

Concepts are Containers

ROBERT O'SHAUGHNESSY and MARK SPREVAK
University of Edinburgh, Edinburgh, UK

In this paper, we propose and defend a theory of concepts. According to Machery (2009), psychologists and philosophers mean different things by 'concept'. Psychologists mean bodies of knowledge used to categorise and infer; philosophers mean constituent of propositional thought. Machery's conclusion would drive a wedge between contributions by psychologists and philosophers on concepts. Theories about the former would have no clear role to play in, and cast no light on, the latter, and vice versa. We argue that, on the contrary, 'concept' has a single core meaning: a container of stored knowledge pertaining to a single category. This single meaning satisfies both the theories of psychologists and philosophers. The divergence in use of the term 'concept' on which Machery focuses arises because words for containers are often used to refer to (a) what is contained by the container and (b) the label of a container. Our account explains what a concept is, and how one might be misled by Machery's challenge.

Keywords: Concepts; mental files; pointers; language of thought; eliminativism; Machery; Fodor.

1. Introduction

Machery (2009) claims that philosophers and psychologists mean different things when they use the term 'concept'. Machery identifies two desiderata for something to be a concept:

1. *Judgement desideratum:* A concept must permit us to make categorisation, typicality, and inferential judgements.
2. *Propositional desideratum:* A concept must be capable of being used as a constituent in propositional thought.

Machery claims that the term 'concept' means something different when it is used to satisfy the first desideratum than it does when it is used to satisfy the second. Our claim is that 'concept' has a single core meaning that satisfies both desiderata: *a container of stored knowledge pertain-*

ing to a single category.¹ If we are right, then philosophers and psychologists are talking about the same thing when they talk about concepts.

To explain the divergence in use of the term 'concept' that Machery describes, we claim that psychologists and philosophers extend a single core meaning in different ways. Unlike Machery, we do not conclude that there are two entirely distinct entities that stand behind our concept talk. Rather, we argue that we have a case of polysemy patterned on forms of polysemy familiar when containers are in play. 'Concept', in this respect, is like 'DVD': the word for the container may be used to refer to the contents of the container (for example, "that DVD was boring" meaning *the movie on that DVD was boring*) or the word for the container ('DVD') may be used to refer to the title or label of the container (for example, handing someone a list of movie titles and asking her to "pick one of those DVDs" meaning *pick one of those titles of, or labels for, DVDs*). 'DVD' may be used to refer to the *movie* or its *associated title*, but no one would think that a DVD *is* a movie or that a DVD *is* a title. What a DVD *is* is a container of stored information.

Our claim is that a concept is a container of stored knowledge that pertains to a single category. Psychologists tend to refer to the *contents* of the container. Philosophers tend to refer to the *label* of the container. No one should conclude from this that a concept *is* the contents or that it *is* the label, any more than they would for 'DVD'.²

2. Machery's claim

Machery (2009) claims that philosophers and psychologists refer to two entirely distinct entities by 'concept'. Philosophers are disposed to favour models in which, for example, RED is an atomic concept linked to other related concepts (such as COLOUR) which do not leave a role for the exemplars and prototypes that psychologists think of as concepts and which facilitate categorization and inference. Psychologists favour

¹ We use the term 'knowledge' the way Machery and psychologists do: as a term for "any contentful state that can be used in cognitive processes" (Machery 2009: 8). States of knowledge in this sense do not have to be true or justified, nor do they have to be explicit or propositional: they may be imagistic (perceptual) or procedural (sensorimotor).

² What we are proposing could be described as an elaborated version of the "different aspects" response to Machery's claim (Machery, personal communication). That is to say, philosophers and psychologists are talking about different aspects of the same thing, not about different things (for responses along this line, see Margolis and Laurence 2010; Piccinini 2011). Machery replies to these objections that the aspects that satisfy the propositional desideratum play no role in theories about the aspects that satisfy the judgement desideratum, and vice versa. This suggests that there are two different entities, not one, involved in concepts. Our account differs in that we take there to be a *third* entity—a container of stored knowledge—and we hold that psychologists are interested in one aspect of this entity (its contents) and philosophers in another aspect (its label). This single third entity unifies the (otherwise puzzlingly different) multiple aspects of concepts and satisfies both Machery's desiderata.

models in which concepts are pieces of information (exemplars, prototypes and theories) with rules of engagement. But there do not seem to be good rules of engagement that allow such information to be constituents of propositional thought. If correct, Machery's claim would create a puzzle. Philosophers' concepts and psychologists' concepts would, in principle, have nothing in common. When you form the propositional thought THE DOG CHASES THE CAT, your understanding of the proposition could be entirely unrelated to your ability to make the kinds of judgements in which psychologists are interested about the constituent concepts DOG, CAT, CHASE, and vice versa. But divorcing the interests of philosophers and psychologists here seems too strong. Rich connections exist between philosophers' concepts and psychologists' concepts. These connections are obscured and relegated to purely contingent happenstance on Machery's account.

Chalmers (2011) makes a distinction between verbal disputes and substantive disputes. A verbal dispute arises when parties take themselves to be using an expression for which they both have the same proposition in mind when, in fact, they have different propositions in mind. The parties disagree, failing to realise that they are talking about different things. Once it is discovered that they are using the same term in different ways the dispute may vanish. For a substantive dispute, the parties use the expression in the same way but they disagree about the underlying facts.

Machery (2009) is, in effect, claiming that disputes between philosophers and psychologists over the nature of concepts are verbal disputes. Philosophers and psychologists use the term 'concept' and they might assume that they are talking about the same thing. But, according to Machery, they are talking about different things. For example, when Fodor says, "most of what contemporary cognitive science believes about concepts is radically, and practically *demonstrably*, untrue" (Fodor 1998: viii), Fodor (according to Machery) is engaged in a verbal, rather than a substantial, attack on cognitive science. Our claim, *contra* Machery, is that philosophers and psychologists mean the same thing by 'concept'. There is scope for substantive, not merely verbal, dispute between philosophers and psychologists about concepts.

3. *Polysemy*

Polysemy arises when a term has multiple meanings that are semantically related. For example, the term 'bank' is polysemous: 'bank' means both *financial institution* and *physical building* where this institution offers services—the words have different meanings but they are semantically related. Polysemy differs from mere homonymy. Homonymy arises when a term has multiple meanings that may be semantically unrelated. The term 'bank' also functions as a mere homonym: it means both *financial institution* and *side of a river*—different meanings that are semantically unrelated.

A polysemous term that will prove instructive to us later is 'table'. Murphy (2002) lists fourteen semantically related meanings of 'table' including: what furniture makers mean by 'table' (*four-legged piece of furniture*), what geographers mean by 'table' (*a flat or level area of land*), and what jewellers mean by 'table' (*a facet of a cut precious stone*). Murphy claims that these meanings are related because they share a common etymological origin. Some original meaning of 'table' (according to Murphy, *four-legged piece of furniture*) was extended to generate the rich palette of meanings now associated with 'table'.

Machery observes that philosophers and psychologists use the term 'concept' in different ways. Our claim, and Machery's view, is not that 'concept' is a mere homonym. Machery's proposal is that 'concept' is polysemous: it has multiple, semantically related meanings that pertain separately to philosophy and psychology. What is distinctive about Machery's proposal is that 'concept' is polysemous *in a particular way*: 'concept' is polysemous in the mouths of philosophers and psychologists in roughly the same way that 'table' is polysemous in the mouths of furniture makers and geographers. In both cases, *entirely distinct and independent entities* are denoted by the same linguistic expression. 'Table' as used by furniture makers refers to a *piece of furniture*. 'Table' as used by geographers refers to a *flat piece of land*. Machery claims that 'concept' as used by philosophers refers to a *constituent of propositional thought*; 'concept' as used by psychologists refers to a *body of knowledge used for classification*. There would be no point in a furniture maker and geographer entering into a dispute about whether "tables have four legs" or whether "tables in South Africa are sedimentary deposits". Similarly, there would be no point in a philosopher and psychologist entering into a dispute about whether "concepts govern typicality judgements" or whether "concepts are involved in linguistic thought".

We agree that 'concept' is polysemous, but we disagree about the kind of polysemy involved. The polysemy involved in 'concept' is not the same as that exhibited by 'table' as used by furniture makers and geographers. A referring term, such as 'table' or 'concept', may be polysemous without denoting entirely distinct and independent entities. If we are correct, then despite the divergence in use of 'concept' that Machery describes, both philosophers and psychologists can and should agree about what a concept is.

4. *Polysemy without proliferation*

Consider two other ways in which 'table' is polysemous (Murphy 2002: 404):

- a. *the company of people eating at a table*: "The entire table shared the plate."
- b. *a painting, sculpture or photograph of a table*: "The table is painted very soulfully."

Our claim is that 'concept', as used by psychologists and philosophers, is polysemous in roughly the same way that 'table' is polysemous in the preceding sentences. Psychologists follow the pattern exemplified by (a). Philosophers follow the pattern exemplified by (b).

Consider (a). Here, 'table' is used to refer to a *person or group of people*. For example, the waiter might say, "Table six wants a coffee refill". Observe that although 'table' means *person or people*, a second table (a piece of furniture) is also involved. One cannot use 'table' in isolation to mean *person or people*. One cannot point to a person walking down the street and say, "Look at that table". In contexts where 'table' means *person or people*, there must be also a table (a *piece of furniture*) present. This is because what is usually meant by 'table' is *person or people seated at the table*. The listener must be able to identify a table (a piece of furniture) in order to know which person or people are intended.

Compare this to a geographer's 'table'. In that case, no table (piece of furniture) is involved. Except from at the origin of the geographer's term, tables as pieces of furniture are irrelevant to the intended meaning. In contrast, for the waiter's 'table' (person or people), a table (piece of furniture) is essential to our present use and understanding. If one did not know what a table (piece of furniture) is, one would not know which person or people are being referred to. And it is because the table (piece of furniture) serves the function of grouping people together that the term can be used to refer to those people.

For psychologists, the polysemy involved in 'concept' is not that of the geographer's 'table', for which the term refers to two entirely distinct and otherwise unrelated entities (*piece of furniture* and *flat piece of land*). It is instead like that of the waiter's 'table', for which the original referent is used to identify something linked to, contained by, or grouped together by, that item. When psychologists use 'concept' what they mean, we claim, is not *body of knowledge* but *container of body of knowledge (pertaining to a single category)*.³

³ Note that Machery cannot make the same response to our claim as his (2010) reply to Margolis and Laurence (2010). Margolis and Laurence suggest that we should think that concepts are mental symbols akin to words and that psychologists are interested in one aspect of these concepts (the exemplars, prototypes and theories linked to such concepts). Machery responds along the following lines. Mental symbols that are constituents of thought play no role in, and cast no light on, the theories of concepts that occupy psychologists. Indeed, Margolis and Laurence's characterization makes most of what psychologists say about concepts literally false. On the other hand, if 'concept' has the meaning set out by Machery then most of what psychologists say about concepts comes out true, and everything else being equal, this is to be preferred. Machery cannot make a similar response to our proposal. First, concepts conceived as containers do play the required role in psychological theories (see Section 7). Second, our proposal does not require what psychologists say about concepts to be false, provided it is understood as we claim: as a linguistic shorthand for *bodies of knowledge contained in (or stored in) a concept* where 'concept' means *container of stored knowledge pertaining to a single category*.

Now consider (b). The polysemy arises here because a single term is used to refer to both the object being represented and the representation of the object. We might say of a painting of a table, "That table is painted very soulfully". It may even take a moment's reflection to realise that this is a case in which 'table' deviates from the furniture makers' intended meaning. But what we mean is not that a table (a piece of furniture) is painted soulfully (perhaps the table has not been painted at all). We mean that *the image of the table or the representation of the table* has been painted soulfully.

Imagine that a restaurant owner has a graphical computer program that allows her to arrange representations of tables on screen to match the bookings for that evening. When she moves the outlines of tables on her screen, perhaps composing outlines into greater wholes, it is natural for her to speak of the representations on the screen as 'tables', even though they are merely proxies for tables. In the same way, we suggest, when philosophers use 'concept' what they mean is *label of a concept or proxy for a concept* (where 'concept' means the same as it did for psychologists: *a container of stored knowledge pertaining to a single category*).

What is distinctive about both kinds of polysemy is that both involve reference to entities that are *related to single common entity*. 'Table' refers to *person/people seated at a table* or to a *representation of a table*. Just as one cannot use 'table' to mean *person/people* but only *person/people at a table*, so one cannot use 'table' to mean *representation* but only *representation of a table*. This differs from the polysemy of the geographer's 'table' and the kind that Machery claims is involved in 'concept'. The waiter and the restaurant owner use 'table' to mean different things (*person/people seated at a table* and *representation of a table*) but they can do this only because a third entity, a table (piece of furniture), is related to both. So, we claim, psychologists and philosophers use 'concept' to mean different things (*body of knowledge for categorisation* and *constituent of thought*) but only because a third entity, a concept (*container of stored knowledge pertaining to a single category*), is related to both. A concept, we propose, is not a generic container but one which has the purpose of containing information pertaining to a single category.

5. Container talk

The term 'table' behaves in this way because tables are containers: tables group people together.⁴ The linguistic patterns described above

⁴ Note we are not claiming that the term 'table' is an ideal analogy for the term 'concept' in all respects ('table' has many more meanings than 'concept'). Rather we make use of Murphy's (2002) discussion to distinguish between what we call polysemy with proliferation (which Machery is arguing for) and polysemy without proliferation (which we are arguing for).

are characteristic of container talk. If one has a term for a container, or a thing that groups other things together, that term can also be used to refer to the *contents* grouped together or contained. 'Table' can be used to mean *person/people sitting at the table* because the table groups the people into a single category. Consider the expression, "the *x* is boring" (cf. Murphy 2002: 438). For *x* we may substitute words for *containers* when we mean *the contents of the container*. We may substitute any of the following: 'DVD', 'newspaper', 'CD', 'video', 'TV', 'file', 'website', 'room', 'bottle of wine', 'Christmas stocking', 'book', and so on. In each case, we mean not *x* itself is boring but the contents of *x* are boring. We submit that this is how psychologists use 'concept': they mean the *contents* of the concept—the stored bodies of knowledge.

A similar pattern holds for *x* meaning *representation of x*. We can use 'DVD' to refer to *the title associated with the DVD*. Suppose someone were to hand you a list of titles and say, "Pick one of those DVDs; we'll watch it tonight". No one in their right mind would reply, "What you should have said is 'pick one of those *titles* standing for a DVD containing the movie of that title". Nobody thinks that a printed title is a DVD. Similarly, if given a catalogue and asked to "underline your favourite *x*s", where *x* could be 'newspaper', 'CD', 'video', 'TV', 'file', 'website', 'room', 'bottle of wine', 'Christmas stocking', 'book', and so on, what is meant is not *underline x itself* (how could one underline a website?) but *underline the label of, or proxy for, x*. We submit that this is how philosophers use 'concept': they mean the *label of or proxy for* the container of stored information. A DVD is a container of stored knowledge that may have an associated label. A concept is a container of stored knowledge that may have an associated label.

Container talk can rapidly switch between referring to the *containers*, referring to the *contents of the containers*, and referring to the *labels of the containers*. In normal communication these switches rarely cause a problem, but pitfalls lurk if one attempts to read off from this practice which entities stand behind the talk. Thus far, our intention has been to observe that container talk sanctions two forms of linguistic shorthand: container talk may refer to the contents of containers or to the labels of containers. With this point in mind, one should not be surprised that, if concepts *were* containers, philosophers and psychologists would use 'concept' in these two different ways even if they could agree that a concept really *is* just one thing: *a container of stored information pertaining to a single category*.

6. *A concept is a container*

But why think that a concept is such a container? We believe that the view has much in its favour. We introduce the view in this section, and in Sections 7 and 8 we argue that under this view concepts satisfy both of Machery's desiderata.

Concepts are often described as the entities that furnish the mind.⁵ Our proposal is that concepts are not the furnishings but the rooms: the containers of the furnishings. What furnishes the mind are pieces of information. Concepts contain and group together those pieces of information in pertinent ways. Rooms contain furnishings. Concepts contain pieces of information or knowledge. Rooms do not contain types of room. Concepts do not contain types of concepts. Concepts contain pieces of information (encoded in the form of exemplars, prototypes, theories), but they are not identical to those pieces of information. Rooms contain pieces of furniture but they are not identical to those pieces of furniture. Inside rooms, we gather together pieces of furniture that belong together. Inside concepts, we gather together pieces of information that belong together. Information, such as “cows have udders” is not anyone’s COW concept, but it could be contained within someone’s COW concept. The same goes for an exemplar (for example, a visual image) of a cow, or any other type of information one might associate with cows.

Notice that the furniture inside a room is important for determining the kind of room it is. The furniture in a bedroom determines the kind of room it is. Similarly, the information inside a concept determines the kind of thing that the concept refers to. Imagine that one labels the doors of one’s rooms so that someone can tell which room they are without having to look inside. The label might be a useful proxy, but it does not determine the kind of room it is. If we label the dining room ‘bedroom’, that does not turn it into a bedroom. However, if we were to swap all the furniture from the dining room with the furniture in the bedroom, even if we leave their labels intact, the dining room would be the bedroom. Rooms are individuated functionally by their contents, not by their labels. As we will see, concepts are individuated functionally by the information that they contain.

A drawback to the room analogy is that buildings do not contain upwards of ten thousand rooms each dedicated to a different purpose. Files are a better model for concepts.⁶ Files are containers of informa-

⁵ Locke famously says that the mind is furnished with ideas: “Let us then suppose the Mind to be, as we say, white Paper, void of all Characters, without any *Ideas*; How comes it to be furnished? [...] Whence has it all the materials of Reason and Knowledge?” (Locke 1975: 104). Ideas, for Locke, are the constituents of propositional thought and permit us to make categorisation, typicality, and inferential judgements. Later, Hume and Reid also describe ideas as the “furniture of human understanding” (Hume 1976: 180; Reid 1983: 116). For Locke and his contemporaries, ‘furniture’ meant *that which furnishes* in the sense of stocking or equipping some container; it was not restricted to tables and chairs (Lewis 1967; Pasanek 2015). The furniture in the quotation, for example, are written characters. Our point is that a concept should always be identified with the container (the room, the page) not with its contents (tables, chairs, written marks). The mind contains rooms (concepts) which contain furniture (ideas).

⁶ Margolis (1998); Prinz (2005); Papineau (2006); Fodor (2008); Recanati (2013) suggest mental files are a helpful model when discussing concepts.

tion. A file is suited to gathering together information pertaining to a single topic. One might imagine having thousands of files that contain pieces of information on specific topics. Files sometimes have labels that help them to be retrieved or referenced more easily. So, we argue, do concepts: concepts gather together pieces of information that pertain to a single category and they may have associated labels that provide an easy way for the rest of the cognitive system to get hold of them.

Unlike both rooms and files, concepts are functional, not spatial, containers. The pieces of information inside a concept are not located within some specific spatial boundary. The pieces of information inside a concept are grouped together, and distinguished from other information in the cognitive system, by a functional relation. To see the contrast between a spatial and a functional container, compare how a movie is stored on a DVD with how it is stored over the internet using BitTorrent. In the case of a DVD, the information that comprises the movie is spatially contained inside (it is "on") a physical container: a discrete, spatially bounded, storage disc. In the case of BitTorrent, the container is a functional, not a spatial, container of the same information. Torrents divide up the information of the movie into small chunks and each chunk is stored on a different computer or host. These host computers may be spatially scattered around the world and they may change rapidly over time. A tiny '.torrent' file contains data that indicates how to retrieve all the various chunks of information so that they can be viewed as a coherent movie. The '.torrent' file specifies a functional relationship between the pieces of information that groups them together and distinguishes them from other pieces of information on the internet. A torrent contains a movie, but it does not spatially contain it, or at least not in a sense that has any specific spatial boundary. In the same way, a concept contains pieces of information but it does not spatially contain them. A concept contains information because that information satisfies some functional relationship that allows it to be found and coherently deployed by the brain under the right circumstances. We will see some proposals for this functional containment relation in the next section.

An issue we can discuss here is that the idea of an empty concept or container of stored information might seem to be difficult to make sense of. There are two issues here. First, one might argue that a DVD can exist with nothing stored on it, but it seems odd to say that one has a concept that does not pertain to any category—that does not contain any inferential or recognitional information. We would argue that this oddness is because concepts are functional containers like torrent files. It makes no sense to open a torrent file without information to store in it and similarly it makes no sense to "open" a new concept until there is a category it is targeting (e.g. some entity in the world that the creature has encountered). However, we submit, that if concepts were spatial containers like little DVDs in the brain it would not be

odd for the brain to pre-fabricate empty ones awaiting allocation to a particular category. It does not affect our argument if one prefers to call such empty containers, for example, 'proto-concepts' and reserve the term 'concept' for when they have a certain amount of information contained in them. The second issue is related. When do we possess a concept or grasp a concept? Can we have a concept when there is only a small amount of information in the container or must we be able to fully recognize, infer, and form associated propositions? We prefer to distinguish between fledgling and fully-fledged concepts (they are both types of concept) but as with the first issue it does not affect our argument that concepts are containers. One may, if one prefers, distinguish between proto-concepts and concepts depending on how much information is in the container. To sum up, on both issues, there can be differences of opinion about when a concept (truly called) comes into existence but this does not harm our proposal that concepts are containers.

We now turn to Machery's two desiderata: the judgement desideratum and the propositional desideratum. Machery claims that no single entity, a concept, meets both the judgement desideratum and the propositional desideratum. We disagree: concepts as containers meet both. Let us consider each desideratum in turn.

7. Meeting the judgement desideratum

The judgement desideratum says that a concept should permit us to make categorisation, typicality, and inferential judgements. If you have a (fully-fledged) concept that satisfies the judgement desideratum, all else being equal, you should be able to do two things: (T1) identify to which category a relevant object belongs or is typical; (T2) apply information you have stored about that category to the current instance. To do this, you need to draw on two sorts of information from somewhere inside your cognitive system: (I1) information pertaining to recognition/identification of instances as belonging to, or typical of, that category; (I2) information pertaining to that category that is relevant for making inferences about and interacting with that instance. Either I1 or I2 in the absence of the other would be insufficient for you to have a concept that meets the judgement desideratum. It would be of no use to you to be able to identify a cow if you could not bring to bear information when you have identified a cow. And it would be of no use to you to draw an inference about cows if you do not know how to identify a cow. To satisfy the judgement desideratum, you need both I1 and I2.

Consider now that you likely have *many* pieces of (sometimes incompatible) information about cows under the headings I1 and I2 inside your cognitive system. You are likely to have many pieces of information that are relevant to identifying a cow. You are also likely to have many pieces of information that are relevant to drawing inferences about cows. What distinguishes the pieces of information that fall in-

side your COW concept from those that fall outside is some functional condition. Machery proposes that this is something like the default, or preferentially available, information you use for solving tasks T1 and T2. Which pieces of information get recruited, normally, rapidly, by default, to solve T1 and T2? This condition draws a functional boundary around certain pieces of information inside your cognitive system. That functional boundary unites certain pieces of information and distinguishes those pieces from others in your cognitive system.

Whether this is the right functional containment relation for concepts may be questioned. Machery proposes that the functional containment relation draws a boundary based around the pieces of information in the cognitive system that the agent uses “by default” in solving T1 and T2 (Machery 2009: 11). This condition is explained in terms of the idea of default inference in artificial intelligence: an inference that is normally drawn by the agent, except when some specific additional information is provided that defeats it. Machery equates this with information that, for the agent, is preferentially available, presumptively taken to be relevant, and spontaneously comes to mind (Machery 2009: 11–12). Prinz suggests that the relevant functional containment relation for concepts is that of information being “under organismic control”: pieces of information inside a concept should be capable of being retrieved and manipulated intentionally (Prinz 2004: 45). Dennett suggests that the functional containment relation is that of being available to “call to mind”: the relevant pieces of information should be capable of being objects of the agent’s second-order, personal-level thoughts (Dennett 1996: 157). We do not wish to argue for the advantages of one specific functional containment relation for concepts. Our claim is merely that any theory of concepts will invoke some functional containment relation or another: it must distinguish between those pieces of information under I1 and I2 in your cognitive system that fall inside your concept from those that fall outside. Concepts are by their nature in the business of containment.

Notice that your COW concept is not, and cannot be, *identical to* the pieces of information that we claim are inside your COW concept. Those pieces of information need to be grouped together to distinguish them from other pieces of information about cows in your cognitive system. Merely enumerating *those specific pieces of information* would not suffice to specify your COW concept. And nor would it be necessary: your COW concept could involve any number of different specific pieces of information. Indeed, the specific pieces of information associated with your COW concept are likely to change over time as you learn more about cows. The pieces of information inside a COW concept are neither sufficient nor necessary for having the concept. Why those pieces of information are important is that they are grouped together by a functional relation that separates them from other pieces of information and hooks them up to behaviour in the right way (for example,

to drive your response in solving T1 and T2). The information inside your COW concept is not your COW concept. Your COW concept is a container that holds specific pieces of information about cows. A more basic concept such as RED will contain pieces of information such as exemplars and prototypes which permit inference and categorization. These pieces of information are not concepts and do not need to be relationships to anything else. Of course, as a creature becomes more sophisticated concepts may contain relationships to other concepts (e.g. RED is a COLOUR).

Interestingly, support for this view comes from Machery himself. Initially, Machery says that for a psychologist, “a concept of x is a body of information about x that is stored in long-term memory” (Machery 2009: 4, emphasis ours). Later, however, he says revealingly:

[...] the knowledge that is *stored in a concept* x is preferentially available when we think reason and so on about x . So to speak it spontaneously comes to mind. By contrast, the knowledge about x that is *not stored in a concept of* x is less available—it does not spontaneously come to mind. (Machery 2009: 11–12, emphasis ours)

Note a shift from saying that a concept *is* a body of knowledge to saying that a concept *contains* a body of knowledge. This is precisely the kind of shift we would expect with container talk. Someone might use ‘concept’ to refer to the container or to its contents. This may not cause confusion in some quarters, but the difference between the two matters. Merely *having a body of information* does not suffice for having that concept. This motivated us, and is motivating Machery here, to switch to a container view about concepts. A concept is not, and cannot be, *a body of information pertaining to a single category*. That information must in addition satisfy functional constraints on use that distinguish it from other information in the cognitive system that pertains to the same category. Only information that satisfies those functional criteria is *stored inside* the concept. The concept is a container and the bodies of information are stored inside it. The container permits the categorisation, typicality, and inferential judgements that interest psychologists. Those mere bodies of information do not. Concepts as containers satisfy the judgement desideratum. Concepts as bodies of information do not.

One might have the concern that when we specify that concepts are containers that pertain to single categories all of the work is being done by the term ‘category’, which we have left undefined. What determines what a single category is? We believe that one of the major advantages of container theory is that it allows us to say very simply what a category is: where there is a single container, there is a single category. It is not the case that the system must first decide what constitutes a single category and then open a container for it. Instead, because containers have boundaries (information is either in the container or not), the intentional content is fixed by the contained information. The general idea is that a new container is “opened” when the system encounters some information that does not fit (according to some similar-

ity metric e.g. in exemplar theory or prototype theory) into any existing container. Information accumulates inside the container when further information that passes the similarity metric is encountered. The fact that our concepts generally accord with what we intuitively consider to be good and useful categories is explained by I1 and I2 above: things you identify as the same are the things that drive successful inferences and action and vice versa.

8. *Meeting the propositional desideratum*

The propositional desideratum is that a concept should be capable of being used as a constituent of propositional thought. Concepts allow us to contemplate and make judgements about complex states of affairs and events rather than about single categories. We claim that the same entity that satisfies the judgement desideratum also satisfies the propositional desideratum.

As Machery observes, psychologists tend to focus on categorisation, typicality and inferential judgements, whereas philosophers tend to focus on how concepts are combined to make (a potentially unbounded number of) complex propositional thoughts. Philosophers have a model of what concepts might be like to enable complex propositional thought: word-like symbols. If concepts are word-like symbols with a fixed meaning, then those symbols could be strung together, with a recursive syntax and compositional semantics, to express an unbounded number of complex propositional thoughts. This is Fodor (1975)'s language of thought hypothesis: we form complex propositional thoughts by combining atomic symbols inside our heads using a recursive syntax and compositional semantics. Fodor claims that the word-like symbols simply *are* concepts and that they directly stand for categories in the world (Fodor 1998; Fodor 2008). Your COW concept is a word-like atom inside your cognitive system that refers to *cows*. In contrast, the view we put forward is that these word-like atoms are not concepts but *proxies for* concepts (i.e. proxies for containers of stored knowledge which refer to categories in the world).

Before getting to this, we first need to show something simpler: that concepts as containers satisfy the propositional desideratum. In other words, concepts as containers *could* be the constituents of our propositional thought. In the course of establishing this, it will become clear that there is no logical *necessity* for labels or word-like atoms to be involved in propositional thought at all.

We can make progress with a simple example. Let us imagine for the moment that the atoms in question are the building blocks that young children play with. Children's building blocks sometimes have letters on them and they can be combined to form words. Suppose that our building blocks have words on them and that they can be combined to form sentences. The crucial idea here is that each building block is an atom in the sense that each block is the smallest unit that can par-

take in a construction. Within a construction, a block retains its identity and it can be manipulated or exchanged as a single item. So, for example, one could remove a block with the label 'cow' and substitute it with a block with the label 'dog' without having to do anything to the rest of the blocks.

Now imagine that our building blocks are containers (for example, crates with lids). For each concept, there is a separate building block. Inside the COW block is stored all of our (preferentially available) pieces of information pertaining to cows: all the information that gives us the ability to recognise instances and make inferences about cows needed in solving T1 and T2. Each block is a container that contains pieces of stored knowledge pertaining to a single category. Observe that each block now satisfies *both* the judgement and propositional desiderata. One can use the stored knowledge in the blocks to make categorisation, typicality, and inferential judgements. One can also use the atomic blocks as the constituents (smallest units) of propositional thoughts: just line them up so that they compose a sentence.

The fact that our blocks contain pieces of information does not affect their ability to be atoms that compose into larger wholes, but it does mean that the blocks' labels are no longer necessary. As we first considered them, each building block is labelled in a way that distinguishes it from the other blocks. It would be pointless to form a composition with indistinguishable building blocks—how would you know which complex thought was expressed? But if the blocks contain pieces of information pertaining to a single category, this removes the need for labels. The contents can be used to tell the blocks apart. You can look inside a block to see exemplars (along with other pieces of information used for categorical judgement) of the kind of thing that a particular block refers to. Note that when you form a complex of blocks, you are not composing each individual piece of information inside each block (you do not remove and compose every piece of information inside the COW block with every piece of information inside the RUN block when you form COWS RUN). Rather, you compose the entire containers (the discrete blocks), each of which pertains to a category. You may use the contents of the blocks to justify or interpret or illustrate the resulting statements. A container satisfies both Machery's desiderata: a concept contains information that can be used to identify instances of that category and a concept can be used as a constituent in complex propositional thought.

In our example, the containers were imagined to be spatial containers: they contain by having items placed inside a crate. The composition relation was also imagined to be spatial: place the blocks one after another—line them up in a row—to express a complex proposition. But concepts are not spatial containers. And as Fodor (1975) observes, the composition relation in propositional thought is unlikely to be spatial. Brains do not place information inside little crates in the head and

they do not move those crates around to make a propositional thought. Brains use functional properties for both concept containment and concept composition. We have seen a number of proposals for the brain's functional containment relation. The concept composition relation is also unknown. Current thinking is that complex propositional thought involves individual concepts being tokened in working memory or some similar central workspace.⁷ The neurocomputational properties involved are unclear.⁸ But the fact that there is more work to be done here does not affect our specific point. No matter how one composes those containers—be it arranging them in a spatial row or via some functional relation—those containers can also be used to express a structured proposition in which the containers are constituents. Moreover, the containers do not need to be individuated by labels: they are already individuated by, and have their intentional content fixed by, what is inside them. As discussed above, concepts/containers will only come into being when there is some information to put in them but just how precisely the intentional content is fixed may depend on whether the concept is a fledgling one or a fully-fledged one.

Containers (with or without associated labels) can be the constituents of propositional thought. But some reflection shows that using concepts as containers to express propositional thought would probably not be an efficient way for a cognitive system to operate. Let us return to the example of the building blocks. To form our building blocks into complex wholes requires a space in which to order them. If we want to form a thought with the concept COW in it, we need to transport the COW block or a copy of it (with all of its stored contents, the many pieces of information pertaining to the category COW) into that workspace. Similarly in the case of the brain, to form a propositional thought would require bringing each concept with all its associated contents into working memory or some similar central workspace. One thing we know about working memory or central workspaces in human thought is that it has limited capacity.⁹ Copying or transporting an entire concept with all its associated contents would likely be inefficient as an information processing strategy. A more efficient solution would be to token in the workspace *labels of* (or, to borrow a notion from computer science, *pointers to*) the container. One could compose the labels of, or pointers to, concepts as proxies for the real concepts. That would do just as well for the purpose of forming complex propositional thoughts. Note that these labels or pointers stand proxy for *concepts* (containers

⁷ Baars (1988); Baars (1997); Carruthers (2014); Carruthers (2015); Dehaene and Changeux (2011); Dehaene and Naccache (2001); D'Esposito and Postle (2015); Fodor (2008); Oberauer and Hein (2012); Penn, Holyoak, and Povinelli (2008); Shanahan and Baars (2005).

⁸ Although see Piantadosi, Tenenbaum, and Goodman (2016).

⁹ Baddeley (2010); Baars (1997); Cowan (2000); Ma, Husain, and Bays (2014); Miller (1956).

of stored knowledge) not for *categories in the world* as Fodor and other LOT theorists propose.

There is no logical necessity that the constituents of complex propositional thoughts be labels or word-like atoms. Containers can play the role of conceptual atoms in propositional thought. However, containers are bulky: they do not have the desirable features of being easily transportable or easily copyable. For that reason, it is likely that labels of, or pointers to, concepts are composed to form propositional thoughts.¹⁰ The labels or pointers are proxies for the container of stored information pertaining to a single category. As we saw above, the labels or proxies by themselves give no understanding. Understanding comes from what is contained in the container. Mental words or conceptual labels should be seen as cues for accessing those contents. Using labels or proxies, rather than containers packed with information, makes composing concepts into complex constructions, and taking those complexes as the subject matter of further thought, easier.

Concepts as containers satisfy the propositional desideratum. They can either satisfy the desideratum directly by being the entity composed in propositional thought. Or they can satisfy it by having associated labels or pointers which are composed in place of the concepts themselves. Notably, the labels or pointers are not, as Fodor has it, concepts. They are proxies for concepts as containers. It is concepts as containers (which may or may not have associated labels) that have the content-fixing properties and that satisfy the propositional desideratum.

8. Conclusion

According to Machery, two distinct and independent types of entity stand behind concept talk in philosophy and psychology. We have argued that this is not the case. A single entity stands behind this talk: *a container of stored knowledge pertaining to a single category*. This entity (which has associated content and may have an associated label) satisfies both the judgement desideratum and the propositional desideratum for concepts. The linguistic divergence between philosophers and psychologists in their use of 'concept' that motivates Machery's view is

¹⁰ There is much empirical evidence that monkeys use pointers in working memory areas to keep track of conceptual information in associative areas in delayed match to sample tasks (Miller et al. 1996; Fuster 1995; Goldman-Rakic 1995 amongst others). In these tasks the subject is shown a number of colours on a screen with one colour being indicated as being the "reward" colour. The colours then disappear from the screen during a delay period. They reappear this time without an indicator of the "reward" colour. The monkey must retain information about which colour to select during the delay period in order to select the right one. The consensus amongst the researchers is that the monkeys use loops between cells in working memory and cells in associative areas (that previously had been found to activate in the presence of e.g. the colour red) to track which colour it needed to indicate after the delay. In other words, the conceptual information was not imported into working memory; instead, cells in working memory pointed to where the information was.

explained by a general property of container talk: container talk can rapidly switch between referring to the *contents* of the container and to the *label* of the container. Despite the pattern of linguistic use that Machery describes, philosophers and psychologists should agree that a concept is just one thing: a container of stored information. They might disagree about specific features of that thing: about the functional containment relation (is “preferential availability” the right relation?) or about the functional composition relation and its implementation (how is composition done in the brain?). But these disagreements are substantive disagreements, not verbal disagreements. Treating concepts as containers untangles Machery’s bind. We arrive at a desirable outcome: philosophers and psychologists share a common, rationally explainable, interest in concepts.

Acknowledgements

We are grateful for helpful comments on earlier versions of this paper from Edouard Machery, Jesse Prinz, and Andy Clark.

References

- Baars, B. 1988. *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.
- Baars, B. 1997. *In the Theater of Consciousness*. Oxford: Oxford University Press.
- Baddeley, A. 2010. “Working Memory.” *Current Biology* 20: R136–R140.
- Carruthers, P. 2014. “On Central Cognition.” *Philosophical Studies* 170: 143–62.
- Carruthers, P. 2015. *The Centered Mind*. Oxford: Oxford University Press.
- Chalmers, D. J. 2011. “Verbal Disputes.” *Philosophical Review* 120: 515–566.
- Cowan, N. 2000. “The Magical Number 4 in Short-Term Memory: A Reconsideration of Mental Storage Capacity.” *Behavioral and Brain Sciences* 24: 87–185.
- Dehaene, S., and J.-P. Changeux. 2011. “Experimental and Theoretical Approaches to Conscious Processing.” *Neuron* 70: 200–227.
- Dehaene, S., and L. Naccache. 2001. “Towards a Cognitive Neuroscience of Consciousness: Basic Evidence and a Workspace Framework.” *Cognition* 79: 1–37.
- Dennett, D. C. 1996. *Kinds of Minds*. New York: Basic Books.
- D’Esposito, M., and B. R. Postle. 2015. “The Cognitive Neuroscience of Working Memory.” *Annual Review of Psychology* 66: 115–142.
- Fodor, J. A. 1975. *The Language of Thought*. Sussex: The Harvester Press.
- Fodor, J. A. 1998. *Concepts*. Oxford: Blackwell.
- Fodor, J. A. 2008. *LOT2: The Language of Thought Revisited*. Oxford: Oxford University Press.
- Hume, D. 1976. *The Natural History of Religion and Dialogues Concerning Natural Religion*. Edited by A. W. Colver and J. V. Price. Oxford: Oxford University Press.

- Lewis, C. S. 1967. *Studies in Words*. 2nd ed. Cambridge: Cambridge University Press.
- Locke, J. 1975. *An Essay Concerning Human Understanding*. Edited by P. H. Nidditch. Oxford: Oxford University Press.
- Ma, W. J., Husain, M. and Bays, P. M. 2014. "Changing Concepts of Working Memory." *Nature Reviews Neuroscience* 17: 347–356.
- Machery, E. 2009. *Doing Without Concepts*. Oxford: Oxford University Press.
- Machery, E. 2010. "The Heterogeneity of Knowledge Representation and the Elimination of Concept." *Behavioral and Brain Sciences* 33: 231–244.
- Margolis, E. 1998. "How to Acquire a Concept." *Mind and Language* 13: 347–369.
- Margolis, E., and Laurence, S. 2010. "Concepts and Theoretical Unification." *Behavioral and Brain Sciences* 33: 219–220.
- Miller, G. A. 1956. "The Magical Number Seven Plus or Minus Two: Some Limits on Our Capacity for Processing Information." *Psychological Review* 63: 81–97.
- Murphy, G. L. 2002. *The Big Book of Concepts*. Cambridge: MIT Press.
- Oberauer, K., and Hein, L. 2012. "Attention to Information in Working Memory." *Current Directions in Psychological Science* 21: 164–69.
- Papineau, D. 2006. "Phenomenal and Perceptual Concepts." In T. Alter and S. Walter (eds.). *Phenomenal Concepts and Phenomenal Knowledge*, Oxford University Press, 111–144.
- Pasanek, B. 2015. *Metaphors of Mind*. Baltimore: Johns Hopkins University Press.
- Penn, D. C., Holyoak, K. J. and Povinelli, D. J. 2008. "Darwin's Mistake: Explaining the Discontinuity Between Human and Nonhuman Minds." *Behavioral and Brain Sciences* 31: 109–178.
- Piantadosi, S. T., Tenenbaum, J. B. and Goodman, N. D. 2016. "The Logical Primitives of Thought: Empirical Foundations for Compositional Cognitive Models." *Psychological Review* 123(4): 392–424
- Piccinini, G. 2011. "Two Kinds of Concept: Implicit and Explicit." *Dialogue* 50: 179–193.
- Prinz, J. 2004. *Gut Reactions*. Oxford: Oxford University Press.
- Prinz, J. 2005. "The Return of Concept Empiricism." In H. Cohen and C. Lefebvre (eds.). *Handbook of Categorization in Cognitive Science*. Amsterdam: Elsevier, 679–699.
- Recanati, F. 2013. *Mental Files*. Oxford: Oxford University Press.
- Reid, T. 1983. *Inquiry and Essays*. Edited by R. E. Beanblossom and K. Lehrer. Indianapolis, IN: Hackett.
- Shanahan, M., and Baars, B. 2005. "Applying Global Workspace Theory to the Frame Problem." *Cognition* 98: 157–76.