Moving Forward (and Beyond) the Modularity Debate: A Network Perspective

Matteo Colombo*†

Modularity is one of the most important concepts used to articulate a theory of cognitive architecture. Over the last 30 years, the debate in many areas of the cognitive sciences and in philosophy of psychology about what modules are, and to what extent our cognitive architecture is modular, has made little progress. After providing a diagnosis of this lack of progress, this article suggests a remedy. It argues that the theoretical framework of network science can be brought to bear on the traditional modularity debate, facilitating our progress in articulating a good theory of the human cognitive architecture.

1. Introduction. At least since Fodor's (1983) *The Modularity of Mind*, the notion of modularity has been one of the most important concepts used to articulate an account of the human cognitive architecture, which provides us with an encompassing theory (a "blueprint") of the nature, arrangement, and form of the structures and processes that are responsible for cognition and adaptive behavior. This should be distinguished from a theory of the origins of cognitive architectures, which is concerned with the evolutionary and developmental history of the structures and processes that are responsible for cognition and adaptive behavior.

Received October 2012; revised February 2013.

*To contact the author, please write to: Tilburg Center for Logic and Philosophy of Science, Tilburg University, PO Box 90153, 5000 LE Tilburg, The Netherlands; e-mail: m.colombo@uvt.nl.

†I thank Filippo Rossi, Andy Clark, and Tim Bates for discussion of some of the ideas that motivated this article. Charles Rathkopf, Liz Irvine, Carlos Zednik, and Jan Sprenger read and commented on earlier drafts, for which I am sincerely grateful. Rhiannon Bolos was so kind to assist me with copyediting. A special thank you goes to Clark Barrett and an anonymous reviewer for detailed comments and suggestions. This work was supported by a grant from the Deutsche Forschungsgemeinschaft (DFG) as part of the priority program New Frameworks of Rationality (SPP 1516).

Philosophy of Science, 80 (July 2013) pp. 356–377. 0031-8248/2013/8003-0002$10.00
Copyright 2013 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/670331 Published online by Cambridge University Press

There has been little agreement on how the concept of modularity should be characterized or on how we should study the extent to which the human cognitive architecture is modular. This disagreement has arguably hindered advancement in our understanding of the human cognitive architecture. There are two main reasons why researchers have often argued at cross-purposes in debates about modularity. First, different researchers have often used the term 'modularity' in significantly different ways (for discussion of different notions of modularity, see Segal [1996] and Samuels [2000]). Second, questions about the nature, arrangement, and form of the structures and processes responsible for cognition often have been conflated with questions about the value and problems of (some form of) adaptationism as a view about the role of natural selection in the evolution of our cognitive traits, in the construction of explanations of our cognitive faculties, and in the definition of the goal of research on the evolution of such faculties (cf. the history of modularity in Barrett and Kurzban [2006, 628–29] or the exchanges between Fodor [2000] and Pinker [2005], between Lickliter and Honeycutt [2003a, 2003b] and Tooby, Cosmides, and Barrett [2003], and between Buller [2005] and Machery and Barrett [2006]).

Because of mere terminological disputes,[1] because of vagueness surrounding putative central features of modularity such as functional specialization, domain specificity, and informational encapsulation, and especially because of little agreement about the proper empirical methods for discovering and justifying the existence of candidate modules (cf. the controversy around the cheater-detection module: e.g., Fodor 2000, 2008; Sperber and Girotto 2003; Cosmides and Tooby 2008a, 2008b), the modularity debate in the cognitive sciences and philosophy of psychology has often been frustratingly fruitless. If the debate could be moved forward and refocused on substantive issues about our cognitive architecture, that would mark significant progress.

The goals of this article are twofold. The first goal is to diagnose why the modularity debate risks remaining stagnant. The second goal is to suggest an effective remedy. Such a remedy is the theoretical framework of network science, which provides the basis of a more productive research program on the human cognitive architecture. Network science rests on bottom-up methods, which use large amounts of available and reliable data about brains, rather than on highly controversial methods, which partly draw on speculation about the evolutionary origins of cognitive mechanisms.

The article is in five sections. In sections 2 and 3, I provide a critical overview of the two main accounts of modularity. These accounts have been

1. Twyman and Newcombe write: "Given this lack of agreed-upon definition, the modularity position becomes analogous to the Hydra, the many-headed monster that Heracles found difficult to combat because there were too many heads to take on simultaneously, and, worse, because other heads grew while he addressed a specific one" (2010, 1317).

https://doi.org/10.1086/670331 Published online by Cambridge University Press

developed by Jerry Fodor and by post-Fodorian researchers such as Clark Barrett, Peter Carruthers, Leda Cosmides, Robert Kurzban, Dan Sperber, and John Tooby. I explain why the notions of domain specificity and functional specialization as characterized within these two accounts do not facilitate theoretical and empirical progress. In section 4, after having introduced the basics of network science, I explain what modularity is from a network science perspective. I outline current methods for discovering modularity in systems of interest, and I argue that conceptual resources and methods from network science will help us move the modularity debate forward. In section 5, I concisely review some of the evidence recently gathered within network neuroscience about the modularity of the human cognitive architecture. Section 6 concludes.

A disclaimer about the scope of this article is in order before moving on. Philosophers and many cognitive scientists have paid little attention to network science, and particularly to network neuroscience. It is not my aim to give a formal treatment of this field (see, e.g., Rubinov and Sporns 2010). Nor is my aim to discuss in detail the theoretical and methodological foundations of network science. The article will raise many questions—I will be barely scraping the surface of this exciting field. Nonetheless, I hope to say enough about network science to convey to philosophers and cognitive scientists a sense of the theoretical opportunities it affords and of the problems it raises.

2. What Is a Module? Current philosophical and psychological literature includes two prominent types of characterizations of modularity: one put forward by Fodor (1983, 2000) and the other shared by post-Fodorian accounts such as Carruthers's (2006) and Barrett and Kurzban's (2006). According to Fodor (1983), there are nine features that collectively individuate modules: domain specificity, mandatory operation, limited central accessibility, fast processing, informational encapsulation, "shallow" outputs, fixed neural architecture, characteristic and specific breakdown patterns, and characteristic ontogenetic timetables.

For my purposes, three points are noteworthy about Fodor's account. First, Fodor is not so much concerned with the evidential status of the modularity of cognition: "This monograph [he writes] is about the current status of the faculty psychology program; not so much its evidential status (which I take to be, for the most part, an open question) as what the program is and where it does, and doesn't, seem natural to try to apply it" (1). So, Fodor's focus is on the nature and scope of a particular research program in cognitive science, which he calls "faculty psychology." This is "the view that many fundamentally different kinds of psychological mechanisms must be *postulated* in order to explain the facts of mental life" (1; emphasis added). Second, Fodorian modules are intended to characterize a fragment of the hu-

https://doi.org/10.1086/670331 Published online by Cambridge University Press

man cognitive architecture. Fodor maintains that only peripheral (i.e., perceptual and motor) and linguistic systems are modular; the architecture of central systems supporting such cognitive capacities as reasoning and decision making is instead nonmodular (1983, 2000). Third, according to Fodor, a cognitive architecture counts as modular if it is modular "to some interesting extent," that is, if it possesses most of the nine features singled out above to an appreciable degree (1983, 37). Hence, whether a cognitive architecture is modular (in Fodor's sense) involves some vagueness. This vagueness is mitigated by the fact that some of the features of modularity, such as information encapsulation and domain specificity, are more important than others (37). If we find that a mechanism is informationally encapsulated and domain specific to some interesting extent, then we have good grounds to believe that it is a Fodorian module.

However, it is also a matter of degree whether a mechanism is informationally encapsulated or domain specific. For, consider the definitions of an informationally encapsulated mechanism and of a domain specific mechanism. A mechanism is informationally encapsulated to the extent that it is less open or permeable to information stored elsewhere in the system over the course of its processes. A mechanism is domain specific to the extent that it has a restricted subject matter. "Domain specificity has to do with the range of questions for which a device provides answers (the range of inputs for which it computes analyses)" (Fodor 1983, 103). From these characterizations, it is unclear to what extent a mechanism should be impermeable to information stored elsewhere in order to count as informationally encapsulated, and it is unclear how narrow the range of inputs that a mechanism can process should be for it to count as domain specific. Hence, both informational encapsulation and domain specificity are vague notions, which do not reliably help us to precisely identify modules.

More recently, many cognitive scientists, evolutionary psychologists, and philosophers have rejected Fodor's notion of modularity. Barrett argues that Fodorian modularity "has been perhaps *too* influential, because it has foreclosed ways of thinking about modularity other than the very specific model [Fodor] proposed in his 1983 book" (2005, 260). Ermer, Cosmides, and Tooby claim that "Fodor's (1983) concept of a module is neither useful nor important" (2007, 153). Carruthers characterizes modularity by dropping some central features of Fodorian modules: "modules might be isolable function-specific processing systems, all or almost all of which are domain specific, whose operations aren't subject to the will, which are associated with specific neural structures (albeit sometimes spatially dispersed ones), and whose internal operations may be inaccessible to the remainder of cognition" (2006, 12).

Although there are differences among particular post-Fodorian accounts of modularity, most post-Fodorian theorists agree that modules should be

https://doi.org/10.1086/670331 Published online by Cambridge University Press

understood as functionally specialized cognitive mechanisms (Tooby and Cosmides 1992; Sperber 1994; Coltheart 1999; Pinker 2005; Barrett and Kurzban 2006; Carruthers 2006). Accordingly, to say that the human cognitive architecture is modular is to say that it consists of a number of isolable mechanisms specialized to carry out some function. A mechanism, in turn, can be understood as "a structure performing a function in virtue of its component parts, component operations, and their organization. The orchestrated functioning of the mechanism is responsible for one or more phenomena" (Bechtel and Abrahamsen 2005, 423). Hence, to say that a cognitive mechanism carries out some function is to say that it transforms some set of inputs into some set of outputs as a result of the organized set of operations performed by its component parts. It should be clear that the structural features of a mechanism—that is, the causal, spatial, and dynamical features of the components it comprises—are relevant to the function it carries out. In order to reliably identify the function carried out by some putative cognitive module, evidence about neuroanatomical and neurophysiological constraints can be relevant.

On this view, the defining feature of modular mechanisms is *functional specialization*. As Barrett and Kurzban put it, "functionally specialized mechanisms with formally definable informational inputs are characteristic of human (and nonhuman) cognition and . . . these features should be identified as the signal properties of 'modularity' " (2006, 630). Functional specialization concerns the types of processes carried out by a mechanism. A cognitive mechanism is functionally specialized to the extent that it carries out a restricted range of types of processes that successfully apply to certain problem domains—where domains are individuated by the formal properties of input representations to the mechanism (630). Accordingly, a module would be a specialized input-output mechanism that accepts inputs of a particular form, performs specialized transformations (or processes) on them, and yields outputs, whose format makes them usable for other cognitive mechanisms (cf. Barrett 2005, sec. 3).

Two points are noteworthy here. First, this characterization involves some vagueness since it does not answer the question of how many functions a mechanism should carry out for it to cease to be functionally specialized. Second, that a module is individuated in terms of its functional specialization does not imply that there is a single distinct anatomical structure in the brain that is responsible for the processes carried out by the module. Modules can correspond to distributed circuits of neural networks that do not neatly map onto any single anatomical structure.

3. Why Current Characterizations of Modularity Have Not Facilitated Progress. There is one important problem shared by Fodorian and post-Fodorian accounts of modularity. They do not offer reliable methods to identify either domain specificity or functional specialization in cognitive

https://doi.org/10.1086/670331 Published online by Cambridge University Press

mechanisms. If domain specificity or functional specialization is essential to modularity, and we lack reliable ways to identify them, then we lack reliable ways to identify modules. If we lack reliable ways to identify modules, then the concept of modularity risks hindering understanding and fruitful debate, rather than being a useful concept for articulating a theory of cognitive architecture.

The difficulty in identifying domain specificity and functional specialization partly depends on the vagueness of these two notions, which is acknowledged by both Fodor and post-Fodorian modularists. It also depends on the type of top-down methodology employed to define these features. This type of methodology—which I explain in a moment—is not directly concerned with using data and information from the brain sciences. To justify disregard for data and information from the brain sciences, it can be argued that modularity concerns a more abstract explanatory level than the level of neuroscience. Modularity would be a concept proprietary of the familiar Marr's (1982) computational and algorithmic levels, which are supposedly unconstrained by evidence about the neural level of implementation. If modularity is a concept proprietary of such levels, then modular architectures in cognitive systems can be identified and studied, while remaining agnostic about how defining features of modularity such as functional specialization are implemented in the brain (cf. Barrett and Kurzban 2006, 642).

In contrast to this conclusion, one of the convictions that motivate this article is that data and information from the brain sciences are relevant to developing an empirically and theoretically fruitful account of modularity and, more generally, a theory of cognitive architecture. A theory of cognitive architecture should draw on available and reliable data about structural, anatomical, and neurophysiological features of the brain, which are likely the bases of a much more productive research program.

Three reasons can be given in support of this conviction. First, the view that an empirically and theoretically fruitful account of modularity can disregard data and information from the brain sciences relies on a misguided understanding of Marr's (1982) three-level framework. Although in some passages Marr emphasizes that "the three levels are only rather loosely related" (25), thereby suggesting some sort of autonomy between them, he also recognizes that accounts of particular cognitive capacities such as human vision should be ultimately assessed in terms of how well the computational and algorithmic specifications they include fit known neurobiological properties and details of the human visual system (chaps. 3 and 4; see also Marr, Ullman, and Poggio 1979, 916). Hence, Marr himself embraced a form of *coevolutionary* research ideology, whereby concepts, models, and theories put forward at one level should be susceptible to correction and reconceptualization in light of discoveries, conceptual refinements, and methods available at some different level (Churchland 1986).

https://doi.org/10.1086/670331 Published online by Cambridge University Press

Second, both the Fodorian and most post-Fodorian accounts of modularity do refer to the brain sciences. For example, Fodor emphasizes "the intimate association of modular systems with neural hardwiring" and draws on "form/function correspondences" between certain patterns of neural connectivity and modularity to argue that input systems processes, but not central processes, are modular (1983, 117–19). Sperber (1994, 40) defines a cognitive module as "a genetically specified computational device in the mind/brain." Cosmides and Tooby (1995) refer to modules as "neurocognitive" adaptations. Carruthers sometimes refers to specific features of modularity as features of the "mind/brain" and draws on anatomical as well as neurophysiological evidence to make a number of arguments about the organization of the mind of nonhuman and human animals (2006, chap. 2). This overview indicates that many cognitive scientists and philosophers working on modularity draw on information from the brain sciences because this information can be relevant to developing and assessing an account of modularity.

Third, and more important, even if developing an account of modularity and, more generally, a theory of cognitive architecture without paying attention to data and information from the brain sciences might not be a hopeless project, it is a significant advantage if such a research program finds data and information from the brain sciences relevant. One should not profess agnosticism about the neural implementation of functional specialization, insofar as neural evidence is available that is relevant to constrain hypotheses about cognitive specializations. In fact, not every functional analysis of a system is equally adequate to play an explanatory role. And one way to assess the explanatory grip of a given functional analysis is to consider the extent to which it is informed by known facts about the human brain. This might not be the only way, but it is useful and congenial to the development of a fruitful account of modularity.

I now turn to the methodology widely adopted by many cognitive scientists and philosophers to identify modularity. Fodor (1983) does not elaborate on it, partly because—as noted above—he is not much concerned with the evidential status of the modularity of cognition. Most of the post-Fodorian literature, instead, addresses the methodological issue of how we should study modularity within an adaptationist framework. A type of top-down, evolutionary/functional methodology is generally adopted. Accordingly, the basis for studying the architecture of the human cognitive system, and the extent to which it is modular, consists in the identification of the types of computational problems that the human cognitive system faced over its evolution (Pinker 1997; Tooby and Cosmides 2005).[2] These kinds of problems, called

2. Buller explains this methodology thus: "Reverse engineering is a process of figuring out the design of a mechanism on the basis of an analysis of the tasks it performs. Evolutionary functional analysis is a form of reverse-engineering in that it attempts to

https://doi.org/10.1086/670331 Published online by Cambridge University Press

adaptive problems, consist of situations such that different variants of some phenotypic trait contribute differently to reproductive fitness. The search for the structures and processes that can plausibly constitute our cognitive architecture is deferred until an adaptive problem—which could be solved by the organized operations of those structures—is identified (see, e.g., Cosmides and Tooby 1995; Barrett and Kurzban 2006). Pursuing this approach, domain specificity is defined in terms of sets of input stimuli characteristic of the adaptive problems faced by our ancestors (Sperber 1994; Carruthers 2006; Ermer et al. 2007; Barrett 2009). Functional specialization is defined similarly, in terms of adaptive problems solved by the cognitive systems of our ancestors. Accordingly, to say that a mechanism is functionally specialized is to say that it carries out a restricted range of types of processes that could solve the types of adaptive problems "that caused the propagation of its genetic basis relative to that of alternative mechanisms" (Ermer et al. 2007, 153).

If domain specificity and functional specialization are to be understood within an adaptationist framework, then some reliable way to individuate adaptive problems is necessary in order to identify modules in our cognitive architecture. However, the reliable individuation of adaptive problems is far from straightforward. It involves problematic methods, which have given rise to controversy about the status of evidence in evolutionary psychology used for establishing that a given cognitive mechanism is an adaptation or is adaptive. Although the research program in evolutionary psychology is clearly not doomed a priori, it is currently very challenging to reliably test adaptationist hypotheses about our past and about the origins of our cognitive architecture (cf. Sterelny and Griffiths 1999; Kaplan 2002; Lloyd and Feldman 2002; Atkinson and Wheeler 2004; Buller 2005; Richardson 2007; Schulz 2011; Machery [forthcoming] provides a more positive assessment of the status of evidence in evolutionary psychology). In the face of such methodological challenges, if the search for modularity mainly depends on our ability to individuate adaptive problems and disregards the structural constraints and neurophysiological dynamical patterns that more apparently shape our cognitive architecture, then the worry that researchers are too unconstrained to suppose any module for any imaginary adaptive task becomes serious.

This worry should not suggest that an adaptive understanding of domain specificity and functional specialization is unimportant to articulating a full account of our cognitive architecture. My point is about the counterproductiveness of adopting a top-down evolutionary/functional analysis as the basis of a research program on modularity and cognitive architecture.If "the language of modularity affords useful conceptual groundwork in which

---

reconstruct the mind's design from an analysis of the problems the mind must have evolved to solve" (2005, 92).

https://doi.org/10.1086/670331 Published online by Cambridge University Press

productive debates surrounding cognitive systems can be framed"—as claimed by Barrett and Kurzban (2006, 644)—then the utility of such a language in moving the debates forward will be greater if the notion of a module is not cashed out in adaptationist terms but instead is understood within an evolutionarily neutral framework.

Network science offers such a framework. Understanding modularity from a network science perspective can enable us to integrate neurophysiological and structural information about cognitive mechanisms, while freeing the notion of functional specialization from adaptationism. Concepts and methods from network science can make the language of modularity more conducive to productive theoretical and empirical debates about the form and organization of the architecture of cognitive systems such as the human brain.

4. What a Module Is: A Network Science Perspective. The brain is a complex system composed of intricately interconnected, interacting elements. The organized activity of such elements is responsible for cognitive capacities and behavior. Network science offers a conceptual framework, methods, and a set of mathematical and statistical techniques, whereby we can understand the architectural organization of brain networks and how such organization is responsible for cognition.

The explanatory targets of network science are complex networks and the phenomena they display. Complex networks can be engineered, biological, and social systems. In order to understand the organization of the structures and informational relationships that characterize complex systems, network science relies on concepts and methods from fields such as graph theory, statistics, information theory, machine learning, measurement theory, and dynamical system theory (Börner, Sanyal, and Vespignani 2007).

Graph theory is the primary analytical framework used in network science. Graphs are sets of nodes and edges, which allow us to represent complex network systems such as the World Wide Web, mafia gangs, or the brain. Nodes represent elements or components of the system. Edges represent connections between pairs of nodes. Edges can be directed or undirected, and they can be binary (i.e., they are either present or absent) or weighted (i.e., they can take on fractional values). Nodes can be connected directly by single edges or indirectly by intermediate nodes and edges. Network measures include degree, strength, and centrality. Degree and strength of a node measure the extent to which the node is connected to the rest of the network. The degree of a given node is the number of connections that link the node to the rest of the network; a node with high strength makes strong connections—where strength is equal to the sum of connection weights. The minimum number of edges that must be traveled to go from one node to the other defines the path length between two nodes. The centrality of a node measures how many shortest paths between other parts of

https://doi.org/10.1086/670331 Published online by Cambridge University Press

the network pass through that node. Nodes with high degree and high centrality are called "hubs." Hubs are crucial for efficient communication between different nodes, as they facilitate global connectedness and integration of information within the network. All these measures refer to the topology of the network. So, two nodes of a network can be physically distant but topologically close, as short path length between two nodes does not imply close physical proximity between them.

In cognitive science, network science is the basis of two main projects. The first project aims to describe the patterns of connectivity displayed by the multiscale networks that constitute nervous systems. The second aims to reveal the organizational principles behind the architecture of the networks of the brain understood as a type of information-processing system. More generally, network neuroscience asks how brain networks, spanning the microscale of individual cells and synapses and the macroscale of embodied and embedded cognitive systems, are responsible for cognition (Sporns 2011).

If we consider a graphical description of a brain network, network nodes can represent neural elements such as cells, populations of neurons, or cortical and subcortical regions, while network edges can represent structural connections between nodes such as synapses or axonal pathways. There are three modes of connectivity between the nodes comprising a brain network. Apart from structural connectivity, there is functional and effective connectivity (Sporns 2007).

Structural connectivity refers to the pattern of physical or anatomical connections linking neural elements. A representation of the structural connectivity of the brain corresponds to the "wiring diagram of the brain," or *connectome*, which provides a map of the anatomical connections between neural structures (Sporns, Tononi, and Kötter 2005). *Functional connectivity* refers to patterns of symmetrical statistical association between activities in different neural elements. Measured in terms of correlation or covariance, mutual information, or spectral coherence between activities in neural elements—regardless of whether they are structurally connected—functional connectivity captures neurophysiological dynamics (Friston 1994). Information about functional connectivity tells us how activity in one neural node affects the correlations between activities in all the other nodes over time. A third mode of connectivity is *effective connectivity*, which refers to patterns of causal effects among neural elements. Measured through structural equation modeling, Granger causality, and other methods, it describes directed relationships between neurophysiological events (Friston 2011).

Information about these three modes of brain connectivity constitutes an important part of a theory of cognitive architecture. This should be obvious if we recognize that a theory of the architecture of a cognitive system should provide us with information about how the system is structured

https://doi.org/10.1086/670331 Published online by Cambridge University Press

and how the organized activity of its elements is responsible for cognitive phenomena and behavior. Information about structural, functional, and effective connectivity helps us describe the anatomical organization of the brain. It enables us to precisely identify the topological, statistical, and information-theoretical principles that might lie behind the architecture of a cognitive system such as the brain. Furthermore, information about these three types of connectivity advances our understanding of how certain causal transactions that take place among certain neural elements produce specific cognitive phenomena.

Bullmore and Sporns (2009, 187) outline four basic steps that are common to most approaches to identify brain networks:

1. Define network nodes.
2. Define edges.
3. Represent the complete set of nodes and edges with a "connection matrix" that specifies which pairs of nodes are directly connected by an edge.
4. Analyze the resulting network with appropriate network-theoretical tools.

Each step requires researchers to make choices that will affect their results. This does not mean, however, that researchers' choices are arbitrary or that most results across studies are incoherent. Appropriate choices depend on the research question under investigation, on the researchers' pragmatic ends and interests, on background knowledge of the system, and on available techniques. For example, if we consider the first step, appropriate parcelling of a given system into neural nodes can rely on histological, anatomical, electrophysiological, receptor distribution, or imaging data, depending on the research question and scale at which the researchers intend to investigate the system. At the level of a large-scale brain system, nodes are better defined as coherent anatomical regions on the basis of prior anatomical criteria and functional imaging data (e.g., Cohen et al. 2008). Unsupervised, data-driven methods can be used to define or validate previously identified boundaries between coherent brain regions (Golland et al. 2008). The idea behind such data-driven methods is—very roughly—that nodes can be defined by searching sets of anatomical or imaging data for clusters of data, which emerge as coherent regions.

Analyses of several different types of brain networks, carried out with different techniques, have shown that some characteristic features of the human cognitive architecture can be detected with high reliability and robustness. One such feature is modularity. *Modularity* is topologically defined in terms of specific patterns of connectivity between neural nodes. Modules are "communities of nodes that share greater numbers of mutual

https://doi.org/10.1086/670331 Published online by Cambridge University Press

connections within each community and fewer connections between them" (Sporns 2011, 113).

If you suspect that this notion is ridden with vagueness as much as the notions of modularity of the Fodorian and post-Fodorian accounts, this suspicion will be dispelled in a moment. Before turning to this issue, two aspects of modularity as characterized within network science are worth pointing out. First, in many complex systems, modularity is a property of architectures that span multiple scales. Many complex systems, including the brain, display the fractal property of *hierarchical* modularity in which "roughly the same kind of community structure is expressed repeatedly at different hierarchical levels or topological scales of the network" (Meunier, Lambiotte, and Bullmore 2010, 2). Second, the distinction among structural, functional, and effective connectivity allows us to distinguish among structural, functional (i.e., statistical), and effective (i.e., causal) modularity. Based on patterns of structural connections in the system, structural modules consist of groups of densely anatomically interconnected nodes that are only sparsely connected with the rest of the network. Similarly, based on the patterns of statistical (and causal) dependencies displayed by the system, modules are individuated by clusters of densely functionally (and effectively) interconnected nodes.

Let us now examine how modules are identified in complex systems and whether the network-scientific notion of modularity is as vague as the traditional notions. A system displays a modular architecture, if the nodes comprising the system cluster in communities such that the nodes in each community are densely connected and the nodes between communities are sparsely connected. How dense? How sparse? There are a number of methods to answer these questions, each of which uncovers in an automated way the modules comprising the system, quantifying how modular its architecture is (see Fortunato 2010, for a review). Each of these methods gives a precise answer to the questions about the number and size of modules in the system. These answers are not necessarily the same. Despite differences, however, most available methods yield coherent results that validate each other.

Currently, one of the most prominent measures of modularity has been developed by Newman (2006; see also Newman and Girvan 2004). The basic idea is that the modularity of a network is identified on the basis of "the number of edges falling within groups [of nodes] minus the expected number in an equivalent network with edges placed at random" (Newman 2006, 8578). Modularity corresponds to "statistically surprising arrangement of edges" in the network (8578). Network-scientific measures of modularity such as this, unlike the notions of modularity underlying Fodorian and post-Fodorian accounts, have received a mathematical formulation. Each of these measures is an objective method, whose assumptions are explicitly laid out and can be independently evaluated in light of empirical results.

https://doi.org/10.1086/670331 Published online by Cambridge University Press

A characterization of modularity within network neuroscience has several advantages over the Fodorian and post-Fodorian accounts. First, the bottom-up and data-driven methods used to identify and measure modularity make it a precise concept and a quantifiable property, which can help us integrate information about different levels of organization in the brain. Second, the focus is on architectural features of our cognitive system, instead of on their origins or on their being adaptations. Third, features such as functional specialization are not matters of stipulation, nor are they labeled in an intuitive manner; rather, they can be rigorously characterized and discovered, drawing on network-theoretical analyses of patterns of brain connectivity. I conclude this section by articulating each claim in turn.

From a network-scientific perspective, modularity "is the result of an objective analysis of network connectivity and not based on intuitive or subjective classification criteria for network elements or on their intrinsic characteristics" (Sporns 2011, 113). Pursuing bottom-up, data-driven methodologies, network science eschews the use of concepts such as domain specificity and functional specialization to define modularity, whose vagueness makes it relatively more problematic to reliably identify modules. Bottom-up, data-driven methods deliver measurements of modularity, which quantify to what extent a given cognitive system is modular. Furthermore, a modularity analysis of connection data sets can provide us with a precise, quantitative characterization of the relationship between structural properties and neurophysiological dynamics of a target brain network. While facilitating us to connect different levels of structural, functional, and causal organization in the brain, modularity analyses can also uncover whether certain organizational principles (e.g., topological efficiency, global efficiency of information transfer, robustness) are conserved "over different scales and types of measurement, across different species and for functional and anatomical networks" (Bullmore and Sporns 2009, 196).

Second, from a network science perspective, whether the properties endowed by a modular architecture are adaptive or evolved because of their fitness benefits to our ancestors are separate, secondary questions (Bullmore and Sporns 2012). The question "How modular is the brain?" is not the same as "How many adaptations does the brain contain?" or "How functionally differentiated is the brain?" And these questions should not be conflated.

Although there is growing appreciation that a network analysis of the topological properties of a biological system is congenial to the study of the "evolvability" of the system (Wagner, Pavlicev, and Cheverud 2007), network science by itself does not answer questions about the evolutionary forces by which modular architectures could be selected (Sporns 2011, chap. 7).[3]

3. The "evolvability" literature does not have exactly the same definition of modularity as the network definition. According to this literature, a system is modular to the degree that

https://doi.org/10.1086/670331 Published online by Cambridge University Press

Modular architectures, for example, have been shown to display robustness. Their topological properties tend to be resilient to perturbations such as the removal or lesion of nodes or edges. Since systems that display a degree of robustness have been shown to possess a fitness advantage, the evolution of modularity might be linked to a topological property such as robustness (Wagner, Mezey, and Calabretta 2005). However, network modularity and "evolvability" can diverge in their empirical implications. It is possible that a highly modular neural network does not exhibit a high degree of within-network evolvability, in the sense that evolutionary forces did not shape different parts of the neural network separately.

Similarly, within network neuroscience, questions about the degree of modularity of the brain are distinguished from questions about the degree to which modules have been selected to carry out distinct functions. This is how it should be since the answers to those questions are not necessarily the same. Barrett's (2012) discussion of the relationship between increasing brain size, network modularity, and the evolution of brain specializations illustrates this point nicely. Barrett considers two distinct (not necessarily mutually exclusive) hypotheses (10737–38). According to one hypothesis, the evolutionary force behind increasing brain size was selection for increased neural processing power. This could have led to increasing modularity purely as an architectural by-product of increasing network size. The resulting more modular brain is not necessarily more functionally differentiated, as the more modular architecture did not undergo selection for new specializations. According to another hypothesis, the evolutionary force behind brain increase was selection for increased specialization. "If the best way to produce new specialized regions is to increase brain size . . . then selection for specialization could have favored mutations that increased overall brain volume, thereby increasing modularity" (10737). In this case, but not in the former, a more modular brain is evidence for more functional differentiation/increased specialization. Examples such as this bear out that, while network neuroscience can be relevant to some questions about the origins of our cognitive architecture (by, e.g., identifying organizational principles of complex brain networks, or neural wiring rules), it does not substitute for other approaches for those interested in those questions.

Finally, from a network perspective, features such as functional specialization, informational integration, and segregation are not matters of stipulation or of intuitive labeling. Especially within post-Fodorian/evolutionary psychologists' accounts of modularity, intuitive labels such as "tool-use

different parts of it can be shaped independently by selection. Hence, modularity here is similar to the inverse of pleiotropy. Nonetheless, also this notion of modularity—in which the relevant networks are genetic and epigenetic networks—can be understood within network science. Thanks to Clark Barrett for drawing my attention to this point.

https://doi.org/10.1086/670331 Published online by Cambridge University Press

module" or "friendship module" have been used to characterize the function of many putative modules (Tooby and Cosmides 1992, 113). One problem with this type of characterization is that it does not allow us to predict structures from functions, or functions from structures. Intuitive labeling does not allow us to predict what type of structural organization or what type of topology a target mechanism will display, given knowledge of the cognitive functions it performs, and it does not help us understand what types of cognitive functions a target mechanism can perform, given knowledge of its structural organization and topology. Part of the reason is that this type of labeling ignores all modes of neural connectivity, which can constrain and bootstrap functional ontologies for cognition (Price and Friston 2005; Bilder et al. 2009).

Let me single out a number of approaches in which network-theoretical analyses of connectivity or measures such as modularity bear on our understanding of cognitive function. First, modularity and connectivity analyses provide grounds for distinguishing networks of the brain, which, in the investigation of structure-function relationships, can be regarded as more appropriate structural units than brain regions considered in isolation from the broader neural context (McIntosh 2000; Klein 2012). Hampshire et al. (2012), for example, addressed the question of whether human intelligence is a single unitary general ability or a set of multiple independent abilities, by relating different brain networks to population differences in performance in a range of cognitive tasks that could yield a measure of general intelligence (IQ). A data-driven method (exploratory factor analysis) was used to identify functional networks from whole-brain-imaging data of participants who performed a battery of cognitive tasks. Reasoning, short-term memory, and verbal processing were respectively found to load most significantly on three distinct functional networks, which best explained cross-task correlations in performance in a larger population sample. The relationships between dissociable functional brain networks, specific cognitive abilities, and general intelligence scores could then be used to support the view that intelligence is not unitary.

According to a second approach, descriptions of structure-function relationships can be validated and constrained by examining the extent to which a given structural network exhibits similar neural dynamics (i.e., functional connectivity) in similar cognitive tasks. Passingham, Stephan, and Kötter (2002), for example, showed that each cytoarchitectonic area has unique patterns of cortico-cortical connections that reliably indicate differences in neural activity during distinct cognitive tasks. This type of evidence suggests that specific patterns of structural connectivity partly determine the types of processes carried out by a given structure and, thereby, can afford constraints on the types of cognitive functions that the structure can perform.

Furthermore, information about functional connectivity can provide an independent test of whether a cognitive function is supported by a target struc-

https://doi.org/10.1086/670331 Published online by Cambridge University Press

ture. Different functional modular arrangements may be observed during different cognitive tasks, suggesting that "the flow of cognition is a result of transient and multiscale neural dynamics, of sequences of dynamical events that unfold across time" (Sporns 2011, 206). Changes in patterns of functional modularity are highly sensitive to the perturbations caused by sensory input, task-specific demands, or lesion. Functional connectivity of a target network can show dynamic changes in modularity predicted by experimentally controlled manipulations in some task. The evidence provided by the extent to which those dynamic changes track specific experimental manipulations bears on two questions. First, in which classes of tasks is the target structure involved? Second, to what extent do the processes carried out by the structure make a difference to whether some cognitive capacity is displayed? These questions are obviously important to describe structure-function relationships underlying functional ontologies for cognition (Friston and Price 2011).

While particular measures in a task can be associated with patterns of functional modularity (i.e., with sequences of dynamical events in the brain), functional networks maintain stable, global topological characteristics (Bassett et al. 2006). The persistence of global topological properties along with the peculiar dynamical-functional changes associated with certain sensory input or particular measures in a cognitive task afford insights about how a variety of cognitive phenomena are determined by the architecture of the networks of the brain (Sporns 2011, chap. 8).

Further insights about structure-function relationships, and, particularly, about how modules relate to cognitive function, are provided by the organizational principles characteristic of certain types of architectures. Two such principles are "functional segregation" and "functional integration" (cf. Tononi, Sporns, and Edelman 1994). Modular architectures display a higher degree of locally *segregated* processing (or functional segregation), according to which different modules selectively capture different statistical regularities in their inputs—while their activities are statistically independent from each other. If a cognitive architecture displays a high degree of segregation, it will tend to display specialized local processing carried out by different modules and little "cross talk" between modules, which reduces both wiring costs and the propagation of noise in the global processing of the system. In light of segregation, a module's functional specialization appears not only to be the result of the intrinsic biophysical properties of the module but also of its extrinsic interactions within the network. "Specialization is not an intrinsic property of any region, but depends on both forward and backward connections with other areas" (Friston and Price 2001, 275). Functional specialization, therefore, becomes meaningful only in the context of global features of networks' connectivity. It cannot be identified by considering a target mechanism in isolation from the processes carried out by other mechanisms with

https://doi.org/10.1086/670331 Published online by Cambridge University Press

which it interacts and from the global features of the system in which it is embedded. Globally *integrated* processing (or functional integration) is another organization principle displayed in varying degrees by different types of architectures. If a cognitive system generates unified cognitive phenomena and coherent behavior, then relevant information processed by segregated modules must be integrated. Globally integrated processing can be achieved in a modular system courtesy of hub nodes, which can enable efficient communication between modules and integrate information processed locally, as well as of topological arrangements such as *small-worldness* in which high local clustering is combined with short paths that connect all nodes of the network facilitating direct cross talk among several modules (Gallos, Makse, and Sigman 2012).

Finally, information about modularity and modes of connectivity can be used to build neural models, which can offer testable hypotheses about structure-function relationships, while performing cognitive functions. Eliasmith et al. (2012), for example, built a 2.5 million neuron model of the brain that could perform almost as well as humans at a number of tasks. One of the added values of models such as this is that they offer a "set of hypotheses regarding the neural mechanisms and organization that may underlie basic cognitive functions" (1205). Part of the explanation for the success of this model is that its structural and functional connectivity embodies constraints motivated by information about real brains' modularity and connectivity. The model captured several aspects of neuroanatomy and neurophysiology and showed how they could give rise to different cognitive functions and adaptive behavior.

As highlighted by this discussion of network-centered approaches to understanding cognitive function, it may well be misguided to ask which module is necessary (or sufficient) for a particular cognitive task. The relationship between modularity and cognitive function should be studied in a context-sensitive way. More attention should be paid to the dynamic patterns of (structural, functional, and effective) connectivity of modules engaged in particular cognitive tasks, rather than to the cognitive function of individual modules.

5. To What Extent the Human Cognitive Architecture Is Modular. A growing number of studies of structural and functional connectivity in the human brain agree that our cognitive architecture comprises "a set of interconnected communities of structural and functionally related elements, arranged on multiple scales from cells to systems" (Sporns 2011, 114). Among such communities, modular structures are prominent. Modules identified in the architecture of neural systems present common features: they are likely to share pathways, if they are spatially close, and they are typically connected through hub nodes and naturally tend to form small-world arrangements. Here is a brief survey of the evidence.

https://doi.org/10.1086/670331 Published online by Cambridge University Press

Hagmann and colleagues (2008) analyzed the structural connectivity in the human cerebral cortex at high spatial resolution. They considered cortical networks of 998 brain regions of interest (with average size of 1.5 cm^2) in six human subjects. Their modularity analyses revealed a set of regions of cortex that are highly central and highly connected. This set of regions appears to form a structural core of the human brain. Six structurally distinct modules interconnected by highly central hubs were also identified. While the structural core was located predominantly within the posterior medial cortex, the six modules consisted of densely connected, spatially contiguous structures spanning frontal, temporoparietal, and medial cortical regions. The connector hubs that linked these modules were located along the anterior-posterior medial axis of the cortex, including the rostral and the caudal anterior cingulate, the paracentral lobule, and the precuneus. More fine-grained analyses revealed additional hierarchically nested modular arrangements: two segregated clusters corresponding to the dorsal and ventral pathway were found in the visual cortex.

Chen and colleagues (2008) carried out a different modularity analysis of the structural connectivity of the human cortex at high spatial resolution, considering 45 regions. They identified six modules, each of which comprised between four and 10 cortical regions. These modules were located in anatomically distinct areas known to carry out processes supporting auditory/linguistic, strategic/executive, sensorimotor, visual, and mnemonic capacities. Chen and colleagues also identified several connector hubs, predominantly located in multimodal or association areas, which receive convergent inputs from multiple cortical regions.

Findings from studies on functional connectivity, adopting different methodological approaches, confirm the results concerning the multiscale modular architecture of structural brain networks. Meunier et al. (2009), for example, studied the modular organization of functional networks under resting (i.e., no-task) conditions at several hierarchical levels. Their results show that functional networks in the human brain have a hierarchical modular organization. At the highest level of the hierarchy, they identified fewer and larger modules, including a somatosensory motor module, a parietal module, and occipital modules. Hubs were identified in the association and cortical areas. At lower levels of the hierarchy, each of these larger modules was found to comprise several smaller submodules and sub-sub-modules.

He et al. (2009) confirmed that functional networks of the human brain present a highly organized modular architecture. Consistent with previous findings, the modules they identified under resting conditions comprised somatosensory, motor, auditory, occipital, and parieto-frontal regions. Highly connected and highly central hubs were identified in the association and limbic/paralimbic area. He and colleagues also found that each of the modules they identified presented unique patterns of internal organization, which sug-

https://doi.org/10.1086/670331 Published online by Cambridge University Press

gests that different modules at the same hierarchical level can present idiosyncratic network properties.

What does this brief survey indicate? Different network-scientific methods consistently show that both structural and functional brain networks display modularity. Modular networks are likely to involve nested hierarchies, ranging from coarse modules to more fine-grained ones, where different modules can present characteristic topological properties. Modularity appears to be a highly conserved organizational feature of the architecture of complex brain networks.

6. Conclusion. I hope to have convinced you that network science provides a more rigorous and tractable concept of modularity than the Fodorian or post-Fodorian accounts, which portends to lead to significant insight into the topological organization of the human cognitive architecture, brain function, and possibly brain evolution. The theoretical framework of network science can be the basis of a more productive research program on the human cognitive architecture.

REFERENCES

Atkinson, Anthony P., and Michael Wheeler. 2004. "The Grain of Domains: The Evolutionary-Psychological Case against Domain-General Cognition." *Mind and Language* 19:147–76.

Barrett, H. Clark. 2005. "Enzymatic Computation and Cognitive Modularity." *Mind and Language* 20:259–87.

———. 2009. "Where There Is an Adaptation, There Is a Domain: The Form-Function Fit in Information Processing." In *Foundations in Evolutionary Cognitive Neuroscience*, ed. Steven M. Platek and Todd K. Shackelford, 97–116. Cambridge: Cambridge University Press.

———. 2012. "A Hierarchical Model of the Evolution of Human Brain Specializations." *Proceedings of the National Academy of Sciences of the USA* 109:10733–40.

Barrett, H. Clark, and Robert Kurzban. 2006. "Modularity in Cognition: Framing the Debate." *Psychological Review* 113:628–47.

Bassett, Danielle S., Andreas Meyer-Lindenberg, Sophie Achard, Thomas Duke, and Edward Bullmore. 2006. "Adaptive Reconfiguration of Fractal Small-World Human Brain Functional Networks." *Proceedings of the National Academy of Sciences of the USA* 103:19518–23.

Bechtel, William, and Adele Abrahamsen. 2005. "Explanation: A Mechanist Alternative." *Studies in History and Philosophy of Biological and Biomedical Sciences* 36:421–41.

Bilder, Robert M., Fred W. Sabb, D. Stott Parker, Donald Kalar, Wesley W. Chu, Jared Fox, Nelson B. Freimer, and Russell A. Poldrack. 2009. "Cognitive Ontologies for Neuropsychiatric Phenomics Research." *Cognitive Neuropsychiatry* 14:419–50.

Börner, Katy, Soma Sanyal, and Alessandro Vespignani. 2007. "Network Science." *Annual Review of Information Science and Technology* 41:537–607.

Buller, David J. 2005. *Adapting Minds: Evolutionary Psychology and the Persistent Quest for Human Nature*. Cambridge, MA: MIT Press.

Bullmore, Edward, and Olaf Sporns. 2009. "Complex Brain Networks: Graph-Theoretical Analysis of Structural and Functional Systems." *Nature Reviews Neuroscience* 10:186–98.

———. 2012. "The Economy of Brain Network Organization." *Nature Reviews Neuroscience* 13:336–49.

Carruthers, Peter. 2006. *The Architecture of the Mind*. Oxford: Oxford University Press.

Chen, Zhang J., Yong He, Pedro Rosa-Neto, Jurgen Germann, and Alan C. Evans. 2008. "Revealing Modular Architecture of Human Brain Structural Networks by Using Cortical Thickness from MRI." *Cerebral Cortex* 18:2374–81.

https://doi.org/10.1086/670331 Published online by Cambridge University Press

Churchland, Patricia S. 1986. *Neurophilosophy: Toward a Unified Science of the Mind-Brain*. Cambridge, MA: MIT Press.

Cohen, Alexander L., Damien A. Fair, Nico U. F. Dosenbach, Francis M. Miezin, Donna Dierker, David C. Van Essen, Bradley L. Schlaggar, and Steven E. Petersen. 2008. "Defining Functional Areas in Individual Human Brains Using Resting State Functional Connectivity MRI." *Neuroimage* 41:45–57.

Coltheart, Max. 1999. "Modularity and Cognition." *Trends in Cognitive Sciences* 3:115–20.

Cosmides, Leda, and John Tooby. 1995. "From Function to Structure: The Role of Evolutionary Biology and Computational Theories in Cognitive Neuroscience." In *The Cognitive Neurosciences*, ed. Michael S. Gazzaniga, 1199–1210. Cambridge, MA: MIT Press.

———. 2008a. "Can a General Deontic Logic Capture the Facts of Human Moral Reasoning? How the Mind Interprets Social Exchange Rules and Detects Cheaters." In *Moral Psychology*, vol. 1, *The Evolution of Morality: Adaptations and Innateness*, ed. Walter Sinnott-Armstrong, 53–119. Cambridge, MA: MIT Press.

———. 2008b. "When Falsification Strikes: A Reply to Fodor." In *Moral Psychology*, vol. 1, *The Evolution of Morality: Adaptations and Innateness*, ed. Walter Sinnott-Armstrong, 143–64. Cambridge, MA: MIT Press.

Eliasmith, Chris, Terrence C. Stewart, Xuan Choo, Trevor Bekolay, Travis DeWolf, Yichuan Tang, and Daniel Rasmussen. 2012. "A Large-Scale Model of the Functioning Brain." *Science* 338:1202–5.

Ermer, Elsa, Leda Cosmides, and John Tooby. 2007. "Functional Specialization and the Adaptationist Program." In *The Evolution of Mind: Fundamental Questions and Controversies*, ed. Steven W. Gangestad and Jeffry A. Simpson, 153–59. New York: Guilford.

Fodor, Jerry A. 1983. *The Modularity of Mind*. Cambridge, MA: MIT Press.

———. 2000. *The Mind Does Not Work That Way*. Cambridge, MA: MIT Press.

———. 2008. "Comment on Cosmides and Tooby." In *Moral Psychology*, vol. 1, *The Evolution of Morality: Adaptations and Innateness*, ed. Walter Sinnott-Armstrong, 137–41. Cambridge, MA: MIT Press.

Fortunato, Santo. 2010. "Community Detection in Graphs." *Physics Reports* 486:75–174.

Friston, Karl J. 1994. "Functional and Effective Connectivity in Neuroimaging: A Synthesis." *Human Brain Mapping* 2:56–78.

———. 2011. "Functional and Effective Connectivity: A Review." *Brain Connectivity* 1:13–36.

Friston, Karl J., and Cathy J. Price. 2001. "Dynamic Representations and Generative Models of Brain Function." *Brain Research Bulletin* 54:275–85.

———. 2011. "Modules and Brain Mapping." *Cognitive Neuropsychology* 28:241–50.

Gallos, Lazaros K., Hernán A. Makse, and Mariano Sigman. 2012. "A Small World of Weak Ties Provides Optimal Global Integration of Self-Similar Modules in Functional Brain Networks." *Proceedings of the National Academy of Sciences of the USA* 109:2825–30.

Golland, Yulia, Polina Golland, Shlomo Bentin, and Rafael Malach. 2008. "Data-Driven Clustering Reveals a Fundamental Subdivision of the Human Cortex into Two Global Systems." *Neuropsychologia* 46:540–53.

Hagmann, Patric, Leila Cammoun, Xavier Gigandet, Reto Meuli, Christopher J. Honey, Van J. Wedeen, and Olaf Sporns. 2008. "Mapping the Structural Core of Human Cerebral Cortex." *PLoS Biology* 6:e159.

Hampshire, Adam, Roger R. Highfield, Beth Parkin, and Adrian M. Owen. 2012. "Fractionating Human Intelligence." *Neuron* 76:1225–37.

He, Yong, Jinhui Wang, Liang Wang, Zhang J. Chen, Chaogan Yan, Hong Yang, Hehan Tang, Chaozhe Zhu, Qiyong Gong, Yufeng Zang, and Alan C. Evans. 2009. "Uncovering Intrinsic Modular Organization of Spontaneous Brain Activity in Humans." *PLoS ONE* 4: e5226.

Kaplan, Jonathan M. 2002. "Historical Evidence and Human Adaptations." *Philosophy of Science* 69 (Proceedings): S294–S304.

Klein, Colin. 2012. "Cognitive Ontology and Region- versus Network-Oriented Analyses." *Philosophy of Science* 79:952–60.

Lickliter, Robert, and Hunter Honeycutt. 2003a. "Developmental Dynamics and Contemporary Evolutionary Psychology: Status quo or Irreconcilable Views? Replies to Critics." *Psychological Bulletin* 129:866–72.

https://doi.org/10.1086/670331 Published online by Cambridge University Press

———. 2003b. "Developmental Dynamics: Toward a Biologically Plausible Evolutionary Psychology." *Psychological Bulletin* 129:819–35.
Lloyd, Elisabeth A., and Marcus W. Feldman. 2002. "Evolutionary Psychology: A View from Evolutionary Biology." *Psychological Inquiry* 13:150–56.
Machery, Edouard. Forthcoming. "Discovery and Confirmation in Evolutionary Psychology." In *Oxford Handbook of Philosophy of Psychology*, ed. Jesse J. Prinz. Oxford: Oxford University Press.
Machery, Edouard, and Clark H. Barrett. 2006. "Debunking *Adapting Minds*." *Philosophy of Science* 73:232–46.
Marr, David. 1982. *Vision*. San Francisco: Freeman.
Marr, David, Shimon Ullman, and Tomaso Poggio. 1979. "Bandpass Channels, Zero-Crossings and Early Visual Information Processing." *Journal of the Optical Society of America* 69:914–16.
McIntosh, A. Randall. 2000. "Towards a Network Theory of Cognition." *Neural Networks* 13: 861–70.
Meunier, David, Renaud Lambiotte, and Edward Bullmore. 2010. "Modular and Hierarchically Modular Organization of Brain Networks." *Frontiers in Neuroscience* 4:200.
Meunier, David, Renaud Lambiotte, Alex Fornito, Karen D. Ersche, and Edward Bullmore. 2009. "Hierarchical Modularity in Human Brain Functional Networks." *Frontiers in Neuroinformatics* 3:37.
Newman, Mark E. J. 2006. "Modularity and Community Structure in Networks." *Proceedings of the National Academy of Sciences of the USA* 103:8577–8696.
Newman, Mark E. J., and Michelle Girvan. 2004. "Finding and Evaluating Community Structure in Networks." *Physical Review* E 69:026113.
Passingham, Richard E., Klaas E. Stephan, and Rolf Kötter. 2002. "The Anatomical Basis of Functional Localization in the Cortex." *Nature Review Neuroscience* 3:606–16.
Pinker, Steven. 1997. *How the Mind Works*. New York: Norton.
———. 2005. "So How *Does* the Mind Work?" *Mind and Language* 20:1–24.
Price, Cathy J., and Karl J. Friston. 2005. "Functional Ontologies for Cognition: The Systematic Definition of Structure and Function." *Cognitive Neuropsychology* 22:262–75.
Richardson, Robert C. 2007. *Evolutionary Psychology as Maladapted Psychology*. Cambridge, MA: MIT Press.
Rubinov, Mikail, and Olaf Sporns. 2010. "Complex Network Measures of Brain Connectivity: Uses and Interpretations." *Neuroimage* 52:1059–69.
Samuels, Richard. 2000. "Massively Modular Minds: Evolutionary Psychology and Cognitive Architecture." In *Evolution and the Human Mind: Modularity, Language and Meta-cognition*, ed. Peter Carruthers and Andrew Chamberlain, 13–46. Cambridge: Cambridge University Press.
Schulz, Armin W. 2011. "Heuristic Evolutionary Psychology." In *Philosophy of Behavioral Biology*, ed. Kathryn S. Plaisance and Thomas A. C. Reydon, 217–34. Berlin: Springer.
Segal, Gabriel. 1996. "The Modularity of Theory of Mind." In *Theories of Theories of Mind*, ed. Peter Carruthers and Peter K. Smith, 141–57. Cambridge: Cambridge University Press.
Sperber, Dan. 1994. "The Modularity of Thought and the Epidemiology of Representations." In *Mapping the Mind: Domain-Specificity in Cognition and Culture*, ed. Lawrence A. Hirschfeld and Susan A. Gelman, 39–67. Cambridge: Cambridge University Press.
Sperber, Dan, and Vittorio Girotto. 2003. "Does the Selection Task Detect Cheater Detection?" In *From Mating to Mentality: Evaluating Evolutionary Psychology*, ed. Kim Sterelny and Julie Fitness, 197–226. New York: Psychology.
Sporns, Olaf. 2007. "Brain Connectivity." *Scholarpedia* 2 (10): 4695.
———. 2011. *Networks of the Brain*. Cambridge, MA: MIT Press.
Sporns, Olaf, Giulio Tononi, and Rolf Kötter. 2005. "The Human Connectome: A Structural Description of the Human Brain." *PLoS Computational Biology* 1:245–51.
Sterelny, Kim, and Paul E. Griffiths. 1999. *Sex and Death: An Introduction to Philosophy of Biology*. Chicago: University of Chicago Press.
Tononi, Giulio, Olaf Sporns, and Gerald M. Edelman. 1994. "A Measure for Brain Complexity: Relating Functional Segregation and Integration in the Nervous System." *Proceedings of the National Academy of Sciences of the USA* 91:5033–37.

https://doi.org/10.1086/670331 Published online by Cambridge University Press

Tooby, John, and Leda Cosmides. 1992. "The Psychological Foundations of Culture." In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, ed. Jerome H. Barkow, Leda Cosmides, and John Tooby, 19–136. New York: Oxford University Press.

———. 2005. "Conceptual Foundations of Evolutionary Psychology." In *The Handbook of Evolutionary Psychology*, ed. David M. Buss, 5–67. Hoboken, NJ: Wiley.

Tooby, John, Leda Cosmides, and Clark H. Barrett. 2003. "The Second Law of Thermodynamics Is the First Law of Psychology: Evolutionary Developmental Psychology and the Theory of Tandem, Coordinated Inheritances; Comment on Lickleiter and Honeycutt (2003)." *Psychological Bulletin* 129:858–65.

Twyman, Alexandra D., and Nora S. Newcombe. 2010. "Five Reasons to Doubt the Existence of a Geometric Module." *Cognitive Science* 34:1315–56.

Wagner, Günter P., Jason Mezey, and Raffaele Calabretta. 2005. "Natural Selection and the Origin of Modules." In *Modularity: Understanding the Development and Evolution of Natural Complex Systems*, ed. Werner Callebaut and Diego Rasking-Gutman, 33–50. Cambridge, MA: MIT Press.

Wagner, Günter P., Mihaela Pavlicev, and James M. Cheverud. 2007. "The Road to Modularity." *Nature Review Genetics* 8:921–31.

https://doi.org/10.1086/670331 Published online by Cambridge University Press