

# The Explanatory Power of Network Models

Carl F. Craver\*†

---

Network analysis is increasingly used to discover and represent the organization of complex systems. Focusing on examples from neuroscience in particular, I argue that whether network models explain, how they explain, and how much they explain cannot be answered for network models generally but must be answered by specifying an explanandum, by addressing how the model is applied to the system, and by specifying which kinds of relations count as explanatory.

---

**1. Introduction.** Network analysis is a field of graph theory dedicated to the organization of pairwise relations. It provides a set of concepts for describing kinds of networks (e.g., small-world or random networks) and for describing and discovering organization in systems with many densely connected components.

Some philosophers describe network models as providing noncausal or nonmechanistic forms of explanation. Lange (2013) and Huneman (2010), for example, argue that network models offer distinctively mathematical (and topological) explanations (sec. 3). Levy and Bechtel (2013) argue that network models describe abstract causal structures in contrast to detailed forms of causal explanation (sec. 4). Rathkopf (2015) distinguishes network explanations from mechanistic explanations because network models apply to non-decomposable systems (sec. 6).

My thesis is that whether network models explain, how they explain, and how much they explain cannot be answered for network models generally but must be answered by fixing an explanandum phenomenon, considering how the model is applied to a target, and deciding what sorts of variables and

\*To contact the author, please write to: Department of Philosophy, University of Washington in St. Louis, 1 Brookings Drive, St. Louis, MO 63130; e-mail: ccraver@wustl.edu.

†Thanks to Steve Petersen and the McDonnell Center for Systems Neuroscience for funding leave time, and to Petersen lab members, especially Becky Coalson, Tim Laumann, and Haoxin Sun. Thanks to Andre Ariew, Mike Dacey, Eric Hochstein, David Kaplan, Anya Plutynski, Gualtiero Piccinini, Mark Povich, and Felipe Romero for comments.

Philosophy of Science, 83 (December 2016) pp. 698–709. 0031-8248/2016/8305-0005\$10.00  
Copyright 2016 by the Philosophy of Science Association. All rights reserved.

relations count as explanatory (see Craver 2014). Uses of network models to describe mechanisms are paradigmatically explanatory, even if the model is abstract, and even if the system is complex, because such models represent how a phenomenon is situated in the causal structure of the world (Salmon 1984). Network models might also provide distinctively mathematical explanations (those that derive their explanatory force from mathematics rather than, e.g., causal or nomic relations; see Lange 2013). Using examples from network science, I illustrate two puzzles for any adequate account of such mathematical explanations—the problem of directionality (sec. 3) and the puzzle of correlational networks (sec. 5)—that are readily solved in paradigm cases by recognizing background ontic constraints on acceptable mathematical explanations.

**2. Basic Concepts of Network Analysis.** Network models are composed of nodes, standing for the network's relata, and edges, standing for their relations. A node's *degree* is the number of edges it shares with other nodes. The *path length* between two nodes is the minimum number of edges required to link them. In a *connected graph*, any two nodes have a path between them.

In *regular networks*, each node has the same degree. In *random networks*, edges are distributed randomly. Few real networks are usefully described as regular or random. In contrast to these special cases, *small-world networks* often have an underlying *community structure* because nodes are connected into clusters. The small-worldness of a network is the ratio of its *clustering coefficient* (the extent to which nodes cluster more than they do in a random graph) to its average path length. This clustering yields networks with nearly decomposable *modules* (Simon 1962; see also Haugeland 1998), collections of nodes that have more and stronger connections with each other than they have with nodes outside the group. Nodes with a high *participation coefficient* have edges with nodes in diverse modules. Such nodes can serve as connecting *hubs* linking modules together. These and other concepts are mathematically systematized and have been used to represent epidemics, academics, indie bands, and airports, to name just a few.

**3. Directions of Mathematical Network Explanation.** Lange (2013) argues that some explanations of natural phenomena are distinctively mathematical. Such explanations work by showing that the explanandum is mathematically (vs. causally or nomically) necessary. Dad cannot evenly distribute 13 cookies between two kids because 13 is indivisible by 2. The sandpile's mass has to be 1,000 g (in a Newtonian world) because it is made of 100,000 grains of 0.01 g each. In each case, the explanandum follows with mathematical necessity in the empirical conditions.

Network models might also provide distinctively mathematical explanations. Take the bridges of Königsberg: seven bridges connect four landmasses;

nobody can walk a path crossing each bridge exactly once (an Eulerian path). An Eulerian path requires that either zero or two landmasses have an odd number of bridges. In Königsberg, all four landmasses have an odd number of bridges. So it's mathematically impossible to take an Eulerian walk around town.

Perhaps some facts about the brain have distinctively mathematical explanations. For example, a network's mean path length is more robust to the random deletion of nodes in small-world networks than it is in regular or random networks. This might explain why brain function is robust against random cell death (see Behrens and Sporns 2011). Signal propagation is faster in small-world networks than in random networks; oscillators coupled in small-world networks readily synchronize (Watts and Strogatz 1998). Perhaps these are mathematical facts, and perhaps they carry explanatory weight.

Huneman (2010) represents mathematical (specifically topological) explanations as arguments. They contain an empirical premise, asserting a topological (or network) property of a system, and a mathematical premise, stating a necessary relation. In our example of Königsberg's bridges:

**Empirical Premise.** Königsberg's bridges form a connected network with four nodes. Three nodes have three edges; one has five.

**Mathematical Premise.** Among connected networks composed of four nodes, only networks containing zero or two nodes with odd degree contain Eulerian paths.

**Conclusion.** There is no Eulerian path around the bridges of Königsberg.

And in our example for the brain:

**EP.** System S is a small-world network. . . .

**MP.** Small-world networks are more robust to random attack than are random or regular networks.<sup>1</sup>

**C.** System S is more robust to random attack than random or regular networks.

On this reconstruction, distinctively mathematical explanations are like covering law explanations (Hempel 1965) except the "law statements" are mathematically necessary (Lange 2013).

1. See Albert, Jeong, and Barabási (2000). This claim must be circumscribed to make it a necessary truth.

The covering law model struggled famously with the asymmetry of causal explanations. The flagpole's height and the sun's elevation explain the shadow's length, and not vice versa. Scientific explanations have a directionality trigonometry lacks.

Similar cases arise for distinctively mathematical explanations. Because each kid got the same number of cookies, it follows necessarily that Dad started with an even batch, but the kids' cookie count doesn't explain the batch's evenness. From the mass of the Newtonian sandpile and the number of identical grains, the mass of each grain follows necessarily, but the mass of the whole does not explain the masses of its parts. Examples also arise for network explanations:

**EP.** Königsberg's bridges form a connected network with four nodes. Marta walks an Eulerian path through town.

**MP.** Among connected networks composed of four nodes, only networks containing zero or two nodes with odd degree also contain an Eulerian path.

**C.** Therefore, either zero or two of Königsberg's landmasses have an odd number of bridges.

Yet Marta's walk doesn't explain Königsberg's layout.<sup>2</sup>

Huneman's and Lange's models are thus incomplete by their own lights; legitimate and illegitimate explanations fit the form (see Lange 2013, 486). The accounts are thus incomplete as descriptions of the defining norms that sort good mathematical explanations from bad.

Furthermore, ontic commitments appear to readily account for this directionality. The evenness of the batch explains the equal distribution (and not vice versa) because the distribution is drawn from the batch. Properties of parts explain aggregate properties (and not vice versa) because the parts compose the whole. Network properties are explained in terms of nodes and edges (and not vice versa) because the nodes and edges compose and are organized into networks. Paradigm distinctively mathematical explanations thus arguably rely for their explanatory force on ontic commitments that determine the explanatory priority of causes to effects and parts to wholes.

**4. Network Models in Mechanistic Explanations.** To explore this ontic basis further, consider three uses of network models—to describe structural con-

2. Perhaps the mayor ordered bridge construction to prevent anyone from ever again taking the Eulerian walk he shared with Marta the night before she left. That's a causal explanation.

nectivity, causal connectivity, and functional connectivity. Models of structural and causal connectivity are sometimes used to represent causally relevant features of a mechanism (Levy and Bechtel 2013; Zednik 2014). When they do, they explain just like any other causal or constitutive explanation (Woodward 2003; Craver 2007). Network models of functional connectivity (secs. 5 and 6) are not designed to represent explanatory information as such.

*4.1. Structural Connectivity.* Network models sometimes represent a mechanism's spatial organization. For example, researchers have mapped the structural connectivity (the cellular connectome) of the central nervous system of *C. elegans*, all 279 cells and 2,287 connections (see White et al. 1986; Achacosa and Yamamoto 1991; see also [www.wormatlas.org](http://www.wormatlas.org)). The nodes here are neurons; the edges are synapses. Algorithms applied to this network reveal a "rich club" structure in which high-degree hubs link to one another more than they do to nodes of lower degree (see, e.g., Towlson et al. 2013). Another example: Sebastian Seung is using automated cell-mapping software and crowdsourcing to map the connectome of a single mouse retina (see [www.eyewire.org](http://www.eyewire.org)). The goal of these projects is to map every cell and connection in the brain (Seung 2012). Network analysis supplies basic concepts (e.g., rich club) for discovering and describing organization in such bewilderingly complex structures.

These structural models contain a wealth of explanatorily relevant information about how brains work. But for any given explanandum (e.g., locomotion), that information is submerged in a sea of explanatorily irrelevant information. Brute-force models of this sort describe all the details and do not sort out which connections are explanatorily relevant to which phenomena.

Suppose, then, we fix an explanandum and filter the model for constitutive explanatory relevance. The resulting model would describe the anatomical connections relevant to some higher-level explanandum phenomenon. But it would only describe structures and would leave out most of how neural systems work. Cells generate temporal patterns of action potentials and transmitter release. Dendrites actively process information. Synapses can be active or quiescent, vary in strengths, and change over time. Two brains could have identical structural connectomes and work differently; after all, dead brains share structural connectomes with their living predecessors. One could model those anatomical connections without understanding the physiology of the system; anatomy is just one aspect of its organization.

*4.2. Causal Connectivity.* Directed graphs are also used to represent causal organization, how the parts (or features) in a mechanism interact (e.g., Spirtes, Glymour, and Scheines 2000). Bechtel and Levy emphasize the importance of causal motifs, abstract patterns in a network's causal organization: autoregulation, negative feedback, coherent feed-forward loops, and

so on. Such motifs have been studied in gene regulatory networks (e.g., Alon 2007) and in the brain (Sporns and Kötter 2004). Motifs can be concatenated to form more complex networks and to construct predictive mathematical models of their behavior.

Causal motifs are mechanism schemas containing placeholders for relevant features of the mechanisms and arrows showing how they interact. The causal (or active) organization of a mechanism can be explanatorily relevant to the mechanism's behavior independently of how the motif is instantiated. One can intervene to alter the details without changing the mechanism's behavior so long as the abstract causal structure is still preserved through the intervention (Woodward 2003; Craver 2007, chap. 6). If so, one can justifiably say that the causal organization, rather than the gory details, is the relevant difference maker for that explanandum.

Like structural networks, abstract network motifs often contain only the thinnest relevant information. For example, Alon's type 1 coherent feed-forward network is shown in figure 1. The arrows stand for "activation." Suppose X is a regulator of promoter Y for gene Z (adding considerable content to the motif). This motif clearly contains information explanatorily relevant to Z's expression. And if we constrain the explanandum phenomenon sufficiently, the motif might describe the most relevant features of the mechanism. For different explananda (Craver 2007, chap. 6) or in different pragmatic contexts (Weisberg 2013), one's explanation will require different degrees of abstraction.

That said, network models of structural and causal connectivity are mechanistic in the sense that they derive their explanatory force from the fact that they explain in virtue of representing how the phenomenon is situated in the causal structure of the world. This general conclusion is supported by the fact that network models can be used to describe things that are not, and are not intended to be, explanations of anything at all. Just as causal-mechanical theories of etiological explanation solved many of the problems confronting Hempel's covering law model by adding a set of ontic constraints that sort good explanations from bad, a causal-mechanical theory of network explana-



Figure 1. Type 1 coherent feed-forward network.

tion clarifies (at least in many cases) why some network models are explanatory and others are not.

**5. Nonexplanatory Evidential Networks.** Consider the use of network models to analyze resting-state functional connectivity (RSFC) in the brain (Power et al. 2011; Wig, Schlaggar, and Petersen 2011; Smith et al. 2013). This work mines magnetic resonance imaging (MRI) data for information about how brain regions are organized into large-scale systems. It offers unprecedented scientific access to brain organization at the level of systems rather than the level of cells (Churchland and Sejnowski 1989). The term “functional connectivity,” however, misleadingly suggests that FC models represent the brain’s working connections as such. They do not (see also Vincent et al. 2007; Buckner, Krienen, and Yeo 2013; Li et al. 2015).

The ground-level data of this project are measures of the blood-oxygen-level dependent (BOLD) signal in each ( $3 \text{ mm}^3$ ) voxel of the brain. The term “resting state” indicates that these measures are obtained while the subject is (hopefully) motionless in the scanner and not performing an assigned task.<sup>3</sup> The scanner records the raw time course of the BOLD signal for each voxel. This noisy signal is then filtered to focus on BOLD fluctuations oscillating between 0.1 and 0.01 Hz.<sup>4</sup> Neuroscientists use this range because it generates the highest-powered signal, not because it is functionally significant. It is likely too slow to be relevant to how the brain works during behavioral tasks.

To model these data as a network, you start by defining the nodes. Power et al. (2011) use two methods. The results mostly agree with each other. *Voxel-wise analysis* treats each voxel of the brain as a node location, yielding at least 10,000 nodes per brain (Power et al. used 44,100 nodes). *Areal analysis*, a less arbitrary approach, uses meta-analyses of task-based functional MRI (fMRI) studies, combined with FC-mapping data described below, to identify 20 mm spheres in regions throughout the brain. This approach yields 264 node locations per brain. Neither approach starts with working brain parts. Voxelwise analysis dices the brain into uniformly sized cubes. Areal analysis reduces it to 264 spheres. In fact, the nodes are not even locations; rather, the nodes are time courses of low-frequency BOLD fluctuations in those locations.

The edges represent Pearson correlations between these time courses. The strength (or width) of an edge reflects the strength of the correlation. This is all represented in a matrix ( $10,000 \times 10,000$  or  $264 \times 264$ ) showing the correlation between the BOLD time course in each voxel (or area) and that in every other voxel (or area). This matrix can then be analyzed with network tools, as described below.

3. Nothing here turns on the specialness of rest relative to activity or on whether the command to rest is an intervention.

4. For an overview of data processing stages, see Van Dijk et al. (2011).

It is now clear why the term “functional connectivity” is misleading. The nodes in the network do not (and need not) stand for working parts. They stand for time courses in the BOLD signal, which are merely indicators of brain activity. These time courses are (presumably) too slow to be part of how the brain performs cognitive tasks, they are present when the brain is not involved in any specific task, and they are measured in conveniently measurable units of brain tissue rather than known functional parts. Likewise, the edges do not necessarily represent anatomical connections, causal connections, or communications. There are, for example, strong functional connections between the right and left visual cortices despite the fact that there is no direct anatomical connection between them (Vincent et al. 2007). These correlations in slow-wave oscillations in blood oxygenation, in short, underdetermine the causal and anatomical structures that presumably produce them (Behrens and Sporns 2011).

We have here a complex analog of the barometer and the storm: a correlational model that provides evidence about explanatory structures in the brain but that is not used to (and would not) explain how brains work. FC matrices are network models. They provide evidence about community structure in the brain. Community structure is relevant to brain function. But the matrices do not explain brain function. They don't model the right kinds of stuff: the nodes aren't working parts, and the edges are only correlations. As for the barometer and the storm, A is evidence for B, and B explains C, but A does not explain C. In my view, network analysis is interesting to the philosopher not primarily because it offers nonmechanistic explanations but because of the role it might play in discovering and describing complex mechanisms. Consider some examples.

*5.1. System Identification.* Like structural networks, FC networks can be analyzed with community detection tools to find clusters of nodes that are more tightly correlated with one another than with nodes outside the cluster. The assumption that clustered nodes in a correlational network form a causally functional unit can be given a quasi-empirical justification: things that fire together wire together, things that wire together synchronize their activities, and this synchronicity is reflected in temporal patterns in the BOLD signal (see Wig et al. 2011, 141).

Using these methods, researchers have discovered several large-scale systems in the cortex, some of which correspond to traditional functional divisions (Power et al. 2011). For example, classical visual and auditory areas show up as clusters. Yet there are also surprises: for example, the classic separation of somatosensory from motor cortices is replaced by an orthogonal separation of hand representations from mouth representations.

*5.2. Brain Parcellation.* Changes in functional connectivity can be used to map cortical boundaries (e.g., Cohen et al. 2008; Wig et al. 2014; Gordon



et al. 2016). An abrupt change in the correlational profile from one voxel to the next indicates that one has moved from one cortical area to another. This approach identifies some familiar boundaries of traditional anatomy but can also be extended to areas, such as the angular gyrus and the supramarginal gyrus, for which anatomical boundaries and parcellations are not currently settled (see, e.g., Nelson et al. 2010).

*5.3. Comparing Brains.* Differences in functional connectivity might be associated with neurological disorders and insults, such as Alzheimer's disease (Greicius et al. 2004), schizophrenia (Bassett et al. 2008), multiple sclerosis (Lowe et al. 2008), and Tourette syndrome (Church et al. 2009). These differences might indicate detectable changes that are etiologically relevant or that are independently useful indicators for diagnosis or prognosis.

*5.4. Lesion Analysis.* Neuropsychologists of a strict localizationist bent prize "pure" cases of brain damage involving complete and surgical damage to single functional regions of cortex. "Impure cases" are notoriously hard to interpret because the symptoms combine and cannot be attributed with any certainty to the damage to specific loci. FC analysis, however, might offer an alternative perspective. Power et al. (2013), for example, identify a number of local "target hubs" in the cortex that are closely connected to many functional systems and have high participation coefficients. Damage to these areas produces wide-ranging deficits out of proportion to the size of the lesion (Warren et al. 2014). The "impure cases" of classical neuropsychology might look more pure through the lens of network science.

To conclude, FC-network models are correlational networks. They are used to discover causal systems, not to represent how causal systems work. Whether a network model explains a given phenomenon depends on how that phenomenon is specified and on whether the nodes and edges represent the right kinds of things and relations. Philosophical debates about how models refer (and what they refer to) are central to understanding how some models have explanatory force (Giere 2004; Frigg 2006).

**6. Near Decomposability and Random Walks.** Rathkopf (2015) argues that network models are distinct from mechanistic explanations because mechanistic explanations apply only to nearly decomposable systems (see sec. 2) whereas network analysis applies to systems that are not nearly decomposable. In drawing a "hard line" between these, Rathkopf's useful discussion both undersells the resources of network analysis and artificially restricts the domain of mechanistic explanation.

Network analysis is not restricted to nondecomposable systems. The primary goal of FC analysis is to reveal community structure in complex networks. The community detection algorithms applied to FC matrices are an-

imated by Simon's (1962) concept of near decomposability: a component in a nearly decomposable system is a collection of nodes more strongly connected to one another than with nodes outside that collection (see Hage-land 1998). Consider random walk algorithms, such as InfoMap (Rosvall and Bergstrom 2008).<sup>5</sup> A "bot" starts at an arbitrary node and moves along edges. It tends to get "trapped" in nearly decomposable clusters precisely because there are more and stronger edges that keep it in the cluster than ones that afford escape. Escape routes are interfaces. The places it gets trapped are nearly decomposable communities. Such algorithms do not deliver binary results (decomposable/not decomposable); rather, they represent a full *spectrum* of organization. Nondecomposable systems (with no community structure) are rare, idealized special cases. Rathkopf's hard line is a blur.

Rathkopf enforces this line by restricting the concept of mechanism. Each network he considers is a causal network, composed of nodes, such as people, and interactions, such as contagion. (Network models are useful in epidemiology because they model community structure). Even in so-called non-nearly decomposable causal networks, the base level of nodes and edges is still a set of causally organized parts: a mechanism. In characterizing that base level, we learn how network behavior is situated in the causal structure of the world. The line between mechanisms (organized causal interactions among parts) and nonmechanisms (e.g., correlations) is much harder than that between nearly decomposable and non-nearly decomposable mechanisms.

**7. Conclusion.** Network analysis is transforming many areas of science. It is attracting large numbers of scientists. It is generating new problems to solve and new techniques to solve them. It is revealing phenomena invisible and unthinkable with traditional perspectives. Yet it does not seem to fundamentally alter the norms of explanation. The problem of directionality and the puzzle of correlational networks signal that, at least in many cases, the explanatory power of network models derives from their ability to represent how phenomena are situated, etiologically and constitutively, in the causal and constitutive structures of our complex world.

#### REFERENCES

- Achacoso, Theodore B., and Willaim S. Yamamoto. 1991. *AY's Neuroanatomy of C. elegans for Computation*. Boca Raton, FL: CRC.
- Albert, Réka, Hawoong Jeong, and Albert-Lászlo Barabási. 2000. "Error and Attack Tolerance of Complex Networks." *Nature* (London) 406:378–82.
- Alon, Uri 2007. "Network Motifs: Theory and Experimental Approaches." *Nature Reviews Genetics* 8:450–61.

5. There are many community detection algorithms. See Lancinichinetti and Fortunato (2009) and Fortunato (2010). Power et al. (2011) use InfoMap.

- Bassett, Danielle S., Edward Bullmore, Beth A. Verchinski, Venkata S. Mattay, Daniel R. Weinberger, and Andreas Meyer-Lindenberg. 2008. "Hierarchical Organization of Human Cortical Networks in Health and Schizophrenia." *Journal of Neuroscience* 28:9239–48.
- Behrens, Timothy E. J., and Olaf Sporns. 2011. "Human Connectomics." *Current Opinion in Neurobiology* 22:1–10.
- Buckner, Randy L., Fenna M. Krienen, and B. Thomas Yeo. 2013. "Opportunities and Limitations of Intrinsic Functional Connectivity MRI." *Nature Neuroscience* 16:832–37.
- Church, Jessica A., Damien A. Fair, Nico U. F. Dosenbach, Alexander L. Cohen, Francis M. Miezin, Steven E. Petersen, and Bradley L. Schlaggar. 2009. "Control Networks in Paediatric Tourette Syndrome Show Immature and Anomalous Patterns of Functional Connectivity." *Brain* 132 (1): 225–38.
- Churchland, Patricia S., and Terry Sejnowski. 1989. *The Computational Brain*. Cambridge, MA: MIT Press.
- Cohen, Alexander L., D. A. Fair, N. U. Dosenbach, F. M. Miezin, D. Dierker, D. C. Van Essen, B. L. Schlaggar, and S. E. Petersen. 2008. "Defining Functional Areas in Individual Human Brains Using Resting Functional Connectivity MRI." *NeuroImage* 41:45–57.
- Craver, Carl F. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Clarendon.
- . 2014. "The Ontic Conception of Scientific Explanation." In *Explanation in the Biological and Historical Sciences*, ed. A. Hütteman and M. Kaiser. Dordrecht: Springer.
- Fortunato, Santo. 2010. "Community Detection in Graphs." *Physics Reports* 486:75–174.
- Frigg, Roman. 2006. "Scientific Representation and the Semantic View of Theories." *Theoria* 55:37–53.
- Giere, Ronald. 2004. "How Models Are Used to Represent Reality." *Philosophy of Science* 71 (Proceedings): 742–52.
- Gordon, Evan M., Timothy O. Lauman, Babatunde Adeyemo, Jeremy F. Huckins, William M. Kelley, and Steven E. Petersen. 2016. "Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations." *Cerebral Cortex* 26 (1): 288–303.
- Greicius, Michael D., Gaurav Srivastava, Allan L. Reiss, and Vinod Minon. 2004. "Default-Mode Network Activity Distinguishes Alzheimer's Disease from Healthy Aging: Evidence from Functional MRI." *Proceedings of the National Academy of Sciences* 101:4637–42.
- Haugeland, John. 1998. *Having Thought: Essays in the Metaphysics of Mind*. Cambridge, MA: Harvard University Press.
- Hempel, Carl G. 1965. *Aspects of Scientific Explanation*. Princeton, NJ: Princeton University Press.
- Huneman, Philippe. 2010. "Topological Explanations and Robustness in Biological Sciences." *Synthese* 177:213–45.
- Lancinichetti, Andrea, and Santo Fortunato. 2009. "Community Detection Algorithms: A Comparative Analysis." *Physical Review* 80:056117.
- Lange, Marc. 2013. "What Makes a Scientific Explanation Distinctively Mathematical?" *British Journal for the Philosophy of Science* 64:485–511.
- Levy, Arnon, and William Bechtel. 2013. "Abstraction and the Organization of Mechanisms." *Philosophy of Science* 80:241–61.
- Li, Jingfeng M., William J. Bentley, Abraham Z. Snyder, Marcus E. Raichle, and Lawrence H. Snyder. 2015. "Functional Connectivity Arises from a Slow Rhythmic Mechanism." *Proceedings of the National Academy of Sciences* 112 (19): E2527–35.
- Lowe, Mark J., Erik B. Beall, Ken E. Sakaie, Ken A. Koenig, Lael Stone, Ruth A. Marrie, and Michael D. Phillips. 2008. "Resting State Sensorimotor Functional Connectivity in Multiple Sclerosis Inversely Correlates with Transcallosal Motor Pathway Transverse Diffusivity." *Human Brain Mapping* 29 (7): 818–27.
- Nelson, Steven M., et al. 2010. "A Parcellation Scheme for Human Left Lateral Parietal Cortex." *Neuron* 67 (1): 156–70.
- Power, Jonathan D., Bradley L. Schlaggar, Christina N. Lessov-Schlaggar, and Steven E. Petersen. 2013. "Evidence for Hubs in Human Functional Brain Networks." *Neuron* 79:798–813.
- Power, Jonathan D., et al. 2011. "Functional Network Organization of the Human Brain." *Neuron* 72 (4): 665–78.
- Rathkopf, Charles. 2015. "Network Representation and Complex Systems." *Synthese*. doi:10.1007/s11229-015-0726-0.

- Rosvall, Martin, and Carl T. Bergstrom. 2008. "Maps of Random Walks on Complex Networks Reveal Community Structure." *Proceedings of the National Academy of Sciences* 105:1118–23.
- Salmon, Wesley. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press.
- Seung, Sebastian. 2012. *Connectome: How the Brain's Wiring Makes Us Who We Are*. New York: Houghton Mifflin Harcourt.
- Simon, Herbert A. 1962. *The Sciences of the Artificial*. Cambridge, MA: MIT Press.
- Smith, Stephen M., et al. 2013. "Functional Connectomics from Resting-State fMRI." *Trends in Cognitive Sciences* 17 (12): 666–82.
- Spirites, Peter, Clark N. Glymour, and Richard Scheines. 2000. *Causation, Prediction, and Search*. Cambridge, MA: MIT Press.
- Sporns, Olaf, and Rolf Kötter. 2004. "Motifs in Brain Networks." *PLoS Biology* 2:e369.
- Towilson, Emma K., Petra E. Vèrtes, Sebastian E. Ahnert, William R. Schafer, and Edward T. Bullmore. 2013. "The Rich Club of the *C. elegans* Neural Connectome." *Journal of Neuroscience* 33: 6380–87.
- Van Dijk, Koene R. A., Trey Hedden, Archana Venkataraman, Karleyton C. Evans, Sara W. Lazar, and Randy L. Buckner. 2011. "Intrinsic Functional Connectivity as a Tool for Human Connectomics: Theory, Properties, and Optimization." *Journal of Neurophysiology* 103:297–321.
- Vincent, Justin L., Gaurav H. Patel, Michael D. Fox, Abraham Z. Snyder, Justin T. Baker, David C. Van Essen, John M. Zempel, Lawrence H. Snyder, Maurizio Corbetta, and Marcus E. Raichle. 2007. "Intrinsic Functional Architecture in the Anaesthetized Monkey Brain." *Nature* 447:83–86.
- Warren, David E., Jonathan D. Power, Joel Bruss, Natalie L. Denburg, Eric J. Waldron, Haoxin Sun, Steven E. Petersen, and Daniel Tranel. 2014. "Network Measures Predict Neuropsychological Outcome after Brain Injury." *Proceedings of the National Academy of Sciences* 111: 14247–52.
- Watts, Duncan J., and Steven H. Strogatz. 1998. "Collective Dynamics of 'Small World' Networks." *Nature* 393:440–42.
- Weisberg, Michael. 2013. *Simulation and Similarity: Using Models to Understand the World*. Oxford: Oxford University Press.
- White, J. G., Southgate Eileen, J. N. Thomson, and Sydney Brenner. 1986. "The Structure of the Nervous System of the Nematode *Caenorhabditis elegans*." *Philosophical Transactions of the Royal Society London B* 314:1–340.
- Wig, Gaugan, Bradley Schlaggar, and Steven Petersen. 2011. "Concepts and Principles in the Analysis of Brain Networks." *Annals of the New York Academy of Sciences* 1224:126–46.
- Wig, Gaugan S., Timothy O. Laumann, Alexander L. Cohen, Johathan D. Power, Steven M. Nelson, Matthew F. Glasser, Francis S. Miezin, Abraham Z. Snyder, Bradley L. Schlaggar, and Steven E. Petersen. 2014. "Parcellating an Individual Subject's Cortical and Subcortical Structures Using Snowball Sampling of Resting-State Correlations." *Cerebral Cortex* 24 (8): 2036–54.
- Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.
- Zednik, Carlos. 2014. "Are Systems Neuroscience Explanations Mechanistic?" In *Preprint Volume for Philosophy Science Association 24th Biennial Meeting*, 954–75. Chicago: Philosophy of Science Association.