

# Phenomena and Objects of Research in the Cognitive and Behavioral Sciences

Uljana Feest\*†

---

It is commonly held that research efforts in the cognitive and behavioral sciences are mainly directed toward providing explanations and that phenomena figure into scientific practice qua explananda. I contend that these assumptions convey a skewed picture of the research practices in question and of the role played by phenomena. I argue that experimental research often aims at exploring and describing “objects of research” and that phenomena can figure as components of, and as evidence for, such objects. I situate my analysis within the existing literature and illustrate it with examples from memory research.

---

**1. Introduction.** In current philosophy of cognitive science, the following two assumptions are fairly widespread: first, research efforts are mainly directed toward explaining phenomena, and second, the explanations in question are likely to be mechanistic. In this article, I do not challenge the importance of the contributions that the mechanistic approach has made to philosophical debates about explanation. However, I believe that the literature about mechanistic explanations has resulted in a skewed picture of research practices in the cognitive and behavioral sciences. In this vein, the main thesis of my essay is that the primary aim of experimental research in these fields is not explanation but rather the description and exploration

\*To contact the author, please write to: Leibniz Universität Hannover, Institut für Philosophie, Im Moore 21, 30167 Hannover, Germany; e-mail: feest@philos-uni-hannover.de.

†Previous versions of this article were presented at the Universities of Hannover, Bielefeld, Wuppertal, and Graz, as well as the 2015 meeting of the European Philosophy of Science Association in Düsseldorf (Germany). A more recent version was presented at the Technische Universität Berlin in July 2016. The author would like to thank the members of these audiences, as well as colleagues who attended my talk at the 2016 meeting of the Philosophy of Science Association in Atlanta, for their very helpful questions and comments.

Philosophy of Science, 84 (December 2017) pp. 1165–1176. 0031-8248/2017/8405-0030\$10.00  
Copyright 2017 by the Philosophy of Science Association. All rights reserved.

of the relevant subject domain, or as I prefer to put it, of “objects of research.” In a nutshell, philosophers who focus on explanation typically not only take it for granted that we already know what the relevant explananda are but also tend to gloss over the problem that the very notion of a phenomenon (especially in its relation to experimental data) is not very clear.

In this article, I address these two problems, arguing that the exploration and conceptualization of both explanantia and explananda are subject to ongoing processes of investigation. Providing an analysis of this process is a relevant desideratum for philosophy of science. Focusing on psychology, I will develop such an analysis, which takes into account the experimental practices of researchers in these fields. In particular, I will argue that experimental effects play an important role in these practices but that this role does not consist in exemplifying explanandum phenomena. Rather, they function as evidence that is drawn on in constructing descriptive accounts of the relevant objects of research. I will defend my use of the term “object of research” over the (more common) term “phenomenon” to indicate that objects of research are more fluid and can in principle be composed of more than one phenomenon.

In section 2, I begin by explaining what I mean by an object of research. Drawing on a case study about working memory, I will illustrate a central feature of such “objects”: that they are (as I will call it) “epistemically blurry.” By this, I mean that from the perspective of practicing scientists, their nature and identity conditions are not well understood. Section 3 will distinguish the expression “object of research” from the more commonly used “phenomenon,” arguing that in the existing literature the latter term is used in an ambiguous fashion, which can result in the conflation of several distinct philosophical issues, thereby detracting from some of the roles that phenomena play in research. Section 4 will turn to a critique of the mechanists’ juxtaposition of phenomena and mechanisms, with regard to both explanation and discovery. I will argue that this juxtaposition does not provide a satisfactory gloss of the role of phenomena in the investigative practices of psychology and that my own account is more promising in this regard.

**2. On the Notion of an Object of Research.** I use the expression “object of research” as distinct from “phenomenon” to describe the fact that psychologists often take their research to be relevant to some overarching subject matter (e.g., memory, emotion, attention). By contrast, I use the term “phenomenon” to refer to more confined empirical patterns.<sup>1</sup> In this section, I explain my contention that scientific activities in psychology are often directed toward some (presumed) overarching research object, where this in-

1. Below, I will provide a more in-depth account of my use of “phenomena,” building upon analyses by Hacking (1983), Bogen and Woodward (1988), and Feest (2011b).

volves the recognition that there is still a great deal of conceptual openness with regard to the contours of this object.

*2.1. Objects of Research as Epistemically Blurry.* We typically think of objects as three-dimensional things in the world, clearly delineated from their environments. Objects of research in psychology are not like that. Take memory, for example. First, while there is little doubt among researchers that memory is a legitimate object of scientific investigation, memory is not a physical object in the usual sense. Second, I claim that psychologists study memory in part because they are not entirely sure about its shape and contours (or even, perhaps, about whether at the end of the day this will be a viable scientific kind). Put differently, when we put ourselves in the epistemic position of practicing scientists, it is clear that they are not (or at least not exclusively or even primarily) in the business of explaining some already clearly delineated phenomenon. Rather, much of their empirical work is directed toward figuring out what they are actually investigating. Following Feest (2011a), I therefore refer to such objects of research as “epistemically blurry.”

With this thesis I do not mean to deny that ultimately the point of categories like “memory” is to explain behavior. In this vein, my approach is somewhat compatible with more traditional philosophy of psychology, according to which memory is a theoretical concept (an explanans), put into place to explain certain types of behaviors (explanandum phenomena). However, as I hope to make clearer in the course of this article, I think that this way of carving up the field (by dividing it into explananda, which are clearly delineated, and explanantia, which are to be determined) does not lend itself to an adequate analysis of the research process itself, which is importantly concerned with making taxonomic constructs and attempting to describe distinctive features of the (purported) objects in their extension. When I say that objects of research are epistemically blurry, I do not mean that scientists lack a theoretical explanation for an already clearly delineated empirical explanandum but that the very question that empirical data are even descriptively relevant to the object in question is part of the investigative project. Thus, it is my contention that often explanation is not the most immediate concern of experimental research in cognitive science.

I use the expression “object of research” to indicate the presumption of a research target that is independent of our own conceptualizations. Notice, though, that I am not committed to the claim that this presumption will pan out in every case. There may well be cases in which the “object” in question will turn out not to possess the inductive projectability required of a scientific kind. In that case, we will probably see the terms being eliminated in the long run. At the same time, this expression is also supposed to indicate that such (presumed) objects are research targets only by virtue of the fact that we di-

rect our epistemic attention to them and hence that we tentatively use some concepts to describe our subject matter. The research process, as I construe it here, is thus one of simultaneously exploring a specific subject domain and of applying, revising, and extending existing concepts. In this respect, the outlook taken here is not so dissimilar from Wilson's (2006) study of the often unpredictable behavior of concepts when confronted with, or extended to, previously unimagined physical domains or circumstances.

*2.2. Short-Term Memory as an Object of Research.* I believe that the rough gloss just provided can make sense of research practices we encounter in psychology, where a great deal of time and effort are spent trying to arrive at adequate descriptions of the presumed research objects. Consider the example of working memory, which is typically thought of, or modeled, as a stage of information processing before items get committed to long-term memory (see, e.g., Baddeley and Hitch [1974] for a classic formulation).<sup>2</sup> I claim that the point of this model is that of delineating and empirically describing a particular object of research (working memory) in relation to long-term memory. To be sure, such work is theoretical, but it does not follow that its immediate aim is explanation. Rather, the empirical research in question aims at gaining a better descriptive understanding of the very objects under investigation.

Let me quickly illustrate the nature of such research (and thereby also of the epistemic blurriness of the research object in question) by looking at debates about two descriptive features long taken to be characteristic of short-term memory: duration and capacity, where the duration was thought to be about 10 to 20 seconds and the capacity  $7 \pm 2$  (Miller 1956). The term "working memory capacity" refers to the number of items that can be stored in working memory at any given time. When talking about efforts to determine working memory capacity as illustrating my notion of an epistemically blurry research object, what I have in mind is not the fact that there is some variation between different individuals' working memory capacities but rather that it is not a trivial (i.e., straightforwardly empirical) question how to determine a figure like  $7 \pm 2$ . When scientists run a psychological experiment, by what rationale do they determine how many memory items are indicated by any given subject's experimental response? My claim here is that these are questions that are quite central to researchers' debates about their subject matter.

2. The authors introduced the expression "working memory," indicating a conceptual shift away from the earlier "short-term memory." The details of this shift cannot be covered here. This article will mainly use "working memory," except when referring to Miller (1956).

One reason why the question of how to determine working memory capacity is not trivial has to do with a particular phenomenon thought to be relevant to short-term (or working) memory: the phenomenon of chunking (Miller 1956). According to the chunking hypothesis, working memory capacity is fixed. However, once chunking takes place, the amount of information that can be kept in working memory radically increases, because the items are now composed of chunks of smaller items. If we want to describe working memory capacity, chunking will inevitably pose obstacles, since the capacity limit may well be inflated by the effects of chunking. Thus, in order to produce an adequate measure of capacity, we need to ensure that we have a way of empirically identifying the chunks (as opposed to their parts) that subjects keep in working memory. None of these are explanatory questions. They are, rather, theoretical and methodological questions about how to investigate and describe a given object of research.

Cowan (2000), in an influential paper (in which he argues that Miller's famous  $7 \pm 2$  estimate should be corrected downward to  $4 \pm 1$ ), tries to develop a method for distinguishing between memory chunks and the total number of items stored in memory, arguing that chunks are associations between concepts and that one needs to control for items that are already "chunked" in long-term memory. However, as the peer responses to Cowan's article show, there are also problems with this approach. For example, one question that comes up is how it can be established that higher-order coding has in fact been prevented (Beaman 2000, 118), suggesting that Cowan is begging the question by presupposing a particular understanding of chunking. Others argue that Cowan's theoretical definition of a chunk is too weak (Halford, Phillips, and Wilson 2000; Milner 2000; Pascual-Leone 2000). Yet others worry whether empirical findings about chunking can be generalized across modalities (Schubert and Frensch 2000, 146).

I take the debate to be one about how to adequately describe a particular feature (storage capacity) of a specific research object (working memory). My aim here is not to take sides in the debate but to use it to illustrate my thesis that scientists are often concerned with providing descriptions of research objects and that there is frequently a great deal of epistemic uncertainty about the objects in question.

**3. Phenomena versus Objects of Research.** In contrast to my talk of objects of research, philosophers often refer to things like memory, attention, emotion, and so forth, as "phenomena." While I certainly do not want to police the way this term is used, I argue that the term "phenomenon" is being used ambiguously in the philosophical literature and that for the purposes of an analysis of the research process, my notion of an object of research is a useful complement.

3.1. *Two Meanings of “Phenomena.”* According to Bogen and Woodward, phenomena are “relatively stable and general features of the world” (Woodward 1988), which “have stable, repeatable characteristics” (Bogen and Woodward 1989, 393). In this vein, one might suppose that memory is a phenomenon that has certain repeatable characteristics, such as the characteristic of chunking. However, when we look at the examples provided, it is clear that Bogen and Woodward regard the chunking effect itself as a phenomenon, not simply as a characteristic of some other phenomenon. So, it seems that they use the term in two distinct senses: (1) as a stable feature or kind that has repeatable characteristics and (2) as simply a repeatable regularity.

I contend that by running these two senses together, we give up an important analytical distinction, a distinction that will help us—among other things—clarify how objects of research are conceptualized in psychology, why they are epistemically blurry, and (most importantly) how they are investigated. In this vein, I suggest restricting the use of the term “phenomenon” to things like the chunking effect, the melting point of lead, and other events that occur regularly under specified circumstances. Now, it is certainly conceivable that such types of phenomena can become objects of research in the sense outlined above; that is, scientists can turn their epistemic attention to them, explore them, probe them by varying the experimental conditions under which they occur, and so forth. However, I contend that objects of research in psychology are often more variegated than that. For example, I suggest that scientists think of working memory as a *cluster of phenomena*, such as chunking, encoding, retrieval, duration. Thinking of objects of research in this way highlights once again what makes them epistemically blurry: they are epistemically blurry both because the phenomena they are thought to be composed of are ill understood in their own right and because the ways in which those phenomena are related to one another are ill understood as well. Figuring out the exact contours of an object of research involves providing descriptions of the various phenomena tentatively thought to constitute the object of research and providing an account of how they hang together.

3.2. *Phenomena versus Effects?* My use of the term “phenomenon” includes regularities both at the experimental level and at the level of processes presumed to exist independently of experimental creation. In this respect, my account differs from Bogen and Woodward’s (1988), who argued that phenomena are mind-independent regularities in the world, to be distinguished from experimental data, which are idiosyncratic to specific experimental contexts. My account also differs from Hacking’s (1983) contention that phenomena often only exist under the very specific experimental circumstances of their creation. Responding to statements such as this

one, in turn, Bogen and Woodward suggested that what Hacking referred to as phenomena were really data in their sense.

The distinction is not so clear, however. I argue that experimental effects should be considered phenomena. Consider the chunking effect. Bogen and Woodward use this as an example of a phenomenon by virtue of its stability and repeatability and by virtue of the assumption that it exists “in the world.” Nonetheless, it is also clearly an experimental effect. Experimental effects, insofar as they are stable and repeatable, are phenomena in the sense I use the term here. Hence, I argue that phenomena can occur on a scale between the context-specific and local and the context-transcendent (see Feest [2011b] for a similar point), although it is typically assumed that local effects instantiate context-transcendent ones. If, as I contend here, objects of research (such as the presumed referents of the taxonomic categories in psychology) are often composed of multiple phenomena, the question is how can experimentally produced phenomena contribute to the ways in which objects of research are described and constructed. I argue that experimentally created phenomena often function as evidence in the course of scientific attempts to provide empirical descriptions of objects of research. In this vein, I suggest a shift away from the prevailing concern with phenomena as explananda toward a consideration of phenomena as experimental evidence (see also Cummins 2000).

At first sight, this latter claim (that phenomena can function as evidence) may be seen as flying in the face of the very distinction at the heart of Bogen and Woodward’s account, namely, that between phenomena (which are the explananda of general theory) and data (which are thought to provide evidence in favor of specific phenomena). While this distinction is very instructive, I think the notion of an object of research puts into sharper relief what purposes experimental evidence serves in the first place. And while I agree with Bogen and Woodward that experimental data are typically not explananda, this does not mean that they are not phenomena. It just means that in the context of research practices in psychology, phenomena are not always “objects of discovery and explanation,” as Bogen and Woodward initially thought. They are, rather, evidence in relation to particular descriptive questions about objects of research.

I also wish to highlight that the local and context-specific phenomena that are created as part of the investigation of specific objects of research are—at least for the duration of the research process—conceptually tied to the research object in question. Put differently, experimentally produced phenomena can function as evidence in relation to specific objects of research, because of an (however preliminary and defeasible) assumption that they are relevant to this object of research.

In this regard, consider, again, the chunking effect. Qua experimental effect it is possible to use it to gather data about an object of research precisely

because of the assumption that the empirical effect of chunking provides epistemic access to some presumed cognitive process of chunking, which in turn is thought to be a central descriptive feature of working memory. Notice that these assumptions, however tentative, already need to be in place in order to even investigate issues like the capacity of working memory, thereby highlighting what I have in mind when I speak of working memory as an example of an epistemically blurry research object.

**4. Phenomena versus Mechanisms?** Having outlined my account of objects of research as the targets of scientists' investigative efforts, as well as the role played by phenomena, I will now argue that mechanistic accounts in cognitive science are too narrowly fixated on explanations and moreover operate with an unclear and ambiguous conception of "phenomenon." Reading these approaches against the grain, this section will argue that my account of phenomena and objects of research promises to provide a more adequate analysis of the investigative process.

*4.1. Phenomena as the Explananda of Mechanistic Explanations.* Carl Craver, Bill Bechtel, and other philosophers of neuroscience, when using the term "phenomenon," often do so in the context of talking about mechanistic explanation, where by "phenomena" they typically mean explanandum phenomena. There are three problems with this. First, the notion of phenomenon, as it is used in this literature, is not entirely clear. Second (and relatedly), it is typically not explicated what the relevant explanandum phenomena of a given mechanistic explanation are. Third, even with a reasonable construal of the explananda of psychological research, I will argue in a moment that the research process is still not very well portrayed as being focused on explanation.

Consider Craver's (2007) well-known discussion of spatial memory as an object of explanation in cognitive neuroscience. Central assumptions about spatial memory are that it enables us to navigate in space, that it involves spatial maps, and that such maps are formed in the hippocampus by means of the mechanism of long-term potentiation. Now, given Craver's claim that this is a good example of how mechanistic explanations function in cognitive neuroscience, we may ask what the explanandum phenomenon is in this case. The answer Craver and Darden (2001, 121) give to this question is that "the phenomenon to be explained" is "spatial memory," suggesting that it will be explained by a multilevel mechanism. However, this explication of the explanandum phenomenon does not fit well with their notion of phenomenon a few pages earlier, where they suggest that they use the term "phenomenon" in the sense of Bogen and Woodward (1988), "as relatively stable and repeatable properties or activities that can be produced, manipulated or detected in a variety of experimental arrangements" (Craver and Darden 2001, 114). As

examples, they name “the acquisition, storage, and retrieval of spatial memories; the release of neurotransmitters, and the generation of action potentials.” Notice that all of these “phenomena” are more specific and confined than “the phenomenon” of spatial memory, thus raising the question what precisely the authors take to be the explanandum phenomenon of this research.

It seems to me that the authors run together two meanings of “phenomenon” outlined above, that is, on the one hand, specific regularities presumed to belong to a research object (acquisition, storage, retrieval, etc., and perhaps also patterns of search behavior of rats in mazes) and, on the other hand, more complex objects (e.g., spatial memory), conceivably composed of many specific regularities. Related to this, they also do not clearly distinguish between two possible (kinds of) explananda of psychology, namely, (1) the behavioral regularities triggered by particular experimental setups and (2) the capacity of an organism to execute a variety of behaviors in more complex environments (see Chow [2015] for a similar point). I will now briefly consider each of these two candidate explanandum phenomena, arguing that the former are often used as empirical data in support of specific descriptions of the latter.

We begin with the first interpretation of explanandum phenomenon. There are several possible readings of Craver’s (2007, 165) claim that the navigational behavior of mice is “the spatial memory phenomenon.” First, scientists might actually want to explain a particular experimental mouse’s behavior. Second, they want to explain the behaviors of mice more generally, extrapolating from an explanation of experimental mice behavior to an explanation of the behavior of mice in the wild. Third, they want to explain human navigational behavior, extrapolating from the explanation of experimental mice behavior to the explanation of human behavior. Now, again, I am not questioning the assumption that we ultimately want the concept of spatial memory to be an explanatory construct. Nor do I (in this article) question mechanistic accounts of explanation. What I do question, however, is the underlying assumption that experimentally generated phenomena, such as the stable and repeatable regularities of navigational behavior exhibited by mice in mazes, are typically intended to instantiate explananda. Instead, I suggest that they are created to provide evidence in favor of particular descriptive accounts of spatial memory, understood here as an object of research in the sense outlined above.

This brings me to my second interpretation of “the explanandum phenomenon of spatial memory,” according to which the explananda in cognitive science are often not isolated behavioral phenomena but the complex behavioral capacities themselves. Craver (2007) distances himself from Cummins’s use of the term “capacity,” but it strikes me as fairly innocuous, as long as we appreciate that the term can be understood as either (1) referring to an explanandum phenomenon (a complex behavioral capacity of a

system) or (2) a placeholder for a detailed descriptive account of a multi-level mechanism (i.e., an explanans). Here, my interest is in the former. I argue that once we are asking what a given behavioral capacity really is, we are inquiring into the empirical characteristics of an elusive research object, and the answer is not simply going to be a matter of filling in the mechanical details of an already clearly delineated explanandum. Again, this draws our attention to the fact that what is at stake here are taxonomic and descriptive questions. In the context of ongoing research, I claim, experimentally generated instances of phenomena often function as evidence with regard to accounts of specific behavioral capacities, not as explananda of mechanistic explanations.

*4.2. Phenomena and Mechanisms as Objects of Discovery.* It is to the credit of mechanist philosophers that they have put not only explanation but also discovery on the agenda of philosophy of science (e.g., Bechtel and Richardson 1993; see also Darden 2006; Bechtel 2008; Craver and Darden 2013). In this context, Bechtel and Richardson have early on highlighted the conceptual openness of the research process, suggesting that phenomena frequently get “reconstituted” in the course of their investigation. It is noticeable, though, that these works are very much focused on the discovery of (explanatory) mechanisms, thus leaving intact the basic idea that the explanation of phenomena, rather than the description of research objects, is the main focus of research.

A widely held assumption in the literature is not only that phenomena are the explananda of mechanistic explanations but also that the discovery of mechanisms typically proceeds by way of constructing “higher level” (phenomenological) descriptions of the relevant explanandum phenomena, with the aim of later supplying the mechanistic explanation, which in turn can “reconstitute” the explanandum phenomenon (e.g., Bechtel and Richardson 1993). In a similar vein, Craver and Darden (2001, 2013) describe the process of discovery as a process of “going down one level” to search for the mechanisms and “going up one level” to try to describe again precisely what the mechanisms do (i.e., to describe again the phenomenon).

This seems like a sensible idea, and the authors provide some nice illustrations of how this can work (pointing to top-down experiments that manipulate rats/mice at the behavioral level and observe effects on the level of the brain and, vice versa, bottom-up experiments). However, I argue that this account at best captures a small part of research practices that play a role in discovery in memory research, which puts a lot of effort into obtaining empirical descriptions of the research objects and the phenomena that pertain to them. As an example of such descriptive questions, recall the discussion in section 2.2 about duration and capacity as descriptive features of working memory.

I suggest that, in spatial memory research, too, there are a host of unresolved descriptive questions of the following kinds: How literally should the metaphor of the spatial map be interpreted? Are spatial maps more like images or like architectural models? Are they accurate? Are they consistent? Insofar as there are distortions in spatial memory, what are these distortions caused by? What are cognitive reference points for spatial information, and how can they vary by context? Do spatial maps encode perspective, or do they represent a bird's-eye view? These are questions that have been raised in the spatial knowledge literature in the past 20 years or so. In this vein, one psychologist has argued that "despite its considerable appeal, as traditionally used, the 'cognitive map' metaphor does not reflect the complexity and richness of environmental knowledge. That knowledge comes in a variety of forms, memory snippets of maps we've seen, routes we've taken, areas we've heard or read about, facts about distances or directions" (Tversky 1993, 21).

This quotation expresses the sentiments I have tried to argue for in this article, as it does not talk about mechanisms or explanations. Rather, the author raises descriptive questions about a research object that is not very well understood or that is, as I have called it in this article, epistemically blurry. While I cannot pursue this case study here, I suggest that it offers a glimpse into the kinds of descriptive questions psychologists often concern themselves with, raising philosophical issues about the ways in which experimentally produced phenomena are used in their research.

**5. Conclusion.** I have questioned the common assumption that the investigative practices in cognitive science are directed toward (the discovery of) mechanistic explanations of phenomena and have instead drawn attention to the fact that (1) the exact contours of the relevant explanandum phenomena are often not clearly defined and (2) the very notion of a phenomenon remains contested and ambiguous in the philosophy of science literature. In this vein, I have argued that the research process is better analyzed as one that tries to construct adequate descriptions of epistemically blurry objects of research (such as working memory or spatial memory), and I have provided the outline of an analysis of the ways in which experimentally created phenomena figure in this process.

In conclusion, let me emphasize that I do not wish to downplay the importance of explanations for science or for philosophy of science. However, if we aim for philosophical accounts of scientific practice, we are well advised to take a closer look at some other concerns and questions researchers in fact struggle with. I suggest that my account of the investigative process as focused on what I call epistemically blurry objects of research offers a fruitful novel account of the types of questions that arise at the intersection of theoretical and methodological debates in psychology.

## REFERENCES

- Baddeley, Alan, and D. J. Hitch. 1974. "Working Memory." In *The Psychology of Learning and Motivation: Advances in Research and Theory*, ed. Gordon Bower, 8:47–78. New York: Academic Press.
- Beaman, C. P. 2000. "The Size and Nature of a Chunk." *Behavioral and Brain Sciences* 24:118.
- Bechtel, William. 2008. *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. London: Routledge.
- Bechtel, William, and Robert Richardson. 1993. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton, NJ: Princeton University Press.
- Bogen, James, and James Woodward. 1988. "Saving the Phenomena." *Philosophical Review* 97 (3): 303–52.
- Chow, Sheldon. 2015. "The Gap between Psychological Explanation and Mechanistic Explanation." Paper presented at the 2015 meeting of the European Philosophy of Science Association.
- Cowan, Nelson. 2000. "The Magical Number 4 in Short-Term Memory: A Reconsideration of Mental Storage Capacity." *Behavioral and Brain Sciences* 24:87–115.
- Craver, Carl. 2007. *Explaining the Brain: The Mosaic Unity of the Cognitive Neurosciences*. Oxford: Clarendon.
- Craver, Carl, and Lindley Darden. 2001. "Discovering Mechanisms in Neurobiology: The Case of Spatial Memory." In *Theory and Method in the Neurosciences*, ed. Peter Machamer, Rick Grush, and Peter McLaughlin, 112–37. Pittsburgh: University of Pittsburgh Press.
- . 2013. *In Search of Mechanisms: Discovery across the Life Sciences*. Chicago: University of Chicago Press.
- Cummins, Robert. 2000. "'How Does It Work?' versus 'What Are the Laws?': Two Conceptions of Psychological Explanation." In *Explanation and Cognition*, ed. Frank Keil and Robert Wilson, 117–44. Cambridge, MA: MIT Press.
- Darden, Lindley. 2006. *Reasoning in Biological Discoveries*. New York: Cambridge University Press.
- Feest, Uljana. 2011a. "Remembering (Short-Term) Memory: Oscillations of an Epistemic Thing." *Erkenntnis* 75:391–411.
- . 2011b. "What Exactly Is Stabilized When Phenomena Are Stabilized?" *Synthese* 182:57–71.
- Hacking, Ian. 1983. *Representing and Intervening*. Cambridge: Cambridge University Press.
- Halford, Graeme, Steven Phillips, and William Wilson. 2000. "Processing Capacity Limits Are Not Explained by Storage Limits." *Behavioral and Brain Sciences* 24:123–25.
- Miller, George. 1956. "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information." *Psychological Review* 101 (2): 343–52.
- Milner, P. 2000. "Magical Attention." *Behavioral and Brain Sciences* 24:131.
- Pascual-Leone, Juan. 2000. "If the Magical Number Is 4, How Does One Account for Operations within Working Memory?" *Behavioral and Brain Sciences* 24:136–38.
- Schubert, Torsten, and Peter A. Frensch. 2000. "How Unitary Is the Capacity-Limited Attentional Focus?" *Behavioral and Brain Sciences* 24:146–47.
- Tversky, Barbara. 1993. "Cognitive Maps, Cognitive Collages, and Spatial Mental Models." In *Spatial Information Theory: A Theoretical Basis for GIS*, ed. Andrew U. Frank and Irene Campari, 14–24. Berlin: Springer.
- Wilson, Mark. 2006. *Wandering Significance: An Essay on Conceptual Behavior*. Oxford: Oxford University Press.
- Woodward, Jim. 1989. "Data and Phenomena." *Synthese* 79:393–472.