Synthese
https://doi.org/10.1007/s11229-020-02816-8

UNREALISTIC MODELS

Mechanist idealisation in systems biology

Dingmar van Eck[1,2] · Cory Wright[3]

Received: 28 February 2019 / Accepted: 28 July 2020

Abstract
This paper adds to the philosophical literature on mechanistic explanation by elaborating two related explanatory functions of idealisation in mechanistic models. The first function involves explaining the presence of structural/organizational features of mechanisms by reference to their role as difference-makers for performance requirements. The second involves tracking counterfactual dependency relations between features of mechanisms and features of mechanistic explanandum phenomena. To make these functions salient, we relate our discussion to an exemplar from systems biological research on the mechanism for countering heat shock—the heat shock response (HSR) system—in *Escherichia coli* (*E. coli*) bacteria. This research also reinforces a more general lesson: ontic constraint accounts in the literature on mechanistic explanation provide insufficiently informative normative appraisals of mechanistic models. We close by outlining an alternative view on the explanatory norms governing mechanistic representation.

Keywords Idealisation · Mechanistic model · Heat shock response · Systems biology · Mechanistic explanation

✉ Dingmar van Eck
Dingmar.vanEck@Ugent.be; d.vaneck@uva.nl

Cory Wright
cory.wright@csulb.edu

[1] Department of Philosophy and Moral Sciences, Centre for Logic and Philosophy of Science, Ghent University, Blandijnberg 2, 9000 Ghent, Belgium

[2] Department of Philosophy, Institute for Logic, Language and Computation, University of Amsterdam, Oude Turfmarkt 143, 1012 GC Amsterdam, The Netherlands

[3] Department of Philosophy, California State University Long Beach, Long Beach, USA

Published online: 17 September 2020

1 Explanatory idealisation

1.1 Introduction

Idealisations abound in scientific practice, and there are many accounts of them available. One such account, which is predominant in the literature on scientific explanation, may be called the 'misrepresent irrelevancies' account (MIA) of idealisation (see Mäki 1992, 2009; Nowak 1980). According to MIA, the main function of idealizations is to highlight explanatorily irrelevancies. While MIA is one among many accounts of idealisations, there is a paucity of specifically mechanistic analyses of idealisation in scientific modelling and explanation. Among the few on offer is a recent one from Glennan (2017), who harnesses MIA because of its putative use in constructing general mechanistic explanations. Glennan develops this application as follows: (1) features that differ from case to case or from token system to system are misrepresented, such that (2) these idealisations make salient that those (misrepresented) features have no significant explanatory import in the general explanations that are sought; (3) deploying idealisations is then a means to 'help us find generality in a world of mechanisms that are ultimately particular, localized, and heterogeneous' (Glennan 2017: p. 83).

While the application of MIA to mechanistic contexts certainly finds support, it cannot be the final analysis; for there are numerous other ways in which idealised mechanistic models are used to promote explanatory ends. In this paper, we elaborate two other ways in which idealised mechanistic models serve explanatory functions. The first function involves explaining the presence of structural/organizational features of mechanisms by reference to their role as difference-makers for performance requirements. The second involves tracking counterfactual dependency relations between features of mechanisms and features of mechanistic explanandum phenomena.

To elaborate these two explanatory functions of idealized mechanistic models, we distinguish between two types of explanation that are frequently constructed (often in tandem) in biological practice: mechanistic explanations and design explanations. Whereas mechanistic explanations articulate the mechanisms that constitute, produce, or are otherwise responsible for explanandum phenomena, design explanations articulate why mechanisms have the structure or organization that they have (Wouters 2007, 2013; Braillard 2015; Green 2015; van Eck and Mennes 2018). To make these functions more salient, we discuss systems biological research on the mechanism(s) for countering heat shock—the Heat Shock Response (HSR) system—in *Escherichia coli* (*E. coli*) bacteria. HSR systems are endogenous cyto-protective mechanisms found in virtually all organisms and crucial to their survival (Wong 2005), and the specific HSR system in *E. coli* bacteria is one of the most extensively studied HSR systems in biology (Richter et al. 2010). In the case studies we discuss, comparisons between more realistic and less realistic mechanistic models are key to articulating design explanations.

1.2 Preliminaries

The term *idealisation* is often used to refer to deliberate distortions or misrepresentations of facts done in the service of explaining or understanding some target phenomenon.[1] Mäki (2009) calls them 'strategic falsehoods'. In modelling contexts, the introduction of a strategic falsehood makes part of the model a misrepresentation of the part of the target system it represents. While there is a live question about how to comprehend the very conceptualisation of models as truth-bearers, we will not quibble over the semantics of the terms *false representation* and *misrepresentation*, so as to use these terms interchangeably.

A model that represents features that its target system does not have, or that does not represent features that its target system does have, may be called an *unrealistic* model (see also Mäki 2009). All scientific models are unrealistic in some way—such is the nature of the endeavour. Generally, this is because of trade-offs between representational realism and practicality: scientific models are designed to represent only selected features of the target, about which scientists want to learn, and their lack of completeness is part of what makes them useful. By contrast, a model that perfectly represented all of the features of its target system while always neglecting to represent all features not had by its target system, would be no more useful to modellers than the target system itself. Instead, the features of a system that scientists think matter to their interests should be relevantly similar to features of the model they construct, since that's what makes it possible to get accurate information about a target from studying a model; and the models they construct should be similar to their targets in relevant ways and dissimilar in irrelevant ways. But the perfection of similitude as a constraint on realistic modelling is an ideal the utility of which does not reside in forcing the actual practice of developing model-based scientific explanations to comport with it.

What is the relationship between unrealistic and idealized models? Not every model is idealized. Sometimes the distortions are not deliberate; not every falsehood is strategic. There are models whose badness owes to their unwitting misrepresentations of the facts, where those misrepresentations are not done in the service of explaining or understanding a given phenomenon. Idealized models, it might be said, have a distinctively pragmatic dimension; by contrast, the dimension of analysis of unrealistic models is alethic.

To say that all scientific models are unrealistic is to acknowledge the role of practice in constructing them. Yet, models can be more or less unrealistic. One idealisation that we will permit ourselves in this paper will be to use the terms *unrealistic* and *idealized* model roughly interchangeably, so as to pair the terms *more realistic* and *less idealized* and the terms *less realistic* and *more idealized*, and to do so in ways that ignore the psychological issues of deliberation that come with analyses of distortion. So, sometimes we speak about more versus less realistic and sometimes about less versus more idealized models. 'More versus less'-talk concerns a comparative relation

[1] Contrast this with the term *abstraction*, which is often used to refer to the omission of details without such intentional misrepresentation (Jones 2005). If omissions result in misrepresentations of a target system, and if these misrepresentations are intended, argues Godfrey-Smith (2006), they also counts as an idealisation. We will follow this usage as well as the qualification, which looms large in the case studies we discuss. More generally, such idealisations by omission are frequent in scientific practice (Potochnik 2017).

between models. A model is more realistic (or less idealized) compared to another model. Let model m_1 be more realistic than model m_2 if a strategic falsehood in m_2 is de-idealised in m_1, i.e., if the part that is misrepresented in m_2 is not misrepresented in m_1 (and m_1 does not contain a strategic falsehood that is not misrepresented in m_2). Comparisons of un/realisticness may be more complex than this, but this comparative scheme suffices for the cases we investigate.

In their research on heat shock, systems biologists use differential comparisons between unrealistic mechanistic models and more realistic ones to explain why mechanisms of a certain kind k have feature F. More specifically, simulations run with those models are used to explain why these mechanisms have the structure/organization that they have. Comparing the results of such simulations helps establish that those mechanisms have F because it is a difference-making factor for specific performance requirements, such as robustness and speed of response. Additionally, these comparisons are instrumental for helping to answer 'what-if-things-had-been-different' questions about mechanisms (Woodward 2003). Idealisations in mechanistic models are thus shown to serve a second explanatory function: they enable scientists to track counterfactual dependency relations between features of mechanisms and features of mechanistic explanandum phenomena.

The structure of the paper follows straightforwardly from these preliminaries. We start by presenting systems biological research on the mechanism(s) for countering heat shock—the Heat Shock Response (HSR) system—in *Escherichia coli* (*E. coli*) bacteria. After discussing some 'landmark' studies (Arkin and Schaffer 2011) on the *E. coli* HSR system and the usage of idealized mechanistic models therein (El-Samad et al. 2005; Kurata et al. 2006), we situate our analysis in extant philosophical work on mechanistic idealisation, showing how it extends and enriches this important (but underdeveloped) line of research. Additionally, our analysis of idealisation in mechanistic and design explanations in the HSR case in systems biology provides for a more general philosophical lesson about explanatory norms. According to 'ontic constraint' accounts, accuracy is the overriding normative constraint on mechanistic model-based explanations. In somewhat different ways, proponents of ontic constraint accounts (Craver 2007, 2014; Kaplan and Craver 2011; Illari 2013) defend the claim that the more accurate a model represents a mechanism, the better it explains. While accuracy is important, we argue that ontic constraint accounts of mechanistic explanation are not gripping, because they provide insufficiently informative normative appraisals of mechanistic models. In the cases we analyse, idealised mechanistic models are vital for identifying difference-makers and tracking counterfactual dependency relations. Insisting on accuracy as an inviolable constraint on each and every representation of difference-makers and difference-making relations in a mechanistic model gives modellers less traction in addressing their explanatory concerns. Ontic constraint accounts of mechanistic explanation, which comprise the most popular account of mechanistic explanation, are thus hard to uphold in the light of idealized mechanistic models. Finally, we offer conclusions.

2 Mechanist idealisation in systems biology

2.1 Mechanistic explanation and design explanation

Mechanisms are organized hierarchical systems of finitely many component parts, each of whose operations are coordinated in such a way as to compose a single unified repetitive higher-level activity. Mechanistic explanations are invoked to identify their target explanandum phenomenon with these activities, and to show how these systems' activities are constituted so as to produce the explanandum phenomenon. Whereas mechanistic explanations explain how systems work the way they do, design explanations explain why systems have the structure and organization they have. They are procured to account for why these systems (organisms, etc.) have certain traits, e.g., specific organizational features, rather than alternative ones, or why different systems or organisms have the same features or traits.

These are complementary explanatory projects. As Wouters remarks, '[e]xplanations of both kinds are needed to understand systems whose existence depends on an organized ability' (2013: p. 462; see also Green 2015). In biology, both explanatory projects are regularly pursued in parallel. Research on the HSR system in *E. coli* is one biological case among many illustrating the complementary nature and construction of mechanistic and design explanations in the analysis of complex biological systems (see also Wouters 2013; van Eck and Mennes 2018). For example, consider research on the presence of emperor penguins' two-voice system:

> The how-questions [about how the explanandum phenomenon is produced] are answered by [mechanistic explanations] describing features of the system that produce the relevant abilities, activities or characteristic (the emperor penguin's ability to recognize their partners is brought about by the interaction of two voices that produces a temporal pattern of beats characteristic of the individual). The why-questions [about why this feature is present in this system] are answered by [design explanations] describing features of the organism, its environment, or its way of life due to which the characteristics to be explained [the two-voice system of emperor penguins] are advantageous to the organisms that have it. (Wouters 2013: p. 463)

Oftentimes, such design explanations hinge on counterfactual comparisons between extant systems or organisms with a specific trait, e.g., a specific organisational feature of a mechanism, and hypothetical alternatives lacking the trait to be explained, to highlight the advantageous character of having the trait in question and thus to explain its presence. For instance, biologists explained the presence of the penguins' two-voice system in terms of a counterfactual (and theoretical) comparison with hypothetical emperor penguins having a one-voice system. With a one-voice system, mates and young cannot be detected efficiently in the harsh Antarctic conditions in which emperor penguins live; having a two-voice system gives the means to do so efficiently (Aubin et al. 2000). Such counterfactual comparisons also loom large in research on the HSR system in *E. coli* bacteria (discussed next). But rather than mere theoretical comparison, as in the emperor penguin case, design explanations for the presence of specific organisational features of *E. coli's* HSR system there hinge on counterfactual

comparisons between more and less realistic mechanistic models. (Prior knowledge of mechanisms in this case is thus a key driver for the comparisons.) This research on the HSR system in *E. coli* bacteria also shows that less and more idealized models of mechanisms are used to procure both mechanistic explanations and design explanations. Mechanist idealisation—the construction and use of idealized models of mechanisms—in this case is thus more encompassing than the standard picture of New Mechanism that focusses on just idealized models of mechanisms in relation to mechanistic explanation.

2.2 The HSR system in *E. coli*

The term *heat shock response* (HSR) names an endogenous cyto-protective mechanism that is found in virtually all organisms, which is crucial to their survival. A central activity of the HSR mechanism is the rapid expression of a class of proteins—heat shock proteins—following exposure of cells, tissues, and organisms to elevated temperatures.[2] Thermal stress can damage proteins by breaking down their (tertiary) structures. Heat shock proteins act either as molecular chaperones by refolding (repairing) damaged intracellular proteins or as proteases by degrading denatured proteins. HSR thereby protects against irreversible injury or cell death and life-threatening effects that may follow from it, such as elevated body temperature due to failed thermoregulation (hyperthermia) and inflammatory reaction causing injury to organs and tissues (sepsis) (Wong 2005). In sum, HSR serves a crucial cyto-protective role in virtually all organisms.

The HSR system in *E. coli* is among the most extensively studied HSR systems (Kaufmann 1990). *Escherichia coli* is a bacterium found in the intestine of warm-blooded animals, playing roles in the digestion of food and the production of vitamin K—the latter being required for blood clotting.

In *E. coli*, HSR gene expression is regulated by mechanisms inducing changes in concentrations of the transcription factor σ^{32}, which is a core component in *E. coli*'s HSR system. These σ^{32} molecules promote the transcription of heat shock proteins by initiating the transcription of specific genes that in turn encode specific heat shock proteins (Straus et al. 1987). In steady-state conditions (30–37 °C), the intracellular concentration of σ^{32} is low, with a relatively short half-life in vivo, thus limiting the extent of gene transcription. As temperature increases and cells are exposed to thermoregulatory stress, gene expression is induced, which transiently increases the concentration of σ^{32} (*induction phase*). The increase in σ^{32} requires several transcriptional, translational, and post-translational regulatory components working in concert. σ^{32} swiftly produces a protein that binds to RNA polymerase (RNAP), and the response effectively allows this protein to control HSR by allowing it to bind to DNA and recognize heat shock promoters (Arsene et al. 2000; Guisbert et al. 2008). Return to a steady state involves rapid decreases in the concentration of σ^{32} and reduction in the synthesis of heat shock proteins (*shut-off phase*).

[2] HSR can also be triggered by non-thermal stress and pharmacological agents. For present purposes, focusing on thermal stress-induced HSR is sufficient.

Past a specific temperature (42 °C), escalation of rpoH gene transcription and synthesis in *E. coli* has less effect, although it can be modulated by other transcription factors. As mentioned, heat shock proteins are classified according to the ways in which they mitigate damaging effects: molecular chaperones do so by refolding denatured proteins, while proteases do so by degrading denatured proteins. The main chaperone system involved in σ^{32} regulation is DnaK, with co-chaperones DnaJ/GrpE. If the response is sufficiently swift and massive, cell death can be prevented by protein repair and/or removal of damaged proteins. Of course, the response needs to be tightly controlled, so that it is only activated in case of heat shock. This is in part because the response is highly energy consuming, and would otherwise make excessive energy demands were heat shock proteins being constantly produced. Cells must therefore maintain a delicate allostatic balance between the protective effect of heat shock protein production on one hand, and the metabolic cost of overproducing these proteins on the other.

In *E. coli*, the mechanism regulating heat shock protein expression following heat shock stress (temperature increase)—i.e., the HSR system—uses both feedforward and feedback loops that process information about temperature and the folding state of proteins in the cell. For instance, there is a loop for sensing temperature and controlling σ^{32} production, while another loop senses the folding level of proteins and governs the degradation of σ^{32}. These regulatory loops are also crucial for ensuring that heat shock protein production returns their population to a normal baseline after a sufficient number of proteins have been produced and the threat to cell damage or death is averted.

2.3 Modelling the HSR system: mechanistic 'virtual mutant' models

There have been many experimental studies of HSR in *E. coli*. Consequently, the biochemistry involved is well known. These details also make it possible for biologists to produce mathematical models of *E. coli's* HSR system and study the regulatory (design) principles involved. This is what El-Samad et al. (2005), and Kurata et al. (2006) did in several studies on the regulatory architecture of the mechanisms involved (Arkin and Schaffer 2011: p. 845; Guisbert et al. 2008). To uncover this architecture, they adopted a control engineering perspective on the HSR system, decomposing it in several functionally defined modules interconnected by multiple feedback and feedforward loops. The deployment of control engineering principles and techniques is now common in systems biology, because they offer tools to decipher the organization of complex biological mechanisms (Csete and Doyle 2002; Stelling et al. 2004; Braillard 2015; van Eck 2017).

In their review of research on *E. coli's* HSR system, Guisbert et al. (2008) succinctly described this control engineering-inspired analysis, the (more and less realistic) mechanistic models constructed in terms of it, and the results obtained through the comparison of these models as advanced by El-Samad et al. (2005, 2006):

> Temperature-regulated translation is considered to be a feedforward module allowing the system to respond to change in temperature before cellular processes are altered. Additionally, two feedback loops, one mediating activity control and

> a second mediating degradation control, report on cellular conditions, allowing a homeostatic response. To examine the function of each of these modules, the response properties of 'virtual mutants' that had various combinations of these modules were modeled. This analysis revealed that regulation is not redundant; instead, each module contributes different features to the response. The simplest system is one in which control is exerted solely by a direct sensor of temperature […] Such a system is inefficient because it utilizes many chaperones to accomplish folding at elevated temperatures even if the level of unfolded proteins is low […] The addition of activity control improves the efficiency of the system and makes it less sensitive to parameter variation. The further addition of degradation control improves the kinetics of the response, increases its efficiency, and reduces cell-to cell variation. Analyses of this type rationalize the complexity of biological control mechanisms. (Guisbert et al. 2008: pp. 549–550)

Let us elaborate. El-Samad et al. (2005, 2006) constructed a more realistic model of the actual *E. coli* HSR system and several less realistic models—so-called 'virtual mathematical mutants' (Kurata et al. 2006: p. 668)—in which specific feedforward and feedback loops were omitted. After running simulations with those models, they compared the simulation results to generate a design explanation of why the actual HSR system has the organizational complexity it has. It turned out that this complexity is necessary to achieve, inter alia, robust performance. These models are depicted in Fig. 1.

One reason why El-Samad, Kurata and co-workers sought a design explanation for the complex architecture of *E. coli's* HSR system (mechanistic organization) is that the system 'shows a level of complexity not justified by the basic functionality demanded from an operational heat shock system' (2005: p. 2738). From an engineering control perspective, a simple operational HSR system consisting solely of a temperature sensor and a transcriptional/translational device would do the trick. However, as simulation results revealed, it is 'alarmingly sensitive to parameter variations' (El-Samad et al. 2005: p. 2738). These results make clear that both feedforward and feedback loops are required to 'achieve robustness, noise rejection, speed of response, and economical use of cellular resources' (2005: p. 2738). So, *E. coli's* HSR system is constrained by specific performance requirements, such as robustness, timing, and sufficiency and efficiency of cellular resources. And the system has specific features that are difference-making factors for those requirements. By comparing the results of less and more realistic simulations, the researchers were able to rationally connect the possession conditions of the specific features of *E. coli's* HSR system with the requirements that those specific features are adapted to meet.

Importantly, note that these 'virtual mutant' models are produced such that they deliberately misrepresent the mechanistic activities and organization of the wild type *E. coli* HSR system in important ways, and are not to be thought of as (merely) abstract models in which details are omitted. The 'virtual mutant' models do so to assess the contribution of the loops omitted from the mutant models to *E. coli*'s HSR. So, here is a case in which the omission of details results in an intended misrepresentation, and thus by our lights counts as a case of idealisation (Godfrey-Smith 2006). Classifying this

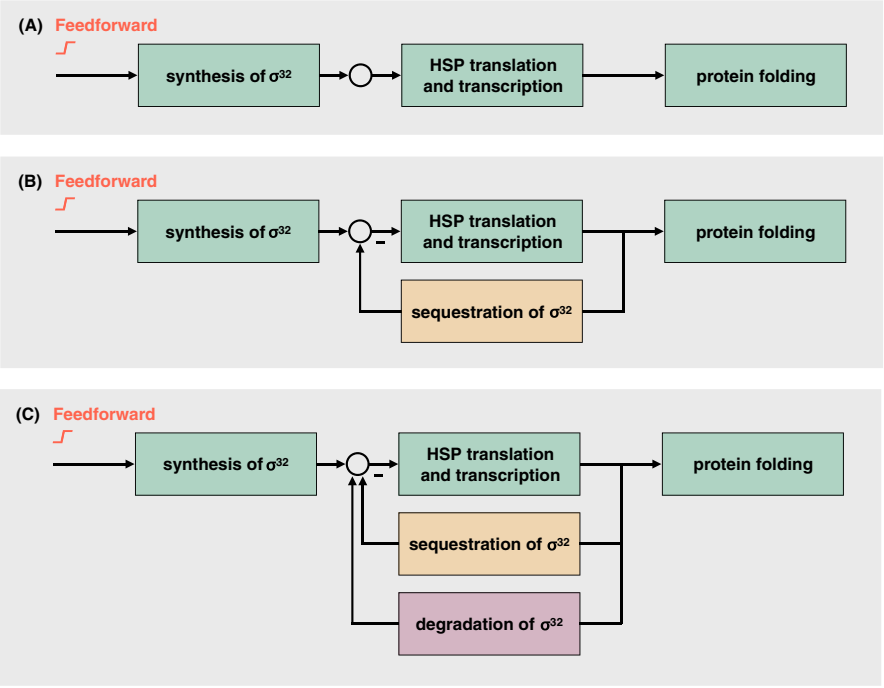

Fig. 1 More realistic and less realistic models of HSR system in *E. coli* (El-Samad et al. 2005)

as a case of abstraction blurs the fact that these scientists intentionally misrepresented the workings of *E. coli*'s HSR.[3]

To support these claims further, consider that Kurata et al. (2006) set the kinetic output parameter values for each model the same and assessed the performance of both more and less realistic 'virtual mutant' mechanistic models relative to these values. The mutant models were, inter alia, much slower and much less efficient in reaching these values than the more realistic model. By setting the kinetic output parameters of the mutant models to the same values appearing in wild type systems, so as to compare different mechanistic organizations and the functions of the loops therein, additional idealizations are introduced in the mutant mechanistic models, viz. their kinetic output parameters are also set to unrealistic values. That is, in the mutant models, the organization of *E. coli's* HSR mechanism is misrepresented, and the output indices are unrealistically adjusted such that these altered organizations can 'virtually'

[3] Let us stress that the mutant models in this research should not be understood as representations of alternative ways in which the *E.coli* HSR system could be organized, and thus have different non-existing targets. The mutant models do misrepresent wild type *E.coli* HSR systems because they are used to say something about that target, viz. that certain characteristics are present in the target system because these improve functionality. The researchers' aim is not to explore possible design variants so as to come up with a story about what the actual design looks like and how it functions. That is known from the start in their research; rather, they want to know why the target system functions the way it does. Given this request for explanation, it seems clear that the *E.coli* HSR wild type system is the target system driving the misrepresentations in the more unrealistic models.

obtain. Importantly, this strategy—i.e., the unrealistic adjustment of output indices in the mutant models and the comparison with the more realistic model—offers insights into the presence of specific organizational features of *E. coli's* HSR system that are difficult to attain with 'ordinary gene knockout mutants' (Kurata et al. 2006: p. 668). When working with ordinary knockouts, the system's organization is altered but kinetic output parameters are not adjusted. Rather, researchers assess how knockout mutants perform with respect to output parameters. In contrast, as Kurata et al. remark:

> Our virtual mathematical mutants, however, compensate for the mutation, by readjusting the kinetic parameter values so as to conserve such properties of the wild-type as yield and efficiency. This allows for the direct comparison between architectures that generate an equivalent output, a difficult task in the wet lab (2006: p. 668)

So, both mechanistic organizations are intentionally misrepresented, and the kinematic parameter values unrealistically adjusted, in order to provide a design explanation for the complex organization of the feedback loop found in wild type *E. coli* HSR systems. This research uses comparisons between more and less realistic mechanistic models to explain why mechanisms of kind k have feature F, and uses simulations run with those models to explain why these mechanisms have the structure and organization they have.

In the context of mechanistic explanation, idealized models serve a second explanatory function, viz. they enable tracking counterfactual dependency relations. And this tracking, in turn, offers control over the behaviour of the target systems that are modelled: counterfactual knowledge of this sort indicates what happens if the target system is manipulated in certain ways. Idealisations thus serve what is considered a key aim of explanatory mechanistic models, viz. control over target systems (Craver 2006). In the HSR case, comparisons between less and more realistic models allowed the researchers to track counterfactual dependency relations between features of mechanisms and features attributed to mechanistic explanandum phenomena. For example, they tracked what would happen to the speed of the heat shock response if the loop feeding information about folding states of the cells back to the components governing heat shock protein production were absent in actual systems (El-Samad et al. 2005: p. 2739). In this example, speed of response would be severely compromised. In Woodwardian terms, these comparisons provide a source to answer 'what-if-things-had-been-different questions' (Woodward 2003). And as elaborated below, they do so in different fashion than other experimental and modelling techniques.

There are a variety of experimental and modelling techniques in use in scientific practice for tracking counterfactual dependency relations between features of mechanisms articulated in mechanistic models and features attributed to mechanistic explanandum phenomena. For instance, intervention techniques in gene knockout experiments are routinely used in different fields of biology (and elsewhere) to track dependency relations. Simulation studies, in which parameter and variable values are adjusted in models, are also commonplace in many branches of biology (and elsewhere) to track how specific values of a phenomenon of interest depend upon specific values of component variables. Research on *E. coli*'s HSR system is a case in point in which both sorts of techniques are employed (e.g., El-Samad et al. 2005; Kurata et al.

2006). We submit that the idealizations just discussed are employed precisely because they offer additional modelling resources for tracking counterfactual dependencies—a point overlooked in the idealization and explanation literature. Importantly, this idealization modelling strategy offers (complementary) evidence for dependency relations that are not to be had with other experimental and modelling techniques. For example, as mentioned, by setting kinetic output parameters to unrealistic values in virtual mutant models, Kurata et al. (2006: p. 668) were able to directly compare different architectures that generate an equivalent output—something that is very difficult, and perhaps just not feasible, to establish with ordinary gene knockout mutants. Likewise, comparisons of the architectures of less realistic (virtual mutant) models with those of more realistic models provide a kind of evidence for dependency relations that cannot be procured solely with more realistic models. Of course, such more realistic models can be used to track and provide evidence for dependencies, but not evidence acquired through comparing more and less realistic models. In sum, idealizations offer an additional, complementary source for tracking counterfactual dependencies, and thus an additional source to assess whether evidence for dependencies procured with different techniques aligns or converges.

What this also indicates is that, to the extent that idealizations are widespread in (biological) modelling and serve important explanatory functions, then always and everywhere insisting on accurate representation of difference-making relations as the overriding norm by which to evaluate mechanistic models, as quite a few authors have it (Craver 2007; Kaplan and Craver 2011; Strevens 2008; Woodward 2003), is counterproductive and too demanding. We return to this issue in §4. For now, let us clarify how our approach relates to other work on mechanist idealisation.

3 Mechanist idealisation: what we have and what we don't have

As mentioned, specifically mechanistic analyses of idealisation in scientific modelling and explanation are few and far between. One recent discussion has it that idealisations are useful for constructing general mechanistic explanations (Glennan 2017). The underlying thought is that features differing from case to case or from token system to system are misrepresented, such that these idealisations make salient that those (misrepresented) features have no explanatory import in the general explanations being sought. The function attributed to idealisations thus concerns highlighting which features are explanatorily irrelevant for explaining mechanistic phenomena at the type level. This perspective on the function of idealisation is a general one, and consonant with the perspective that Strevens (2008) advanced almost a decade earlier in the context of idealisations in causal explanation.

The modelling strategy pursued in the HSR case bears similarities to the 'misrepresent irrelevancies' account of idealization (MIA) in the sense of pursuing a similar goal, viz. identifying difference-makers that should not be misrepresented in a more realistic model; but the identification of difference-makers in the modelling strategy we elaborated works differently. In contrast to MIA, the deployment of mutant models misrepresents actual, explanatorily relevant, mechanistic activities and organization precisely for the purpose of understanding the manner in which those features are

difference-makers for specific performance requirements. That is a remarkable difference, and one that shows that our understanding of idealizations has to be broader than the view provided by MIA.

Love and Nathan (2015) also stressed the point: in addition to the 'misrepresent irrelevancies' function of idealisation, idealisations also get used to misrepresent what is relevant. As they illustrate with several examples of the modelling of features of gene expression, 'in molecular biology, the causal relations responsible for the explanandum phenomenon are deliberately misrepresented on a regular basis' (2015: p. 764). Given that the causal relations that are responsible for an explanandum phenomenon (difference-makers) are explanatorily relevant features on any account of mechanistic explanation, mechanistic idealisations are thus used to misrepresent explanatorily relevant mechanistic features. Like us, Love and Nathan inquired as to why scientists use idealisations to 'misrepresent what is relevant', and they provided a different (complementary) answer than what we presented in §2. They endorse Weisberg's (2007, 2013) account of multiple-models idealization, arguing that the cases of mechanistic modelling they focus on are to be understood as ones in which different mechanistic models, idealised in different, incompatible ways, are used to highlight different features of the mechanistic explanandum phenomena targeted for explanation: '[t]he goal of mechanistic explanation is not an all-inclusive single model but a series of many complementary diagrams and descriptions comprising different idealisations' (Love and Nathan 2015: p. 772). Similarly, Hochstein argued that this kind of multiple models idealisation helps with generating explanations that represent different features of the same causal system: '[i]n these situations, there is no single model that can be used to provide a mechanistic explanation, since different models must be employed which adopt conflicting idealisations in order to represent different features of that mechanism needed for the explanation' (2016: p. 1398). The general point is that, by comparing and contrasting such different idealized mechanistic models, researchers can gain a better understanding of how a given target explanandum is produced. The particular HSR case analysed here illuminates this point (although in different fashion; the HSR case concerns comparisons between more and less idealized models, and the models are not idealised in incompatible ways). Being constrained by specific performance requirements, the HSR system has specific features that are difference-making factors for those requirements; and researchers increased their understanding of the design explanation only by comparing the more realistic and less realistic virtual mutant mechanistic models.

Regarding the question as to why difference-makers (explanatorily relevant features) are routinely misrepresented in mechanistic modelling practices, our analysis provides a complementary answer: such distortions serve (at least) two different, additional explanatory functions: viz. explaining the presence of structural/organizational features of mechanisms by reference to their role as difference-makers for performance requirements and tracking counterfactual dependency relations between features of mechanisms and features of mechanistic explanandum phenomena.[4]

[4] We do not claim completeness for our analysis: likely, there are other explanatory functions served by intentional misrepresentations of mechanistic difference-makers than the ones we and Love and Nathan (2015) identified. Our aim here was just to elaborate two salient and important ones.

Our analysis also compares with a recent treatment of the (central) roles of idealisation in explanation from Potochnik (2017). Potochnik identifies numerous intertwined reasons or motivations scientists may have to idealize, an important one being that '[i]dealizations provide a way to set aside complicating factors to discern a causal pattern of interest' (2017: p. 43; see also Mäki 2009). For instance, the frictionless plane assumption in physics, the rational agent assumption in economics, and the infinite population size assumption in biology are all idealizations that enable scientists to disregard complicating factors and focus on specific features of interest. Initially, Potochnik's observation appears to conform to MIA, also elaborated by Strevens (2008) and Glennan (2017); she claims that idealisations can 'eliminate […] noise or non-central influences' (2017: p. 50), facilitating 'the neglect of features not central to the patterns of immediate interest' (2017: p. 57). However, Potochnik is quick to note that she also assigns other representational roles to idealizations in scientific models (and other sorts of scientific representations). One is that they can also play a positive representational role by making the relevance of a factor salient by misrepresenting it; for in some cases, 'an idealization asserts the nature of a factor's relevance by saying something false about that factor—by representing it as if it were some way it isn't. In doing so, idealizations are fulfilling a subtle, but positive, representational purpose' (2017: p. 52). The same idealization may fulfil both roles simultaneously. Potochnik uses the Bay Area transit system as an example. Stops in transit system maps are often depicted as being in a straight line and equidistant from one another. Such a representation misrepresents spatial characteristics of actual transit systems but conveys a relevant functional similarity between representation and actual system, viz. the spatial and temporal order between stops is similar and it is those features that often matter to users of transit systems.

While the modelling strategies discussed here are not considered by Potochnik (2017), our analysis resonates with her treatment of idealized model-based explanation. It extends Potochnik's framework by elaborating two specific representational roles by which idealizations (can be used to) make salient the nature of a factor's relevance, viz. explaining the presence of structural/organizational features of mechanisms by reference to their role as difference-makers for performance requirements, and tracking counterfactual dependency relations between features of mechanisms and features of mechanistic explanandum phenomena.

Just as mechanist idealisation has not received sustained analysis, neither have the ramifications of mechanist idealisation for ontic constraint accounts of (the explanatory force of) mechanistic explanation been explored. We turn to this issue next.

4 Idealisation and the explanatory force of mechanistic models

We start with mechanistic modelling and the question of what makes mechanistic models explanatory. We consider one influential answer, often associated with ontic constraint accounts of mechanistic explanation: that explanatory models should accurately represent as many explanatorily relevant features of ontic structures (mechanisms) in the world as possible. On this account, the more accurate a model represents explanatorily relevant features of its target, the more explanatory it is. Our analysis

concurs with others in showing that, in the idealisation cases we analysed, this perspective is not gripping. Insofar as ontic constraint accounts emphasize accuracy as the key constraint on explanatory mechanistic models, ontic constraint accounts are not gripping either.

4.1 Mechanistic modelling

In his analysis of the gap between models that are explanatory and those that are not, Craver (2006) concludes that—their usefulness for making predictions, summarizing data, and designing experiments notwithstanding—non-explanatory models do not afford control over the behaviour of target systems. Control is gained by deploying and refining explanatory models; and such models, he argues, are mechanistic. Craver's account is not a general account of the conditions under which models explain, however, and is restricted to those concerned with mechanisms. This is because mechanistic systems are not the only type of phenomena targeted for explanation. One consequence is that not all explanation is mechanistic explanation. Another is that explanations implicate mechanistic modelling only to an extent, and it is unclear what that extent is (See, e.g., Craver 2006: p. 367; Kaplan 2017: p. 70 fn. 1).

Also unclear is the exact nature of the relationship between mechanistic models and mechanisms. Traditionally, it has been construed as one of representation; but scientific representation is fraught with problems: e.g., whether there is any specific problem to begin with (Callender and Cohen 2005), whether it can be analysed as similarity, analogy, morphism of varying degrees of stringency, or other such relations. However, Knuuttila (2011) has convincingly argued that conceptual analyses of the representational relationship, regardless of adicity, are often either too demanding or too minimal.

Some New Mechanists have suggested that the representations constitutive of mechanistic models are descriptions. For instance, Glennan claimed that 'a mechanical [mechanistic] model consists of (1) a description of the mechanism's behavior (the behavioral description); and (2) a description of the mechanism that accounts for that behavior (the mechanical description)' (Glennan 2005: p. 446). Others have suggested that the representations constitutive of mechanistic models are depictions. For instance, Perini (2005) and Ratti (2019) claim that mechanistic models involve visuospatial diagrams, and the more precisely they depict how entities and activities are organized to fully account for the phenomenon, the more they are explanatory.

The relationship between models and mechanisms is complicated in other ways. Some models can be mechanisms themselves without necessarily having mechanisms as their targets, and hence do not qualify as mechanistic models. One example would be the 'Phillips/Newlyn Machine' or MONIAC, which was a hydrodynamic model of the British economy (Newlyn 1950). And some mechanisms themselves can also serve as models, such as the giant squid axon, for representing or depicting features of other mechanisms (Matthewson 2017). We have focused here on models of mechanisms that represent target mechanisms, such as *E. coli's* HSR system. Such models come in different forms, such as diagrams that visually represent entities and activities and organizational relations between them, or equations and formulæ that represent

features of mechanisms in terms of mathematical relations between variables. Sometimes such models are superimposed on one another, as in the work of El-Samad et al. (2005, 2006) on *E. coli's* HSR system.

So what makes models of mechanisms explanatory, and what makes some models better than others? One widely discussed perspective comes from Kaplan and Craver, who wrote that a model of a target phenomenon explains that phenomenon to the extent that:

> the variables of the model correspond to identifiable components, activities, and organizational features of the target mechanism that produces, maintains, or underlies the phenomenon, and the (perhaps mathematical) dependencies posited among these (perhaps mathematical) variables in the model correspond to causal relations among the components of the target mechanism. (2011: p. 611)

The heavy lifting in this passage is done by the term *correspondence*, which Kaplan & Craver leave unanalysed. Since the concept it expresses has been understood in various ways, there occurs the exegetical question. They probably don't have correspondence-as-isomorphism in mind, as Craver and many others have noted that 'few models are actually isomorphic with the phenomenon, given that models typically abstract away from the precise details of the system being modelled, that they typically are only approximate, and that they make simplifying assumptions in order to apply a particular formalism' (2006: p. 357). A plausible interpretation is just that they intend something simpler: e.g., the more accurate a mechanistic model represents its target mechanism(s), the better it explains.[5] This interpretation is reinforced by Craver, who later wrote that 'the norms of [mechanistic] explanation fall out of a commitment by scientists to describe as accurately and completely as possible the relevant ontic structures in the world' (2014: p. 48; see also Kaplan and Craver 2020).

The completeness constraint has been the subject of debate recently, often with an eye to the role of abstraction in mechanistic models (Levy and Bechtel 2013; van Eck and Weber 2014; van Eck 2015a; Boone and Piccinini 2016). Yet, as a commitment out of which the norms of mechanistic explanation fall, Kaplan and Craver explicitly suppose that completeness is unworkable: 'the idea of an ideally complete how-actually model, one that includes all of the relevant causes and components in a given mechanism, no matter how remote, negligible, or tiny, without abstraction or idealisation, is a philosopher's fiction. Science would be strikingly inefficient and useless both for human understanding and for practical application if it dealt in such painstaking minutiæ' (2011: pp. 609–610). So, pace Craver, commitments to completeness are counterproductive because attempts to force mechanistic explanations to satisfy that norm preclude scientists from developing efficient and genuine understanding. And so, if constraints of accuracy and completeness are said to enjoy a sort of global priority over all others, it is ultimately accuracy that is fundamental.

[5] Determining which variables and dependencies ought to be articulated in a given mechanistic model is a different matter. Levy and Bechtel (2013) pitch their account of explanatory relevance against the views of e.g., Craver (2007). One of us argues (van Eck 2015a, 2017) that these sets of authors in fact endorse compatible rather than competing positions; they subscribe to different notions of difference-making, which are suitable for different explanatory requests.

These commitments imply that idealised mechanistic models in which difference-making features are distorted explain less well than more accurate models, in so far as idealised models are explanatory at all.[6,7] This negative view might itself explain why idealisations have been largely ignored in the mechanistic explanation literature. But is it really the case that idealised mechanistic models (always) make for worse explanations? This negative view is difficult to square with the fact that scientists make such extensive explanatory use of idealisations in their models, mechanistic or otherwise.

A look at scientific practice undermines the assumption that scientists are always committed to modelling, as accurately as possible, the structures the workings of which they attempt to explain. To the contrary, as we have seen, they have good reasons for modelling such structures inaccurately in certain contexts: idealisations, instead of making for worse explanations, can serve important explanatory functions. One of them—helping to track counterfactual dependency relations—is especially important for mechanistic explanatory purposes: explanatory mechanistic models (should) afford control over the behaviour of target systems, and knowledge of counterfactual dependencies helps gain control by articulating what happens if the system is manipulated in certain ways (Craver 2007). For instance, the comparisons between the (simulation results of) more and less idealised models in the discussed research on *E. coli's* HSR system tell you that if certain regularity loops were shut down, such as the σ^{32} inactivation loop, system performance is affected, such as speed of the HSR in case of inactivating the σ^{32} inactivation loop.

There is more at stake here than merely the (alleged) normative priority of accuracy on mechanistic explanatory models. Since ontic constraint accounts on mechanistic explanation suppose that accuracy is the key constraint on explanatory mechanistic models, these accounts fall through when accuracy falls through. In the next subsection we outline an alternative to ontic constraint accounts on explanatory norms governing mechanistic representations.

[6] It might be that theorists committed to this accuracy perspective sanction the use of idealizations that distort irrelevant features; and it might also be that some of those theorists concede that idealizations may distort explanatorily relevant features as long as this is required to accurately represent more important explanatorily relevant features. But the concession contradicts the norm that mechanistic explanations should describe as accurately as possible the relevant ontic structures in the world; what is then minimally required is a theory of 'relative' explanatory importance that explains when, and when not, the distortion of explanatorily relevant features is sanctioned. As long as such a theory is not on offer, the second concession seems to entail giving up on the accuracy constraint. Also, given that organization is key to the operation of mechanisms, key features are distorted in the cases analysed here, not relatively minor details. Thus, the accuracy perspective would still be hard pressed to account for the cases of idealization discussed here.

[7] The term *explain* here refers both to design explanations that hinge on the comparison of mechanistic models and to the articulation of counterfactual dependencies using mechanistic models. On a side note, models that represent target mechanisms can be used both for mechanistic explanatory purposes (as in the latter case) and for non-mechanistic explanatory purposes (as in the former case). Also, when one is inclined to think that both explanatory purposes are mechanistic (cf. Matthiessen 2017), *explain* still has a twofold sense.

4.2 An alternative to ontic constraint accounts of mechanistic explanation

In the literature, the current debate between advocates of ontic versus epistemic conceptions of mechanistic explanation revolves around the question of the nature of explanation, i.e., whether explanations are things in the world or representations of them (Craver 2007; Wright 2012, 2015; Bokulich 2016; Wright and van Eck 2018). This is a different debate than the one over which explanatory constraints (should) govern judgements about the goodness of mechanistic representations (Illari 2013; van Eck 2015a, b; Sheredos 2016).

Advocates of ontic constraint accounts consider accuracy to be the key constraint on explanatory mechanistic models. Prima facie, such views are hard to reconcile with pervasive practices of idealized modelling. The ill fit between endorsing accuracy as the key constraint on mechanistic models and the abundant use of idealisations in scientific mechanistic practice was presciently described by Love and Nathan:

> The idealisation of causal relations demonstrates that these models do not depict how the mechanism actually works. If actual difference-makers are represented in such a way that they are not difference-makers, according to what is already known about the mechanism, mechanistic explanations appear to fail according to their own criteria. (2015: p. 768)

The problem they signal is clear: models that misrepresent difference-making causes do not, ipso facto, show how their effects are actually produced; in contradistinction, accurate models of difference-making causes do not involve deliberate distortions of the relations between these difference-making causes and their effects.

Overemphasis on accuracy fits poorly with scientific practice and turns a blind eye to important explanatory functions of idealisation, such as explaining the presence of structural/organizational features of mechanisms by reference to their role as difference-makers for performance requirements, and the tracking of counterfactual dependency relations between features of mechanisms and features of mechanistic explanandum phenomena. We do not mean to imply that accuracy is immaterial to mechanistic modelling, of course; no proper explanatory realism can get by without it. Instead, what our analysis shows—like Love and Nathan's (2015)—is that accounts of the explanatory force of mechanistic models that assign central importance to accuracy are uninformative in those contexts in which explanatory mechanistic models idealize explanatorily relevant features. Ipso facto, ontic constraint accounts lack support there as well.

What, then, are appropriate explanatory norms that should govern the evaluation of mechanistic representations? Limitations on space prevent elaborating such an account in detail; but we do want to offer some suggestions. Two widely endorsed explanatory norms on explanations in the literature are that (1) explanations should identify difference-makers and (2) capture or track dependency relations (Woodward 2003; Craver 2007; Glennan 2017); and another widely endorsed connection and precisification of these norms is in terms of the constraint (3) that explanations should track counterfactual dependencies between difference-makers cited in the explanans and an explanandum phenomenon. For instance, explanatory mechanistic models (should) afford control over the behaviour of target systems and one important way to gain

such control is by knowledge of counterfactual dependencies (Craver 2007). In other words, one important, widely recognized, aspect of explanatory force is the ability to pose and answer 'what-if-things-had-been-different' questions on the basis of (explanatory) information concerning difference-makers expressed in a (mechanistic) model (Woodward 2003; Ylikoski and Kuorikoski 2010). Our positive proposal thus is that the tracking of counterfactual dependencies between difference-makers and explanandum phenomena is a key constraint for evaluating mechanistic models.

In light of this constraint, we see that misrepresentations of explanatory relevant (difference-making) features in mechanistic models can—and often do—play a positive explanatory role. In the systems biology case investigated here, meeting this key constraint operates over comparisons between a variety of more idealized models that misrepresent mechanistic difference-makers in different ways on the one hand, and a less idealized model wherein these difference-makers are not distorted on the other. Such comparisons give insight into the nature of difference-makers' relevance and show how explanandum phenomena depend on them. In this case, the heavy lifting in explanation is not done by accurate representation alone: idealisations are a crucial source of explanatory power.

We suggest that this perspective is more broadly applicable than just the cases discussed in this paper. We take it that this perspective also gives a yardstick to evaluate mechanistic multiple-models idealization practices as discussed by Love and Nathan (2015): when, due to a target system's complexity, different models misrepresent explanatorily relevant mechanistic features in different ways, but the set of those models, as revealed by comparing them, enables identifying difference-makers for a target explanandum phenomenon, and tracking counterfactual dependencies between them, the set of models gains explanatory power.[8]

5 Conclusion

This paper charts two related explanatory functions of idealisation in mechanistic models in systems biology, viz. (1) explaining the presence of structural/organizational features of mechanisms by reference to their role as difference-makers for performance requirements, and (2) tracking counterfactual dependency relations between features of mechanisms and features of mechanistic explanandum phenomena. Our discussion of systems biological research on the mechanism(s) for countering heat shock—the Heat Shock Response (HSR) system—in *Escherichia coli* (*E. coli*) bacteria illustrates these functions. We closed by elaborating a general lesson from our research on idealised mechanistic models: ontic constraint accounts of mechanistic explanation provide inferior normative appraisals of mechanistic models.

This work was motivated by the observation that, whilst idealisations are routinely deployed in scientific practice, there are preciously few specifically mechanistic analyses of idealisation. We seek to add momentum to this important but underdeveloped

[8] We hasten to say that we do not take our suggestions to be the only way to spell out explanatory norms on (mechanistic) explanations. We do take our suggestions to embody widely endorsed commitments in the explanation literature and to enable accounting for the positive role of idealizations in mechanistic modeling, something which ontic constraint views have a hard time accommodating.

line of research. Whereas some extant work on mechanist idealisation clarified why idealisations are used to misrepresent (and thus make salient) explanatorily irrelevant features, other work clarified why idealisations are also used to misrepresent explanatorily relevant features. We charted additional reasons for the latter use of idealizations, clarifying further why idealisations are used to misrepresent features that are explanatorily relevant: such misrepresenting serves the two functions mentioned above. No doubt, mechanist idealisation serves other functions as well. As such, we hope that mechanist idealisation will receive more sustained analysis.

Acknowledgements Funding was provided to Dingmar van Eck by the Research Foundation Flanders (FWO). We thank our reviewers and the guest editors of this special issue for very useful comments and suggestions.

References

Arkin, A., & Schaffer, D. (2011). Network news: Innovations in 21st century systems biology. *Cell, 144,* 844–849.

Arsene, F., Tomoyasu, T., & Bukau, B. (2000). The heat shock response of *Escherichia coli*. *International Journal of Food Microbiology, 55,* 3–9.

Aubin, T., Jouventin, P., & Hildebrand, C. (2000). Penguins use the two-voice system to recognize each other. *Proceedings Royal Society London, 267,* 1081–1087.

Bokulich, A. (2016). Fiction as a vehicle for truth: Moving beyond the ontic conception. *The Monist, 99,* 260–279.

Boone, W., & Piccinini, G. (2016). Mechanistic abstraction. *Philosophy of Science, 83,* 1–13.

Braillard, P. A. (2015). Prospects and limits of explaining biological systems in engineering terms. In P. Braillard & C. Malaterre (Eds.), *Explanation in biology* (pp. 319–344). Berlin: Springer.

Callender, G., & Cohen, J. (2005). There is no special problem about scientific representation. *Theoria, 55,* 67–85.

Craver, C. (2006). When mechanistic models explain. *Synthese, 153,* 355–376.

Craver, C. (2007). *Explaining the brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Oxford University Press.

Craver, C. (2014). The ontic conception of explanation. In M. Kaiser, O. Scholz, D. Plenge, & A. Hüttemann (Eds.), *Explanation in the special sciences* (pp. 27–52). Berlin: Springer.

Csete, M., & Doyle, J. (2002). Reverse engineering of biological complexity. *Science, 295,* 1664–1669.

El-Samad, H. J., Kurata, H., Doyle, J. C., Gross, C. A., & Khammash, M. (2005). Surviving heat shock: control strategies for robustness and performance. *PNAS, 102,* 736–2741.

Glennan, S. (2005). Modeling mechanisms. *Studies in History and Philosophy of Biological and Biomedical Sciences, 36,* 443–464.

Glennan, S. (2017). *The new mechanical philosophy*. Oxford: Oxford University Press.

Godfrey-Smith, P. (2006). The strategy of model-based science. *Biology and Philosophy, 21,* 725–740.

Green, S. (2015). Revisiting generality in biology: Systems biology and the quest for design principles. *Biology and Philosophy, 30,* 629–652.

Guisbert, E., Yura, T., Rhodius, V., & Gross, C. (2008). Convergence of molecular, modeling, and systems approaches for an understanding of the of *Escherichia coli* heat shock response. *Microbiology and Molecular Biology Reviews, 72,* 545–554.

Hochstein, E. (2016). One mechanism, many models: A distributed theory of mechanistic explanation. *Synthese, 193,* 1387–1407.

Illari, P. (2013). Mechanistic explanation: Integrating the ontic and epistemic. *Erkenntnis, 78,* 237–255.

Jones, M. (2005). Idealisation and abstraction: a framework. In M. Jones & N. Cartwright (Eds.), *Idealisation XII: Correcting the model* (pp. 173–217). Bulgaria: Rodopi.

Kaplan, D. (Ed.). (2017). *Explanation and integration in mind and brain science*. Oxford: Oxford University Press.

Kaplan, D., & Craver, C. (2011). The explanatory force of dynamical and mathematical models in neuroscience: A mechanistic perspective. *Philosophy of Science, 78,* 601–627.
Kaplan, D., & Craver, C. (2020). Are more details better? On the norms of completeness for mechanistic explanations. *British Journal for the Philosophy of Science, 71,* 287–319.
Kaufmann, S. (1990). Heat shock proteins and the immune response. *Immunology Today, 11,* 129–136.
Knuuttila, T. (2011). Modelling and representing: An artefactual approach to model-based representation. *Studies in History and Philosophy of Science, 42,* 262–271.
Kurata, H., El-Samad, H., Iwasaki, R., Ohtake, H., Doyle, J. C., Grigorova, I., et al. (2006). Module-based analysis of robustness tradeoffs in the heat shock response system. *PLoS Computational Biology, 2,* 663–675.
Levy, A., & Bechtel, W. (2013). Abstraction and the organization of mechanisms. *Philosophy of Science, 80,* 241–261.
Love, A., & Nathan, M. (2015). The idealisation of causation in mechanistic explanation. *Philosophy of Science, 82,* 761–774.
Mäki, U. (1992). On the method of isolation in economics. *Poznań Studies in the Philosophy of the Sciences and the Humanities, 26,* 319–354.
Mäki, U. (2009). Realistic realism about unrealistic models. In H. Kincaid & D. Ross (Eds.), *Oxford Handbook of the philosophy of economics*. Oxford: Oxford University Press.
Matthewson, J. (2017). Models of mechanisms. In S. Glennan & P. Illari (Eds.), *Routledge Handbook of mechanisms and mechanical philosophy* (pp. 225–237). London: Routledge.
Matthiessen, D. (2017). Mechanistic explanation in systems biology: Cellular networks. *British Journal for the Philosophy of Science, 68,* 1–25.
Newlyn, W. (1950). The Phillips/Newlyn hydraulic model. *Yorkshire Bulletin of Economic and Social Research, 2,* 111–127.
Nowak, L. (1980). *The structure of idealization*. Kufstein: Reidel.
Perini, L. (2005). Explanation in two dimensions: Diagrams and biological explanation. *Biology & Philosophy*, *20*, 257–269.
Potochnik, A. (2017). *Idealization and the aims of science*. Chicago: University of Chicago Press.
Ratti, E. (2019). 'Models of' and 'models for': on the relation between mechanistic models and experimental strategies in molecular biology. *British Journal for the Philosophy of Science, 71,* 773–797.
Richter, K., Haslbeck, M., & Buchner, J. (2010). The heat shock response: Life on the verge of death. *Molecular Cell, 40,* 253–266.
Sheredos, B. (2016). Re-reconciling the epistemic and ontic views of explanation (or, why the ontic view cannot support norms of generality). *Erkenntnis, 81,* 919–949.
Stelling, J., Sauer, U., Szallasi, Z., Doyle, F., & Doyle, J. (2004). Robustness of cellular functions. *Cell, 118,* 675–685.
Straus, D., Walter, W., & Gross, C. (1987). The heat shock response of *E. coli* is regulated by changes in the concentration of σ^{32}. *Nature, 329,* 348–351.
Strevens, M. (2008). *Depth: An account of scientific explanation*. Harvard: Harvard University Press.
van Eck, D. (2015a). Mechanistic explanation in engineering science. *European Journal for Philosophy of Science, 5,* 349–375.
van Eck, D. (2015b). Reconciling ontic and epistemic constraints on mechanistic explanation, epistemically. *Axiomathes, 25,* 5–22.
van Eck, D. (2017). Mechanisms and engineering science. In S. Glennan & P. Illari (Eds.), *Routledge handbook of mechanisms and mechanical philosophy* (pp. 447–461). London: Routledge.
van Eck, D., & Mennes, J. (2018). Mechanism discovery and design explanation: Where role function meets biological advantage function. *Journal for General Philosophy of Science, 49,* 413–434.
van Eck, D., & Weber, E. (2014). Function ascription and explanation: Elaborating an explanatory utility desideratum for ascriptions of technical functions. *Erkenntnis, 79,* 1367–1389.
Weisberg, M. (2007). Three kinds of idealisation. *Journal of Philosophy, 104,* 639–659.
Weisberg, M. (2013). *Simulation and similarity: Using models to understand the world*. Oxford: Oxford University Press.
Wong, H. (2005). Endogenous cytoprotective mechanisms. In L. Bertok & D. Chow (Eds.), *Natural immunity* (pp. 49–65). Amsterdam: Elsevier.
Woodward, J. (2003). *Making things happen*. Oxford: Oxford University Press.
Wouters, A. (2007). Design explanations: Determining the constraints on what can be alive. *Erkenntnis, 67,* 65–80.

Wouters, A. (2013). Biology's functional perspective: roles, advantages, and organization. In K. Kampourakis (Ed.), *Philosophy of biology: A companion for educators* (pp. 455–486). Berlin: Springer.
Wright, C. (2012). Mechanistic explanation without the ontic conception. *European Journal of Philosophy of Science, 2,* 375–394.
Wright, C. (2015). The ontic conception of scientific explanation. *Studies in the History and Philosophy of Science, 54,* 20–30.
Wright, C., & van Eck, D. (2018). Ontic explanation is either ontic or explanatory, but not both. *Ergo, 5*(38), 997–1029.
Ylikoski, P., & Kuorikoski, J. (2010). Dissecting explanatory power. *Philosophical Studies, 148,* 201–219.