Multiple Realizability and Biological Modality

Rami Koskinen*†

Critics of multiple realizability have recently argued that we should concentrate solely on actual here-and-now realizations that are found in nature. The possibility of alternative, but unactualized, realizations is regarded as uninteresting because it is taken to be a question of pure logic or an unverifiable scenario of science fiction. However, in the biological context only a contingent set of realizations is actualized. Drawing on recent work on the theory of neutral biological spaces, the article shows that we can have ways of assessing the modal dimension of multiple realizability that do not have to rely on mere conceivability.

1. Introduction. In a series of critical works culminating in their recent *The Multiple Realization Book*, Thomas Polger and Lawrence Shapiro draw the distinction between things being multiply realiz*able* and multiply realiz*ed* (e.g., Shapiro 2008; Polger and Shapiro 2016). Furthermore, they claim, it is only the latter half of this conceptual distinction that is a well-motivated topic of study for the philosophers of science and mind. It is not that the former class of entities is empty or somehow ill defined. To the contrary. The problem, according to them, lies in the fact that there are too many multiply realizable things, and we even possess a perfectly good algorithm for recognizing them: elementary logic. According to their criticism, the question whether something is multiply realizable is uninteresting precisely because it can be answered on the basis of pure analytical reasoning: "There appears to be no logical contradiction in the possibility that the same psychological property, for example, being hungry, might be realized in various physical ways. Moreover, there is no obvious nomological constraint on multiple realization" (Shapiro 2008, 514).

*To contact the author, please write to: Department of Philosophy, University of Vienna, Universitätsstraße 7, 1010 Vienna, Austria; e-mail: rami.koskinen@univie.ac.at.

†I would like to thank Tarja Knuuttila as well as the audience at the PSA 26th Biennial Meeting, Seattle, November 2018 for helpful comments. Thanks also to Tero Kuparinen for helping with the preparation of fig. 1.

Philosophy of Science, 86 (December 2019) pp. 1123–1133. 0031-8248/2019/8605-0026$10.00
Copyright 2019 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/705478 Published online by Cambridge University Press

It is no wonder, then, that it is the realiz*ed* part of the distinction that Polger and Shapiro take to be the more substantial. This shows in Shapiro's (2008, 514) earlier discussion on Hilary Putnam's original treatise on alternative physical realizations for mental properties: "although Putnam [1975], in his groundbreaking discussion of multiple realizability, mentions only pain and hunger as examples of mental properties that are likely to be realized differently in mammals and mollusks, he clearly believes multiple realization to extend to many other psychological properties as well. This would be unremarkable if it were only the multiple realiz*ability* of psychological states at issue, but Putnam took himself to be offering something much more, namely, that psychological properties are actually multiply realized."

Shapiro and Polger have not been the only ones with worries concerning sloppy multiple realization talk. Indeed, in their highly influential critique of the Multiple Realizability Thesis (MRT), Bechtel and Mundale (1999, 176) seem to use the same kind of distinction. They also note that it is not the mere logical possibility that philosophers usually take to be the main thrust of MRT but rather some more concrete belief about things actually being multiply realized out there in the real biological world. What seems to be troubling, they argue, is that this belief is often based on just pure intuition, not on any actual empirical research. The moral seems to be that, instead of speculating whether a thing could be multiply realizable by means of, say, yet-to-be-discovered forms of alien life or the methods of a future science fiction technology, philosophers should rather concentrate on the firm empirical findings that tell us whether things, as they now stand, actually are multiply realized.[1]

Although MRT seems to be stating an empirical hypothesis, Putnam and other early advocates of the thesis relied heavily on conceptual intuitions and thought experiments. Moreover, if philosophical case studies of various natural phenomena provide no evidence of actual multiple realization, MRT seems to be under a serious threat (Bechtel and Mundale 1999; Shapiro 2000, 2008; Polger 2009). As Polger and Shapiro (2016, 53) nicely summarize the situation, "Lacking evidence of actual multiple realization, what kind of evidence for multiple realizability could we have?"

The central aim of this article, in contrast, is to argue that regardless of the evidence of actual multiple realization, we can have ways of assessing possible multiple realizability. What is more, I will show that this does not have to rely on mere conceivability. I focus on the biological context in which only a contingent subset of possible realizers will ever become actualized. Recent work by evolutionary systems and synthetic biologists does not only address how things have been realized in naturally evolved systems; it also explores possible realization bases. For researchers in these fields the question whether some realizer is biologically possible is not just a matter of logic. Rather, it is

1. For a similar point, see also Wimsatt (2007, 384–85).

https://doi.org/10.1086/705478 Published online by Cambridge University Press

an important empirical-cum-hypothetical question. Furthermore, if these new fields can provide scientific knowledge on the modal dimension of multiple realizability, it seems that we can have evidence for it that is relatively independent of actual multiple realization in nature.

The outline of the article is as follows. Section 2 provides a short background for the argument by discussing evolutionary accounts of MRT. Rosenberg (2001, 2006) has argued that the truth of MRT should follow from the fact that natural selection selects for effects, not realizers. Although a useful starting point for inquiry, my modal analysis of MRT shows why Rosenberg's argument ultimately fails for an important class of biological systems. In section 3, it is shown how according to a recent evolutionary systems biological theory of neutral spaces (Wagner 2005, 2014) there exists a huge space of equivalent solutions to most biological functions, only a tiny portion of which will ever be realized. Through the careful modeling of the topological structure of this space, we can gain knowledge of the unactualized solutions and also explain why some of them have remained unreachable. The situation will then be highlighted by an example of a neutral space of genetic alphabets in section 4. Section 5 concludes the article by suggesting how some of these solutions might nevertheless become realized by rendering them targets for biological engineering.

2. Multiple Realizability and Biological Evolution. Although the paradigmatic arguments for multiple realizability have often come from the philosophy of mind, biological examples have also played a central role in the assessment of MRT. Indeed, most of the cases that are discussed tend to come from the functional life sciences broadly understood, be they in the context of genes or neurocognitive architectures or organs like eyes (e.g., Hull 1974, 116, 121; Kincaid 1990; Bechtel 2008, 137–43; Weiskopf 2011; Polger and Shapiro 2016). Furthermore, ever since Putnam's original essay on multiple realizability in the context of the philosophy of mind, evolutionary considerations have had a distinctive part in arguments concerning MRT. This is also true of many of the more recent case studies on the topic (e.g., Bechtel and Mundale 1999; Keeley 2000; Shapiro 2000; Richardson 2008; Weiskopf 2011). Rosenberg (2001, 2006) has even argued that MRT is in fact a pretty straightforward consequence of evolution by natural selection and that this is the fundamental mechanism that explains why we should expect it to hold across all the other special sciences as well, since they are ultimately based on biological entities and processes.

The main idea behind Rosenberg's claim is that natural selection is what can be termed a "structure-blind" process (2001, 2006, 22, 30). This means that it cannot discriminate between two different solutions to the same problem, given that these differences do not bear on the system's fitness. As such, we should expect selection to routinely find multiple solutions to the same

https://doi.org/10.1086/705478 Published online by Cambridge University Press

functional problems (Rosenberg 1985, 2001, 2006, 22, 41–42; for a similar claim, see also Wimsatt [2007], 29). Moreover, in an evolving biological world that is in a constant flux, different solutions to the same problem may be favorable in different times and under different circumstances. According to Rosenberg, it is almost guaranteed that different realizers for most biological functions will be tried out sooner or later in the course of evolution: "If the 'space' of adaptational 'moves' and countermoves is very large, in the long run every functional kind will at least sometimes be realized by multiple causally distinct realizations" (2001, 367).[2] As candidates of multiply realized biological functions, Rosenberg mentions familiar concepts from Mendelian genetics, like allelic dominance, but also some less discussed ones, like the secondary structure of hemoglobin molecules and the mechanisms of cellular energy production (Rosenberg 1985, 61, 89, 95, 102).

However, as an empirical fact, evolution often seems to be rather conservative in its use of different designs, and many biological functions often share realizers even across species (Bechtel and Mundale 1999; Raerinne and Eronen 2012). Because of various tacit evolutionary constraints and the prevalence of functional homologies, it seems that the only way to go is to evaluate the disparity of biological realizers on a case-by-case basis. Plausible instances of multiple realization include, among others, the location mechanisms of weakly electric fish produced by convergent evolution (Keeley 2000) and the lateral inhibition of the eyes of horseshoe crabs and different mammalian species (Weiskopf 2011). However, with other cases, like the cortical areas of different mammalian brains (Bechtel and Mundale 1999) or the optical solutions of human and cephalopod eyes, multiple realization seems not to be empirically vindicated (Shapiro 2000). In line with Wimsatt's (2007, 29) remark that "evolution feeds opportunistically on outcomes, manipulating probabilities by diverse paths so that *multiple realizability and multiple exceptions* are inevitable," it does indeed seem that there are also multiple exceptions to multiple realizability.

More illuminating tools are needed if one wants to critically evaluate the scope of multiple realizability in nature. In the next section, I argue that the concepts of biological robustness and neutral biological spaces (Wagner 2005, 2014) provide such tools. They can save the intuitively appealing idea that there is a link between multiple realizability and biological evolution. However, they have a more restricted scope than conceptual arguments and also point toward cases in which multiple realizability is not to be expected. Moreover, by providing a rigorous framework for assessing the differential modality of function-equivalent solutions in nature, the distinction between

2. See also Rosenberg (1985, 48): "There are many ways to skin a cat, and nature has at one time or another employed most of them."

https://doi.org/10.1086/705478 Published online by Cambridge University Press

what is realizable and what is not becomes more nuanced than what recent critics of MRT have assumed.

3. Robustness, Functional Equivalence, and Neutral Biological Spaces. Most biological systems are characterized by their ability to maintain their functioning in the face of considerable environmental changes, genetic mutations, and metabolic fluctuations. This is often expressed by saying that biological systems are *robust*. According to an influential view in systems biological theorizing, a key element behind robustness seems to be that biological functions are often realizable in multiple different structures, mechanisms, or processes (e.g., Edelman and Gally 2001; Greenspan 2001; Kitano 2007; Chouard 2008).

A recent, theoretically powerful way to conceptualize the phenomenon of biological robustness comes from the idea of *neutral spaces* (Wagner 2005, 2014). The general concept of a neutral space is defined as follows:

> **(NS)** A neutral space is a collection of equivalent solutions to the same biological problem. It can also be thought of as a set of alternative configurations of a biological system, configurations that solve the same problem. (Wagner 2005, 195)

Neutral spaces are always relativized to a particular biological domain or class of biological functions. For example, a gene can be neutral for the substitution of another gene, but it generally does not make sense to speak of a gene being neutral to some completely other kind of macromolecule, say, a hemoglobin. In what follows, I focus on the concept of a neutral sequence space, although much of what I say would hold in the case of other kinds of spaces as well (e.g., like the space of neutral metabolic networks).[3]

As an example, I use a schematic presentation of neutral RNA/protein structure spaces presented originally in Wagner (2005, 201). In figure 1, there are three different types of sequences that fold into three different types of structures, respectively, with each structure in turn serving a specific biological function. Sequences folding into the most frequent structure form a single neutral space that is spread across the entire sequence space in a network-like manner (*white circles*). There is another moderately frequent structure that is also realized by several sequences (*gray circles*). However, these sequences do not form a single network but rather come in three separate island chains. Finally, there is a rare structure that is realized by only two isolated

3. Topologically, all the neutral spaces that Wagner (2005) considers are multidimensional metric spaces. Some are discrete, while some can be treated as continuous. In order to represent these spaces in a convenient way, one must necessarily strip away some of their complexities, like multidimensionality.

https://doi.org/10.1086/705478 Published online by Cambridge University Press

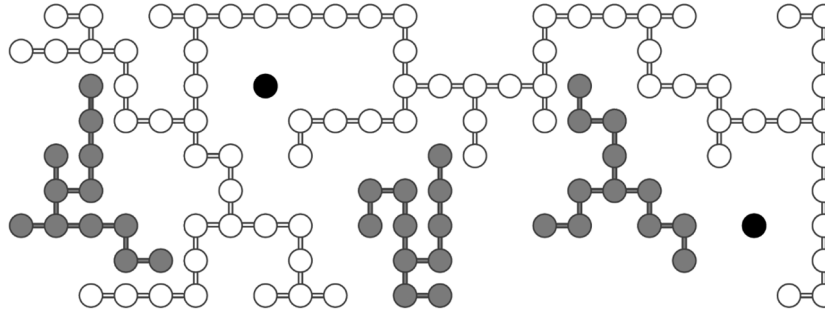


Figure 1. Schematic presentation of neutral RNA/protein structure spaces. Rectangular area represents sequence space, with individual circles corresponding to different sequences. Circles with the same shading correspond to sequences with the same RNA secondary structure/protein fold, thus realizing the same biological function. Unlike the image, the actual sequence space is multidimensional. Adapted from Wagner (2005, 201).

sequences (*black circles*). The sequences for the first structure form a single neutral space that is *connected*, while the neutral space for the second structure is only *partially connected*. The neutral space for the third structure is *not connected* at all.

Two sequences are connected only if there exists a stepwise path of function-preserving mutations between them. For example, the neutral space of the two black circles in figure 1 is not connected because no series of mutations would be able to transform one to the other without a loss of function in the process. Thus, although each function is realizable by multiple different sequences, not all functions are alike in how "easy" it is to find or lose a respective sequence for them. In general, the more neutral solutions there exist and the more connected these solutions are, the higher the robustness of that particular function.

It is natural to identify multiple realizability in biology with neutral solutions in Wagner's sense (see also Wimsatt 2012). It is because most biological functions seem to be realizable in multiple different function-equivalent ways that makes biological systems robust. This of course also depends on what we count as a genuine case of multiple realizability in the philosopher's sense.[4] However, the purpose of this article is not primarily to address the analytical criticism of MRT but rather to answer the empirical challenge regarding the modal dimension of multiple realizability. Similar remarks hold for the concept of a neutral solution as for multiple realizability: they are best

4. See Shapiro (2000, 2008) and Polger and Shapiro (2016) for what has become known as the "official recipe" for multiple realization.

https://doi.org/10.1086/705478 Published online by Cambridge University Press

interpreted as relative concepts. Like neutral biological solutions, multiple realizability is never absolute but comes in degrees and, most importantly, is always perspectival and anchored to a particular biological function (Wimsatt 2006, 459). As Wagner notes, "A subtle philosophical question is what constitutes different solutions to the same problem. A chemist might argue that two proteins differing in their amino acid sequence but cleaving a small molecule with the same reaction mechanism are similar solutions, whereas two proteins that use a different reaction mechanism are different solutions. From an evolutionary perspective, however, it is sensible to view all genotypes that serve the same function as different solutions to the same problem, because each of these phenotypes can, in principle, be discovered independently from other such genotypes" (2005, 240).

It is clear that in the case of most neutral spaces, with any two neighboring nodes, the structural difference between them is not necessarily great. For example, in the case of many sequence spaces, the difference is typically that of a point mutation or a single change of a resulting amino acid. This may seem rather modest, especially given how many mutable positions a sequence can have. However, taken collectively, an entire neutral space of function-equivalent sequences can form quite a heterogeneous set. Although there might not be a significant difference in structure between nearest neighbors, the same does not typically hold for nodes that are farther apart.[5] Also, it is often not known a priori what different sequences will realize the same function or how many such sequences exist.

However, for the purpose of this article, the most important aspect of the theory of neutral spaces comes from the tools it provides for reasoning about the differential modal nature of various function-equivalent solutions. I highlight this next by an example from the study of the topology of a neutral space of genetic alphabets. This simple example will show on a conceptual level how it is possible to reason about the modal dimension of multiple realizability without reverting to arguments based on mere conceivability.

4. A Neutral Space of Genetic Alphabets. Life on earth is based on the DNA and RNA molecules with their four letter, two base-pair system (G-C, A-T/U), but is this structure necessary, or could there be an alternative realizer for genetic functions? Wagner (2005, 196–97) presents an elegant case to show that, once a particular genetic alphabet, say A_1, is fixed, it might be impossible for it to change anymore. What is striking in his example is that it holds even if there in theory exists an alternative alphabet A_2 that is just as optimal as A_1.

5. There are even cases of some enzymes with the same fold but that do not have a single amino acid in common (Wagner 2014, 124).

https://doi.org/10.1086/705478 Published online by Cambridge University Press

The alphabet that Wagner (2005, 18–21) considers is the pair (κ-X, α-Γ). He chooses this particular alphabet because it has been shown to express a robust Watson-Crick base-pairing pattern, comparable to that of the natural (G-C, A-U) alphabet. Basically, this means that binding between its complementary bases, and repulsion between noncomplementary ones, is strong. It is also chemically stable. Because of this, it can be treated as an alternative solution to the functional problem of faithful genetic replication, even if no known organism is based on it: "[the (κ-X, α-Γ) alphabet] is not realized in life as we know it, perhaps because prebiotic conditions were not conductive to the synthesis of its bases, perhaps because of other chemical instabilities, or perhaps because of a mere historical accident" (196).

To transition from one genetic alphabet to the other, a set of intermediates is needed that makes the change chemically attainable. In the case of the alphabets considered here, the following chain of five different alphabets is formed:

$$\langle(\boldsymbol{\kappa}\textbf{-X}, \boldsymbol{\alpha}\textbf{-}\boldsymbol{\Gamma}), (\kappa\text{-X}, \alpha\text{-}\Gamma, \text{G-C}), (\kappa\text{-X}, \text{G-C}), (\kappa\text{-X}, \text{G-C}, \text{A-U}), (\textbf{G-C}, \textbf{A-U})\rangle$$

Because the alphabets at the endpoints share no bases, a total of four steps (an addition or a deletion of a base pair) is required to get from one to the other. However, the problem now arises that all the intermediates between (G-C, A-U) and (κ-X, α-Γ) are suboptimal; they fare worse than alphabets at either endpoint in realizing the function of faithful genetic replication. Because of this, natural selection would never favor a stepwise transition from either endpoint to the other. As Morange (2009, 370) notes, "No variation can be selected if it transiently decreases fitness, even if additional variation going in the same direction or complementary variation induces a later increase in fitness—natural selection always acts in the present."[6] What makes the example remarkable is that it essentially depends only on the topological structure of the neutral space in question (i.e., whether nodes belonging to it are connected or not). In other words, in the above chain of alphabets, only those indicated in bold belong to the neutral space of optimal genetic alphabets. However, as can be seen, this neutral space is not connected; the alphabets at the ends can be likened to the isolated black circles in figure 1. This means that once either alphabet is fixed, it cannot revert and transform itself anymore without a loss of function in the process (Wagner 2005, 197).

I take this example as having two take-home messages regarding multiple realizability. First, it provides evidence against Rosenberg's claim that natural selection can freely pick between alternative solutions to the same functional problem, given only the simple requirement that the nature of these

6. This is even more prominent in the case of such a highly generatively entrenched feature of living systems as the genetic alphabet (see Wimsatt 2007, 135–36).

https://doi.org/10.1086/705478 Published online by Cambridge University Press

solutions does not bear on the system's overall fitness. Although theoretical reasons suggest that both (G-C, A-U) and (κ-X, α-Γ) alphabets should have an equivalent capacity to realize the function of faithful genetic replication, natural selection does not have the power to wander freely between them. Ironically for Rosenberg, it is exactly natural selection's unyielding demand for fitness that makes the transition through suboptimal solutions biologically impossible. Thus, selection cannot provide a universal algorithm for attaining multiply realized solutions.

Second, it shows how coarse-grained the simple realized-realizable dichotomy of Polger and Shapiro is in evaluating the scope and limits of MRT, at least in the case of many biological functions. Even though the alternative alphabet considered here is not found in the genetic makeup of any known organism—that is, it is not realized in nature—there is a perfect sense in which it is nevertheless realizable. It is a biochemically feasible way to achieve the same function as that of the natural DNA/RNA alphabet. Furthermore, the potential to realize faithful genetic replication cannot surely be just a matter of pure a priori logic. Rather, it is a matter of distinctly biological possibility.

The skeptic can still of course question whether the kind of chemical alphabet considered here is ultimately adequate or even realizable by actual biophysical means. It could rely on, say, unrealistic modeling assumptions. The case that I have presented here certainly does not constitute a proof of MRT when it comes to biological systems, not to mention special science kinds in general. However, I argue that it gives us compelling reasons to regard the multiple realizability of many important biological functions as a live possibility worthy of scientific and philosophical investigation. In the next and final section, I conclude by briefly pointing to some future avenues in contemporary biological engineering that can be used to further refine our understanding of this modal character of MRT.

5. Conclusion: An Empirical Claim or an Engineering Hypothesis? In their new book-length treatment on multiple realization, Polger and Shapiro (2016, 53) ask their readers: "Lacking evidence of actual multiple *realization*, what kind of evidence for multiple *realizability* could we have?" And they answer, provocatively: "This is where Strawperson usually steps into the debate. Strawperson doesn't think that we need to consider the evidence for multiple realizability. Strawperson doesn't care whether we have observed any actual cases of multiple realization, because what matters is only the *modal* claim that multiple realization is possible" (53). Here they of course interpret possibility simply as noncontradictoriness; this is what makes the supposed modal argument so obviously lacking.

In this article I have argued for a new perspective that helps to refine our understanding of the modal dimension of multiple realizability without having to revert to the role of the Strawperson. More precisely, I have shown that

https://doi.org/10.1086/705478 Published online by Cambridge University Press

recent critics of MRT have been too restrictive on their insistence to limit our focus solely on actual realizations and to claim that whatever goes beyond this must be just a matter of logical speculation or science fiction. Now, it has been surely a welcome trend to bring the debate on multiple realizability in closer contact with the empirical content of actual science. When Hilary Putnam originally formulated MRT, he was clear that what he was proposing is in fact a bold empirical hypothesis (Polger and Shapiro 2016, 42). Naturally evolved systems may or may not turn out to be such that their functional features are multiply realized (Putnam 1975, 436). The best way to test an empirical hypothesis is surely by the means of empirical scientific investigation.

However, there is also a further sense in which Putnam and other early advocates of MRT were putting forth a hypothetical claim. This is based on the fact that many of the reasons that were used to support MRT came from the engineering sciences. Especially important were computer science, artificial intelligence, and robotics—booming areas of applied engineering in the 1960s. Although the main motivation for MRT came from conceptual considerations stemming from the theory of computation, there is surely an empirical component in here, too: computers, robots, and mindful machines are all concrete material artifacts that can be investigated empirically. However, what is noteworthy is that these are all things that are not readily found in nature. Rather, they are products of human engineers and require lots of nontrivial design work for their realization. Regarding thinking machines and the like, it is also the case that we arguably still do not have any proper examples of them. That is, at the moment, they are still unactualized engineering hypotheses. They may or may not turn out to multiply realize anything. However, according to Putnam at least, they are often designed so as to multiply realize some interesting phenomenon or function.

As critics of MRT have noted, it is certainly "unusual to adduce possible future evidence in support of a hypothesis" (Polger 2009, 458). However, in the biological context at least, synthetic biologists with a proof-of-concept mind-set (Kendig 2015) are now in a position to concretely study alternative biological realizations that have not been realized by natural selection. These can be likened to the unconnected solutions in Wagner's neutral spaces. Cutting-edge tools from biological engineering help researchers transcend mutational gaps that would be deleterious for natural systems (Morange 2009). For example, many labs around the world have already done experimental studies on artificial genetic alphabets. Thus far, the results have been promising, and multiple alternatives to the natural DNA molecule have already been probed (Kwok 2012; Karalkar and Benner 2018). It is of course still an oftentimes radically open question whether there ultimately exist alternative ways to implement any given biological function. However, valuable information is gained even in the case when multiple realizability is

https://doi.org/10.1086/705478 Published online by Cambridge University Press

not attainable. One should not expect MRT to be an either-or matter. As with all good hypotheses, what makes it even more interesting is that we cannot know its modal scope and empirical prevalence a priori.

REFERENCES

Bechtel, William. 2008. *Mental Mechanisms*. New York: Taylor & Francis.

Bechtel, William, and Jennifer Mundale. 1999. "Multiple Realizability Revisited: Linking Cognitive and Neural States." *Philosophy of Science* 66:175–207.

Chouard, Tanguy. 2008. "Beneath the Surface." *Nature* 456:300–303.

Edelman, Gerald M., and Joseph A. Gally. 2001. "Degeneracy and Complexity in Biological Systems." *Proceedings of the National Academy of Sciences* 98:13763–68.

Greenspan, Ralph J. 2001. "The Flexible Genome." *Nature Reviews Genetics* 2:383–87.

Hull, David. 1974. *Philosophy of Biological Science*. Englewood Cliffs, NJ: Prentice-Hall.

Karalkar, Nilesh B., and Steven A. Benner. 2018. "The Challenge of Synthetic Biology: Synthetic Darwinism and the Aperiodic Crystal Structure." *Current Opinion in Chemical Biology* 46:188–95.

Keeley, Brian L. 2000. "Shocking Lessons from Electric Fish: The Theory and Practice of Multiple Realization." *Philosophy of Science* 67:444–65.

Kendig, Catherine E. 2015. "What Is Proof of Concept Research and How Does It Generate Epistemic and Ethical Categories for Future Scientific Practice?" *Science and Engineering Ethics* 22:735–53.

Kincaid, Harold. 1990. "Molecular Biology and the Unity of Science." *Philosophy of Science* 57:575–93.

Kitano, Hiroaki. 2007. "Towards a Theory of Biological Robustness." *Molecular Systems Biology* 3:137

Kwok, Roberta. 2012. "DNA's New Alphabet." *Nature* 491:516–18.

Morange, Michel. 2009. "Synthetic Biology: A Bridge between Functional and Evolutionary Biology." *Biological Theory* 4:368–77.

Polger, Thomas W. 2009. "Evaluating the Evidence for Multiple Realization." *Synthese* 167:457–72.

Polger, Thomas W., and Lawrence A. Shapiro. 2016. *The Multiple Realization Book*. Oxford: Oxford University Press.

Putnam, Hilary. 1975. "The Nature of Mental States." In *Mind, Language and Reality*, ed. Hilary Putnam, 429–40. Cambridge: Cambridge University Press.

Raerinne, Jani, and Markus Eronen. 2012. "Multiple Realizability and Biological Laws." *History and Philosophy of the Life Sciences* 34:521–38.

Richardson, Robert C. 2008. "Autonomy and Multiple Realization." *Philosophy of Science* 75:526–36.

Rosenberg, Alexander. 1985. *The Structure of Biological Science*. Cambridge: Cambridge University Press.

———. 2001. "On Multiple Realization and the Special Sciences." *Journal of Philosophy* 98:365–73.

———. 2006. *Darwinian Reductionism*. Chicago: University of Chicago Press.

Shapiro, Lawrence A. 2000. "Multiple Realizations." *Journal of Philosophy* 97:635–54.

———. 2008. "How to Test for Multiple Realization." *Philosophy of Science* 75:514–25.

Wagner, Andreas. 2005. *Robustness and Evolvability in Living Systems*. Princeton, NJ: Princeton University Press.

———. 2014. *Arrival of the Fittest*. New York: Current.

Weiskopf, Daniel A. 2011. "The Functional Unity of Special Science Kinds." *British Journal for the Philosophy of Science* 62:233–58.

Wimsatt, William C. 2006. "Reductionism and Its Heuristics: Making Methodological Reductionism Honest." *Synthese* 151:445–75.

———. 2007. *Re-engineering Philosophy for Limited Beings*. Cambridge, MA: Harvard University Press.

———. 2012. "Evolution and the Stability of Functional Architectures." In *Functions: Selection and Mechanisms*, ed. Philippe Huneman, 19–42. Dordrecht: Springer.

https://doi.org/10.1086/705478 Published online by Cambridge University Press