categories (i.e., categories for which subjects have a word). From this he infers that concepts of lexicalized categories share many properties with concepts of non-lexicalized categories – therefore, contrary to Piccinini and Scott, there is no difference in kind between the two.⁹

This objection seems to be based on a misunderstanding, which I will now attempt to clear up. The objection assumes that according to Competence Pluralism, for each category, subjects possess either an explicit or an implicit concept, but not both. If that were the case, it is at least plausible that the same kind of concept explains the phenomena pointed at (typicality effects, etc.) for both lexicalized and non-lexicalized categories. If there were no other phenomena left to explain, it would be at least plausible that – contrary to Piccinini and Scott – lexicalized and non-lexicalized categories are represented by the same kind of concept.

But the assumption underlying the objection is inconsistent with Piccinini and Scott's Competence Pluralism. Our hypothesis is that, with respect to lexicalized categories, people possess two kinds of concept: explicit and implicit. Implicit concepts of lexicalized categories work similarly to implicit concepts of nonlexicalized categories; they are employed for the same tasks that implicit concepts of nonlexicalized categories are employed for. Thus, they explain the same phenomena that implicit concepts of nonlexicalized categories explain, including typicality effects. So, contrary to Machery's suggestion, the evidence marshalled in his objection is naturally predicted by our Competence Pluralism. People possess implicit concepts of a category whether or not the category is lexicalized. By contrast, explicit concepts are posited to explain *other* phenomena – phenomena that implicit concepts are unable to explain, such as language understanding and linguistic inference.¹⁰

With Competence Pluralism thus clarified, I now turn to the positive reasons to conclude that people possess both explicit and implicit concepts, and that explicit concepts are different in kind from implicit concepts.

The Argument from Two Minds

The following argument begins with the observation that there are two kinds of cognitive process, both of which can control behaviour. The two kinds of process can be doubly dissociated; thus, the systems that exhibit them can operate independently of one another. Since both kinds of process rely on

concepts and they are doubly dissociable, there must be at least two distinct conceptual systems – one for each kind of process. This conclusion takes us part of the way toward Competence Pluralism, but it doesn't get us there yet. For Competence Pluralism maintains that there are two *different* kinds of concept. The argument from two minds concludes that there are two conceptual systems and suggests that they are likely to contain concepts of different kinds, but it doesn't rule out that the concepts in the two conceptual systems are of the *same* kind. That further step requires a separate argument, to be given in the next section. For now, let me briefly flesh out the argument from two minds.

Cognitive psychologists distinguish between explicit (or declarative) and implicit (or procedural) cognitive processes. Properties often attributed to explicit processes include being conscious, controlled, sequential, fast-learning, and slow-acting. By contrast, properties often attributed to implicit processes include being unconscious, automatic, parallel, slow-learning, and fast-acting. For present purposes, the only important difference between the two kinds of process is that explicit processes are the only ones whose contents the language faculty can access: subjects can give verbal reports on the contents of explicit processes, but not on the contents of implicit processes. Aside from this, the precise characterization of explicit versus implicit processes does not matter much for present purposes. What matters is that there is broad consensus that these are two different kinds of process, each of which explains different phenomena.

Machery discusses the distinction between implicit and explicit cognition under the banner of “multi-process theories” (*DwC*, pp. 140-50, 185-6). He argues that the multi-process theories he reviews are insufficiently detailed, rely on unclear distinctions, and are based on artificial tasks that might lack ecological validity. For present purposes these criticisms are inconsequential, because they pertain only to some aspects of some specific theories. They do not impugn the general distinction between explicit and implicit cognition, which is supported by decades of experiments in areas that include learning and memory, reasoning, decision making, and social cognition.

Here are some results that support the distinction between implicit and explicit cognitive processes:

- People who can describe how to do something (e.g., drive a car or play tennis) may be unable to do that thing, and vice versa (e.g., touch typists generally cannot verbally recall how the keys are arranged on the keyboard) (cf. Jilk *et al.*, 206).
- People with bilateral damage to the hippocampus can no longer recall what they experience (deficit in explicit learning and memory) but can still learn new skills and habits (preserved implicit learning and memory).
- People may verbally express beliefs and desires that are in conflict with other aspects of their behaviour. Furthermore, people may find it difficult to change their behavioural habits on the basis of their verbally expressed

attitudes. The consensus in social psychology is that people's behaviour is influenced by both implicit and explicit beliefs and desires, and that implicit and explicit attitudes are not always in agreement.

- People exhibit nonverbal behavioural sensitivity to certain categories considerably earlier in development than when they exhibit language-based understanding of the same categories. For example, although children do not show linguistic understanding of the distinction between true and false beliefs (e.g., by verbally solving a false belief task) until the age of about four years, infants as young as fifteen months old exhibit habituation and dishabituation behaviours that suggest they are sensitive to the difference between true and false beliefs. A plausible explanation is that while it takes children until the age of about four to acquire an explicit understanding of false beliefs, they possess an implicit understanding of false beliefs much earlier than that (Onishi and Baillargeon).

These are just a few examples of phenomena that support the distinction between explicit and implicit cognition. There is no room here to review the extensive literature that bears on this topic (see Evans 2008, Evans and Frankish 2009). Suffice it to say that throughout psychology and neuroscience it is widely accepted that explicit and implicit cognitive processes are distinct kinds of process that can control behaviour, are dissociable, and are implemented in distinct kinds of neural system (Lieberman 2007, 2009; Samuels).

Both explicit and implicit cognitive processes are assumed by mainstream cognitive psychologists to involve the manipulations of representations. As we saw in section 1, concepts are nothing but the primitive representational components of such representations. Therefore, since explicit and implicit processes are implemented in different kinds of systems, there are at least two distinct conceptual systems – one for implicit cognition and one for explicit cognition.¹¹

Implicit and explicit cognition differ in important respects. First and foremost, only the contents of explicit cognition can be verbalized. The differences between implicit and explicit cognition strongly suggest that the concepts belonging to the two systems are of different kinds – kinds whose differences would explain the differences between explicit and implicit cognition. I will now argue that this is indeed the case: there are two different *kinds* of concept.

The Argument from Explicit Cognition

The present argument is a refinement of Piccinini and Scott's Argument from Language from "Splitting Concepts." In a nutshell, the argument is that explicit cognition and implicit cognition require concepts with different properties. Therefore, there are two kinds of concept.

Language marks a significant discontinuity in kind between cognizers. On one hand are non-human animals, whose cognitive systems appear unable to

acquire and process language. This is not to say that non-human animals are unable to communicate, of course. They have sophisticated communication systems. But none of them can acquire a language in the strict sense of the term – a communication system based on a recursive syntactic system capable of producing infinitely many structures out of finitely many primitive components.

On the other hand are human beings. Human beings share many aspects of their cognitive endowment with other animals. Thus, it is overwhelmingly plausible that human beings share with other animals whatever concepts allow other animals to categorize and draw (implicit) inferences about categories. These are what I call *implicit* concepts. Implicit concepts may well represent statistical information about categories as well as some causal information. In addition to implicit concepts and the mechanisms that process them, human beings have a special cognitive apparatus that allows them to acquire and process language.

It is almost commonplace that explicit cognition – cognition that is capable of processing language (strictly so called) – confers cognitive advantages. The range of messages, inferences, and behaviours that explicit cognition makes possible appears to be of a different kind than the kind of messages, inferences, and behaviours that implicit cognition allows. The cognitive power of explicit cognition is due, to a large extent, to the productive recursive structure that language possesses. Assuming that non-human animals are as incapable to acquire a language as they appear to be, we may conclude that acquiring and processing a language requires a special cognitive apparatus that is unique to human beings – a “faculty of language in the narrow sense” (FLN) (Hauser, Chomsky, and Fitch).¹²

Given that language is unique to human beings and that explicit cognition enhances a creature’s cognitive power (more on this below), it would be surprising if the same kind of concept that people share with non-human animals were enough for human beings to support explicit cognition. On the contrary, I will now list several considerations pointing in the opposite direction: concepts different in kind from and cognitively more powerful than implicit ones.

Dissociation between Recognition and Naming

A traditional assumption about the relation between concepts and words is that there is a simple one-to-one mapping between the two. To each concept there corresponds a word (except for non-lexicalized concepts), and to each word there corresponds a concept (except for ambiguous words). (Opinions differ as to whether grouping objects into categories causes or is caused by naming patterns, but this disagreement is irrelevant here.) There is growing evidence that this assumption of a one-to-one mapping between words and concepts is incorrect.

Specifically, a number of studies have demonstrated a partial dissociation between recognition – as measured by perceived similarity between objects – and naming. In one study, English, Chinese, and Spanish speakers were asked both to cluster 60 containers into groups based on their similarity and to name each

container. Notice that the similarity measures that were used included not only structural but also functional features of the objects. When naming the stimuli, the three groups of speakers classified the stimuli according to three different classificatory schemes, which depended on the language each spoke. The three naming schemes often cross-classified the stimuli – that is, they were mutually inconsistent. But when sorting the objects based on their similarity, all subjects sorted the objects in approximately the same way, regardless of which language they spoke (Malt *et al.* 1999, 2003). Similar results were obtained on the classification of motion events (Gennari *et al.* 2002; see also Malt and Sloman 2004).

Other studies, which Machery reviews (*DwC*, pp. 185-6, 195-6), provide evidence of two dissociable categorization processes: an implicit, similarity-based process, and a categorization process that may rely on explicit rules or definitions (Allen and Brooks; Smith and Sloman 1994, Smith *et al.* 1998).

These results suggest that human beings have two sources of categorization: an implicit concept, which operates on the basis of similarity relations between stimuli (including not only structural but also causal and functional features), and an explicit concept, which is constrained by conventions, definitions, and other linguistic constraints. The resulting two forms of categorization – recognition and naming – are correlated but also somewhat independent, because they have different constraints.

Syntax and the Format of Explicit Concepts

As I already pointed out, understanding and producing linguistic sentences requires a productive, recursive syntax. I assume that human beings use concepts to represent the meaning of words and employ such concepts when processing linguistic structures. Since linguistic structures have syntactic properties and their syntax must be parsed for linguistic structures to be understood or produced, and since syntax interacts heavily with semantics, it follows that the semantic representations that enter language processing must possess a format that allows them to interact with the parsing of syntax.

For instance, one way in which explicit concepts interact with syntax is by fulfilling certain syntactic roles, such as noun, verb, adjective, etc. In order to do so, FLN must be able to sort concepts into syntactic categories. One way to do so would be to store information about syntactic category into the concepts themselves. However FLN accomplishes this task, it is a task that is specific to it and not shared with the cognitive systems involved in implicit cognition. In summary, explicit concepts must have a format that allows them to interact appropriately with syntax – something that implicit concepts need not and presumably do not have. Thus, the format of explicit concepts is likely to be different in kind from the format of implicit concepts.

Explicit Representational Power

As many others have pointed out, linguistic structures can represent things that are difficult or impossible to represent without them. Consider conditional

statements, counterfactual statements, modal statements, negated statements, and quantified statements. It's difficult to see how the contents of any of these types of statement could be represented without language-like resources (cf. Fodor, 175). This is due in part to the already-mentioned syntactic properties of language. In addition, the representational power of language requires concepts that play special roles: negation, the quantifiers, possibility and necessity, etc.

Language also allows human beings to represent entities and properties that are unobservable (e.g., *atom*, *black hole*), non-existent (e.g., *fairy*, *superhero*), arbitrarily abstract (e.g., *justice*, *truth*, *square root*), or ad hoc (e.g., *things to do for weekend entertainment*). Suffice it to say that implicit cognition appears to lack concepts for these types of categories, and that most of these concepts appear to lack the kind of similarity structure that is reliably exhibited by naturally occurring, referring object concepts (such as *apple*, *dog*, *pen*, etc.) (Piccinini and Scott, 404-5).

Both of these observations suggest that explicit cognition requires concepts that have a different structure and a greater representational power than implicit concepts: explicit concepts can represent what implicit concepts cannot.

Explicitly Encoded Information

Although many concepts appear to lack definitions, other concepts have definitional structures. One example is the concept of grandmother. G is x 's grandmother if and only if G is the mother of a parent of x . While the concept of grandmother is relatively simple, its definition is already more complex than anything that is likely to be expressible without explicit representational resources. And definitions can be indefinitely long and complex – requiring an indefinite amount of syntactic structure to be formulated. Definable concepts are strongly associated with their definitions, which provide criteria for their application. At least for simple and commonly used definable concepts, such as *grandmother*, most competent speakers of the language know the definition and use it to apply the concept.¹³

It is likely that formulating, acquiring, storing, and deploying definitions require cognitive resources that are available only through FLN, such as syntax and concepts for the types of categories mentioned in the previous section. Thus, definable concepts are explicit concepts. The existence of definable concepts, and the role that their definitions play in cognitive processes, is another piece of evidence that explicit concepts are different in kind from implicit concepts.

Explicit Inferential Power

Inferences involving linguistic structures are more powerful and versatile than implicit inferences. Human beings, unlike other animals, can manipulate linguistic structures in ways corresponding to deductively valid inferences, such as *modus ponens*, as well as inductive inferences. Perhaps most relevantly, people can perform a wide range of material inferences, such as the inferences

from “*x* is blue” to “*x* is coloured,” from “*x* is above *y*” to “*y* is below *x*,” or from “*x* is a raccoon” to “*x*’s parents were raccoons.” By contrast, non-linguistic cognizers lack the ability to perform these kinds of inference. Patently, performing inferences is not merely a matter of acquiring a language; one must also acquire the ability to generate appropriate conclusions from the available premises. Such an ability presumably is grounded on information stored either in the conceptual system or in a system of semantic memory that is closely associated with the conceptual system. Given that linguistic inferences are manifested through language, and hence involve FLN, and that they are not exhibited by non-linguistic cognizers, the conceptual system behind linguistic inferences must be the explicit one. Therefore, explicit concepts – unlike implicit ones – are capable of participating in cognitive structures that support linguistic inference. Once again, this is evidence that explicit concepts and implicit concepts constitute different kinds.

Conclusion

As Machery argues, people possess different kinds of concept. This is a plausible and important conclusion, which paves the way for a new generation of theories of cognition. While I agree that concepts split into different kinds, I disagree about which kinds they split into.

To be sure, the above points deserve to be elaborated at much greater length. Even so, they do suggest that contrary to Machery, concepts do not split into exemplars, prototypes, and theories. Instead, concepts split into explicit and implicit concepts. Implicit concepts encode statistical and some causal information about categories; they subsume the implicit aspects of both prototypes and theories. (As to exemplars, they exist but are not concepts of general categories.) Explicit concepts may encode statistical and causal information, but more importantly, they may encode syntactic information, definitional information, and whatever else is needed for the language faculty (in the narrow sense) to process them. Thus, explicit concepts subsume the explicit aspects of what psychologists call theories (and perhaps prototypes). Explicit concepts are necessary for explicit cognition – the distinctively human ability to use language, to represent unobservable, nonexistent, abstract, and ad hoc aspects of the world, and to perform linguistic inferences.¹⁴

Notes

- 1 Unless otherwise noted, references are to Machery 2009.
- 2 Machery offers a different programmatic functional characterization of concepts in psychology. He defines concepts as bodies of “knowledge” (i.e., information) stored in long-term memory and used by default in the processes underlying most higher cognitive competences when these processes result in judgments (*DwC*, p. 12). This account has several problems. First, contrary to common usage in both philosophy and psychology, Machery’s account rules out structures that are not

stored in long-term memory. Second, it relies on the vague notion of “higher” cognitive process. Third, it’s unclear why cognitive competences that do not result in judgments should be ruled out – or else, it’s not clear whether the notion of judgment Machery employs is adequate to the task of demarcating concepts. For these reasons, I will stick with the more traditional account given in the main text.

- 3 What Machery offers is a programmatic characterization of concepts in philosophy that differs from the one he gives for psychology. According to him, in philosophy having a concept of x is being able to have propositional attitudes about x as x (*DwC*, p. 32). But even leaving aside that many philosophers would reject this characterization, Machery offers no reason to believe that his two functional characterizations are mutually exclusive, i.e., about different things.
- 4 The term “exemplar” is not always used in its strict sense. For example, Machery (*DwC*, p. 172) discusses Smith’s (2002) argument that the typicality gradients observed during dot-distortion category tasks disconfirm exemplar theories in favor of prototype theories. During the dot-distortion category task, subjects learn to classify patterns of dots on a screen created by distorting a “prototypical” pattern. Smith refers to specific patterns of dots, and representations thereof, as exemplars. But a pattern of dots is not a particular object: a pattern of dots is repeatable, whereas an object is unrepeatable. Since the stimuli in dot-distortion category task experiments are not particular physical objects but random patterns of dots, nothing in these experiments shows that subjects store exemplars, i.e., representations of particular objects, as opposed to representations of patterns of dots. This equivocal use of “exemplar” may contribute to the belief that exemplars belong in a theory of concepts. My argument in the text applies only to exemplars properly so called.
- 5 People with “photographic” memory or at least exceptionally good memory might be a partial exception here.
- 6 Objection: Your second reason relies on the assumption of a classical, serial architecture. What about doing the processing in parallel? Wouldn’t parallel processing get around the objection from computational feasibility? Reply: Maybe, although even with parallel processing, computing similarity based on exemplars would be computationally more costly than doing it based on prototypes.
- 7 In some experimental paradigms, subjects are presented with stimuli in such a way that statistical information is available, but no causal information is available. In such cases, a concept can only store statistical information. This doesn’t show that, when causal information about a category is available, the causal information is represented by a separate (“theoretical”) concept, as opposed to being stored together with the statistical information within the same concept.
- 8 I’m now calling *implicit* and *explicit* concepts what Piccinini and Scott called, respectively, *linguistic* and *nonlinguistic* concepts. Thanks to Machery for requesting a more perspicuous nomenclature.
- 9 Here is how he puts it:

According to Competence Pluralism, every category ... is represented by several concepts, one for each cognitive competence. ... Moreover, the concepts of a given

category (substance, etc.) belong to kinds of concept that have little in common with each other ...

There is a wealth of evidence that lexicalized concepts (that is, concepts that are expressed by a word), which can be recruited by the linguistic competences, and non-lexicalized concepts, which cannot be recruited by these competences, have very similar properties. Consider, for example, typicality. It is well-known that typical objects are categorized more quickly and more accurately than atypical objects. The membership of typical objects in a given category is also learned more quickly than the membership of atypical objects in this category. For present purposes, the important point is that these properties are common to lexicalized concepts and to non-lexicalized concepts 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 non-lexicalized figures into different categories and are then asked to classify new figures into these categories, typical figures are classified more quickly and more accurately than atypical figures. The similarity between lexicalized and non-lexicalized concepts shows that it is not the case that very different kinds of concepts are used in linguistic tasks and in non-linguistic tasks. Thus, Piccinini and Scott's defense of Competence Pluralism fails (*DwC*, pp. 58-60).

In my rendition of Machery's argument in the main text, I reformulate the objection without using the expression "lexicalized concepts" because I find it tendentious.

- 10 More precisely, typicality effects with respect to lexicalized categories may be explained in one of two ways, consistent with Competence Pluralism. Either typicality effects are always explained by implicit concepts (for both lexicalized and nonlexicalized categories), or explicit concepts (like implicit concepts) also have a similarity structure that can explain typicality effects. I'd rather not take a stance on which of these explanations is correct.

- 11 Two caveats:

In section 1 I rejected Machery's alternate account of concepts in psychology, but the present argument does not depend on that rejection, for the primitive representations manipulated by explicit and implicit cognitive processes count as concepts under Machery's account too.

Some philosophers may be reluctant to posit implicit concepts, preferring to say that implicit processes manipulate not concepts but rather "nonconceptual" representations. But I am here trying to explain phenomena that psychologists explain in terms of what they call concepts, so I will keep using the term "concept."

- 12 I am not hereby committing myself to any particular theory of the language faculty, including Hauser *et al.*'s hypothesis that the language faculty in the narrow sense includes "only the core computational mechanisms of recursion as they appear in narrow syntax and the mappings to the interfaces" (Hauser, Chomsky, and Fitch, 1573). For some recent discussion of the language faculty, see also Pinker and Jackendoff; Fitch, Hauser, and Chomsky; and Jackendoff and Pinker. My argument in this paper is intended to be neutral between the positions in this debate. It's also intended to be neutral as to whether explicit cognition is made possible by the acquisition of language or vice versa.

- 13 Machery appears to be an exception. He maintains that “Tina Turner is a grandmother” is “true under one reading and false under another reading” (*DwC*, p. 72). On the contrary, “Tina Turner is a grandmother” is just true.
- 14 Thanks to Dan Weiskopf for comments on a previous version and to Jim Virtel for editorial assistance. Special thanks to Edouard Machery for discussion and comments.

References

- Allen, S. W. and L. R. Brooks
1991 “Specializing the Operation of an Explicit Rule.” *Journal of Experimental Psychology: General* 120: 3–19.
- Evans, J. S. B. T.
2008 “Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition.” *Annual Review of Psychology* 59: 255–78.
- Evans, J. S. B. T. and K. Frankish, eds.
2009 *In Two Minds: Dual Processes and Beyond*. Oxford: Oxford University Press.
- Fitch, W. T., M. D. Hauser, and N. Chomsky
2005 “The Evolution of the Language Faculty: Clarifications and Implications.” *Cognition* 97: 179–210.
- Fodor, J. A.
2008 *LOT 2: The Language of Thought Revisited*. Oxford: Oxford University Press.
- Gennari, S. P., S. A. Sloman, B. C. Malt, and W. T. Fitch
2002 “Motion Events in Language and Cognition.” *Cognition* 83: 49–79.
- Hauser, M. D., N. Chomsky, and W. T. Fitch
2002 “The Faculty of Language: What Is It, Who Has It, and How Did It Evolve?” *Science* 298: 1569–79.
- Jackendoff, R. and S. Pinker
2005 “The Nature of the Language Faculty and Its Implications for Evolution of Language (Reply to Fitch, Hauser, and Chomsky).” *Cognition* 97: 211–25.
- Jilk, D. J., C. Lebiere, R. C. O’Reilly, and J. R. Anderson
2008 “SAL: An Explicitly Pluralistic Cognitive Architecture.” *Journal of Experimental and Theoretical Artificial Intelligence* 20(3): 197–218.
- Lieberman, M. D.
2007 “Social Cognitive Neuroscience: A Review of Core Processes.” *Annual Review of Psychology* 58: 259–89.
2009 “What Zombies Can’t Do: A Social Cognitive Neuroscience Approach to the Irreducibility of Reflective Consciousness.” In *Two Minds: Dual Processes and Beyond*, ed. J. S. B. T. Evans and K. Frankish, 293–316. Oxford: Oxford University Press.
- Machery, E.
2009 *Doing without Concepts*. Oxford: Oxford University Press.

- Malt, B. C., S. A. Sloman, S. Gennari, M. Shi, and Y. Wang
1999 "Knowing versus Naming: Similarity and the Linguistic Categorization of Artifacts." *Journal of Memory and Language* 40: 230–62.
- Malt, B. C., S. A. Sloman, and S. P. Gennari
2003 "Universality and Language Specificity in Object Naming." *Journal of Memory and Language* 49: 20–42.
- Malt, B. C. and S. A. Sloman
2004 "Conversation and Convention: Enduring Influences on Name Choice for Common Objects." *Memory and Cognition* 32(8): 1346–54.
- Onishi, K. H. and R. Baillargeon
2005 "Do 15-Month-Old Infants Understand False Beliefs?" *Science* 308: 255–58.
- Piccinini, G. and S. Scott
2006 "Splitting Concepts." *Philosophy of Science* 73(4): 390–409.
- Samuels, R.
2009 "The Magical Number Two, Plus or Minus: Dual-Process Theory as a Theory of Cognitive Kinds." In *Two Minds: Dual Processes and Beyond*, ed. J. S. B. T. Evans and K. Frankish, 129–46. Oxford: Oxford University Press.
- Smith, E. E. and S. A. Sloman
1994 "Similarity- versus Rule-Based Categorization." *Memory and Cognition* 22: 377–86.
- Smith, E. E., A. L. Patalano, and J. Jonides
1998 "Alternative Strategies of Categorization." *Cognition* 65: 167–96.
- Smith, J. D.
2002 "Exemplar Theory's Predicted Typicality Gradient Can Be Tested and Disconfirmed." *Psychological Science* 13(5): 437–42.
- Weiskopf, D. A.
2009 "The Plurality of Concepts." *Synthese* 169: 145–73.