# Neural Redundancy and Its Relation to Neural Reuse

John Zerilli*†

Evidence of the pervasiveness of neural reuse in the human brain has forced a revision of the standard conception of modularity in the cognitive sciences. One persistent line of argument against such revision, however, cites the evidence of cognitive dissociations. While this article takes the dissociations seriously, it contends that the traditional modular account is not the best explanation. The key to the puzzle is neural redundancy. The article offers both a philosophical analysis of the relation between reuse and redundancy as well as a plausible solution to the problem of dissociations.

**1. Introduction.** Cognitive science, linguistics, and the philosophy of psychology have long been under the spell of the "modularity of mind"—or the idea of the mind as a modular system (Fodor 1983). In contemporary psychology, a modular system is generally understood to be "one consisting of functionally specialized subsystems responsible for processing different classes of input (e.g. for vision, hearing, human faces, etc.), or at any rate for handling specific cognitive tasks" (Zerilli 2017, 231). According to this theory, "human cognition can be decomposed into a number of functionally independent processes, [where] each of these processes operates over a distinct domain of cognitive information" (Bergeron 2007, 176). What makes one process distinguishable from another is its "functional independence, the fact that one can be affected, in part or in totality, without the other being affected, and vice versa" (176). Furthermore, given that functional processes are realized in the brain, a functionally specialized process is one that presumably occupies a distinctive portion of neural tissue, although not necessarily a small, closely circumscribed and contiguous region. So fruitful and

*To contact the author, please write to: Leverhulme Centre for the Future of Intelligence, University of Cambridge; e-mail: jz303@cam.ac.uk.

†I gratefully acknowledge the suggestions of those who either read or heard this article, in particular Kim Sterelny, Andy Barron, and Richard Menary.

Philosophy of Science, 86 (December 2019) pp. 1191–1201. 0031-8248/2019/8605-0032$10.00
Copyright 2019 by the Philosophy of Science Association. All rights reserved.



https://doi.org/10.1086/705529 Published online by Cambridge University Press

influential has this model been that it is safe to say that in many quarters of the cognitive sciences—and most especially in cognitive psychology, cognitive neuropsychology, and evolutionary psychology—modularity is essentially the received view.

Developments in cognitive neuroscience over the past 30 years, however, have discomfited the modular account. More evidence than ever before points to the pervasiveness of neural reuse in the human brain—the "redeployment" or "recycling" of neural circuits over widely disparate cognitive domains (Anderson 2010, 2014). As the terminology suggests, theories of "reuse" posit the "exaptation" of established and diachronically stable neural circuits over the course of evolution or normal development without loss of original function, so that the functional contribution of a circuit is preserved across multiple task domains. As Anderson (2010, 246) explains, "rather than posit a functional architecture for the brain whereby individual regions are dedicated to large-scale cognitive domains like vision, audition, language and the like, neural reuse theories suggest that low-level neural circuits are used and reused for various purposes in different cognitive and task domains." According to the theory, just the same circuits exapted for one purpose can be exapted for another, provided sufficient intercircuit pathways exist to allow alternative arrangements of them. Indeed, the same parts put together in different ways will yield different functional outcomes, just as "if one puts together the same parts *in the same way* one will get the same functional outcomes" (247, my emphasis). The evidence here converges from heterogeneous sources and research paradigms, including neuroimaging, computational, biobehavioral, and interference paradigms, and exempts practically no area of the brain, including areas long regarded as specialized hubs for certain types of sensory processing (e.g., visual and auditory pathways). Among other things, this means that one of the hallmark features of a module—its domain specificity—looks too stringent a requirement to prove useful. For neural reuse demonstrates that any one module will typically be sensitive to more than one stimulus, including—most importantly—those channeled along intermodal pathways. Meanwhile efforts to salvage a computational or "software" theory of modularity, which carries no commitments regarding implementation, have met with skepticism (Anderson 2010), if not outright opposition (Zerilli 2017). And while the brain could still be modular in some other sense, what is clear is that the strict domain-specific variety of modularity can no longer serve as an appropriate benchmark.

And yet there is a persistent line of argument against this conclusion that draws from a large body of experimental literature attesting to the existence of cognitive dissociations, in which a cognitive ability (say language) is either selectively impaired (linguistic ability is compromised, but no other cognitive ability seems to be materially affected) or selectively spared (general intelligence is compromised, while linguistic abilities function more or

https://doi.org/10.1086/705529 Published online by Cambridge University Press

less as they should). This literature, most vividly exemplified in lesion studies, is frequently cited in support of classical modularities of mind—be they inspired by the likes of Jerry Fodor (1983), evolutionary psychology, or some variation thereof. While numerous rejoinders to this line of thinking have been offered over the years, few have grappled seriously with the phenomenon, either dismissing the dissociations as noisy or reasoning from architectural considerations that even nonmodular systems can generate dissociations (Plaut 1995). The aim of this article is to offer a fresh perspective on this vexed topic. I take the dissociation evidence seriously, on the one hand, while affirming that traditional modularities of mind do not do justice to the evidence of neural reuse, on the other. I do this by invoking neural redundancy, an important feature of cortical design that ensures we have various copies of the same elementary processing units that can be put to alternative (if computationally related) uses in enabling diverse cognitive functions. In the course of the discussion I offer a philosophical explication of the relationship between neural reuse and neural redundancy.

**2. What Is the Problem? Cognitive Dissociations and Neural Reuse.** Let us take an especially contentious question to underscore the nature of the problem we are dealing with and how redundancy might assist in its illumination. The question is this: Does language rely on specialized cognitive and neural machinery, or does it rely on the same machinery that allows us to get by in other domains of human endeavor? The question is bound up with many other questions of no less importance, questions concerning the uniqueness of the human mind, the course of biological evolution, and the power of human culture. What is perhaps a little unusual about this question, however—unusual for a question whose answer concerns both those working in the sciences and the humanities—is that it can be phrased as a polar interrogative, that is, as a question that admits of a yes-or-no response. And indeed the question has divided psychologists, linguists, and the cognitive science community generally for many decades now more or less into two camps. I would like to sketch the beginnings of an answer to this question—and others like it—in a way that does not pretend it can receive a simple yes-or-no response.

First of all, let me stress again that neural reuse is as well-verified a phenomenon as one can expect in the cognitive sciences and that it has left virtually no domain of psychology untouched. Neural reuse suggests that there is nothing so specialized in the cortex that it cannot be repurposed to meet new challenges while retaining its capacity for meeting old ones. In that regard, to be sure, what I am proposing is unapologetically on the side of those who maintain that language, as well as many other psychological capacities, is not cognitively special (e.g., that there is no domain-specific "language organ"; cf. Chomsky 2002, 84–86).

https://doi.org/10.1086/705529 Published online by Cambridge University Press

And yet I would like to carefully distinguish this claim from the claim that there are no areas of the brain that subserve exclusively linguistic functions. The neuropsychological literature offers striking examples of what appear to be fairly clean dissociations between linguistic and nonlinguistic capacities, that is, cases in which language processing capacities appear to be disrupted without impeding other cognitive abilities and cases in which the reverse situation holds (Poeppel 2001). An example would be where the ability to hear words is disrupted, but the ability to recognize nonword sounds is spared. Although dissociations are to some extent compatible with reuse—indeed there is work suggesting that focal lesions can produce specific cognitive impairments within a range of nonclassical architectures (Plaut 1995)—and it is equally true that often the dissociations reported are noisy, still their very ubiquity needs to be taken seriously and accounted for in a more systematic fashion than many defenders of reuse have been willing to do (see, e.g., Anderson 2010, 248; 2014, 46–48). After all, a good deal of support for theories of reuse comes from the neuroimaging literature, which is somewhat ambiguous taken by itself. As Fedorenko, Behr, and Kanwisher (2011, 16428) explain: "standard functional MRI group analysis methods can be deceptive: two different mental functions that activate neighbouring but non-overlapping cortical regions in every subject individually can produce overlapping activations in a group analysis, because the precise locations of these regions vary across subjects, smearing the group activations. Definitively addressing the question of neural overlap between linguistic and nonlinguistic functions requires examining overlap within individual subjects, a data analysis strategy that has almost never been applied in neuroimaging investigations of high-level linguistic processing."

When Fedorenko and her colleagues applied this strategy themselves, they found that "most of the key cortical regions engaged in high-level linguistic processing are not engaged by mental arithmetic, general working memory, cognitive control or musical processing," and they think that this indicates "a high degree of functional specificity in the brain regions that support language" (2011, 16431). While I do not believe that claims of this strength have the least warrant—as I will explain, functional specificity cannot be established merely by demonstrating that a region is selectively engaged by a task—these results do at least substantiate the dissociation literature in an interesting way and make it more difficult for those who would prefer to dismiss the dissociations with a ready-made list of alternative explanations.

**3. How Might Redundancy Feature in a Solution?** With rare exceptions (e.g., Friston and Price 2003; Barrett and Kurzban 2006; Jungé and Dennett 2010), redundancy has passed almost unnoticed in the philosophical and cognitive science literature. This is in stark contrast to the epigenetics literature, where redundancy and the related concept of degeneracy have been

https://doi.org/10.1086/705529 Published online by Cambridge University Press

explored to some depth (e.g., see Edelman and Gally 2001; Whiteacre 2010; Iriki and Taoka 2012).[1] The idea behind neural redundancy is that, for good evolutionary reasons (see below), the brain incorporates a large measure of redundancy of function. Brain regions (such as cortical columns and similar structures) fall in an iterative, repetitive, and almost lattice-like arrangement in the cortex. Neighboring columns have similar response properties: laminar and columnar changes are for the most part smooth—not abrupt—as one moves across the cortex, and adjacent modules do not differ markedly from one another in their basic structure and computations (if they really differ at all when taken in such proximity). Regional solitarity is therefore not likely to be a characteristic of the brain (Anderson 2014, 141).[2] That is to say, we do not possess just one module for X, and one module for Y, but in effect several copies of the module for X, and several copies of the module for Y, all densely stuffed into the same cortical zones. As Jungé and Dennett (2010, 278) conjecture: "It is possible that specialized brain areas contain a large amount of structural/computational redundancy (i.e., many neurons or collections of neurons that can potentially perform the same class of functions). Rather than a single neuron or small neural tract playing roles in many high-level processes, it is possible that distinct subsets of neurons within a specialized area have similar competencies, and hence are redundant, but as a result are available to be assigned individually to specific uses. . . . In a coarse enough grain, this neural model would look exactly like multi-use (or reuse)." This is plausibly why capacities that are functionally very closely related, but that for whatever reason are forced to recruit different neural circuits, will often be localized in broadly the same regions of the brain. For instance, first and second languages acquired early in ontogeny settle down in nearly the same region of Broca's area, and even when the second language is acquired in adulthood the second language is still represented within Broca's area (while artificial languages are not). The neural coactivation graphs of such composite networks must look very similar. Indeed these results suggest—and a redundancy model would predict—that two very similar tasks that are forced to recruit different neural circuits should exhibit similar patterns of activation. And this is more or less what we find (see below).

One might be tempted to think that redundancy and reuse pull in opposite directions. This is because whereas reuse posits that neural circuits get reused

1. Redundancy occurs when items have the same structure and function (i.e., are both isomorphic and isofunctional). Degeneracy occurs when items having different structures can perform the same function (i.e., are heteromorphic but isofunctional). Degeneracy implies genuine multiple realization (see Zerilli 2019).

2. The term "solitarity" is Anderson's (2014), but while he concedes that solitarity will be "relatively rare," he does not appear to believe that anything particularly significant follows from this. See also Anderson (2010, 296).

https://doi.org/10.1086/705529 Published online by Cambridge University Press

across different tasks and task categories, redundancy accommodates the likelihood of diverse cognitive functions being activated by structurally and computationally equivalent circuits running in parallel: instead of a single circuit being reused across domains, two, three, or more copies of that same circuit may be recruited differentially across those domains, such that no single circuit gets literally "reused." But there is no substantive tension here. The redundancy account in truth supplements the reuse picture in a way that is consistent with the neuroimaging data, faithful to the core principle of reuse, and compatible with the apparent modularization and separate modifiability of technical and acquired skills in ontogeny. Evidence of the reuse of neural circuits to accomplish different tasks has, in fact, been adduced in aid of a theory that posits the reuse of the same neural tokens to accomplish these different tasks. Redundancy means we must accept that at least some of the time what we may actually be witnessing is reuse of the same types to accomplish these tasks. This does not diminish the standing of reuse. Let me explain.[3]

To the extent that a particular composite reuses types and is dissociable pro tanto—residing in segregated brain tissue that is not active outside the domain in question—it is true that to that extent its constituents will appear to be domain specific. But in this case looks will be deceiving. The classical understanding of domain specificity in effect assumes solitarity—that a module for X does something that no other module can do as well, or that even if another module can do X as well, taken together these X-ing modules do not perform outside the X-domain (see, e.g., Bergeron 2007, 176). But a module that appears dedicated to a particular function may not be domain specific in the classical sense. Dedication is not the same as domain specificity, and redundancy explains why. A composite of neural regions will be dedicated without being domain specific if its functional resources are accessible to other domains through the deployment (reuse) of neural surrogates (i.e., redundant or "proxy" tokens). In this case its constituents will be multipotential but single use (Jungé and Dennett 2010, 278), and the domain specificity on display somewhat cosmetic. To take an example, a set of cortical columns that are structurally and computationally similar may be equally suited for face recognition tasks, abstract-object recognition tasks, the recognition of moving objects, and so on. One of these columns could be reserved for faces, another for abstract objects, another for moving objects, and so on. What is noteworthy is that while the functional activation may be indistinguishable in each case, and the same type of resource will be employed on each occasion, a different token module will be at work at any one time. To quote Jungé and Dennett again: "In an adult brain, a given neuron [or set of neurons] would

3. For a developmental twist on the type/token distinction invoked in the context of modular theorizing about the mind, see Barrett (2006).

https://doi.org/10.1086/705529 Published online by Cambridge University Press

be aligned with only a single high-level function, whereas each area of neurons would be aligned with very many different functions" (278). Such modules (and composites) are for all intents and purposes qualitatively identical, although clearly not numerically identical, meaning that while they share their properties, they are not one and the same. The evidence of reuse is virtually all one way when it comes to the pervasiveness of functional inheritance across cognitive domains. It may be that this inheritance owes to reuse of the same tokens (literal reuse) or to reuse of the same types (reuse by proxy), but the inheritance itself has been amply attested. This broader notion of reuse still offers a crucial insight into the operations of cognition, and I dare say it represents a large part of the appeal of the original massive redeployment hypothesis.

It is interesting to note in this respect that although detractors have frequently pointed out the ambiguity of neuroimaging evidence on account of its allegedly coarse spatial resolution, suggesting that the same area will be active across separate tasks and task categories even if distinct but spatially adjacent or interdigitated circuits are involved in each case, this complaint can have no bearing on reuse by proxy. Fedorenko et al. (2011, 16431) take their neuroimaging evidence to support "a high degree of functional specificity in the brain regions that support language," but their results do not license this extreme claim. The regions they found to have been selectively engaged by linguistic tasks were all adjacent to the regions engaged in nonlinguistic tasks. Elementary considerations suggest that they have discovered a case of reuse by proxy involving language: the domains tested (mental arithmetic, general working memory, cognitive control, and musical processing) make use of many of the same computations as high-level linguistic processing, even though they run them on duplicate hardware. Redundancy makes it is easy to see how fairly sharp dissociations could arise—knocking out one token module need disrupt only one high-level operation: other high-level operations that draw on the same type of resource may well be spared.

The consequences of this distinction between literal reuse and reuse by proxy for much speculation about the localization and specialization of function are potentially profound. In cognitive neuropsychology the discovery that a focal lesion selectively impairs a particular cognitive function is routinely taken as evidence of its functional specificity (Coltheart 2011; Sternberg 2011). Even cognitive scientists who take a developmental approach to modularity, that is, who concede that parts of the mind may be modular but stress that modularization is a developmental process, concede too much when they imply, as they frequently do, that modularization results in domain-specific modules (Karmiloff-Smith 1992; Barrett 2006; Guida, Campitelli, and Gobet 2016). This is true in some sense, but not in anything like the standard sense, for redundancy envisages that developmental modules form a special class of neural networks, namely, those that are qualitatively identical but numerically distinct. The appearance of modularization in development

https://doi.org/10.1086/705529 Published online by Cambridge University Press

is thus fully compatible with deep domain interpenetration. In any event, redundancy does not predict that all acquired skills will be modular. The evidence suggests that while some complex skills reside in at least partly dissociable circuitry, most complex skills are implemented in more typical neural networks, that is, those consisting of literally shared parts.[4]

**4. What Else Might Redundancy Explain?** It is generally a good design feature of any system to have spare capacity. For instance, in engineered systems, "redundant parts can substitute for others that malfunction or fail, or augment output when demand for a particular output increases" ( Whiteacre 2010, 14). The positive connection between robustness and redundancy in biological systems is also clear (Edelman and Gally 2001; Whiteacre 2010; Iriki and Taoka 2012). So there are good reasons for evolution to have seen to it that our brains have spare capacity. But in the case of the brain and the cortex most especially, there are other reasons why redundancy would be an important design feature. It offers a solution to what Jungé and Dennett (2010, 278) called the "time-sharing" problem. It may also offer a solution to what I call the "encapsulation" problem.

The time-sharing problem arises when multiple simultaneous demands are made on the same cognitive resource. This is probably a regular occurrence. Here are just a few examples.

- Driving a car and holding a conversation at the same time: If it is true that some of the selfsame motor operations underlying aspects of speech production and comprehension are also required for the execution of sequenced or complex motor functions (Pulvermüller and Fadiga 2010), as perhaps exemplified by driving a manual vehicle or operating complex machinery (e.g., playing the organ), how do we manage to pull this off?
- By reflecting the recursive structure of thought (Christiansen and Chater 2016, 51), the language circuits may redeploy a recursive operation simultaneously during sentence production. This might be the case during the formation of an embedded relative clause—the thought and its encoding may require parallel use of the same sequencing principle. Again, how do we manage this feat?
- If metarepresentational operations are involved in the internalization of conventional sound-meaning pairs, and also in the pragmatics and mind reading that carry on simultaneously during conversation, could this not simply be another instance of time-sharing? The example is contentious, but it still raises the question: How does our brain manage to do things like this?

4. This seems to be true regardless of whether the complex skills are innate or acquired.

https://doi.org/10.1086/705529 Published online by Cambridge University Press

- Christiansen and Chater's (2016) Chunk and Pass model of language processing envisages multilevel and simultaneous chunking procedures. As they put it, "the challenge of language acquisition is to learn a dazzling sequence of rapid processing operations" (116). What must the brain be like to allow for this dazzling display?

Explaining these phenomena is difficult. Indeed when dealing with clear (literal) instances of reuse, results from the interference paradigm show that processing bottlenecks are inevitable—true multitasking is impossible. Redundancy offers a natural explanation of how the brain overcomes the time-sharing problem. It explains, in short, how we are able to walk and chew gum at the same time.

Redundancy might also offer a solution to what I have called the encapsulation problem. The neural networks that implement cognitive functions are not likely to be characterized by informational encapsulation if they share their nodes with networks implementing other cognitive functions. This is because in sharing their nodes with these other systems they will prima facie have access to the information stored and manipulated by those other systems (Anderson 2010, 300). If, then, overlapping brain networks must share information, it would be reasonable to suppose that central and peripheral systems do not overlap. For peripheral systems, which are paradigmatically fast and automatic, would not be able to process inputs as efficiently if there were a serious risk of central system override—that is, of beliefs and other central information getting in the way of automatic processing. But we know from the neuroimaging literature that quite often the brain networks implementing central and peripheral functions do overlap. This is puzzling in light of the degree of cognitive impenetrability that certain sensory systems still seem to exhibit—limited although it may be. If it is plausible to suppose that the phenomenon calls for segregated circuitry, redundancy could feature in a solution to the puzzle, since it naturally explains how the brain can make parallel use of the same resources. Neuroimaging maps might well display what appear to be overlapping brain regions between two tasks (one involving central information, the other involving classically peripheral operations), but the overlap would not exist—there would be distinct albeit adjacent or interdigitated and nearly identical circuits recruited in each case. Of course there may be other ways around the encapsulation problem that do not require segregated circuitry: the nature and extent of the overlap is presumably important. But clearly redundancy opens up some fascinating explanatory possibilities.

To the extent that acquired skills must overcome both the time-sharing problem as well as the encapsulation problem—for acquired competencies are often able to run autonomously of central processes—we might expect that their neural implementations incorporate redundant tissue. In concluding,

https://doi.org/10.1086/705529 Published online by Cambridge University Press

let me illustrate this point by offering a gloss on a particular account of how skills and expertise are acquired during development elaborated by Anderson (2014) and Guida et al. (2016). The process involved is called "search" (Anderson 2014). Search is an exploratory synaptogenetic process, "the active testing of multiple neuronal combinations until finding the most appropriate one for a specific skill, i.e., the neural niche of that skill" (Guida et al. 2016, 13). The theory holds that in the early stages of skill acquisition, the brain must search for an appropriate mix of brain areas and does so by recruiting relatively widely across the cortex. When expertise has finally developed, a much narrower and more specific network of brain areas has been settled on, such that "as a consequence of their extended practice, experts develop domain-specific knowledge structures" (13). The gloss (and my hunch) is this: first, that repeated practice of a task that requires segregation (to get around time-sharing and encapsulation issues) will in effect force a search into redundant neural territory (Karmiloff-Smith 1992; Barrett 2006; Barrett and Kurzban 2006); second, that search will recruit idle or relatively underused circuits in preference to busy ones as a general default strategy. Guida et al. (2016) cite evidence that experts' brains reuse areas for which novices' brains make only limited use: "Whereas novices use episodic long-term memory areas (e.g., the mediotemporal lobe) for performing long-term memory tasks, experts are able to (re)use these areas also for performing working-memory tasks" (14). Guido and colleagues, in agreement with Anderson (2014), seem to have literal reuse in mind. But the same evidence they cite is consistent with reuse by proxy.

**5. Conclusion.** It is true that language and other cognitive skills frequently dissociate from other skills, but redundancy puts this sort of modularization in its proper context. Redundancy predicates functional inheritance across tasks and task categories even when the tasks are implemented in spatially segregated neural networks. Thus, dissociation evidence alone does not always indicate true functional specificity. In particular, these dissociations provide no evidence that language is cognitively special vis-à-vis other cognitive domains.

## REFERENCES

Anderson, Michael L. 2010. "Neural Reuse: A Fundamental Organizational Principle of the Brain." *Behavioral and Brain Sciences* 33 (4): 245–66; discussion 266–313.

———. 2014. *After Phrenology: Neural Reuse and the Interactive Brain*. Cambridge, MA: MIT Press.

Barrett, H. Clark. 2006. "Modularity and Design Reincarnation." In *The Innate Mind*, vol. 2, *Culture and Cognition*, ed. Peter Carruthers, Stephen Laurence, and Stephen P. Stich, 199–217. New York: Oxford University Press.

Barrett, H. Clark, and Robert Kurzban. 2006. "Modularity in Cognition: Framing the Debate." *Psychological Review* 113 (3): 628–47.

https://doi.org/10.1086/705529 Published online by Cambridge University Press

Bergeron, Vincent. 2007. "Anatomical and Functional Modularity in Cognitive Science: Shifting the Focus." *Philosophical Psychology* 20 (2): 175–95.
Chomsky, Noam. 2002. *On Nature and Language*. New York: Cambridge University Press.
Christiansen, Morten H., and Nick Chater. 2016. *Creating Language: Integrating Evolution, Acquisition, and Processing*. Cambridge, MA: MIT Press.
Coltheart, Max. 2011. "Methods for Modular Modelling: Additive Factors and Cognitive Neuropsychology." *Cognitive Neuropsychology* 28 (3–4): 224–40.
Edelman, Gerald M., and Joseph A. Gally. 2001. "Degeneracy and Complexity in Biological Systems." *Proceedings of the National Academy of Sciences of the USA* 98 (24): 13763–68.
Fedorenko, Evelina, Michael K. Behr, and Nancy Kanwisher. 2011. "Functional Specificity for High-Level Linguistic Processing in the Human Brain." *Proceedings of the National Academy of Sciences of the USA* 108 (39): 16428–33.
Fodor, Jerry A. 1983. *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, MA: MIT Press.
Friston, Karl J., and Cathy J. Price. 2003. "Degeneracy and Redundancy in Cognitive Anatomy." *Trends in Cognitive Sciences* 7 (4): 151–52.
Guida, Alessandro, Guillermo Campitelli, and Fernand Gobet. 2016. "Becoming an Expert: Ontogeny of Expertise as an Example of Neural Reuse." *Behavioral and Brain Sciences* 39:13–15.
Iriki, Atsushi, and Miki Taoka. 2012. "Triadic (Ecological, Neural, Cognitive) Niche Construction: A Scenario of Human Brain Evolution Extrapolating Tool Use and Language from the Control of Reaching Actions." *Philosophical Transactions of the Royal Society* B 367:10–23.
Jungé, Justin A., and Daniel C. Dennett. 2010. "Multi-Use and Constraints from Original Use." *Behavioral and Brain Sciences* 33 (4): 277–78.
Karmiloff-Smith, Annette. 1992. *Beyond Modularity: A Developmental Perspective on Cognitive Science*. Cambridge, MA: MIT Press.
Plaut, David C. 1995. "Double Dissociation without Modularity: Evidence from Connectionist Neuropsychology." *Journal of Clinical and Experimental Psychology* 17 (2): 291–321.
Poeppel, David. 2001. "Pure Word Deafness and the Bilateral Processing of the Speech Code." *Cognitive Science* 21 (5): 679–93.
Pulvermüller, Friedmann, and Luciano Fadiga. 2010. "Active Perception: Sensorimotor Circuits as a Cortical Basis for Language." *Nature Reviews Neuroscience* 11:351–60.
Sternberg, Saul. 2011. "Modular Processes in Mind and Brain." *Cognitive Neuropsychology* 28 (3–4): 156–208.
Whiteacre, James M. 2010. "Degeneracy: A Link between Evolvability, Robustness and Complexity in Biological Systems." *Theoretical Biology and Medical Modelling* 7 (6): 1–17.
Zerilli, John. 2017. "Against the 'System' Module." *Philosophical Psychology* 30 (3): 235–50.
———. 2019. "Multiple Realization and the Commensurability of Taxonomies." *Synthese* 196 (8): 3337–53.

https://doi.org/10.1086/705529 Published online by Cambridge University Press