European Journal for Philosophy of Science (2019) 9: 15
https://doi.org/10.1007/s13194-018-0243-3

PAPER IN THE PHILOSOPHY OF THE LIFE SCIENCES

Multiple Realizability as a design heuristic in biological engineering

Rami Koskinen [1]

Received: 3 April 2018 / Accepted: 28 November 2018 / Published online: 12 December 2018

Abstract

Recently, several critics of the multiple realizability thesis (MRT) have argued that philosophers have tended to accept the thesis on too weak grounds. On the one hand, the analytic challenge has problematized how philosophers have treated the multiple realization relation itself, claiming that assessment of the sameness of function and the relevant difference of realizers has been uncritical. On the other hand, it is argued that the purported evidence of the thesis is often left empirically unverified. This paper provides a novel strategy to answer these worries by introducing a role for multiple realizability in the context of biological engineering. In the field of synthetic biology, bioengineers redesign the evolutionary realizations of biological functions, even constructing artificial chemical surrogates in the laboratory. I show how in the rational design approach to biological engineering, multiple realizability can function as a design heuristic in which the sameness of function and difference of realizers can be controlled. Although practically motivated, this engineering approach has also a theoretical, exploratory component that can be used to study the empirical limitations of multiple realizability. Successful realization of the engineering designs would amount to a concrete demonstration of multiple realizability, taking evidence for MRT beyond what is readily found in nature.

Keywords Multiple realizability · Synthetic biology · Biological engineering · Engineering design · Engineering heuristic

This article belongs to the Topical Collection: EPSA17: Selected papers from the biannual conference in Exeter
Guest Editors: Thomas Reydon, David Teira, Adam Toon

✉ Rami Koskinen
rami.koskinen@helsinki.fi

[1] University of Helsinki, Unioninkatu 40, A522, 00014 Helsinki, Finland

1 Introduction

In its biological version, the *multiple realizability thesis* (MRT) states that biological systems, like genes, metabolic networks, or organs like eyes, can be realized in various physically different ways (e.g., Putnam 1967/1975; Griffiths and Stotz 2013, 58). Although often simply assumed in the philosophical literature, many take MRT to follow from the fact that biological systems are often individuated in functional terms – that is, through reference to their goals, capacities, causal roles, and the like (Rosenberg 2001, 2006, 30–34; Weiskopf 2011). It is argued that because there are no prima facie reasons to regard functions as being subordinate to particular designs, biological systems should be seen as realizable in multiple material ways (e.g. Rosenberg 2006, 30; see also Fodor 1974), with some philosophers even providing tentative empirical evidence for this (e.g. Hull 1974). Like with psychology and cognitive science in the philosophy of mind, MRT is traditionally taken to show that biology cannot be reduced to lower-level physical sciences. Although each biological token corresponds to some physical token, type-level reductionism becomes impossible under multiple realizability because of the disjunctive heterogeneity of underlying realizers.[1]

However, as popular as MRT has become especially among those with an antireductionist alignment, in recent years a number of critics have begun to seriously question its adequacy (e.g. Bechtel and Mundale 1999; Shapiro 2000, 2008; Couch 2005; Polger 2008, 2009; Raerinne and Eronen 2012; Polger and Shapiro 2016). Philosophers like Bechtel and Mundale (1999) and Polger and Shapiro (2016) have forcefully argued that the traditional proponents of MRT have tended to accept the thesis on too weak grounds. The criticism has come on essentially two fronts, which can be labelled respectively as the analytic challenge and the empirical challenge.[2] Assume that x and y are realizers and F is a function. The analytic challenge can be understood as the conjunction of two basic worries:

(1) Do x and y both really realize *F*? (The problem of the sameness of function.)
(2) Do x and y really *multiply* realize F? (The problem of the difference of realizers.)

That is, the analytic challenge deals with the meanings of the basic concepts of MRT, and it asks us to consider what *kinds* of situations count as legitimate cases of multiple realizability. For example, Polger and Shapiro (2016) have convincingly argued that, in purported cases of multiple realizability, philosophers often have not been careful enough to distinguish whether it is actually a profound causal difference in the underlying mechanisms that is taken to be realized differently or just some superficial characteristic of the system. In a similar fashion, many proponents of MRT have also been too eager to count two *similar* natural functions as instances of the exact *same*

[1] For reasons of space, and for the fact that this is not first and foremost a paper about reductionism, I am not going to go into the details of the traditional reductionist-antireductionist debate in the philosophy of biology here. For an excellent summary of the debate, including the central role that multiple realizability plays in it, the reader can consult Brigandt and Love (2017). It is perhaps worth mentioning that there are also advocates of multiple realizability who do not see it as essentially connected to reductionism (e.g. Wimsatt 2007, 276).
[2] My analysis follows that of Polger and Shapiro (2016), which is the most up-to-date and thorough treatment of MRT and its new wave of criticism.

function. According to Polger and Shapiro's analysis, neither situation would count as a legitimate case of multiple realizability in the strict sense.

The empirical challenge, on the other hand, can be formulated as follows: are various scientific target-systems (functions) in nature actually multiply realized, or is MRT just a philosopher's fiction? While the analytic challenge concerns itself with how concepts like "function" and "realization" are usually left vague by the proponents of MRT, the empirical challenge claims that evidence for the thesis is typically based on merely hypothetical cases of what is conceptually possible instead of firm empirical results (Bechtel and Mundale 1999, 176–177; Shapiro 2000, 636).

These challenges are intimately linked and answering one of them often requires also paying close attention to the other. Indeed, especially in practice-oriented philosophy of science, many researchers have begun to provide detailed case studies in order to see whether actual empirical findings and scientists' own accounts can provide support for multiple realization and offer a principled way to understand the nature of the multiple realizability relation (e.g. Bechtel and Mundale 1999; Keeley 2000; Richardson 2008, 2009; Weiskopf 2011; Aizawa 2013). However, it is fair to say that, taken collectively, these studies have been rather inconclusive. Moreover, due to a deliberate avoidance of unwarranted speculation, many have begun to see MRT solely as a claim about actual entities in the here and now and their properties (e.g., Richardson 2008; Shapiro 2008). The problem seems to be that, on a closer look, nature as such does not come with neatly discernible realizations (e.g. Bechtel and Mundale 1999; Shapiro 2008). Furthermore, as nicely summarized by Polger and Shapiro (2016), the few covincing cases of actual multiple realization do not seem to provide enough evidence for the whole thesis as traditionally conceived.

The purpose of this paper is to try to disentangle this situation by introducing a novel role for multiple realizability in the context of biological engineering. More precisely, I will argue that MRT can function there as a fruitful *design heuristic*. In the nascent field of bioengineering called synthetic biology, researchers often describe their work as the modelling and building of alternative ways to realize the functions of living systems. Given the pursuit to rationally streamline the naturally evolved realizations of biological functions, or even construct artificial alternatives in the laboratory, the multiple realizability of biological systems can be seen as a methodological tool for synthetic biologists with both practical, heuristic value as well as theoretical, exploratory value. By examining the rational design method used by synthetic biologists, I show how bioengineers can partially control the sameness of function and difference of realizers for the artificial systems they construct. These artificial systems, in turn, can shed new light on the complicated case of hypothetical evidence of MRT by rendering some of them as concrete material objects.

I begin Section 2 by providing a short introduction to synthetic biology, in order to show how the idea of multiple realizability as a design heuristic fits into its research practice. I then provide two illuminating cases from synthetic biology research, which will exemplify my points. In Section 3, I first analyse the design of minimal genomes where researchers believe that natural gene networks can be materially "rewired" for engineering purposes while maintaining their essential functions as intact. The purported multiple realizability of circuit functions serves the practical purpose of providing standardized modules that can realize the same function while exhibiting rationally minimalized designs. Then, as a second case in Section 4, I examine how the

construction of unnatural biochemical systems can have important effects for our understanding of the multiple realizability of genetic material. Recent studies on xeno nucleic acids (XNAs) and alternative genetic alphabets suggest that DNA and RNA might not be the only molecules capable of carrying genetic information. This kind of research also supports the hypothesis that it is possible to have materially different realization bases for the essential functions of living systems, thus further vindicating MRT, if successful. Section 5 concludes the article, summarizing the main points and drawing some of their implications for general philosophy of biology.

2 Multiple realizability as a design tool in synthetic biology

Synthetic biology is an emerging field of biological engineering whose purpose is to build artificial biological systems. Although the field is characterized by a focus to engineer human-valued products like drugs, biomaterials and other commercial substances, many synthetic biologists can also be seen using engineering methods to study what biology is fundamentally all about. Elowitz and Lim (2010, 889) explain: "Conventionally, biologists have sought to understand life as it exists. Increasingly, however, from stem-cell reprogramming to microbial factories, researchers are describing what is and exploring what could be." This means that a lot of work that is done under the label of synthetic biology can be interpreted as biological basic research (Knuuttila and Loettgers 2013a). The field can be partly seen as a continuation of the systems biology approach that emerged in the 1990s (Cameron et al. 2014; Green 2017). Systems biology tries to overcome some of the limitations associated with more traditional molecular biology methods by modelling emergent phenomena that arise in complex genetic and metabolic networks. By doing so, systems biologists believe that they can shed light on the basic organizing principles hypothesized to underlie organisms' gene expressions and metabolic functions (Powell and Dupré 2009; Green 2017).

Because of the heterogeneity and messiness of most biological systems, this task is notoriously difficult. Instead of embracing this complexity, synthetic biologists would rather try and get rid of it (O'Malley et al. 2008, 63). Biological systems are products of natural evolution. This means that they are well adapted to their environments and often have manifold fine-tuned ways of regulating their behaviour. The downside of this is that they are also extremely complex systems whose functions are often intertwined and messy. They function, but it is far from clear from the perspective of a cognitively limited researcher *how* exactly they function. Even with the most data-intensive methods of modern systems biology, researchers are still far from being able to fully understand the plethora of different biological design principles that govern the natural world.

One key idea of synthetic biology is that if researchers were able to design and build a certain kind of biological system, they would probably also understand its functioning better (Endy 2005, 449). For example, when studying complex gene regulatory networks, synthetic biologists have built their own versions of genetic circuits, which share many common features with their naturally occurring cousins but are nevertheless artificial objects with many constrained features (e.g. Elowitz and Leibler 2000; Gardner et al. 2000; for a philosophical discussion, see Knuuttila and Loettgers

2013b). Indeed, in the context of synthetic biology, people often speak about a particular "build to understand" heuristic to research (Elowitz and Lim 2010). This method of planned (re) construction is sometimes also dubbed *rational design*.[3]

According to a classic characterization by a pioneer of the field, Drew Endy, the approach of rational design is expressed in at least three distinct ways. First, the products of synthetic biology should be of a *standardized* nature. This means that they are easily replaceable and work as expected in a variety of predefined contexts. In principle, this would enable the mass production of useful synthetic products for industrial purposes. It also makes biological design accessible to a wider audience (Endy 2005, 450). Second, complex design process should be ideally *decoupled* into several simpler problems that can be solved independently (Endy 2005, 451). Third, the conceptualization of synthetic biological systems should allow for hierarchical *abstraction*. This means that biological engineers can exchange useful inter-level information without fully fleshed understanding of how everything works at the level below or above one's own expertise. For example, if a group working on a synthetic genetic circuit requires a certain kind of promoter gene, in principle they can simply describe its required function and leave the implementation work to a group who operates at a lower level of organization (Endy 2005, 451–452).[4]

Although synthetic biologists are still far from achieving all the aforementioned engineering ideals, what is interesting for the purposes of this paper is how the engineering approach depicted here often conceptually separates the function of a given system from its material basis or realization. This is especially prominent when one considers how an important part of the actual design work of synthetic biologists nowadays is computer-assisted, often with no direct wet-lab interaction whatsoever. In the philosophy of engineering, design is recognized as an essentially creative form of problem-solving. Engineers typically do not aim for just one all-purpose solution, but assesses alternatives based on pragmatic criteria, like efficiency, reliability, or even aesthetics (Kroes 2012). Biological engineering is no exception here (Koskinen 2017).

Here the idea of multiple realizability looms large. As in the case of software engineering (Calcott 2014), one ideal for synthetic genomics, for example, is to have a number of different modules capable of performing the same capacity or function. Because they do not affect the output of the higher-level system, they can be swapped within that context (Gibson 2014, 525). However, it is important to note that the sameness of function is not absolute but relative to particular interests and problems. For example, synthetic biologists often make use of the fact that two alternative modules are function-equivalent with respect to one function, but not to some other

[3] Another method commonly utilized by synthetic biologists, which is sometimes contrasted with that of rational design, is *directed evolution* (Haseltine and Arnold 2007). In directed evolution, synthetic genetic circuits are made more robust by allowing natural evolution to adjust them. This happens through an iterative process of scanning and selection, where researchers simply let nature do its job. However, the evolving sample is systematically biased at each step, due to selecting for further refinement only mutations with desired functions. For a methodological discussion of the comparative merits of rational design and directed evolution, see Ijäs (2018).

[4] Another partly overlapping characterizing perspective is provided by the concept of *modularity*. A module is something whose own function can be understood in relative isolation from other parts of the system (Cambray et al. 2011: 628; see also Hartwell et al. 1999). However, a module is not necessarily of a standardized nature. Also, whereas decoupling is more of a methodological idea, modularity comes also with some ontological baggage.

ones. This is especially true in the case of so-called *orthogonal* parts, which are capable of filling some familiar capacity but are built in such a way as to prevent them from having unwanted interactions with other components (Slusarczyk et al. 2012, 412). This can be achieved by implementing the function by using, for example, genetic components that are not recognizable elsewhere in the host organism. This is multiple realizability driven by the need of one very specific function. Especially in the case of orthogonal parts, the disparity of realizers is exceptionally clear: to ensure that all unwanted interactions are eliminated, synthetic biologists will have to resort to solutions that are as degenerate, or materially disparate, as possible. Similar cases for the the importance of functionally equivalent designs can be found throughout the field, characterizing in particular the approach favoured by biological redesigners (e.g. Morange 2009; Benner et al. 2011; Slusarczyk et al. 2012).[5]

Implicit ideas of multiple realizability are often evoked in the context of the analysis of natural systems. For example, it has been argued that when scientists consider several different "how-possibly" models for one and the same function, they are in a way assuming the possible multiple realizability of that function (Weiskopf 2011, 243; Koskinen 2017, 499). Of course, this does not say anything about the likelihood of the function actually *being* multiply realiz*ed* in the real world. However, it would border on contradiction to deny the *possible* multiple realiza*bility* of a function while at the same time seriously entertaining two different hypotheses about the nature of its material realizer. However, things get much more concrete in the engineering sciences when researchers are actively pursuing alternative ways to realize a given function. It is then that the multiple realizability of that function becomes an explicit design hypothesis. Although clearly fallible, it nevertheless marks a conscious pursuit to come up with concrete, material, human-engineered alternatives to naturally occurring realizers.

In the engineering of novel biological systems, this fixing of a shared function has a number of beneficial features. First, it imposes an effective constraint that limits the search space for novel designs (see Koskinen 2017). To be sure, there may in theory be a myriad of different ways to implement even a single function. But often in reality only a tiny subset of these abstract possibilities turn out to be practically feasible. In any case, the number of options is much smaller compared to a purely hypothetical, exploratory situation where no function whatsoever is kept fixed. Second, it comes with some important practical uses. As I will show in the following cases, there are times when multiply realizable solutions are desirable from the point of view of biodevice standardization and biosafety. It is thus a valuable design heuristic for synthetic biologists whose rationale goes beyond abstract philosophical theorizing.

From the philosopher's standpoint, one methodological advantage that synthetic biology brings to the table in the study of MRT is the predefined nature of the function being studied. That is, when synthetic biologists propose two different blueprints for

[5] There are also synthetic biologists who explicitly define their goal as that of engineering novel functions *not* found in nature. Indeed, this may be even the more popular way to introduce the field. However, it is good to notice that the difference here is not absolute but relativized to one's level of abstraction. Even if the ultimate function of some synthetic circuit might be new to nature, the same may not hold for all of the sub-functions that its design is based on. However, this paper will consciously focus on the branch of synthetic biology that at heart is more about the redesigning of biological systems. Thanks are owed to an anonymous referee for pressing on this point.

realizers of a given function, it is clear from the start that they are indeed alternatives for *one and the same thing* (compare with the worries of Shapiro and others about the sameness of function). Of course, it is not guaranteed that every design is actually realizable. However, with suitable design settings, researchers can ensure that what they are doing is in fact serving some predefined engineering goal.[6] In the context of naturally evolved systems, the problem is often that, even when there are two apparently different realizers that exhibit some similar capacities, it is not clear that they must serve the *exact same* function (see Shapiro 2000, 647; Couch 2005). This is especially pressing for someone like Rosenberg (2006), who has argued for MRT on the basis of the etiological function concept; it is well-known that the evolutionary function of a given trait is not always easily assessable from our current point of view. In synthetic design, the ultimate purpose of the system and its different sub-capacities can be expected to be more transparent. I should stress that I am not endorsing the absurd claim that synthetic biological systems will not have all kinds of side-effects and contextual peculiarities. However, some independent pragmatic criteria for the sameness of function is obtained by observing whether different material realizers were designed with the same purpose in mind. If two different systems were designed to perform the same Boolean function in the same genetic circuit, we can expect them to be function-equivalent with regard to *that* function (but probably not some other ones).

From a more exploratory perspective, synthetic biology can also study hypothetical realizations that are unrealized in nature. For example, by using tools from contemporary bioengineering to construct artificial systems which do not have to meet the harsh requirements that natural selection imposes on their fitness, synthetic biology can leapfrog some of the mutational gaps that would be deleterious for natural systems (see Morange 2009). If these designs prove successful, they can provide a concrete demonstration of alternative realizations, taking the evidence of MRT beyond naturally evolved systems. On the other hand, if a particular realization seems especially recalcitrant vis-à-vis re-engineering attempts, it may reveal tacit constraints that partly explain its prevalence in nature, providing also valuable information.

In the following two sections, I provide cases from synthetic biology research that exemplify my claims. I first analyse the probing of alternative genetic/metabolic circuits, where researchers believe that genomes of selected microorganisms can be downsized for engineering purposes while still leaving their essential functions intact. As a second case, I examine how the construction of unnatural biochemical systems can have important effects on our understanding of the multiple realizability of the genetic systems. Both of these approaches assume that it is possible to provide function-equivalent solutions to familiar biological problems. In doing so, they make bold empirical bets on the nature of the systems that they study, as well as the biological world in general. However, they also highlight the heuristic role that MRT can have in the context of biological engineering, regardless of whether all of the proposed endeavours turn out to be eventually successful.

[6] See Kitcher (1998, 270): "Entities have functions when they are designed to do something, and their function is what they are designed to do."

3 Engineering of minimal genetic systems

One of the central streams of synthetic biology research is focused on the idea of minimal life. This manifests in the way that all complexly functioning biological systems are thought of as being in principle reducible to a hypothetical minimal set of biologically essential factors – a kind of biological version of Occam's razor. This quest for minimalism is perhaps most clearly concretized in the study of minimal genomes by J. Craig Venter and others. Venter's team's model organism of choice, namely, *Mycoplasma genitalium*, is estimated to have one of the shortest genomes found on Earth (Glass et al. 2006). This little bacterium's genome consists of roughly 580,000 base pairs. For comparison, this is only about one-eighth of that of *E. coli*. With a genome this small, the existence of totally non-functional DNA is very improbable. One might then ask, what are Venter and his team about? Surely one cannot shrink something that is already about as diminutive as there is.

Wrong. By systematically silencing parts of the genome, Venter and his team managed to shut down some of the 482 protein-coding genes of *M. genitalium* while maintaining the organism's vital functions intact. In fact, only 382 of the genes were ultimately identified as "essential" (Glass et al. 2006). How is this possible? Part of the answer to the puzzle lies in the fact that, as a freely living bacterium, *M. genitalium* must exhibit at least some capacity to react to varying environmental conditions. This means that Venter and his team probably ended up disrupting some of the organism's functional responses that manifest only when certain very specific environmental conditions are met, but which do not have phenotypic effects in the highly constrained laboratory environment. Secondly, even with its nearly minimal set of genes, *M. genitalium* apparently has some level of redundancy in its metabolic networks. More precisely, the organism seems to be able to supply some of its vital enzymes in the face of gene deactivation by resorting to alternative metabolic pathways (Glass et al. 2006, 428).

Now, does the new minimal version of the mycoplasma count as a case of multiple realization? It is important to be careful here before drawing any far-reaching conclusions. Many philosophers of science would probably object that, although scientifically impressive, the above example was achieved a bit too cheaply for it to count as genuine evidence for MRT. For even though the total change was dramatic, there was actually not that much of a change in any of the underlying molecular mechanisms of the poor bacterium. Indeed, as Polger and Shapiro (2016) have convincingly argued, in a proper case of multiple realization it is not enough that there is simply *variation* in the underlying realizers of a particular function. What is important is that this variation be somehow *relevant* to the realization of that function; it must have some kind of causal connection to it (Shapiro 2000: 644; Polger and Shapiro 2016). Mere deletion of redundant but identical – or at least closely similar – components does not really seem to cut it.

Here I am inclined to agree with Polger and Shapiro's analysis. Often when providing supposed examples of multiple realization, many philosophers have not been careful enough to distinguish whether it is actually a profound difference in the underlying causal mechanisms that they take to be realized differently or just some superficial characteristic of the system. This does not mean that judging these matters is always easy. To the contrary, to know which factors are causally relevant in a given system under study is a difficult and highly context-sensitive matter, especially in the case of complex biological systems. Nevertheless, the above example of the minimal bacterial genome clearly has its limitations.

However, things get more interesting when one realizes that some of the changes provoked by Venter and his team were due not only to redundant pathways, but completely *degenerate* metabolic mechanisms. That is, they were not only functionally equivalent, but also structurally different underlying solutions (Edelman and Gally 2001). This is explained by the fact that the organism has relaxed substrate specificity in some metabolic situations – different types of enzymes can sometimes be used to catalyse the same reaction (Glass et al. 2006, 428). This suggests at least a modest case of multiple realizability for some of the functions in question, at least at the genetic/metabolic level. Whether these changes are enough to count the whole down-sized genome as a multiply realized instance of *Mycoplasma genitalium* is another question. However, by studying different minimalizing assumptions and metabolic wirings, synthetic biologists can push the evidence for multiply realizable solutions beyond naturally occurring systems. In a sense, they are testing their limits. If vital biological functions can be maintained even under such dramatic circumstances, as in some of Venter's experiments, the multiple realizability of those functions gains considerable plausibility. The only way to test these hypothetical configurations is through hands-on biological engineering.

4 Construction of unnatural biochemical systems

Another dimension of synthetic biology focuses on alternative chemical bases of living systems. Called *xenobiology*, it goes beyond cells, genes and metabolic pathways even. In a sense, this line of research is as close to chemistry as to biology. However, since this area of synthetic biology is concerned with the higher-order functions of living systems, its motivations are completely biological. As Steven Benner, one of the pioneers of the field explains, although some synthetic biologists seek to forward engineer completely new kinds of functions that are not found in nature, among chemists "synthetic biology" means the opposite, namely, "to use *unnatural* molecular parts to do things that are done by natural biology." According to Benner, researchers in this area "believe that if they can reproduce biological behavior *without* making an exact molecular replica of a natural living system, then they have demonstrated an understanding of the intimate connection between molecular structure and biological behavior. If taken to its limit, this synthesis would provide a chemical understanding of life." (Benner et al. 2011, 372; emphasis in original). As in the case of the rewiring of minimal genomes, a hypothetical kind of multiple realizability seems to be often evoked in the case of unnatural biochemical systems.

Although the work of xenobiologists is engaged in very fundamental theoretical questions in biology, it can also provide tools of surprising practical utility. For example, a pressing concern of contemporary biosafety is that efficient tools are needed to ensure that a synthetically created organism cannot have unwanted interactions with natural species (Schmidt 2010; Torres et al. 2016). One such tool is derived from the study of alternative molecular chirality, sometimes called *mirror life* (Church and Regis 2012, 25–29). Many important biomolecules come in two different varieties, which are defined by their molecular handedness. However, living systems usually feature only one type of chirality per any given class of molecule. This is because systems built on one chirality cannot usually accommodate molecules of the opposing chirality in a

functional way. Thus, although there are no important chemical differences between the chiral variants of a single molecule, aggregative chirality can make a huge difference at a higher level. An artificial system built on the same molecules as a natural system, but with the opposing chirality, should in theory have the same macro-level functions as the natural system. However, certain unwanted molecular interactions between these systems could essentially be prevented in this way.

Another way to engineer a molecular safeguard is to base the synthetic organism's genetic code on some other medium than natural DNA or RNA molecules. In this situation, no harmful interbreeding between synthetic and natural organisms would happen, since their genetic codes are orthogonal; they simply cannot recognize the information content that the other one is carrying (Schmidt 2010, 323–324). Orthogonally designed xeno-organisms would still be able to interact with natural species at purely ecological and mechanical levels. This is desirable if, for example, a synthetic organism is designed to extract some unwanted chemical from the environment. It is only at the level of the genetic code that a genetic or semantic firewall makes sure that no unwanted hereditary information is passed into nature (Torres et al. 2016, 399). However, at the same time, these systems are often designed so as to realize many of the same functions as natural systems. The point is that they can operate orthogonally in the relevant biological context, but otherwise being able to function normally – a genetic firewall is not a biological firewall (Schmidt 2010, 327). Because of these considerations, I suggest that there is a sort of social market demand within the bioengineering community for designing solutions that are multiply realizable: besides being theoretically interesting, they could provide useful biosafety applications in bioindustry and therapeutics, to name just a couple areas.

Although research on unnatural genetic systems is still in its infancy, some interesting results have already been reported. These can be divided into roughly three categories: those dealing with alternative genetic coding systems, those dealing with unnatural nucleobases, and those that try to change to backbone of the natural DNA molecule. All known organism store genetic information in a four-letter, two base-pair system. This code is read in triplets of letters called codons. Philosophers of science have already recognized that the natural genetic code provides a prima facie case of multiple realization (Richardson 2008: 531). This is because the mapping relation between the codons in the genome and the resulting amino acids is heavily redundant; there are 64 possible three-letter combinations of the basic genetic letters A, C, G, and T and only 20 amino acids that are used in standard Terran biochemistry. This means it is possible to have two completely syntactically different genomes that code for exactly the same functional phenotype. Furthermore, the particular natural assignment of codons with their corresponding amino acids does not seem to be in any interesting sense necessitated by the underlying chemical properties of the molecules, such as their stereochemical structure. This has lead researchers speculate that there could in principle be alternative ways to establish the genetic code. Recent work in synthetic biology has produced systems that alter the code in many ways, from systems that change the particular codon assignments to codes that are read in quadruplets instead of triplets (Chin 2017; Zhang et al. 2017).

However, the idea of simply altering the assignment relationships of the genetic code is still firmly anchored in the same basic set of chemical constituents that make up the natural system. To really push the evidence for the kind of unnatural biochemical realization of life that Benner and others have called for, researchers have put a lot of

effort into the design and construction of artificial genetic systems that feature also molecules that are alien to life as we know it. Although no full-fledged, self-sustaining xeno-organism exists yet, researchers have managed to construct many interesting variants of these kind of artificial genetic systems, be they genetic molecules with alternative backbone structures or systems of unnatural genetic alphabets (see, e.g., Malyshev et al. 2014; Marlière et al. 2011; Thyer and Ellefson 2014; Anosova et al. 2016; Benner et al. 2011). For example, a class of so-called XNA molecules has been probed where the natural deoxyribose backbone of the DNA molecule has been replaced with some alternative molecule, as in the case of threose nucleic acid TNA, for example (Anosova et al. 2016). Examples of alternative nucleobases include 5-chlorouracil inserted by Marlière et al. (2011) into a laboratory strain of *E. coli*, and the pair d5SICSTP-dNaMTP by Malyshev et al. (2014). Recently, Benner and his team were able to come up with the artificial alphabet (A, T, C, G, P, Z) (Benner et al. 2011). Benner's alphabet is an example of an expanded artificial genetic alphabet. With a mix of natural and unnatural bases, it is possible to use it to study both disparate and higher-cardinality genetic alphabets at the same time.

It is true that to some extent the molecular pairing mechanism of P and Z resemble those of the natural A-T and C-G pairs. This would at first seem to go against Shapiro's (2000, 643) requirement of a causally relevant difference when testing for multiple realizability. However, because of the mixture of unnatural and natural bases the whole causal structure of the system tends to change a lot – small differences tend to lead to catastrophic consequences. This makes synthesis as a way of inquiry such a difficult task (Benner et al. 2011). Furthermore, there is also an epistemic point to be made regarding whether to count a difference between two distinct realizers of a biological function as interesting enough. It is suggested that the lower one goes in the organization of biological systems, the fewer live options there are to realize their functions (cf. Wagner 2014, 215). For example, according to Rosenberg (1985, 169), there must be some level of fundamental biological organization "where the disjunction of vast and heterogeneous set of alternatives gives way […] leaving only one way the cat is skinned". As such, alternative genetic systems are in no way trivial achievements. Indeed, according to previous scientific theories, the informational specificity of the genetic code was thought to be encoded in the intrinsic chemical properties of the four natural nucleobases (Benner et al. 2011). In particular, there was no input from the lower-level physical/chemical sciences to inform otherwise. However, using exploratory engineering methods guided by the requirements of a higher-level function, synthetic chemists have now begun to produce results that give reason to believe otherwise. I argue that this provides considerable empirical support for MRT as a pragmatic hypothesis. At the very least, it shows the heuristic importance that assumptions about multiple realizability can have in a real scientific context.

5 Conclusions

In the nascent bioengineering field of synthetic biology, researchers are consciously pursuing alternative ways to realize familiar biological functions. This is especially true among those synthetic biologists who start from the idea of redesigning naturally occurring systems. Indeed, some take this even further by

explicitly announcing that their main aim is to *better* some of Mother Nature's designs (see, e.g., Endy 2005; Kwok 2012). In any case, when designing a different version of the same system – be it simpler, more rational, or simply a structurally alternative configuration – there must be something that is kept invariant in the process; otherwise it would just be a case of two totally different systems. In the cases of biological engineering considered in this paper, this something is what philosophers of science would most often call the function of the system (the problem of the sameness of function). This lends itself very naturally to the idea that some of these synthetic systems can be taken to multiply realize biological functions. Although often practically motivated, this approach can provide a novel and fruitful way to study the empirical limitations of multiple realizability. Successful realization of the designs would amount to a concrete demonstration of multiple realizability, providing remedy for issues of purely hypothetical evidence.

I provided two cases from synthetic biology research that further exemplified this heuristic role: the rewiring of minimal genomes and the construction of unnatural biochemical systems. If my arguments are correct, at least in some cases it is indeed possible to provide alternative realizers for particular biological systems by using synthetic design methods. However, this conclusion must be taken with some caution. The field of synthetic biology is still in its early stages and no definite conclusions can be drawn by any responsible researcher studying the field. Nonetheless, what is important is the role that multiple realizability plays in the shaping of the design heuristics in synthetic biology research. Practical matters like rationality assumptions and orthogonality motivate researchers to try and come up with functionally equivalent, but structurally disparate designs (the problem of the difference of realizers). Although these might not be attainable in every actual case, multiple realizability can still be seen as a guiding ideal behind specific theoretical and concrete design choices. Its heuristic utility is at least partly independent from whatever nature as such is taken to exemplify, be it none or few or many multiply realized designs.

Traditionally, the multiple realizability of biological systems has been used as an argument against scientific reductionism. However, because the rational design approach to synthetic biology often employs reductionist heuristics in its design of biological systems, it is questionable how the potential multiple realizability of engineered biological systems serves the supposed autonomy of the life sciences. One possible answer is to see the situation as providing further support for the view that traditional views concerning scientific reductionism are in fact outdated (Bechtel and Richardson 1993; Wimsatt 2006a, b, 2007). For example, it can be argued that the multiple realizability of a given function is not in itself a stumbling block for any particular research strategy, but rather an incentive to utilize different, but mutually complementary, heuristics (see Keeley 2000, 446).

More important than the question concerning reductionism are the new tools that can be used to study the scope and nature of multiple realizability in biological systems. In particular, they show that even if MRT does not hold in the case of all naturally evolved systems, it can nevertheless be fruitful in some cases to search for hypothetical realizers

by means of rational design methods. Aside from being of pragmatic and theoretical biological importance, multiple realizability as a an engineering heuristic can shed light on age-old philosophical problems concerning the nature of unactualized possibilities.[7] By exploring the space of possible designs that go beyond the limits of naturally evolved systems, it might even open doors for completely new kinds of questions by asking, for example, whether life itself could be multiply realizable.

Funding Information Open access funding provided by University of Helsinki including Helsinki University Central Hospital.

References

Aizawa, K. (2013). Multiple realization by compensatory differences. *European Journal for Philosophy of Science, 3*, 69–86.

Anosova, I., Kowal, E. A., Dunn, M. R., Chaput, J. C., van Horn, W. D., & Egli, M. (2016). The structural diversity of artificial genetic polymers. *Nucleic Acids Research, 44*(3), 1007–1021.

Bechtel, W., & Mundale, J. (1999). Multiple realizability revisited: Linking cognitive and neural states. *Philosophy of Science, 66*, 175–207.

Bechtel, W., & Richardson, R. C. (1993). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Princeton: Princeton University Press 2nd ed. 2010. Cambridge, MA: The MIT Press.

Benner, S. A., Yang, Z., & Chen, F. (2011). Synthetic biology, tinkering biology, and artificial biology. What are we learning? *Comptes Rendus Chimie, 14*, 372–387.

Brigandt, I., & A. Love. (2017). Reductionism in biology. In E. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*, URL https://plato.stanford.edu/entries/reduction-biology/

Calcott, B. (2014). Engineering and evolvability. *Biology and Philosophy, 29*, 293–313.

Cambray, G., Mutalik, V. K., & Arkin, A. P. (2011). Toward rational design of bacterial genomes. *Current Opinions in Microbiology, 14*, 624–630.

Cameron, D. E., Bashor, C. J., & Collins, J. J. (2014). A brief history of synthetic biology. *Nature Reviews Microbiology, 12*, 381–390.

Chin, J. W. (2017). Expanding and reprogramming the genetic code. *Nature, 550*, 53–60.

Church, G., & Regis, E. (2012). *Regenesis: How synthetic biology will reinvent nature and ourselves*. New York: Basic Books.

Couch, M. B. (2005). Functional properties and convergence in biology. *Philosophy of Science, 72*, 1041–1051.

Edelman, G. M., & Gally, J. A. (2001). Degeneracy and complexity in biological systems. *Proceedings of the National Academy of Sciences, 98*, 13763–13768.

Elowitz, M. B., & Leibler, S. (2000). A synthetic oscillatory network of transcriptional regulators. *Nature, 403*, 335–338.

Elowitz, M. B., & Lim, W. A. (2010). Build life to understand it. *Nature, 468*, 889–890.

Endy, D. (2005). Foundation for engineering biology. *Nature, 438*, 449–453.

[7] As one of the reviewers pointed out, a fully realized hypothetical system of course becomes part of actual reality. However, what matters here is the epistemic point that before its concretization, a hypothetical realization is just that: a hypothesis. Whether it turns out to be biologically feasible or not is something which cannot be decided a priori.

Fodor, J. (1974). Special sciences (or, the disunity of science as a working hypothesis). *Synthese, 28*, 97–115.
Gardner, T. S., Cantor, C. R., & Collins, J. J. (2000). Construction of a genetic toggle switch in Escherichia coli. *Nature, 403*, 339–342.
Gibson, D. G. (2014). Programming biological operating systems: Genome design, assembly and activation. *Nature Methods, 11*, 521–526.
Glass, J. I., Assad-Garcia, N., Alperovich, N., Yooseph, S., Lewis, M. R., Maruf, M., Hutchison, C. A., III, Smith, H. O., & Venter, J. C. (2006). Essential genes of a minimal bacterium. *Proceedings of the National Academy of Sciences, 103*, 425–430.
Green, S. (2017). Philosophy of systems and synthetic biology. *The Stanford Encyclopedia of Philosophy* (summer 2018 edition), E. N. Zalta (ed.), URL https://plato.stanford.edu/archives/sum2018/entries/systems-synthetic-biology/.
Griffiths, P., & Stotz., K. (2013). *Genetics and philosophy*. Cambridge: Cambridge University Press.
Haseltine, F. A., & Arnold, F. H. (2007). Synthetic gene circuits: Design with directed evolution. *Annual Review of Biophysics and Biomolecular Structure, 36*, 1–19.
Hull, D. (1974). *Philosophy of biological science*. Englewood Cliffs: Prentice-Hall, Inc..
Ijäs, T. (2018). Design methodologies and the limits of the engineering-dominated conception of synthetic biology. *Acta Biotheoretica*. https://doi.org/10.1007/s10441-018-9338-7.
Keeley, B. L. (2000). Shocking lessons from electric fish: The theory and practice of multiple realization. *Philosophy of Science, 67*, 444–465.
Kitcher, P. (1998). Function and design. In D. L. Hull & M. Ruse (Eds.), *The philosophy of biology* (pp. 258–279). Oxford: Oxford University Press.
Knuuttila, T., & Loettgers, A. (2013a). Basic science through engineering? Synthetic modeling and the idea of biology-inspired engineering. *Studies in History and Philosophy of Biological and Biomedical Sciences, 44*, 158–169.
Knuuttila, T., & Loettgers, A. (2013b). Synthetic modeling and mechanistic account: Material recombination and beyond. *Philosophy of Science, 80*, 874–885.
Koskinen, R. (2017). Synthetic biology and the search for alternative genetic systems: Taking how-possibly models seriously. *European Journal for Philosophy of Science, 7*(3), 493–506.
Kroes, P. (2012). *Technical artefacts: Creations of mind and matter*. Heidelberg: Springer.
Kwok, R. (2012). DNA's new alphabet. *Nature, 491*, 516–518.
Malyshev, A., Dhami, K., Lavergne, T., Chen, T., Dai, N., Foster, J. M., Corrêa, I. R., Jr., & Romesberg, F. L. (2014). A semi-synthetic organism with an expanded genetic alphabet. *Nature, 509*, 385–388.
Marlière, P., Patrouix, J., Döring, V., Herdewijn, P., Tricot, S., Cruveiller, S., Bouzon, M., & Mutzel, R. (2011). Chemical evolution of a bacterial genome. *Angewandte Chemie International Edition, 50*, 7109–7114.
Morange, M. (2009). Synthetic biology: A bridge between functional and evolutionary biology. *Biological Theory, 4*, 368–377.
O'Malley, M. A., Powell, A., Davies, J. F., & Calvert, J. (2008). Knowledge-making distinctions in synthetic biology. *BioEssays, 30*, 57–65.
Polger, T. W. (2008). Two confusions concerning multiple realization. *Philosophy of Science, 75*, 537–547.
Polger, T. W. (2009). Evaluating the evidence for multiple realization. *Synthese, 167*, 457–472.
Polger, T. W., & Shapiro, L. A. (2016). *The multiple realization book*. Oxford: Oxford University Press.
Powell, A., & Dupré, J. (2009). From molecules to systems: The importance of looking both ways. *Studies in History and Philosophy of Biological and Biomedical Sciences, 40*, 54–64.
Putnam, H. (1967/1975). The nature of mental states. In Mind, language and reality: Philosophical papers*, vol.* 2, ed. Hilary Putnam, 429–440. Cambridge: Cambridge University Press.
Raerinne, J., & Eronen, M. (2012). Multiple realizability and biological laws. *History and Philosophy of the Life Sciences, 34*, 521–538.
Richardson, R. C. (2008). Autonomy and multiple realization. *Philosophy of Science, 75*, 526–536.
Richardson, R. C. (2009). Multiple realization and methodological pluralism. *Synthese, 167*, 473–492.
Rosenberg, A. (1985). *The structure of biological science*. Cambridge: Cambridge University Press.
Rosenberg, A. (2001). On multiple realization and the special sciences. *The Journal of Philosophy, 98*, 365–373.
Rosenberg, A. (2006). *Darwinian reductionism*. Chicago: The University of Chicago Press.
Schmidt, M. (2010). Xenobiology: A new form of life as the ultimate biosafety tool. *BioEssays, 32*, 322–331.
Shapiro, L. A. (2000). Multiple realizations. *The Journal of Philosophy, 97*, 635–654.
Shapiro, L. A. (2008). How to test for multiple realization. *Philosophy of Science, 75*, 514–525.
Slusarczyk, A. L., Lin, A., & Weiss, R. (2012). Foundations for the design and implementation of synthetic genetic circuits. *Nature Reviews Genetics, 13*, 406–420.
Thyer, R., & Ellefson, J. (2014). New letters for life's alphabet. *Nature, 509*, 291–292.

Torres, L., Krüger, A., Csibra, E., Gianni, E., & Pinheiro, V. B. (2016). Synthetic biology approaches to biological containment: Pre-emptively tackling potential risks. *Essays in Biochemistry, 60*, 393–410.

Wagner, A. (2014). *Arrival of the fittest: Solving Evolution's greatest puzzle*. New York: Current.

Weiskopf, D. A. (2011). The functional unity of special science kinds. *The British Journal for the Philosophy of Science, 62*, 233–258.

Wimsatt, W. C. (2006a). Aggregate, composed, and evolved systems: Reductionistic heuristics as means to more holistic theories. *Biology and Philosophy, 21*, 667–702.

Wimsatt, W. C. (2006b). Reductionism and its heuristics: Making methodological reductionism honest. *Synthese, 151*, 445–475.

Wimsatt, W. C. (2007). *Re-engineering philosophy for limited beings: Piecewise approximations to reality*. Cambridge, MA: Harvard University Press.

Zhang, Y., Ptacin, J. L., Fischer, E. C., Aerni, H. R., Caffaro, C. E., San Jose, K., Feldman, A. W., Turner, C. R., & Romesberg, F. E. (2017). A semi-synthetic organism that stores and retrieves increased genetic information. *Nature, 551*, 644–647.