

# Construct Stabilization and the Unity of the Mind-Brain Sciences

Jacqueline Anne Sullivan\*†

---

This article offers a critique of an account of explanatory integration that claims that explanations of cognitive capacities by functional analyses and mechanistic explanations can be seamlessly integrated. It is shown that achieving such explanatory integration requires that the terms designating cognitive capacities in the two forms of explanation are stable but that experimental practice in the mind-brain sciences currently is not directed at achieving such stability. A positive proposal for changing experimental practice so as to promote such stability is put forward, and its implications for explanatory integration are briefly considered.

---

**1. Introduction.** Debates about the unity of the mind-brain sciences have been reinvigorated in recent years as new accounts of the nature of explanation in psychology and neuroscience have been introduced into the philosophical literature. Whereas previous versions of the debate focused on whether psychological theories could be reduced to neuroscientific theories—a possibility blocked by the argument for the multiple realizability of psychological kinds at the neural level (Fodor 1974)—the new debate concerns whether a unified science of cognition can be achieved via the integration of psychological and neuroscientific explanations. Advocates of mechanistic explanation (Piccinini and Craver 2011) argue that cognitive psychology is not autonomous from neuroscience because explanations of cognitive capacities by functional analysis are simply incomplete mechanistic explanations. Once the structural de-

\*To contact the author, please write to: Department of Philosophy, 3148 Stevenson Hall, 1151 Richmond St., Western University, London, ON N6A 5B8, Canada; e-mail: jsulli29@uwo.ca.

†The author would like to thank Muhammad Ali Khalidi, an anonymous referee, and members of the Rotman Institute of Philosophy's 2015 Summer Writing Workshop, including Frédéric-Ismaël Banville, Danny Booth, Robert Foley, John Jenkinson, Andrew Peterson, Nicholas Slothouber, and Jesse Wright, for very helpful comments on an earlier draft of this paper.

Philosophy of Science, 83 (December 2016) pp. 662–673. 0031-8248/2016/8305-0002\$10.00  
Copyright 2016 by the Philosophy of Science Association. All rights reserved.

tails—the physical entities and activities that realize cognitive capacities—are “filled in,” explanations by functional analysis become “full-blown mechanistic explanations” (Piccinini and Craver 2011, 283). As a consequence of successful explanatory integration, cognitive psychology and neuroscience will come to form a unified science of cognition.

The aim of this paper is to argue that experimental practice in cognitive psychology and neuroscience is not conducive to the type of explanatory integration Piccinini and Craver advocate. In section 2, I outline the main features of the account of explanatory integration they put forward. In section 3, I make the case that the integration of functional analyses and mechanistic explanations requires that components of the two types of explanation, namely, cognitive capacities, are *stable*. I define stability by appeal to conceptual tools on offer in the theoretical literature in psychology and the social sciences and identify certain facts about experimental practice in cognitive psychology and neuroscience that have contributed to the instability of constructs designating cognitive capacities. In section 4, I propose some changes to experimental practice conducive to stabilizing these constructs and consider the implications for explanatory integration.

**2. A Unified Science of Cognition.** Explanations in neuroscience, insofar as they describe the physical entities/components and activities/processes that realize phenomena of interest, have been characterized as *mechanistic* (e.g., Craver 2007; Bechtel 2008). Given the complex nature of the kinds of phenomena mechanistic explanations are intended to explain, their development is taken to require input from multiple different laboratories and areas of neuroscience situated at multiple different levels of analysis. To take a celebrated example from the philosophical literature on mechanistic explanation, activation of N-methyl-D-aspartate (NMDA) receptors in area CA1 of the rat hippocampus is one component in the description of the multilevel mechanism of rodent spatial memory. According to Craver, such explanations arise as findings from many different cellular, molecular, and behavioral neuroscience laboratories are “integrated” into descriptions of multilevel mechanisms (Craver 2007).

In contrast to mechanistic explanations, explanations in cognitive psychology are “explanations by functional analysis” (e.g., Cummins 1983) and are used to explain mental functions or processes without regard for anatomical, structural, biochemical, or physiological facts about brains. According to Jerry Fodor (1968, 107–8), “the psychologist [seeks] functional characterizations of psychological constructs,” and “the criteria employed for individuating such constructs are based primarily on hypotheses about the role they play in the etiology of behavior.” Cognitive psychologists design complex tasks in order to tease apart distinct cognitive processes by appeal to subjects’ behavioral performance on those tasks. The resulting expla-

nations are sometimes depicted by means of box and arrow diagrams where the boxes stand in for psychological capacities (e.g., working memory) and the arrows represent the input-output/feed-forward/feed-backward connections or information flow from stimulus inputs to behavioral outputs. In contrast to the mechanistic explanation of spatial memory provided above, an early explanation of spatial memory by functional analysis described an “‘internal navigation’ system” that received sensory data and “movement feedback” from the motor system and sent information to a “map construction system” (O’Keefe and Nadel 1978, 94).<sup>1</sup>

Although cognitive psychology and neuroscience are regarded as distinct scientific enterprises, Piccinini and Craver (2011) have recently argued that the two fields are not *explanatorily* autonomous.<sup>2</sup> While both areas of science aim to explain cognitive capacities like spatial memory, they claim that only neuroscience is successful insofar as it identifies both the functional and structural details—the activities and the entities—of the systems that realize cognitive capacities. Piccinini and Craver may be described as conceiving of the two forms of explanation as situated at different points on an explanatory completeness continuum. Functional analyses or “mechanism sketches” lie at one end; complete mechanistic explanations of cognitive capacities lie at the other. Once neuroscience fills in “the structural aspects that are missing from a functional analysis,” it “turns into a more complete mechanistic explanation” (Piccinini and Craver 2011, 308). To return to our example, “the cognitive map system,” which was originally a component of an explanation of spatial behavior by functional analysis, may be described as being later “filled in” with a brain structure, namely, the hippocampus (e.g., O’Keefe and Nadel 1978). At that point, the entities and activities of the hippocampus, namely, place cells in area CA1, became relevant to explaining how the hippocampus comes to produce a cognitive map. Craver’s (2007) depiction of the mechanism of spatial memory thus may be regarded as an explanation by functional analysis that has since moved further on down the explanatory completeness continuum.<sup>3</sup> Such examples at first blush appear to support Piccinini and Craver’s idea that “functional analyses can be seamlessly integrated with mechanistic explanations, and psychology can be seamlessly integrated with neuroscience” (2011, 308).

1. O’Keefe and Nadel (1978, 89–101) outline the “psychological basis” of cognitive maps.

2. Craver is amenable to scientists making autonomous decisions to have their mechanistic explanations “bottom out” where they see fit. This does not preclude another investigator locating the bottom somewhere else.

3. While Piccinini and Craver do not appeal to this example to support their argument, it instantiates the kinds of features they have in mind.

In addition to what appear to be successful explanations like that of spatial memory, which support the idea that explanations by functional analysis and mechanistic explanations are being integrated, Piccinini and Craver's argument derives support from methodologically integrative scientific areas like neuropsychology and cognitive neuroscience, whose very existence may be taken to suggest that cognitive psychology cannot advance our understanding of cognition in the absence of neuroscience. Although many neuropsychologists uphold the information processing view of the mind characteristic of cognitive psychology and use behavioral tasks to decompose cognitive processes into their component subprocesses, they regard comparing task performance of normal subjects with that of subjects with brain lesions and neurological disorders essential for such functional decomposition. While many cognitive neuroscientists also endorse an information processing view of the mind and use behavioral tasks designed to individuate cognitive processes, they combine these methods with imaging (e.g., functional magnetic resonance imaging), recording (e.g., electroencephalography), and intervention (e.g., transcranial magnetic stimulation) techniques that are intended to facilitate the localization of such processes in the brain. These two methodologically integrative fields, at least at first blush, provide good grounds for thinking that Piccinini and Craver are right and explanations of cognitive capacities by functional analysis alone are insufficient; knowledge about the structural details of brains that realize those capacities is relevant. However, as I aim to show in the next section, when we look more closely at these areas of science, we realize that they are not currently on a trajectory toward integrating functional analyses with mechanistic explanations because current practice both within and across the relevant areas of science is not directed at stabilizing the meanings of the terms designating cognitive capacities that occur in the two forms of scientific explanation.

**3. Construct Stabilization as Prerequisite for Integration.** Historically, advocates for unity of science have argued for *theory reduction* (e.g., Nagel 1961). Although Piccinini and Craver advocate for unity via explanatory integration, as I will show, at least one of the traditional constraints on inter-theoretic reduction, *connectability*, is presupposed by their account. The basic idea behind the connectability condition is that theories contain terms that have certain referents, and for two theories to be participants in a successful reduction relation, "a bridge law" must be established, which specifies that the referents of the terms in the theory to be reduced are bidirectionally equivalent to the referents of the terms in the reducing theory. The classic example of successful satisfaction of the connectability condition is the reduction of the term "temperature of a gas" in thermodynamic theory to "mean kinetic energy of the molecules" in statistical mechanics.

The connectability condition applies to explanatory integration insofar as the explanations that are candidates for integration must have the same referents. More specifically, the terms designating cognitive capacities in an explanation by functional analysis must have roughly the same referents as the terms designating cognitive capacities in a mechanistic explanation. To refer back to the example in the previous section, an explanation by functional analysis that contains the term *spatial memory* ought to refer to the same phenomenon as a mechanistic explanation that contains the term *spatial memory*. As Piccinini and Craver claim, whereas explanations by functional analysis identify capacities and subcapacities, mechanistic explanations identify capacities, subcapacities, and the structural parts of brains and their activities that realize those capacities. Terms designating cognitive capacities are the common denominator between the two forms of explanation, and satisfying the connectability condition requires that the terms designate the same thing. Otherwise, what we have is not explanatory integration, but elimination and replacement of terms in one area of science for the other.

My aim in the rest of this section is to demonstrate that a prerequisite for connectability—*construct stability*—cannot be met because the terms designating cognitive capacities in cognitive psychology, and particularly in neuroscience, do not have stable referents, and experimental practice in these areas of science currently is not directed at securing such stability. In order to make my case, some conceptual tools for thinking about how cognitive capacities are investigated experimentally and how theoretical constructs attain stability in sciences that study cognitive capacities are relevant.

The starting point for my analysis is the individual laboratory. This choice of starting point is justified by virtue of the fact that Piccinini and Craver identify two ways explanatory integration comes about. The first is described above: mechanistic explanations fill in the structural details of explanations by functional analysis. This kind of integration seems to involve already-developed and stable functional components of functional analyses and/or mechanistic explanations being integrated together. However, Piccinini and Craver also identify another form of explanatory integration that involves “the integration of findings from different areas of neuroscience and psychology into a description of multilevel mechanisms” (2011, 285). Findings about cognitive capacities originate in individual laboratories. So, if we are interested in whether the connectability condition is being met, our analysis should begin with intralab practices for stabilizing constructs designating cognitive capacities and be extended to interlab practices across laboratories. In putting forward this set of conceptual tools, I am interested primarily in those features of experimental practice that those areas of cognitive psychology and neuroscience that study cognitive capacities have in common, so as to use these tools as a basis for identifying differences in these features.

When a cognitive psychologist or neuroscientist goes into the laboratory to investigate a cognitive capacity, she will have likely grouped together instances of what she takes to be the same capacity under a concept or construct. She may rely on how other investigators in her field define the concept, but she may also define it slightly differently. Examples of constructs that designate cognitive capacities in cognitive psychology and neuroscience include *spatial memory*, *working memory*, *attention*, *face recognition*, and *procedural memory* (to name only a handful). Such constructs originate with a concept that investigators associate with certain observations, which serves as a basis for theory building and experimental task/paradigm design and construction.

Once an investigator has selected a cognitive capacity of interest, which is designated by a construct, she then develops an experimental paradigm—a set of procedures for producing, measuring, and detecting an instance of that capacity in the laboratory. For example, an experimental paradigm used to investigate a cognitive capacity like spatial memory will include a set of production procedures that specify the stimuli (e.g., distal and local cues) to be presented, how those stimuli are to be presented/arranged (e.g., spatially, temporally), and how many times each stimulus is to be presented during phases of pre-training, training, and post-training/testing. The paradigm will also include measurement procedures that specify the response variables to be measured in pre-training and post-training/testing phases of the experiment and how to measure them using apparatuses designed for such measurement. Finally, a set of detection procedures specifies what the comparative measurements of the response variables from the different phases of the experiment must equal in order to ascribe the cognitive capacity of interest to the organism and/or the locus of the function to a given brain area or neuronal population.

An investigator will, in the ideal case, aim to design an experimental paradigm that produces an instance of the kind of capacity she intends to detect and measure. She ought to want the match between the effect she produces in the laboratory and the phenomena she takes to be grouped together under the general construct to be valid. Another way to put this is that she aims for the experimental paradigm she selected to have a high degree of “construct validity.” Construct validity “is involved whenever a test is to be interpreted as a measure of some attribute or quality which is not operationally defined” (Cronbach and Meehl 1955, 282). It “involves making inferences from the sampling particulars of a study to the higher-order constructs they represent” (Shadish, Cook, and Campbell 2002, 65). Experimental paradigms or cognitive tasks may have anywhere from a low to high degree of construct validity. The higher the degree of construct validity, the closer the match between the effect under study in the laboratory and the cognitive phenomena designated by the construct.

It is important to note that the experimental process within any given laboratory is rarely one-shot. Oftentimes, an investigator and/or his/her critics wonder whether the investigative procedures she has used in the laboratory satisfy the criterion of construct validity. Such worries prompt the processes of “construct explication” and “construct assessment” (Shadish et al. 2002). These processes may be understood in terms of a series of questions that ideally become a fundamental part of the experimental process. Specifically, an investigator asks the following at the relevant stages of this process: (1) Which instances of worldly phenomena should be grouped together under the concept designating the construct? (2) Which investigative strategies will yield instances that instantiate it? (3) Are the investigative strategies adequate, or should they be modified? (4) Given the data these investigative strategies yield, should the construct be revised to exclude phenomena that do not belong in the category or to include additional phenomena that do?<sup>4</sup>

Returning to Piccinini and Craver’s account of explanatory integration, it is important to note that construct stabilization will involve more than a single lab and more than a single area of science. In other words, stabilizing constructs via processes like construct explication and construct assessment will involve coordination across labs situated in the same and different areas of science to come to specific agreement about (1) how to generally define terms, (2) what are the best experimental paradigms for studying a given cognitive capacity, and (3) the conditions under which two experimental paradigms can be said to measure the same cognitive capacity. Yet, do we encounter such coordination in the form of a consistent emphasis on construct validation/explication/assessment across laboratories and investigators in the same and different areas of cognitive psychology and neuroscience?

A proper answer to this question requires investigating the stability of constructs designating cognitive capacities in the sciences that study cognition on a case-by-case basis,<sup>5</sup> a project that cannot be undertaken in the context of a single paper. Instead, the current approach is to point to facts that are suggestive that the meaning of constructs designating cognitive capacities is not stable in the sciences studying cognition owing in large part to the fact that strategies for stabilizing constructs are not consistently adopted across investigators and research areas.

Let’s begin by considering construct stabilization in cognitive psychology. As a long-standing scientific tradition, one of its paradigmatic features is to educate its members on the importance of engaging in rigorous task analyses to determine the component cognitive processes operative in the production of behavioral data. This should provide us with some confidence

4. Adapted from Shadish et al. (2002, 66).

5. Piccinini and Craver do not provide an example of a psychological explanation by functional analysis successfully integrated with a mechanistic explanation.

that intralab strategies are in place to stabilize constructs designating cognitive capacities. This does not mean, however, that interlab practices are conducive to stability. For example, two investigators may be interested in studying spatial memory in the rodent but disagree about the most suitable task for this purpose. One investigator may use the Morris water maze, and another, the elevated T-maze. Yet, stimuli and task demands differ radically between these two tasks, and it is difficult to tease apart the component cognitive processes involved in each.<sup>6</sup> Investigators also often disagree about which component cognitive processes are involved in the production of a given set of behavioral data, and often the behavioral data are compatible with multiple different explanations by functional analysis.

Piccinini and Craver might respond that the way to overcome such underdetermination is by investigating the brain structures that realize the cognitive processes in question. This is because, as they claim, structure places constraints on function; structure determines the kinds of cognitive processes that can be realized and how. It is at this point that they advocate a move to cognitive neuroscience and toward explanatory integration. Yet, there are certain challenges that this move faces. One concerns the limitations of the method of reverse inference (e.g., Poldrack 2006). A second problem, with which I am concerned here, is that successful explanatory integration requires, at a bare minimum, that the constructs designating cognitive capacities are stable and thus connectable between the two areas of science.

There are good reasons, however, to think that this is not the case. First, cognitive neuroscientists do not agree among themselves about whether achieving construct validity and engaging in construct explication or construct assessment are important. Some investigators do aim to identify the component cognitive processes thought to be engaged in experimental tasks and determine how the variables manipulated in an experiment affect these processes (see Sullivan 2014a, 2014b). However, Russell Poldrack suggests that many cognitive neuroscientists rarely engage in such task analysis at all: “Unfortunately, . . . task analyses are very rarely presented in neuroimaging papers. Whereas formal theories from cognitive psychology could often provide substantial guidance in the design of such tasks, it is uncommon for neuroimaging studies to take meaningful guidance from such theories. Rather, the task comparisons in many studies are based on intuitive judgments regarding the cognitive processes engaged by a particular task” (2010, 149). In other words, task analysis, which is a component of construct explication and assessment, is not something that currently occurs across laboratories or investigators in a consistent, coordinated way.

6. For example, Morris’s “key message” in a recent book chapter on the water maze is that it “is not just one task, but a family of procedures suited to diverse scientific questions” (2015, 73).



Another factor contributing to construct instability in cognitive neuroscience is the far-reaching methodological pluralism. If we look across labs in cognitive neuroscience and do a comparative analysis, we encounter “a multiplicity of experimental protocols” (Sullivan 2009) insofar as investigators often do not agree on which experimental paradigms ought to be used to investigate a given cognitive capacity, and they have freedom to design tasks as they deem most appropriate to their explanatory goals. Carrie Figdor (2011) puts the point nicely in claiming that the terms used to designate kinds of cognitive capacities do not have stable meanings; even if different investigators use the same term to refer to a kind of cognitive function or a kind of experiment, it does not mean that they intended to designate the “same” cognitive function by means of the term. Investigators may also look at the very same task and yet disagree about the component processes involved given either, as Poldrack claims, their intuitive judgments or prior theoretical commitments.

Cognitive neuroscientists, like Poldrack, acknowledge the widespread construct instability in cognitive neuroscience (and lack of a proper cognitive ontology) and have offered solutions that have yet to be broadly implemented in practice. Some claim that to localize cognitive functions we need a coordinated effort to develop a taxonomy of more general constructs (e.g., “sensory-motor integration”) that are more suitable for capturing what particular brain areas do (Price and Friston 2005). Others claim that we need coordinated efforts to develop “process pure” tasks that individuate finer-grained constructs than those on offer in cognitive psychology (see Sullivan 2014b). In addition to Poldrack’s suggestions to develop cognitive tasks more appropriate to functional localization (2006) and to engage in more rigorous task analysis (2010), he advocates the use of meta-analyses and data-mining techniques as a basis for assessing the strength of hypotheses about what functions specific brain areas are performing (2006).

These facts, taken in combination, provide grounds for doubting that the constructs designating cognitive capacities in cognitive neuroscience are stable in the way required for Piccinini and Craver’s explanatory integration. Further, while various investigators working in cognitive neuroscience have begun to acknowledge the problem and to itemize its sources, there currently is no agreed-on panacea. Part of the problem is that cognitive neuroscience, insofar as it is integrative, is eclectic. Investigators do not necessarily share a Kuhnian paradigm in common. However, one important theme that arises is the continued importance of the perspective and methods of cognitive psychology for individuating cognitive capacities and stabilizing constructs in cognitive neuroscience. This suggests that what is needed in integrative areas of neuroscience is the preservation of a plurality of perspectives, as well as the promotion of perspectives likely to aid in the achievement of integrative explanatory goals. Establishing that perspectival

pluralism specifically is necessary for explanatory integration is the aim of the next section.

**4. Perspectival Pluralism and Explanatory Integration.** Cognitive psychologists and cognitive neuroscientists adopt different ontological perspectives on cognitive systems insofar as they appeal to different “set[s] of variables . . . to characterize” and “partition” those “systems . . . into parts”; these perspectives directly inform how these investigators “interact causally with [those] system[s]” (Wimsatt 2007, 227) and impact how they design their experiments (see also Giere 2010). For example, when Richard Morris designed the water maze, he was interested in the construct “place learning”—the cognitive ability to find a hidden target in the absence of local cues. He adopted an information processing view of the mind, and this intimately shaped the experimental design of the water maze (see Sullivan 2010). Data from his experiments originally led him to conclude that the water maze individuated place learning.

However, a separate and later research study (Eichenbaum, Stewart, and Morris 1990) revealed that rats with hippocampal lesions—the structure thought to underlie place learning—could still perform successfully in the water maze. These results, which were obtained when investigators adopted an information processing view of the brain and its structures, suggested, contrary to Morris’s findings, that the water maze does not individuate a discrete cognitive capacity. Rather, other cognitive processes (e.g., nonspatial, associative) are involved.

In contrast, cognitive neurobiologists, who use the water maze to study cellular and molecular activity, are not concerned with these constituent information processes. Their failure to recognize that the water maze involves multiple distinct cognitive processes has likely contributed to the instability of the construct used to designate the phenomenon under study in the water maze (see Sullivan 2010). It has also resulted in mechanistic explanations that lack clear explananda, like the claim that NMDA-receptor activation in the hippocampus is a necessary component of the mechanism of “spatial memory.” As evidence in support of this point, investigators with training in cognitive psychology working in collaboration with Morris raised the question why rats with blocked NMDA receptors fail to perform successfully in the water maze. They employed a battery of cognitive tests designed to identify “what” informational processes are disrupted by NMDA-receptor blockade and “what” information rats actually learn in the water maze. In taking this information processing perspective, they demonstrated that NMDA-receptor activation likely “disrupts non-spatial as well as spatial components of water maze learning” (Bannerman et al. 1995, 185).

The water maze illustrates nicely that stabilizing constructs designating cognitive capacities is an iterative process that requires multiple distinct

perspectives to be operative when experimental paradigms are being designed and implemented in the lab and the resulting data are being interpreted. Further, it shows that ensuring stabilization of constructs used to designate cognitive capacities requires that investigators engage in the process of construct explication. In other words, cognitive psychology does not have a time-limited role to play in explanatory integration; its involvement should be ongoing.

That perspectival pluralism of the form I am advocating is essential for explanatory integration also derives support from two recent initiatives spearheaded by the National Institutes of Mental Health: the Research Domain Criteria (RDoC) project and the Cognitive Neuroscience Treatment Research to Improve Cognition in Schizophrenia (CNTRICS) initiative. Investigators who are involved in these interdisciplinary initiatives include cognitive psychologists, cognitive neuroscientists, experts in animal behavior, cognitive neurobiologists, clinical pharmacologists, and members of industry. They all share in common the aim of developing experimental paradigms to identify the cognitive and behavioral capacities that are disrupted in persons with mental illnesses, so that treatments for these dysfunctions may be identified. They believe that this aim can only be achieved if the different perspectives they represent each play a role in the design, implementation, and revision of experimental paradigms (Sullivan 2014b).

**5. Conclusion.** Explanatory integration requires stable explanatory targets, stable constructs. We do not have such stability in the neurosciences of cognition. Perspectival pluralism of the form advocated here might be a viable means of achieving it. Indeed, recent initiatives in mental health research emphasize the importance of perspectival pluralism for explanatory integration.

#### REFERENCES

- Bannerman, David, Mark Good, Steven Butcher, Mark Ramsey, and Richard Morris. 1995. "Distinct Components of Spatial Learning Revealed by Prior Training and NMDA Receptor Blockade." *Nature* 378:182–86.
- Bechtel, William. 2008. *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. New York: Taylor & Francis.
- Craver, Carl. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Oxford University Press.
- Cronbach, Lee, and Paul Meehl. 1955. "Construct Validity in Psychological Tests." *Psychological Bulletin* 52:281–302.
- Cummins, Robert. 1983. *The Nature of Psychological Explanation*. Cambridge, MA: MIT Press.
- Eichenbaum, Howard, Caroline Stewart, and Richard Morris. 1990. "Hippocampal Representation in Place Learning." *Journal of Neuroscience* 10 (1): 3531–42.
- Figdor, Carrie. 2011. "Semantics and Metaphysics in Informatics: Toward an Ontology of Tasks." *Topics in Cognitive Science* 3:222–26.
- Fodor, Jerry. 1968. *Psychological Explanation: An Introduction to the Philosophy of Psychology*. New York: Random House.

- . 1974. “Special Sciences (or: The Disunity of Science as a Working Hypothesis).” *Synthese* 28:97–115.
- Giere, Ronald. 2010. *Scientific Perspectivism*. Chicago: University of Chicago Press.
- Morris, Richard. 2015. “The Watermaze.” In *The Maze Book: Theories, Practice and Protocols for Testing Rodent Cognition*, ed. Heather Bimonte-Nelson, 73–92. New York: Springer.
- Nagel, Ernest. 1961. *The Structure of Science: Problems in the Logic of Scientific Explanation*. New York: Harcourt, Brace & World.
- O’Keefe, John, and Lynn Nadel. 1978. *The Hippocampus as a Cognitive Map*. Oxford: Clarendon.
- Piccinini, Gualtiero, and Carl Craver. 2011. “Integrating Psychology and Neuroscience: Functional Analysis as Mechanism Sketches.” *Synthese* 183 (3): 283–311.
- Poldrack, Russell. 2006. “Can Cognitive Processes Be Inferred from Functional Imaging Data?” *Trends in Cognitive Sciences* 10 (2): 59–63.
- . 2010. “Subtraction and Beyond: The Logic of Experimental Designs for Neuroimaging.” In *Foundational Issues in Human Brain Mapping*, ed. Stephen Hanson and Martin Bunzl, 147–59. Cambridge, MA: MIT Press.
- Price, Cathy, and Karl Friston. 2005. “Functional Ontologies for Cognition: The Systematic Definition of Structure and Function.” *Cognitive Neuropsychology* 22 (3/4): 262–75.
- Shadish, William, Thomas Cook, and Donald Campbell. 2002. *Experimental and Quasi-experimental Designs for Generalized Causal Inference*. Boston: Houghton Mifflin.
- Sullivan, Jacqueline A. 2009. “The Multiplicity of Experimental Protocols: A Challenge to Reductionist and Nonreductionist Models of the Unity of Science.” *Synthese* 167:511–39.
- . 2010. “Reconsidering Spatial Memory and the Morris Water Maze.” *Synthese* 177 (2): 261–83.
- . 2014a. “Is the Next Frontier in Neuroscience a Decade of the Mind?” In *Brain Theory*, ed. Charles Wolfe, 45–67. New York: Palgrave Macmillan.
- . 2014b. “Stabilizing Mental Disorders: Prospects and Problems.” In *Classifying Psychopathology: Mental Kinds and Natural Kinds*, ed. Harold Kincaid and Jacqueline Sullivan, 257–81. Cambridge, MA: MIT Press.
- Wimsatt, William. 2007. *Re-engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge, MA: Harvard University Press.