

Machery's Alternative to Concepts and the Problem of Content

Bernardo Pino¹  · Bernardo Aguilera^{2,3}

Received: 21 September 2015 / Accepted: 11 May 2017 / Published online: 7 June 2017
© The Author(s) 2017. This article is an open access publication

Abstract Edouard Machery has argued that the notion of concept should be eliminated from scientific theorising about cognition on the grounds that what psychologists call concepts do not form a natural kind and that keeping this notion would encumber scientific progress. His view is that the class of concepts really divides into three distinct yet co-referential kinds of *bodies of knowledge* typically used in distinct cognitive processes. The main purpose of this paper is to challenge Machery's eliminativist conclusion on the grounds that his alternative to concepts fails to be combined with any prominent account of content determination. As a case of philosophical scrutiny of a psychological proposal, the paper also intends to show how philosophical and psychological interests in concepts can be brought together.

1 Introduction

Within psychology and the philosophy of mind, concepts are recognised as essential elements of psychological explanation. Human beings deploy conceptual abilities to coordinate their behaviour with respect to the features of the environment their

Bernardo Pino and Bernardo Aguilera contributed equally to this work.

✉ Bernardo Pino
bpino@ug.uchile.cl
Bernardo Aguilera
baguilera@udd.cl

¹ Department of Philosophy, University of Sheffield, 45 Victoria Street, Sheffield S3 7QB, UK

² Centro de Bioética, Facultad de Medicina, Clínica Alemana Universidad del Desarrollo, Santiago, Chile

³ Facultad de Medicina, Universidad de Chile, Santiago, Chile

concepts represent, and psychologists attempt to formulate generalizations about these conceptual abilities to explain and predict people's behaviour. Likewise, a psychological theory of concepts should explain a myriad of interrelated phenomena such as, for example, the acquisition of concepts, learning, categorization, and reasoning processes in general. Thus, it should come as no surprise that 'concept' is a theoretical term that plays a unifying role in psychological explanation.

Another reason why concepts are so significant for psychological explanation is that they provide for an account of the intentionality of thoughts. What this means is that they provide an explanation for the capacity that thoughts have to represent objects, individuals, properties and relations (hereafter, just objects). This capacity is customarily explained by appealing to the notion of content. The notion of content has been the focus of extensive philosophical debate, but what is clear is that concepts are traditionally understood as content-bearing, to the extent that providing an account of content has been considered a desideratum for a theory of concepts (Fodor 1998; Prinz 2002).¹

However, the very existence of concepts has recently been put into question. In his *Doing without Concepts*,² Machery claims they should be eliminated from the theoretical vocabulary of psychology. His suggestion is that different kinds of *bodies of knowledge* can do the job of accounting for the same cognitive phenomena that those working in the psychology of concepts have tried to explain by appealing to the notion of concept. Machery's central argument is that these classes of bodies of knowledge are so heterogeneous that they cannot be subsumed by the general notion of concept. He concludes that this latter notion should not be considered as a genuine psychological kind.

Our purpose in this paper is to critically examine Machery's notion of bodies of knowledge, especially with respect to content. Given that a theory of concepts has traditionally provided psychology with a notion of content, we should expect that Machery's alternative to concepts could be supplemented with an account of content as well. Rejecting concepts in favour of a kind recognised in psychological theory might only be costly for philosophical interest in concepts. Giving up the explanatory functions subserved by concepts would be a heavier cost, and one which psychologists would have to pay as well as philosophers. So Machery's intention must be to do without concepts, yet still do the required explanatory work, using bodies of knowledge for the purpose. In order to do this, those bodies of knowledge will need to be contentful in something like the way that concepts were supposed to be.

However, in his book he does not provide us with an account of content for bodies of knowledge. In this paper we contend that there is indeed no available notion of content that can be combined with Machery's heterogeneous classes of bodies of knowledge. Thus, in Sect. 2, we begin by presenting Machery's case for the elimination of concepts and his proposed notion of bodies of knowledge. In

¹ In addition to helping to explain how we manage to interact efficiently with external entities, content is essential for an account of the semantics of language. Following Grice (1957), many philosophers maintain that linguistic expressions have meaning only by virtue of deriving their contents from that of the concepts of language users.

² Unless otherwise indicated, references to Machery's work in this paper correspond to this book.

Sect. 3, we discuss what he says about content. Section 4 sets forth the most prominent approaches to content adopted in philosophy. Then, in Sect. 5, we assess whether or not Machery's view of bodies of knowledge is consistent with any of these approaches to content. Section 6 discusses additional challenges and the prospects for Machery's alternative to concepts. Ultimately, we conclude that none of the examined accounts of content can be made to work for Machery's purposes. We discuss some implications of this conclusion for both his case for concept eliminativism and the sharp division of labour he contends for between psychologists and philosophers working on concepts.

2 Machery's Alternative to Concepts: Heterogeneous Bodies of Knowledge

Concept theorists have posited different models of concepts (notably, *prototypes*, *exemplars* and *theories*) to explain a variety of cognitive phenomena (e.g., concept learning, inferences, categorization, language understanding, etc.) and there is significant consensus about the explanatory strengths of the dominant models with respect to some of those phenomena.³ While some cognitive scientists have proposed that different conceptual models may co-exist within a single framework (e.g., Murphy 2002; Laurence and Margolis 1999), the study of concepts has typically aimed at the development of a single theory of concepts that is able to account for all the relevant phenomena. Machery (2009, 2010b) has called the assumption that there must be such single general theory of concepts "the Received View."

An interesting reaction to this view is Machery's *Heterogeneity Hypothesis*, that is, the view that the class of concepts divides into several distinct (or heterogeneous) kinds which Machery takes to be different bodies of knowledge (BoKs, for short). On this view, a BoK is information about a certain category of objects that, out of the whole knowledge about this category that may be stored in our long-term memory, is preferentially retrieved to be used in cognitive processes underlying our higher cognitive competences. Machery distinguishes between *default knowledge* (i.e., knowledge about a given category that is preferentially retrieved) and *background knowledge* (i.e., knowledge about the same category that isn't preferentially retrieved). While background knowledge is said to be context-dependent, default knowledge is said to be context-independent because it can be retrieved in all contexts.

In his book,⁴ Machery looks at the kinds of properties that are said to be scientifically relevant for characterising concepts in psychology. Essentially, these properties involve the nature of the information that is constitutive of concepts (e.g., statistical or causal information); the nature of the processes that use concepts (e.g.,

³ For a general evaluation of main theories of concepts, see Laurence and Margolis (1999), Margolis and Laurence (2014), Prinz (2002), Chap. 3–4 and Machery (2009), Chap. 4. For a detailed account of findings in experimental cognitive psychology, see Murphy (2002).

⁴ See also Machery (2005, 2010b, 2015).

similarity-based processes); the nature of the vehicles of concepts (e.g., perceptual representations); the brain areas involved in processing concepts; and processes of concept acquisition. Given those properties, he provides substantial empirical evidence to support the claim that prototypes, exemplars and theories amount to distinct kinds of BoKs with very little in common. This view is in contrast with the traditional idea that the class of concepts forms a single natural kind with many common properties. Then, by appealing to the methodological doctrine that, in any given science, theoretical terms that fail to pick out a natural kind should be eliminated from the corresponding scientific classificatory scheme, Machery argues that the term ‘concept’ should be displaced from psychological taxonomy because concepts do not form a single psychological kind capable of supporting scientific generalisations.⁵ He therefore defends the view that the existence of these distinct types of BoKs provides pragmatic grounds for the elimination of the theoretical notion of concept.

To emphasise, Machery’s central claim for the heterogeneity hypothesis is that the models of concepts posited by Prototype, Exemplar and Theory theories really describe the role of three different kinds of BoKs in distinct and independent cognitive processes. If this is right, then, contrary to the Received View, the explanatory advantages associated with any one of these theories should not be viewed as evidence against the other theories. Hence, Machery’s contention that, for each category (i.e., a certain type of objects), an individual typically has several concepts, which, in his terminology, correspond to distinct co-referential kinds of BoKs.

Hybrid theories of concepts are the nearest rivals to Machery’s view. Proponents of these theories (e.g., Anderson and Betz 2001; Keil 1989) acknowledge the existence and cognitive relevance of various psychological structures (or BoKs), but do not acknowledge them as distinct kinds of concepts. Instead, or so they argue, concepts are divided into different parts. Machery thinks those who advocate hybrid theories have failed to explain what it means for different BoKs to be parts of the same concept. By contrast, he provides us with two criteria to determine when someone uses distinct concepts (i.e., distinct BoKs) that refer to the same category of objects.

According to the first criterion, ‘connection’, two BoKs are connected when the retrieval of one of them (e.g., *that water is transparent*) from long-term memory in order to be used in a given cognitive process (e.g., categorisation) is *necessarily linked* to the retrieval of the other (e.g., *that water is made of H₂O*) to be used for other purposes. In turn, according to the second criterion, ‘coordination’, two BoKs are coordinated when they do not yield conflicting outcomes, notably inconsistent categorisation judgments (e.g., the judgment that whales are fish and the judgment that whales are not fish). Given these criteria, as Machery contends, two BoKs are parts of the same concept in case they are both connected and coordinated. Otherwise, they belong to two distinct concepts. He goes on to argue that hybrid

⁵ To support the claim that the class of concepts can’t qualify as a natural kind, Machery endorses the relevant view defended by Boyd (1999). According to this view, “natural kinds are classes whose members share many scientifically important properties in virtue of one or several causal mechanisms” (Machery 2010b, p. 200).

theories can't be right about concepts being divided into several BoKs because it is a mistake to assume (as hybrid theory proponents do) that, to the extent that those BoKs are identified with co-referential prototypes, exemplars and theories, these BoKs are both connected and coordinated.

Finally, as Machery points out, his case for the Heterogeneity Hypothesis is fully empirical and it therefore provides room for potential falsification such as, for example, future discoveries regarding brain location of concepts that may challenge the claim that his BoKs have little in common. At the same time, Machery accepts the possibility that prototypes, exemplars and theories do not exhaust the different kinds of BoKs we have about one and the same category.

3 What Does Machery Say About Content?

In this book, I will use the term “knowledge” as psychologists do. By “knowledge”, psychologists mean any contentful state that can be used in cognitive processes (p. 8).

In *Doing Without Concepts* Machery understands his distinct kinds of BoKs as mental states with content and indeed explicitly exploits language of aboutness when characterising each of them as a type of “body of knowledge about x that is stored in long-term memory” (p. 12). However, he does not develop a theory of content that could be suitable for these types of BoKs. One reason for this omission can be extracted from the quote above: in setting forth his view, Machery adopts a psychologist's standpoint regarding the study of concepts, which (among other issues) is concerned with the nature of the knowledge that is constitutive of our contentful BoKs, but is (allegedly) neutral about how it is that these BoKs attain their contents.⁶

What underlies Machery's commitment to this standpoint is his methodological contention that theorising about concepts in philosophy and psychology does not pursue a common agenda. According to him, while philosophers are concerned with the development of a semantic theory for our concepts or thoughts, psychologists are mainly interested in the causal role that concepts play in cognitive processes such as learning, categorizing and reasoning. One way to lay down this contrast is in terms of a division of labour, according to which “while psychologists explain behavior and cognition by ascribing contentful mental states, philosophers explain how people can have contentful states” (p. 31). According to this view, philosophers would have the foundational role of explaining how it is that BoKs have content and psychologists would just take such an account for granted. It is worth mentioning that this view appears to be endorsed by some philosophers working on concepts (cf. Fodor 1994; Peacocke 1992; Rey 1998).

⁶ Throughout his book, Machery identifies the task of accounting for the content of concepts with providing a “semantic theory for our propositional attitudes” (p. 47). This is somewhat misleading, given that he also understands knowledge (and hence any BoK) as not necessarily exhausted by propositional states (e.g., p. 8).

Machery problematizes this view by pointing out that philosophers' theories of content make a feeble appeal to intuitions. Given that philosophers cannot have direct access to the contents of people's thoughts, they proceed by putting forward intuitions about the conditions under which one would find it plausible to ascribe to someone certain propositional attitudes with certain contents. Machery contends that this procedure is unreliable on the basis of some evidence suggesting that people's intuitions regarding the ascription of propositional attitudes are neither uniform (e.g., Hewson 1994) nor free from cultural variation (e.g., Nisbett et al. 2001; Machery et al. 2004). But Machery's conclusion is conservative. He claims that, even though the philosophical task of explaining conceptual content should be taken with caution, it might still be developed successfully in future work. We agree with Machery in this respect, and, in that spirit, the following sections of this paper will attempt to explore what available theory of content could possibly work for Machery's revisionary theory of concepts qua BoKs.

Machery sometimes flirts with a more radical motivation for excluding issues about content in his theory of BoKs. Suppose someone who is persuaded by Machery's eliminativist thesis regarding concepts is also tempted to eliminate the notion of content from psychological explanation as well. After all, if it is possible to make the case that concepts play no role in the scientific study of the mind, it might well be the case that the notion of content is explanatorily idle too. Indeed, some have specifically argued that it is methodologically controversial that reference could count as a constraint or desideratum for psychological theories of concepts (e.g., Machery et al. 2004; Piccinini and Scott 2006).

However, this move would be too fast. Consider that Machery's proposal is to replace concepts with three distinct classes of psychological kinds, namely prototypes, exemplars and theories. These classes of BoKs are supposed to be responsible for the same psychological tasks formerly attributed to concepts, such as making the world intelligible to us by classifying the perceptual input into a collection of meaningful categories. To perform that job, BoKs must have, not just processing properties, but also semantic properties that enable them to refer to external objects. Simply put, as in the case of research on concepts, research on BoKs should not be just concerned with *how* they discriminate and categorise, but also with *what* they discriminate and categorise. To appreciate this point, consider Machery's contention that distinct BoKs can be co-referential, as would be the case when a DOG-prototype, a DOG-exemplar or a DOG-theory are used to categorise dogs. To make sense of this idea, we need to know how each BoK manages to refer to external objects in the first place, viz. how it is that they have contents.

Someone might urge that Machery need not be committed to distinct types of BoKs (about the same category of objects) being entirely co-referential, but, instead, to them having sufficiently overlapping extensions.⁷ Admittedly, psychologists can advance theories about the types of BoKs that an individual deploys when deciding whether certain objects belong in the same category without them having to supply their views of BoKs with a theory of co-referentiality (cf. Machery 2009, Sect. 2). But, at least for the purposes of explaining categorisation, those BoKs would still

⁷ We thank an anonymous reviewer for drawing our attention to this point.

presuppose co-referentiality, and this is something that will have to be spelt out at some point. Since the standard way of explaining what makes certain BoKs (at least partially) co-referential is the fact that they pick up the same contents, we believe that, without an account of content that justifies the presumption that those BoKs are co-referential, Machery's theory would be at least incomplete.

In sum, Machery does not offer an account of content for his proposed BoKs on the main grounds that he is focusing on a psychological theory of concepts, where issues about the nature of the content of our mental states do not figure as explananda. Indeed, he regards psychological and philosophical theories of concepts as separate projects, given that (allegedly) psychologists and philosophers do not talk about the same thing when they use the term 'concept'. On his view, because concept theorists working on each of these two projects are not meant to answer the same questions, they end up talking past each other (p. 51). But, at the same time, in replying to his critics, Machery acknowledges that it is important to have an account of how concepts can have content and how different types of BoKs could be co-referential.⁸ We concur with Machery in this respect and also in recognising that there might be some division of labour between psychologists and philosophers theorising about concepts, such that issues of content are the focus of the latter but not the former.

However, we cannot really be content with the claim that when theorising about concepts philosophers and psychologists are really talking about different things. Since both research fields have common objects of study, their aim should be to keep talking about the same thing when using the same theoretical notion. So, if psychologists suggest that the psychological notion of a concept (or BoK) is inconsistent with what philosophers have traditionally understood by that notion, instead of separating both research projects philosophers should begin to examine whether these new theoretical constructs can fulfil the role concepts have been supposed to play in psychological explanation, such as how it is that we can identify a new object as being a referent of the same sort as objects we have previously encountered (viz. how it is that distinct BoKs can share the same content).

So the question of how BoKs attain their contents is something philosophers would have to address, if not for a strictly psychological account of concepts, at least for a complete and coherent theory of concepts (or BoKs). And if they conclude that no plausible notion of content can be applied to BoKs, then this shortcoming can constitute a reason for resisting Machery's concept eliminativism. At least, this drawback should amount to questioning the plausibility of the notion of BoK as an alternative to that of a concept. Thus, we believe it would be misleading to infer that a more developed theory of BoKs relevant for psychologists is incompatible with an account of conceptual content on grounds of Machery's contention that philosophical and psychological theories of concepts do not answer the same questions. In the following sections, we expect to show that the problem of content invites important interaction between psychology and philosophy.

⁸ See Machery (2010b, pp. 234–235).

4 Notions of Conceptual Content

Concepts are, at a fundamental level, information about the world stored/activated in memory in a format suitable for mental processing. We propose to distinguish two broad approaches to characterize this relation between concepts and the things in the world they are supposed to be about. Call them *Intensionalism* and *Referentialism*.⁹ According to Intensionalism, the content of a concept determines its extension. According to Referentialism, the extension of a concept determines its content.

Broadly characterised, Intensionalism can comprise more than one view of content determination and here we will focus on two. In the first view, content is determined by ‘Fregean senses’ and, in the second, it is determined by ‘conceptual role’. Concerning Referentialism, we will focus on its most popular version according to which the content of concepts is fixed by some sort of informational-cum-nomological relation to its referent (e.g., *Informational Semantics*).

In the following section, we will assess whether or not Machery’s view of BoKs is consistent with any of the two broad approaches to content presented above. Even though Machery takes it for granted that his alternative to concepts amounts to certain contentful states, we want to explore whether his notion of bodies of knowledge could be supplemented by some version of Intensionalism or Referentialism.

5 Is There a Viable Notion of Content for Machery’s View of Distinct Bodies of Knowledge?

5.1 Intensionalism and BoKs

In characterising Intensionalism, we broadly distinguished two main views. One of these views is based on the notion of Fregean senses and the other is based on the idea of conceptual role. Let us examine each of them in turn.

Many have followed the Fregean distinction between *sense* and *referent* and the idea that the expressions of a language can have both of them, such that the sense of an expression contains the *mode of presentation* of its referent. Thus, for example, the terms ‘elephants’ and ‘long-trunked pachyderms’ can be two modes of presentation (i.e., senses) of the same referent in the world. According to Frege, being able to grasp the sense of an expression is what allows us to access its referent. In an attempt to legitimate Frege’s ontology of senses as abstract objects and explain how such objects could intervene causally in people’s cognition, some neo-Fregean philosophers have proposed that grasping a certain sense consists in possessing a particular range of complex abilities or dispositions, which amount to the possession of a concept (Dummett 1991). So to possess the concept of

⁹ This terminology is motivated, though not strictly in accordance, with the relevant distinction proposed by Fodor and Pylyshyn (2015).

ELEPHANT, for example, might consist in having the ability to discriminate elephants from non-elephants and to draw certain inferences about elephants.

Perhaps the most elaborated proposal along these lines is Peacocke's (1992, 1996), who further explored the cognitive capacities that might allow a person to grasp a Fregean sense. He proposed that a concept (i.e., a Fregean sense containing a mode of presentation of a given referent) is individuated by the conditions a thinker must meet to possess that concept, which in turn correspond to tacit knowledge about concept-individuating norms—the norms for making judgements in which the relevant concepts are applied. Thus, since Fregean senses are concepts and they are supposed to be the primary bearers of content, when a thinker satisfies the conditions to possess a concept qua Fregean sense, this concept *ipso facto* has content.

In some passages, Machery appears to be inclined towards something like an intentionalist neo-Fregean view of BoKs. What makes prototypes, exemplars and theories distinct psychological kinds is that each of them is retrieved by default in order to be used in similar yet independent cognitive processes underlying our cognitive competences. Thus, for example, what distinguishes a BIRD-prototype from a BIRD-exemplar is that they are automatically retrieved in order to be used for independent epistemic mechanisms (e.g., prototype-based categorisation and exemplar-based categorisation, respectively). Since these mechanisms explain how thinkers come to satisfy concept individuating-norms, such as those that make it rational to recognise and categorise a concept under certain conditions, the three distinct BoKs about the same referent can then be understood as three distinct modes of presentation. From this perspective, it could be claimed that when using a BoK an individual is grasping a Fregean sense, since entertaining a certain BoK involves satisfying the corresponding possession conditions for entertaining a certain mode of presentation of a given referent.

Prima facie, this strategy would bring in many advantages connected with the neo-Fregean notion of content. First, it would account for the intentionality of BoKs, if it is true that senses are what fix reference. Secondly, it would also explain how BIRD-prototypes, BIRD-exemplars and BIRD-theories could be distinct cognitive mechanisms while, at the same time, they share the same referent. The idea here is that these BoKs could be conceived of as encoding those capacities the neo-Fregean understands as constituting our grasp of different senses of a given referent, such as when WATER and H₂O in the propositions “water is liquid” and “H₂O is liquid” are distinguished by means of grasping two senses of a single referent (water). And finally, since senses can be grasped at different times and by different individuals, they would give stability to BoKs and explain how they can be shared.

Some might find this move awkward on the grounds that neo-Fregean approaches have traditionally had a hard time both trying to account for the ontological status of senses qua non-mental abstract objects and, hence, explaining how they can be grasped by our representational mechanisms. However, the notion of senses could just as well be subsumed by theories that do not need to hold the Fregean commitment that senses are non-mental abstract objects. As has been suggested by Margolis and Laurence (2007) and Prinz (2002), the notion of abstract sense could

be adapted to a representationalist approach to concepts by applying a type/token distinction, where concept types are conceived of as abstract objects and concept tokens as their particular mental realizations. While Prinz points out that the putative abstract ontological status of senses can be reconciled with the abstract status of mental-representation types, Margolis and Laurence propose a Mixed View of concepts which accommodates senses as the putative meanings that mental-representation tokens can have or express. In any case, viewing senses in terms of a type/token distinction regarding concepts would be no more problematic than presupposing a metaphysical type/token distinction regarding properties in general, as, for example, in saying that red objects instantiate the property of *being red*.

However, we believe that this neo-Fregean approach to content does not work for Machery's purposes. Note that if our use of BoKs satisfies the conditions for grasping Fregean senses, then the existence of concepts would not be challenged by the fact that different BoKs are effected by independent cognitive mechanisms. Recall that, according to our initial characterization, grasping a Fregean sense requires that a thinker satisfies the conditions that account for mastering a concept. Hence, insofar as those possession conditions amount to psychological mechanisms that are suited for grasping a Fregean sense (or instantiating a concept type, for that matter), the corresponding BoKs can be said to possess conceptual content. But then there would be no motivation for eliminating concepts, and so no reason for talking about replacing concepts with distinct kinds of BoKs after all.

Someone might want to insist on adapting this neo-Fregean approach to BoKs and claim that combining BoKs with Fregean senses can ascribe content to BoKs without having to identify senses with concepts. If the notion of sense were thus dissociated from that of concept, then the fact that these BoKs can grasp Fregean senses would not threaten Machery's concept eliminativism. However, it would be awkward to stay in the Fregean tradition without assuming some of its fundamental commitments. For example, a theory of the grasping of a sense is normally formulated in terms of knowledge of certain contents, which presupposes the capacity to combine those contents in complete thoughts and assign them truth-conditions (viz. knowing the semantic value of thoughts; cf. Peacocke 1992). If a thinker understands the content F, she has to be capable of knowing what it is for the thoughts Fa and Fb to be true. This involves the capacity to compose new thoughts, often called the *generality constraint* (e.g., Evans 1982). But note that this may be too demanding for Machery's view of BoKs, since (to mention a well-known problem we further explore in Sect. 6) prototypes have long been criticised on the grounds that they cannot compose.

The point is that characterising content in terms of Fregean senses brings along features traditionally associated with the conceptual, such as compositionality, reference determinacy and force independence.¹⁰ Indeed, resorting to this Fregean approach is a standard strategy used for establishing criteria for concept possession (Gunther 2003), the satisfaction of which some think allows for concept

¹⁰ Toribio (2010) makes this point explicit when describing the Neo-Fregean approach as a "full-fledged view" of concept possession that demands "the possession of certain inferential abilities—the kind of inferential abilities the Generality Constraint makes explicit" (p. 14).

individuation (e.g., Peacocke 1992, 1996; Cussins 1990, 1993). By contrast, detaching senses from concepts may imply pruning the notion of sense to the point of losing its significance within the theoretical framework that motivates them. We conclude from this that, if what is at stake is the fact that concepts split into BoKs with distinct inferential structures, then it may be better to avoid engaging with the Fregean tradition and explore the viability of Intensionalism through an alternative account of content determination that is attached to fewer and less demanding theoretical commitments.

As we pointed out at the beginning of this section, Intensionalism can also be linked with the view that content is determined by conceptual role. According to this view, the content of a concept is determined by (at least some of) its inferential connections to other concepts. Standard versions of the Prototype, Exemplar and Theory theories of concepts are good cases in point. According to these theories, a concept is some kind of mental representation that bears a particular type of structural relation to other concepts. Thus, whereas the content of a prototype concept is determined by a set of statistically frequent relations to its constituent concepts, the content of an exemplar concept is determined by its similarity relations to the representation of its exemplars as stored in memory. Likewise, the content of a theory concept is supposed to be determined by its inherent inferential relations with other concepts within a given theory (or belief system, for that matter).

However, these theories raise a common problem related to the very inferential character of content determination that they put forward. Given this character, it becomes evident that, because the corresponding inferential relations are subject to inevitable variation, the contents of the proposed models of concept are inherently unstable (cf. Fodor and Lepore 1992). For instance, if my prototypical guitar changes from wooden-made to made-of-diverse-materials, then, according to Intensionalism based on conceptual roles, my old and new prototypical guitars can be said to be BoKs with different semantic structures and, hence, different contents. Similarly, people whose GUITAR-prototypes differ in some respect would not be able to think or talk about the same thing when using their GUITAR-prototypes. What this example suggests is that BoKs with unstable inferential structures both within and across individuals are unsuited to support generalisations about various phenomena (e.g., how people learn, reason and communicate with one another) in virtue of BoKs that refer to the same objects. In other words, if Machery's proposed classes of BoKs are to attain their contents from their inferential relations to other prototypes, exemplars and theories, then this account of content determination is bound to inherit the problem of content stability stemming from the models of conceptual representation these kinds of BoKs are based on.

Admittedly, the way we have just characterised the problem of the instability of content determined by conceptual role is based on the assumption that conceptual content is constituted by every inference that a concept is involved in. This is what Fodor and Lepore (1992) have described as semantic holism. Some might claim, however, that holism is not a challenge to content determination by means of conceptual role because a limited number of inferences, instead of every one of them, can do the job of providing our mental states with stable content. This view is

normally committed to the idea that content similarity, not content identity, is sufficient for the stability of meaning required for different people (or the same person at different time slices) to be said to share the same concepts (e.g., Harman 1993). Regarding this alternative, Fodor and Lepore (1992) have replied that constancy of content fails to be accounted for in terms of a limited number of inferences without presupposing some sort of identity between inferences that overlap within or between individuals. A related problem pointed out by Fodor and Lepore is that, even if we granted that there is some invariant set of inferences that is constitutive of content, such an account of content would have to appeal to some kind of analytic/synthetic distinction in order to specify that set of inferences. The problem arises because, according to these authors, Quine's critique to this distinction showed the notion of analyticity to be untenable.

In response to Fodor and Lepore (1992), Greenberg and Harman (2006) argue for an ecumenical version of Conceptual Role Semantics (CRS) that is said to be immune to the problems we have just mentioned. In this version of CRS, conceptual roles are identified with several aspects of use of mental representations in thought. These aspects of use do not need to be restricted to internal inferential relations, which, according to Greenberg and Harman, correspond to just one of the several possible aspects of conceptual role (i.e., one of the several ways mental representations are used in thought).¹¹ In these authors' view, whereas conceptual roles are said to be part of what determines that a mental representation has a given content, they should not be deemed constitutive parts of the content of the mental representation. Hence, as Greenberg and Harman claims, CRS is committed to a notion of content stability that is based on content similarity, without presupposing sameness of content. Such possible theory of content determination deserves our attention since it could provide Machery's view of co-referential BoKs with a plausible notion of content stability by means of conceptual role. However, we think there is some important incompatibility between Greenberg and Harman's notion of content similarity determined by conceptual role and Machery's putative co-referential BoKs.

According to Greenberg and Harman, the content of a mental representation derives from use. If this view was to work for Machery's BoKs, then what a given BoK represents would have to be determined by the roles of the BoK in a person's cognitive economy. Thus, for instance, what a prototype (e.g., DOG-prototype) represents is determined by the way it is used to categorise certain things in the world (viz. dogs), and the same applies for exemplars and theories, *mutatis mutandis*. Moreover, given the assumption that content similarity is sufficient condition for content stability, the similar ways different people may represent a given object, say, prototypically, could provide grounds for those people to be said to possess co-referential prototypical representations of that object. And the same can be said about the ways different people may represent a given object by means of exemplars and theories. However, even if we granted that the similarity between

¹¹ According to Greenberg and Harman (2006), other aspects of use may include perceptual representation, recognition of implications, modelling, inference, labelling, categorization, theorizing, planning, and control of action.

two tokens of the same type of BoK can be enough to explain their co-referentiality (e.g., that two DOG-prototypes are both about dogs), things become fuzzier when it comes to different types of BoKs which are also supposed to be co-referential (i.e., DOG-prototypes, DOG-exemplars and DOG-theories). In these cases, we can expect considerable variation between their cognitive roles, and so it is not at all clear whether Greenberg and Harman's notion of content similarity is able to guarantee the co-referentiality of these BoKs.

To clarify the point, note that, even though distinct yet co-referential types of BoKs are said to be used in similar cognitive processes (e.g., categorization, induction, deduction, analogy-making, planning, or linguistic comprehension, etc.), they are also said to be used in very different ways. For example, as Machery (2010b) points out, while prototypes are assumed to be used in processes that compute the similarity between a prototype and other representations in a linear manner, exemplars are assumed to be used in processes that compute the similarity between a set of exemplars and other representations in a nonlinear manner. In turn, theories (e.g., a BoK involving some causal knowledge) are assumed to be used in "cognitive processes that are similar to the algorithms involved in causal reasoning" (p. 201). Thus, if the content of a BoK derives from its use, then prototypes, exemplars and theories associated to the same category are not expected to bear the similar contents that are supposed to guarantee their co-referentiality. We conclude from this that Greenberg and Harman's version of CRS cannot be successfully adapted to Machery's view of BoKs.

5.2 Referentialism and BoKs

As opposed to Intensionalism, Referentialism assumes that the extension of a concept determines its content. A prominent version of this view states that the content of concepts is determined by virtue of them being in a certain causal-cum-nomological relation to their referents. Here, Fodor's informational semantics is a good case in point (Fodor 1987, 1990). According to this theory, the content of a concept qua primitive (or unstructured) mental representation is determined by the concept's standing in an appropriate causal-cum-nomological relation to things in the world. Motivated by the notion of content as information put forward by Dretske (1981), the theory is based on the idea that a concept carries information about a given property if the concept is under the nomological control of that property. More specifically, the idea is that types of concepts are causally connected to the properties their tokens express by means of there being a natural law that prescribes this (actual or counterfactual) causal connection. Thus, contrary to the alternative described in the preceding paragraphs, the content of a concept like APPLE would not be determined by its relations to other concepts (e.g., FRUIT, TREE, etc.) but by a causal-cum-nomological relation that connects the property of *being an apple* with the concept APPLE.

The mention of Fodor's informational semantics is pertinent here since Machery himself resorts to this view when pressed by his critics to explain how his different kinds of BoKs refer and how they can be said to be co-referential (Machery 2010b, p. 235). He suggests that the retrieval from long-term memory of a given BoK (e.g.,

a prototype of squirrel) could be nomologically linked to the presence of a given property (e.g., *being a squirrel*) in a person's environment. Co-referentiality would then be explained by the fact that different BoKs are retrieved by virtue of being under the nomological control of the same property, as it would be the case with a SQUIRREL-prototype, a SQUIRREL-exemplar and a SQUIRREL-theory, for example. In addition, given that those BoKs are supposed to share the same informational content, the problem of instability of content would then not be a threat to Machery's heterogeneous BoKs. However, we think there are good reasons why a causal view of content can't be made to work for Machery's distinct kinds of BoKs. Let us elaborate why.

The main problem is that this causal view does not seem to guarantee the co-referentiality of different BoKs associated to the same category. There are situations where prototypes, exemplars and theories about the same category of objects could be said to be nomologically linked to things in the world that belong to different categories. To use a common example, consider the case of BoKs associated to the category of grandmothers. Prototypical grandmothers are grey-haired, have wrinkled skin and wear glasses. But there are grandmothers that do not satisfy these properties (e.g., Sarah Palin) and thus they fail to trigger¹² a GRANDMOTHER-prototype to be used in cognitive processes that mainly depend on similarity to a prototype (e.g., prototype-based categorization). Instead, we would expect them to be linked to a BoK corresponding to a GRANDMOTHER-theory, the structure of which is likely to include the element of information that grandmothers are mothers of a parent. Consequently, a subject's GRANDMOTHER-prototype and GRANDMOTHER-theory can fail to co-refer because these BoKs can fail to be nomologically linked to the same referents.

To make this point clearer, consider the reverse case: an old woman that looks just like a prototypical grandmother but has never been a mother (e.g., Betty White). People will probably identify her as a grandmother through the activation of their GRANDMOTHER-prototype, while their GRANDMOTHER-theory takes no part in cognitive processes that mainly depend on, say, causal, functional or explanatory knowledge about a category (e.g., theory-based categorization). So we have another situation where a prototype and a theory about the same category are not co-referential (cf. Margolis and Laurence 2010).¹³

Machery might address these worries in two ways and we shall look at each of them in turn. First, suppose Machery responds by adding that BoKs about a given category share a common informational core that guarantees their co-referentiality. Thus, it is possible to claim that such core is a common constituent of different

¹² We use the term 'trigger' to express the same idea as 'retrieving X from long-term memory'.

¹³ Here a Fodorian might reply that this critique does not affect the referentialist view of the content of BoKs since, according to Informational Semantics, content is determined by the information carried by a concept, not by its inferential role (whether prototypical or theoretical). But in the context of Machery's view of BoKs, the information that is constitutive of a BoK is that which is activated by default in order to be used in cognitive processes, rather than by the role of that BoK in cognitive inferences (see Machery 2015). So, if content is identified with information as per Informational Semantics, BoKs whose respective default bodies of information fail to be co-activated by a given referent cannot be said to be under the nomological control of that referent and, therefore, they cannot be co-referential.

BoKs about a given category and that it is this core that is under the nomological control of a property in the environment. In this sense, someone could think that the nomologically-based retrieval of any given BoK about x (e.g., GRANDMOTHER-prototype) does not need to imply the retrieval of a whole different BoK about x (e.g., GRANDMOTHER-theory), provided that the shared core about x is retrieved every time one of those BoKs is used in cognitive processes where they play a central role. However, this solution does not work because it is incompatible with Machery's Heterogeneity Hypothesis according to which co-referential BoKs have very few properties in common, including the information retrieved by default from long-term memory when using them (see Sect. 2, above). Recall Machery's distinction between context-independent default knowledge about x , which is constitutive of a given BoK, and the background knowledge about x that is activated in a context-dependent manner and is not constitutive of that BoK. If we were to propose that distinct yet co-referential BoKs about x (e.g., CAT-prototype, CAT-exemplar and CAT-theory) share a common informational core, we should expect that information to always overlap when thinking about cats in different contexts. But since Machery assumes that each co-referential BoK has its own proprietary default knowledge, if there were some information that overlaps across contexts, it would correspond to some background information that is not constitutive of any of those BoKs and thus cannot provide grounds for individuating their contents. (Additionally, note that even if distinct BoKs about x happened to share common informational cores which are nomologically activated by means of independent retrieval processes, allowing these cores within Machery's theoretical framework would undermine the motivation for questioning the existence of concepts as a single phenomenon. After all, those cores would be doing the work we normally expect concepts to do).

Another way Machery might address the worries mentioned above is by responding that the occurrence of a given BoK about x 's can be sometimes triggered by non- x 's, or else, fail to be activated by atypical x 's, but still carry information about x 's insofar as the relevant nomic link between the occurrence of a given BoK and the presence of its referent is to be understood counterfactually: if the object x were more typical, x would retrieve the correct BoK by default.¹⁴ Indeed, a solution of this sort is proposed by Fodor (1987) with his Asymmetric Dependence Theory to allow for the possibility of misrepresentation within informational semantics. This would explain, for example, why a BoK can be wrongly activated (e.g., a GRANDMOTHER-prototype triggered by Sara Palin) and still count as co-referential with a different BoK about the same category (e.g., a GRANDMOTHER-theory), given that the former's BoK-referent nomic connection is counterfactually dependent on the latter's.

Unfortunately, we believe the suggested solution fails for, at least, two reasons. First, there are cases (e.g., newborn marsupials vs. rodents; so-called "good fakes" such as cubic zirconia vs. diamonds, etc.) where the requirement of a high degree of typicality based on superficial features does not guarantee an account of the misapplication of certain BoKs, notably prototypes and exemplars. The reason is that the information encoded in those BoKs is not enough to guarantee that the

¹⁴ We thank an anonymous reviewer for pointing this out.

failure to identify, say, an atypical diamond is counterfactually supported by a BoK that carries information about *typical diamonds*, since a BoK that carries information about, say, *cubic zirconia that are good fake diamonds* could be said to provide the same counterfactual support. Thus, there has to be more to an account of the misapplication of these kinds of BoKs than just information about typical superficial features of objects in the world, which brings us to the second problem.

Just as Jacobson (2010) has argued, cases of “good fakes” present Machery’s BoKs with a critical problem. She observes that theories encode a type of information (e.g., causal, functional, etc.) that is often essential for distinguishing between genuine objects and good fakes, information which prototypes and exemplars fail to encode. So, it would appear that, in the case of good fakes, the incorrect activation of a prototype (or exemplar) about a given category is dependent on the nomic link that determines the content of theories about that category, as would be the case in the example of an atypical grandmother mentioned above. However, since people’s theories can also contain false information (e.g., the belief that there can’t be mothers who are younger than 55 years old) and, consequently, fail to identify categories correctly, it would be problematic to say that an account of the incorrect activation of theories is counterfactually dependent upon the occurrence of typical instances.¹⁵ The reason is that, as we have seen in the case of good fakes, information about typical instances, which is preferentially encoded by prototypes and exemplars, seems to inevitably rely upon the type of information that is preferentially encoded by theories. Hence, we believe the incorrect activation of Machery’s BoKs can’t be explained in terms of the counterfactual typicality requirement mentioned above without risking circularity.

One final worry regarding the application of a referentialist approach to Machery’s view of distinct BoKs has to do with whether these BoKs are construed at the right level of abstraction. As Schneider (2010) notes, if the semantic properties of co-referring BoKs can be individuated informationally, then the fact that prototypes, exemplars and theories conform to distinct processing properties does not challenge the idea that these BoKs could fall under the same concept type. Concepts could thus be understood as a superordinate—functional—kind (cf. Edwards 2010; Samuels and Ferreira 2010; Rey 2009; Lalumera 2013). Consider, for example, the view that the class of concepts is a natural kind that contains subordinate kinds including, e.g., prototypes, exemplars and theories in the same way that the class of minerals is a natural kind with subordinate kinds such as pyrite, topaz and quartz. In this sense, the causal view of content Machery thinks might work for his BoKs really fixes the content of a superordinate kind, which, in turn, explains how any three different BoKs about a given category are actually co-referential.

At this point Machery might contend the notion of concept as a superordinate kind makes no useful work in psychology and that, as he has argued, keeping this notion would impede scientific progress (Machery 2010b, 2015). But this claim is highly controversial, and many have pointed out that the notion of concept does have many explanatory virtues. Besides providing psychology with a notion to make

¹⁵ Put differently, it would be problematic to say that if grandmother *x* were more typical, *x* would retrieve the correct theory by default.

interesting generalisations about high-level cognitive processes (Scarantino 2010; Weiskopf 2010), concepts have shown to be a useful heuristic in guiding psychological research (Hayes and Kearney 2010; Margolis and Laurence 2010). To elucidate this view, consider Cowie's account of "messy" or imprecise concepts (Cowie 2009).¹⁶ In a related discussion about good and bad concepts, Cowie defends what she calls *vulgar pragmatism*, according to which imperfect and imprecise concepts should be preferred in the absence of better options when those concepts are part of a flourishing research programme. Together with the objections above, Cowie's idea of vulgar pragmatism can thus provide additional reasons to reject Machery's claim that 'concept' hinders scientific progress.

So far, we have explored several approaches to content for Machery's BoKs. Each approach has failed to satisfy an account of how these BoKs could attain their contents, specifically, one that allows them to be both distinct and co-referential. The next section discusses additional issues raised by our critical assessment of Machery's BoKs and presents some concluding remarks.

6 Discussion and Final Remarks

So far, we have shown that Machery's view of BoKs is incompatible with what we consider standard versions of the two main approaches to content determination: Intensionalism and Referentialism. This can be a drawback for Machery's proposal insofar as an account of content has traditionally been considered a desideratum for a theory of concepts.

As we know, the notion of content has a number of explanatory virtues (see Sects. 1 and 3). Importantly, it could help to explain the pervasive capacity of cognitive agents to think differently about the same thing and share their thoughts with others thanks to the co-referentiality of their concepts. Likewise, the notion of content may be connected with people's notable capacity to entertain and form an unbounded number of contentful novel thoughts. This latter capacity has been explained in terms of the combinatorial properties of concepts as the constituents of thoughts. Indeed, it can be considered a received view among philosophers and psychologists that the content of thoughts is *compositional*.¹⁷

However, Machery's heterogeneous BoKs are bad candidates for explaining the combinatorial properties of thoughts. We concur with Hill (2010) in that, while BoKs involved in high level cognitive processes are expected to provide an account of our capacity to form novel thoughts, Machery's BoKs are incapable of explaining the content of complex mental representations in terms of combinatorial properties. Here Hill echoes arguments regarding prototypes and compositionality raised by Fodor and Lepore (1996), who defend the view that, if prototypes can't compose, they can't be identified with concepts.

¹⁶ In her paper, Cowie makes no distinctions of usage between 'concepts' and 'terms'.

¹⁷ Roughly, 'compositionality' refers to the principle according to which the content of a thought is inherited from the contents of its constituent concepts, together with their structural arrangement (see, e.g., Fodor and Lepore 1992).

Machery (2010a) thinks Hill's argument fails because the empirical fact that the prototype associated with some complex expression (e.g., PET FISH) does not prove to be a function of the prototypes associated with the combined expressions (e.g., PET and FISH) does not rule out the possibility that BoKs can combine compositionally. We are not moved by Machery's response because it explicitly appeals to Prinz's contention that prototypes can combine compositionally, which we do not think supports Machery's view of distinct BoKs.

Contrary to Fodor and Lepore (1996), Prinz (2012) claims that concepts are usually (but not exclusively) prototypes and that prototypes can combine compositionally, even though they may sometimes fail to do so. He appeals to evidence suggesting the potential viability of different models of prototype combination (e.g., Smith et al. 1988; Hampton 1991) and he also advances his own model (viz. his RCA or Retrieval, Composition and Analysis Model). However, these models make no distinction between different and independent BoKs as Machery does, and all the models (explicitly or implicitly) rely on theoretical and exemplar knowledge to be involved in an account of prototype integration. Hence, even if we granted that Prinz is right about compositionality, Prinz's proposal and the like are not compatible with Machery's BoKs.

It should be noted that Machery has explicitly claimed that the combinatorial properties of thoughts is not among the explananda of psychological theories of concepts (Machery 2010a). Bracketing concerns about this methodological claim,¹⁸ we could make sense of Machery's view by interpreting his characterisation of concepts qua BoKs as couched in a *subpersonal* level of analysis, where personal-level explanations about the combinatorial and normative properties of thought are not at issue. So without denying that there are such genuine properties of thinkers, one could say that they are simply not within the scope of the theoretical vocabulary of subpersonal psychology, in the same sense as the personal-level notion of 'pain' does not literally figure in subpersonal explanations of pain describing information-processing or neural events in the brain. As Dennett (1969) has pointed out, the fact that the concept of 'pain' (or 'concept', for present purposes) is not part of subpersonal explanations does not mean that a subpersonal account of pain cannot shed light on its nature and associated behavioural properties.¹⁹

An interesting consequence of understanding BoKs within a subpersonal information-processing framework is that their contents could then be regarded as nonconceptual.²⁰ A view of this sort has been developed by Bermudez (1995). He gives the example of representations at different stages of early visual processing

¹⁸ Say, e.g., if that property of our cognitive capacities is not part of the explananda of a scientific theory of concepts, then whose explanandum is it?

¹⁹ In his book, Machery (pp. 25–26) contends that placing the term 'concept' at the personal-level is misleading given that this term is sometimes used by psychologists and cognitive scientists in explanations normally viewed as subpersonal. However, since Machery's purpose is to regiment the use of the term 'concept', the idea that this term should be replaced by that of 'BoK' within subpersonal explanations is worth considering.

²⁰ The notion of nonconceptual content was introduced by Evans (1982) to characterise ways of representing the world independent of thinkers' conceptual abilities. It is often negatively defined in opposition to conceptual content, as having an iconic/pictorial aspect and lacking predicative structure.

according to Marr's theory of vision, such as zero-crossings and 2 1/2 D sketches, and contends that they should be described as nonconceptual. One reason he gives is that, even though these representations carry information from the environment, they have not yet generated a perceptual representation (whether conceptual or not). Another is that the contents of subpersonal states such as zero-crossings are so abstruse and domain specific that it seems implausible that individuals possess the concepts required for specifying them. Even though Bermúdez's examples revolve around visual processing, the suggestion is that the attribution of nonconceptual content can be extended beyond perception, encompassing subpersonal information processing at higher cognitive levels.

It is conceivable that this less constrained notion of content allows for more heterogeneity and thus might contribute to an explanation of how it is that we possess several BoKs referring to the same kind. To avoid the obvious objection that when attributing nonconceptual content to BoKs one has to presume that states with conceptual content exist somewhere else, this view may be further adapted for Machery's BoKs by endorsing what Gunther (2003, p. 16) calls *global nonconceptualism*, the view that concepts are irrelevant for the individuation of contents and that psychology does not require a notion of conceptual content (cf. Stalnaker 2003; see Bermudez 2007 for a skeptical look to global nonconceptualism).

The possibility of combining Machery's BoKs with some version of global nonconceptualism should be explored in a more detailed way, but this is not the place to do it, so we leave it for future work. Nevertheless, it deserves noting that applying global nonconceptualism to BoKs would render the notion of content almost explanatorily idle. Since not just BoKs but all representational states would have the same kind of informational content, how we process information or think about the world would not be a function of how we represent the world, viz. of the content of our representational states. So in case of adopting such a view, Machery would have to be prepared to accept this consequence and embrace the relegation of content to a peripheral role in psychological explanations.

In sum, we believe that the present scrutiny of Machery's view of BoKs, under the light of prominent theories of content, can constitute an example of how philosophers and psychologists can engage in a constructive dialogue regarding their common purpose in elucidating the nature of concepts. Machery says that the psychology of concepts is in a state of disarray because none of the available theories can provide a unifying theoretical framework. If this is right, then looking at how a certain theory of concepts (or BoKs) comports with well settled ontological or semantic issues regarding concepts may well constitute an additional criterion for assessing the plausibility of those theories (cf. Vallejos 2008). In any case, we conclude from the present analysis that the problems to supply Machery's proposal with a viable notion of content can constitute a reason to resist his argument for concept eliminativism.

Acknowledgements We are grateful to Stephen Laurence, George S. Botterill, Manuel Rodriguez for their comments on earlier drafts of this paper. We also benefited from audiences at postgraduate seminars at the University of Sheffield, UK, and the Universidad Nacional de Córdoba, Argentina. This work was supported by the National Commission for Scientific and Technological Research (CONICYT) under Becas Chile scholarship for doctorate programs abroad.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Anderson, J. R., & Betz, J. (2001). A hybrid model of categorization. *Psychonomic Bulletin and Review*, 8, 629–647.
- Bermudez, J. L. (1995). Non-conceptual content: From perceptual experience to subpersonal computational states. *Mind and Language*, 10, 333–369.
- Bermudez, J. L. (2007). What is at stake in the debate about nonconceptual content? *Philosophical Perspectives*, 21, 55–72.
- Boyd, R. (1999). Homeostasis, species, and higher taxa. In R. Wilson (Ed.), *Species: New interdisciplinary essays* (pp. 140–185). Cambridge: MIT Press.
- Cowie, F. (2009). Why isn't Stich an eliminativist? In D. Murphy & M. A. Bishop (Eds.), *Stich and his critics* (pp. 14–74). Hoboken: Wiley.
- Cussins, A. (1990). The connectionist construction of concepts. In M. Boden (Ed.), *The philosophy of artificial intelligence* (pp. 368–440). Oxford: Oxford University Press.
- Cussins, A. (1993). Nonconceptual content and the elimination of misconceived composites. *Mind and Language*, 8(2), 234–252.
- Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT Press.
- Dummett, M. (1991). *Frege and other philosophers*. Oxford: Oxford University Press.
- Dennett, D. (1969). *Content and consciousness*. London: Routledge & Keegan.
- Edwards, K. (2010). Unity amidst heterogeneity in theories of concepts. *Behavioral and Brain Sciences*, 33, 210–211.
- Evans, G. (1982). *The varieties of reference*. Oxford: Oxford University Press.
- Fodor, J. A. (1987). *Psychosemantics: The problem of meaning in the philosophy of mind*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1990). *A theory of content and other essays*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1994). Concepts: A potboiler. *Cognition*, 50, 95–113.
- Fodor, J. A. (1998). *Concepts: Where cognitive science went wrong*. New York: Oxford University Press.
- Fodor, J., & Lepore, E. (1992). *Holism: A Shopper's guide*. Oxford: Blackwell.
- Fodor, J., & Lepore, E. (1996). The red herring and the pet fish: Why concepts still can't be prototypes. *Cognition*, 58, 253–270.
- Fodor, J., & Pylyshyn, Z. (2015). *Minds without meanings: An essay on the content of concepts*. Cambridge, MA: MIT Press.
- Greenberg, M., & Harman, G. (2006). Conceptual role semantics. In E. LePore & B. Smith (Eds.), *Oxford handbook of philosophy of language* (pp. 295–322). Oxford: Oxford University Press.
- Grice, P. (1957). Meaning. *Philosophical Review*, 66, 377–388.
- Gunther, H. Y. (Ed.). (2003). *Essays on nonconceptual content*. Cambridge, MA: MIT Press.
- Hampton, J. (1991). The combination of prototype concepts. In P. J. Schwanenflugel (Ed.), *The psychology of word meanings* (pp. 91–116). Hillsdale: Lawrence Erlbaum Associates.
- Harman, G. (1993). Meaning holism defended. In J. A. Fodor & E. Lepore (Eds.), *Holism: A consumer update* (pp. 163–171). Amsterdam: Rodopi.
- Hayes, B. K., & Kearney, L. (2010). Defending the concept of “concepts”. *Behavioral and Brain Sciences*, 33, 214.
- Hewson, C. (1994). Empirical evidence regarding the folk psychological concept of belief. In A. Ram & K. Eisele (Eds.), *Proceedings of the sixteenth annual conference of the Cognitive Science Society* (pp. 403–406). Hillsdale, NJ: Erlbaum.
- Hill, C. (2010). I love Machery's book, but love concepts more. *Philosophical Studies*, 149, 411–421.
- Jacobson, A. J. (2010). The faux, fake, forged, false, fabricated, and phony: Problems for the independence of similarity-based theories of concepts. *Behavioral and Brain Sciences*, 33, 215.

- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge: MIT Press.
- Lalumera, E. (2013). Concepts exist. More about Eliminativism. *Method—Analytic Perspectives*, 2, 126–133. Retrieved July 9, 2016. <http://www.methode.unito.it/methOJS/index.php/meth/article/view/26/23>.
- Laurence, S., & Margolis, E. (1999). Concepts and cognitive science. In E. Margolis & S. Laurence (Eds.), *Concepts: Core readings* (pp. 3–81). Cambridge, MA: The MIT Press.
- Machery, E. (2005). Concepts are not a natural kind. *Philosophy of Science*, 72, 444–467.
- Machery, E. (2009). *Doing without concepts*. New York: Oxford University Press.
- Machery, E. (2010a). Replies to my critics. *Philosophical Studies*, 149, 429–436.
- Machery, E. (2010b). Précis of doing without concepts. *Behavioral and Brain Sciences*, 33, 195–244.
- Machery, E. (2015). By default. In E. Margolis & S. Laurence (Eds.), *The conceptual mind* (pp. 567–588). Cambridge, MA: The MIT Press.
- Machery, E., Mallon, R., Nichols, S., & Stich, S. P. (2004). Semantics, crosscultural style. *Cognition*, 92, B1–B12.
- Margolis, E., & Laurence, S. (2007). The ontology of concepts—Abstract objects or mental representations? *Noûs*, 41(4), 561–593.
- Margolis, E., & Laurence, S. (2010). Concepts and theoretical unification. *Behavioral and Brain Sciences*, 33, 219–220.
- Margolis, E., & Laurence, S. (2014). Concepts. In E.N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition). Retrieved July 9, 2016. <http://plato.stanford.edu/archives/spr2014/entries/concepts/>.
- Murphy, G. (2002). *The big book of concepts*. Cambridge: MIT Press.
- Nisbett, R. E., Peng, K., Choi, I., & Norenzayan, A. (2001). Culture and systems of thought: Holistic vs. analytic cognition. *Psychological Review*, 108, 291–310.
- Peacocke, C. (1992). *A study of concepts*. Cambridge, MA: MIT Press.
- Peacocke, C. (1996). Can possession conditions individuate concepts? *Philosophy and Phenomenological Research*, 56(2), 433–460.
- Piccinini, G., & Scott, M. (2006). Splitting concepts. *Philosophy of Science*, 73, 390–409.
- Prinz, J. J. (2002). *Furnishing the mind: Concepts and their perceptual basis*. Cambridge, MA: MIT Press.
- Prinz, J. (2012). Regaining composure: A defense of prototype compositionality. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford handbook of compositionality* (pp. 437–453). Oxford: Oxford University Press.
- Rey, G. (1998). Concepts. In E. Craig (Ed.), *Routledge Encyclopedia of Philosophy*. London: Routledge. Retrieved April 21, 2015. <https://www.rep.routledge.com/articles/concepts/v-1/>.
- Rey, G. (2009) Review of E. Machery, *Doing without Concepts*. *Notre Dame Philosophical Reviews*. Retrieved October 10, 2014. <http://ndpr.nd.edu/news/24087-doing-without-concepts>.
- Samuels, R., & Ferreira, M. (2010). Why don't concepts constitute a natural kind? *Behavioral and Brain Sciences*, 33, 222–223.
- Scarantino, A. (2010). Evidence of coordination as a cure for concept eliminativism. *Behavioral and Brain Sciences*, 33, 223–224.
- Schneider, S. (2010). Conceptual atomism rethought. *Behavioral and Brain Sciences*, 33, 224–225.
- Stalnaker, R. (2003). What might non-conceptual content be? In Y. Gunther (Ed.), *Essays on nonconceptual content* (pp. 95–106). Cambridge, MA: MIT Press.
- Smith, E. E., Medin, D. L., Rips, L. J. & Keane, M. (1988). Combining prototypes: A selective modification model. *Cognitive Science*, 12, 485–527.
- Toribio, J. (2010). The animal concepts debate: a metaphilosophical take. *Teorema*, 29(2), 11–24.
- Vallejos, G. (2008). *Conceptos y ciencia cognitiva*. Santiago: Bravo y Allende Editores.
- Weiskopf, D. (2010). The theoretical indispensability of concepts. *Behavioral and Brain Sciences*, 33, 228–229.