

Mechanistic Explanation in Systems Biology: Cellular Networks

Dana Matthiessen

ABSTRACT

It is argued that once biological systems reach a certain level of complexity, mechanistic explanations provide an inadequate account of many relevant phenomena. In this article, I evaluate such claims with respect to a representative programme in systems biological research: the study of regulatory networks within single-celled organisms. I argue that these networks are amenable to mechanistic philosophy without need to appeal to some alternate form of explanation. In particular, I claim that we can understand the mathematical modelling techniques of systems biologists as part of a broader practice of constructing and evaluating mechanism schemas. This argument is elaborated by considering the case of bacterial chemotactic networks, where some research has been interpreted as explaining phenomena by means of abstract design principles.

- 1 *Introduction*
 - 2 *The Strategies of Mechanistic Science in the Study of Cellular Networks*
 - 3 *Are Mechanisms Insufficient?*
 - 4 *Design Explanations and Bacterial Chemotaxis*
 - 5 *Conclusion*
-

1 Introduction

The last two decades have seen a widening embrace of the mechanistic perspective among philosophers of biology. This view, which received early contributions from Bechtel and Richardson ([1993]), Glennan ([1996]), and Machamer *et al.* ([2000]), holds that much of the research being done in the life sciences is best understood as a search for the mechanisms responsible for phenomena. Explaining phenomena using mechanisms is seen as an alternative to the traditional deductive-nomological or unificationist models of scientific explanation, which, because of the scarcity of general, exceptionless

laws in biology, fail to adequately characterize these fields (Beatty [1995]; Bechtel and Abrahamsen [2005]). In contrast, the ‘new mechanistic philosophy of science’ (Skipper and Milstein [2005]) emphasizes the practice of specifying organized complexes of entities and activities that produce phenomena as an explanatory strategy.¹

Yet it is still developing as a perspective and there is disagreement over how broadly this notion of explanation can or ought to be applied.² Mechanistic philosophy draws heavily from examples in the fields of molecular biology, neuroscience, and basic cell biology. While authors acknowledge its adequacy for such fields, it is frequently contended that, as biological systems approach higher levels of behavioural and compositional complexity, the explanatory purchase of standard mechanistic concepts and strategies begins to falter (see, for example, Ramsey [2008]; Weiskopf [2011]; Andersen [2012]; Brigandt [forthcoming]). In particular, the sufficiency of mechanistic philosophy is currently being disputed for research that falls under the banner of ‘systems biology’. This term refers to a collection of research programmes, largely outgrowths of molecular and cell biology but also including studies at the population level, that share a certain methodological overlap. Unlike much of biology, the use of quantitative models predominates, especially those of dynamic systems theory and advanced statistics. This results in a greater emphasis on computational models and experiments *in silico*. Affiliated authors call for a ‘holistic’, ‘synthetic’, ‘anti-reductionist’, or ‘integrative’ conception of biology (Boogerd *et al.* [2007a]). While some hold that systems biology is mechanistic (Boogerd *et al.* [2007b]; Richardson and Stephan [2007]; Bechtel [2011]; Bechtel and Abrahamsen [2013]), others argue that it involves a distinct, non-mechanistic form of explanation. The authors that I will consider (Braillard [2010]; Kuhlmann [2011]; Silberstein and Chemero [2013]) take a similar stance on the status of systems understood through network modeling.³ The sense that such research falls outside of the mechanistic purview is, I think, motivated by three factors:

- (1) A tendency in the mechanistic literature to privilege linear sequential processes that produce a phenomenon as their end state over those

¹ This is the terminology introduced in the Machamer–Darden–Craver variant of mechanistic philosophy (Machamer *et al.* [2000]). Other authors make use of other terms with slightly different connotations and uses, although all retain the basic dualism between entities and activities. Bechtel and Richardson ([1993]) prefer ‘parts and operations’, while Glennan ([1996]) cites ‘parts and their interactions’.

² For this reason, some authors have voiced concern over the perceived excesses of ‘mechanistic imperialism’ (cf. Weiskopf [2011]).

³ It should be noted that each author varies slightly in their terminology. Only Braillard speaks directly of systems biology. Kuhlmann instead describes ‘dynamically complex systems’, and Silberstein and Chemero are more immediately concerned with systems neuroscience. That said, I believe the substantial similarity in the form of their arguments and the suggestion that one can draw implications for systems biology from them justifies this admitted act of lumping.

that are cyclic and underlie or maintain phenomena. The belief that these are the fundamental forms of mechanisms is perhaps exacerbated by the phrasing of the oft-cited Machamer–Darden–Craver definition, according to which they are ‘are productive of regular changes from start or set-up to finish or termination conditions’ (Machamer *et al.* [2000], p. 3).

- (2) The emphasis in the early work of Bechtel and Richardson ([1993]) on the strategies of decomposition and localization, by which one partitions a complex system into its spatial components and divides up the functions of the mechanism in terms of them. This has led some authors to believe that mechanistic explanation is strictly downward-looking: describing the functioning of the whole in terms of the combined properties of its lower-level parts.
- (3) The methodological differences between systems biology and other biological fields. The heavy use of mathematics for modelling purposes, which runs counter to the standard reliance on visual, pictorial diagrams in biology, is thought to conflict with mechanistic reasoning. In particular, the predictions made on the basis of model structures are viewed as explanations that are not strictly mechanistic.

In what follows, I will evaluate these latter claims with respect to a representative programme in systems biological research: the study of regulatory networks within single-celled organisms. I argue that these networks are amenable to mechanistic philosophy without need to appeal to some alternate form of explanation. I intend to show how commonly accepted postulates and procedures of mechanistic philosophy apply to and are referenced within the study of cellular networks. In addressing the critics, I aim to advance a picture of mechanistic science that is in accord with the methodology of those studying cellular networks. For this I will draw on a depiction of the explanatory strategies of mechanistic science, which I owe to Craver and Darden ([2013]); they cover this in greater detail. After briefly arguing for the compatibility of systems biology with these strategies, I will consider and reject an alternative view of systems biology, which I associate with the notion of ‘design explanation’.

2 The Strategies of Mechanistic Science in the Study of Cellular Networks

The search for mechanisms can be presented as three general stages: characterization of the phenomena, construction of schemas, and evaluation of schemas. The aim of this process is to establish a representation of a mechanism that captures all of the features contributing to the phenomenon under

consideration. Once a relatively complete representation has been constructed, it enables one to give a mechanistic explanation of the phenomenon. This is done by referring to the manner in which the phenomenon is constituted by what the schema represents—namely, the productive continuity and organization of the entities and activities that comprise the mechanism.⁴

Making sure one accurately characterizes the phenomenon is a crucial initial step as it serves as a guiding norm for the following stages. For our purposes, a phenomenon is taken to be a regular natural occurrence, ‘a repeatable type of event or product’ (Craver and Darden [2013], p. 54). Phenomena are characterized with the aid of data collected through observation and experimentation. William Bechtel ([2007]; [2011]; Bechtel and Abrahamsen [2013]) has noted that the phenomenon of life, at its most basic level, involves the maintenance of a system that is autonomous from its environment. Living organisms resist entropic forces, which tend towards disorder, by continually drawing energy from their surroundings and directing it towards processes that uphold their internal organization. By maintaining their structure and behaviours with relative consistency within varying environmental or internal conditions, organisms manifest robustness in various ways. Accounting for these and related phenomena is one of the aims of systems biology. The systems studied often involve multiple entities interacting in causal cycles. As a result, the systems as a whole tend to exhibit non-linear dynamics—that is, their behaviour fluctuates at a non-constant rate and may involve complex oscillations or deterministic chaos (Strogatz [1994]). Due to the complexity of these phenomena, which cannot be precisely captured by verbal or diagrammatic reconstruction alone, computational modelling of the system’s dynamics is often employed.

Cells are maintained by multilevel systems composed of different molecular entities and their interactions. These regulatory networks of genes, transcribed RNA, proteins, and metabolites function in highly complex, coordinated ways. They thus fortify basic cell activities against changes in the cell’s environment. The guiding aim of those who study cellular networks is the construction of a complete model of such self-regulating behaviour (Herrgard *et al.* [2004]). While comprehension of relevant intra-cellular processes—transcription, protein modification, ribosome assembly, and so on—provides a conceptual underpinning, the study of cellular networks regularly abstracts from the details in order to precisely describe the quantitative relations between their products. For instance, researchers focus on relative rates at which genes are being expressed by measuring the concentration levels of transcribed mRNA or synthesized proteins within a cell. By mapping

⁴ I sometimes make reference to ‘components’ instead of entities and activities. I think of these as discrete, intra-mechanism complexes of entities and activities, although a single entity could also constitute a component.




Figure 1. Three mechanism–phenomenon relations.

these relations, they aim to determine the corresponding network of interactions by means of which regular cell activity is maintained: gene–gene, gene–protein, protein–protein, and the like. Moreover, this will allow for a more accurate characterization of the specific functional role of each gene within the cell.

In contrast to point (1) above, Craver and Darden ([2013]) distinguish between three relations that a mechanism may have to a phenomenon (Figure 1); the mechanism may produce a phenomenon at the end of a continuous process; it may underlie it as a number of organized working parts that give rise to a phenomenal whole; or it may maintain a phenomenon by counterbalancing forces that shift some process away from a homeostatic steady-state point. I want to suggest that cellular networks are mechanisms that are best thought of as underlying and maintaining phenomena. That is, the observed behaviour of a cell is held to be constituted by an underlying organization of genes, transcription factors, proteins, and metabolites. This underlying organization is given schematically in terms of a network graph. The cell’s stable functioning under shifts in environmental or internal conditions is maintained by means of

the complex regulatory relations existing between network constituents. This functional maintenance is given schematically in terms of the system-level dynamics derived from the network representation.

The task of schema construction works to determine the components and organization of the mechanism that produces a phenomenon. The result is a representation of how a collection of entities and activities are organized in a manner that displays the productive continuity underlying the phenomenon—a mechanism schema. Roughly speaking, schemas can be thought of along similar lines as models in other areas of science; they are objects, more or less abstract, whose parts stand in some kind of representational relation to things in the world and that are not strictly reducible to linguistic description, instead often drawing upon mathematical or diagrammatic presentations (Teller [2001]; Frigg and Hartmann [2012]).⁵

Schema construction is the main stage in the mechanistic logic of discovery. It is orientated by a focused, pragmatic interest in the particular phenomenon under concern and the levels of organization within which its mechanism or mechanisms are individuated. Analogical reasoning may play an important role here (Hanson [1970]): researchers may enter the construction stage with a stock of mechanism types known to operate in structurally or environmentally similar phenomena. These can be used to help generate hypotheses about constitutive entities or activities or, by abstracting from particular details, they may give clues as to the spatial, temporal, or causal organization of the mechanism (Craver and Darden [2013], Chapter 5). The specific advantages of such abstraction have been noted recently by Levy and Bechtel ([2013]), who note that a high degree of abstraction often helps in determining the contributions of causal organization to system-level behaviours.

An initial step in schema construction for systems biology researchers is that of network inference. Given a set of data,⁶ a number of statistical methods are available for the construction of cellular network models, many of which involve building point-and-line graphs in which the nodes represent genes, proteins, or other cell products, and the edges stand for some form of dependency relation (derived algorithmically from the data) between them. Using one

⁵ Due to the predominance of the term ‘model’ in discussions of systems biology, I’ll tend to use it instead of Darden and Craver’s ‘schema’. For my purposes, they are interchangeable.

⁶ Research in cellular networks has been spurred by the increased availability of high-throughput data gathering given by measurements such as gene expression profiling, which employs robotics to gauge the expression activity of prepared mRNA samples. Multiple parallel experiments may be carried out at once, yielding very large quantities of data. So much data are produced that curated databases are used to store the results, but it is important to note that difficulties in accuracy attend to this process. Most databases are not designed to account for context-sensitive gene activity; high-throughput analyses often fail to detect rare events or unstable interactions; and the data available for model organisms usually address a small number of cell processes and experimental conditions (De Backer *et al.* [2010]). The challenge for researchers is finding an appropriate way to infer the actual network of interactions from data drawn from experiments, mined from databases, and sourced from extant publications.

of the most computationally rich means of determining regulatory relations—namely, Bayesian procedures—only those genes that are found to show a strong form of stochastic dependence have their representative nodes linked in the graph.⁷ In this way, network graphs serve to represent the key mechanistic features: nodes stand for the networks components and directed edges stand for their functional interactions while indicating direction of causal influence.

Initial models may not accurately depict an underlying mechanism, but instead present ‘how-possibly’ explanations of a phenomenon. They show, given the available evidence, researchers’ hypotheses for how a phenomenon could arise. There may be multiple competing how-possibly mechanisms for a single phenomenon. During schema construction, constraint-based reasoning helps researchers pare down the space of models by eliminating those that are deemed unrealizable on the basis of accepted physical or biochemical principles. Available models are then evaluated and tested in an effort to better represent the mechanism.

These constructive and evaluative techniques come together in systems biology when network graphs inferred through statistical methods are analysed and subjected to algorithmic or manual selection procedures. Authors attempt to give an even more accurate rendering of the actual network by integrating additional data about cell structures, such as analyses of gene location.⁸ Such a practice is indicative of the merging of top-down approaches (constructing approximate networks based on large-scale data describing functional relations) with bottom-up approaches (building up from detailed data describing molecular components of a system) (De Backer *et al.* [2010]).

Since an inferred network graph extends beyond the observed data, it allows for the prediction of a large number of regulatory interactions between genes. Computational analysis of the graph allows for the identification of numerous

⁷ Bayesian networks are those in which nodes X_i and X_j , which represent molecules (genes, proteins, metabolomes, and so on), are only connected by an edge if the corresponding molecules’ activities are correlated and, knowing the behaviour of all other relevant molecules and subsets thereof in the system, the behaviour of X_i still yields additional information about X_j (Markowitz and Spang [2007]). Formally, the relation of conditional independence (written $X \perp Y | Z$) is given when the following evidential probabilities hold: $P(X=x, Y=y | Z=z) = P(X=x | Z=z) \cdot P(Y=y | Z=z)$. Nodes X_i and X_j in Bayesian networks are thus only connected by an edge when their observed activity is correlated and one has:

$$\sim (X_i \perp X_j) | X_S \text{ for all } S \subseteq V \setminus \{i, j\},$$

where V is the complete set of nodes, S is a subset of nodes, and X_S is the collective activity of this subset. This technique allows for the construction of graphs with directed edges showing determinate pathways of influence.

⁸ These determine the DNA binding sites of proteins, providing physical evidence for regulatory relations between a gene that produces a given protein and those genes bound by the protein. Such information can be incorporated into model selection by giving no weight to those models that fail to include edges required by location data to be present, or vice versa—a form of constraint-based reasoning applied to automated modelling procedures (Hartemink [2002]).

structural features. For example, Bechtel and Abrahamsen ([2013], p. 15) mention topological analyses that indicate that many biological networks are characterized by ‘small-world organization’ in which there are a number of distinct, highly connected clusters with fewer links in between them; these can operate with lower energy costs than highly connected networks and ‘allow for specialized sub-populations that differ from the overall population’ ([2013], p. 720). Likewise, network analysis allows researchers to note the presence of ‘network motifs’—sub-systems such as feed-forward or feedback loops whose recurrence may play a functional role in the overall structure (Alon [2006]; Albert [2007]). These and other global patterns of connectivity among network components contribute to inferences regarding system behaviour.

The fit of the organization that the schema imputes to the mechanism can be further tested by interventions; precisely altering the mechanistic process at one point should yield predictable alterations in those points that are thought to be causally downstream. Such experiments enable one to update a schema—further specifying the organization and functional role of the mechanism’s components—or consider more serious revisions (Craver and Darden [2013], Chapter 8). Obtaining a better representation of the functional organization of the actual network often requires the implementation of dynamic modelling. Models of this sort supplement an interaction network graph with functions describing how the state of a given node depends on the state of those to which it is connected. This may be achieved by assuming that the regulatory network can be represented in terms of a system of differential equations and measuring its response to perturbations (as in Gardner *et al.* [2003]).⁹

As Bechtel and Abrahamsen ([2013]) have argued, there is a close relation between, on the one hand, the parameters and variables of the differential equations and, on the other hand, the mechanistic entities and activities being modelled. In the case of network dynamics, variables can stand for the individual expression rates of the genes (equivalently, the concentration of different mRNAs) and parameter coefficients describe the regulatory interactions between the genes (Gardner *et al.* [2003]). Schematizing the system-level

⁹ In this approach, the dynamics of the system under perturbation are described using a system of ordinary differential equations, $dx/dt = Ax + u$. While the authors recognize that relations between networks components are typically non-linear, they reason that a system near a steady-state point is stable enough to be approximated by linear equations. For example, x is an n -dimensional vector, the n components of which stand for the concentration of different mRNA; u is a vector representing an external perturbation, offsetting the concentration of one or more distinct mRNA sequences; and A is an $n \times n$ matrix of parameter coefficients that weigh how much the presence of each mRNA depends on the presence of the others. By assuming that the system is perturbed from a steady-state point (so $dx/dt = 0$), researchers can make a series of small determinate perturbations (u), measure the resulting levels of mRNA (x), and solve $Ax = -u$ for A . After data has been gathered from numerous perturbations, multiple linear regression is applied to numerically determine the correlation coefficients relating gene concentrations (A).

dynamics generated by a cellular network provides a model of how these underlying mechanisms collectively function to maintain observed states or behaviours of the organism by showing how the combined interactions of network components tend toward certain system-level equilibria.

3 Are Mechanisms Insufficient?

Thus one can detect a methodology among systems biologists that can be subsumed under the general strategies of mechanistic science. A phenomenon is characterized: the robust, self-regulating behaviour of living systems—single cells, in our case—with respect to their environment. This is done by integrating knowledge of distinct intra-cellular processes with the assumption that the interlocking effects of their combined activity regulate the cell. Due to the complexity of the mechanisms and the difficulty in mentally simulating constitutive interactions, it is supposed that any reasonably accurate schema has to incorporate advanced mathematical modelling techniques. A network graph is inferred from data tracking component interactions. Because of relative limitations in data and computational technologies (there are far more genes than measurement samples; there are multiple models that fit the data), the how-possibly schemas must be analysed in terms of correctness, that is, how realistic they are for a biological system. This is done by incorporating informed assumptions (for example, economy of regulation (cf. Albert [2007], p. 3332)) and integrating additional information (for example, genomic location data) into the process of model selection. Finally, schemas are evaluated and further specified by testing predictions derived from network inference analysis and determining causal networks through perturbational interventions.

A dynamic model of the whole system, given by a system of differential equations, can be developed from these interventions. Just as suggested as a strategy for constructing a mechanism schema (Craver and Darden [2013], Chapter 5), systems biologists reason analogically, predicting system-level properties and dispositions from the dynamic systems encountered in other fields, such as engineering and control theory (drawing on (Ljung [1999]), for instance). Thus a representational schema of the underlying aspects (given by a network graph) and maintaining aspects (given by the system-level dynamics developed from this network) of the mechanism or mechanisms is constructed, which can itself be tested by further experimental intervention. Throughout the field, careful researchers do not appear to lose sight of the fact that they are detailing the active organization of specific entities and the interactions between them.¹⁰

¹⁰ Consider Albert ([2007], p. 3336), for whom ‘one needs to know (1) the identity of the components that constitute the biological system; (2) the dynamic behavior of these components

The claims made here and in the previous section constitute a response to point (1) of the introduction. In contrast to the linear, sequential notion of mechanisms, systems biology can be understood as a practice of constructing and evaluating schemas for complex mechanisms that underlie and maintain the phenomena of self-regulating states and behaviours in organisms. I can now address point (2) regarding the downward-looking characterization of mechanistic explanation. Silberstein and Chemero ([2013], p. 961) state, ‘localization and decomposition are universally regarded as the *sine qua non* of mechanistic explanation’. It is asserted that mechanistic explanations take the intrinsic properties of a system’s components as the basis from which the overall behaviour is derived. This view of mechanistic explanation as essentially reductionist allows authors to frame complex dynamics as outstripping its resources; in dynamically complex systems ‘most facts about the nature of these components as well as their initial arrangement have no bearing on the complex system behavior one wants to explain’ (Kuhlmann [2011], p. 4). The reasons for this are two-fold: First, dynamics at the level of the whole system have properties and functions (sometimes called ‘emergent’) that result from the global structure of its components’ interaction network. It follows that these properties and functions are, in a way, distributed throughout the entire system and cannot be located in one place or another. Second, the properties of network components—gene expression rates, rates of enzyme activity, and so on—cannot be adequately described through intrinsic features, but are determined holistically as the result of their place in an entire network. Because mechanistic explanation is based on localization of functions in spatial components, because its explanations depend on the intrinsic properties of these components, and because such a task is not possible in the analysis of dynamically complex systems, it is argued that any explanation based on this method must not be mechanistic.

Drawing on (Craver and Darden [2013]), I have suggested an alternate account of mechanistic explanation as rooted in the elaboration of a schema that adequately describes the manner in which the organization of entities and activities constitute and maintain a phenomenon. Their account fits well with the intuition that there is no need to give a complete characterization of every aspect of every component when one is seeking to explain a higher-level organization; this is the pragmatic element of mechanistic explanation. Against the reductionist notion of mechanisms, Craver ([2007]) has argued that articulating what he calls the ‘active organization’ of systems is a crucial norm of mechanistic explanation. In contrast to spatial and temporal organizations, the active organization of a mechanism is described in terms of

(i.e. how their abundance or activity changes over time in various conditions); and (3) the interactions among these components’—all of which are basic mechanistic details.

the functional role that each entity plays with respect to one another. In particular, the concept indicates the relevance of cooperative or inhibitory interactions between constituents to the net behaviour of the system. Mechanisms, as Craver notes, ‘are not mere static or spatial patterns of relations, but rather patterns of allowance, generation, prevention, production, and stimulation’ (Craver [2007], p. 136).

It is because mechanisms do not just function in a single context but interact with the larger systems in which they are situated that such patterns can be identified at different levels of organization. Observing these patterns may not depend directly on the specification of the activities and entities that comprise a component mechanism. This is clearly the case even in examples that are undeniably mechanistic. For example, there is no need to refer to the properties of every individual cell when explaining how the heart pumps blood (Craver and Darden [2013], Chapter 7). It is readily understood without any apparent change to the character of the explanation that tissues behave in ways that their constituents cannot and that these constituents are affected by this behaviour. Similarly, it is understood that any attempt to explain a phenomenon like the heart pumping blood through the activities of individual cells alone would be overly complicated and serve only to obscure the pattern of interest. As I said before, we must bear in mind that mechanistic explanation is orientated by a pragmatic interest in the particular phenomenon of concern and the levels of organization within which its mechanism or mechanisms are individuated.

Philosophers have noted that certain real patterns of phenomena only become salient under certain conditions of abstraction (Dennett [1991]). Scientific models may employ abstraction in order to more accurately characterize the core causal properties underlying a phenomenon of interest (Cartwright [1989], Chapter 5; Weisberg [2007]). If mechanism schemas are kinds of models, as I have suggested, then we should not be surprised to see similar techniques at work in their construction. This point has been made explicitly by Arnon Levy. With Bechtel he writes that when considering the dynamic behaviour of mechanisms, ‘abstract models, such as models of connectivity [...] highlight the features of that specific system that make a difference in it—namely, its pattern of internal causal connections’ (Levy and Bechtel [2013], p. 259).¹¹ Further, it is hard to see how a component’s properties being affected by its role in a higher level causes trouble for mechanistic explanation. Mechanists can recognize the independence of causal relations

¹¹ In citing their work, I do not mean to wholly endorse some of Levy and Bechtel’s ([2013]) conclusions. In particular, I think there are reasons to be suspicious of the sufficiency of fully general explanations of mechanistic phenomena. As I will attempt to show with the case of Braillard and chemotaxis, explanations relying on general formal principles alone deprive us of the ability to sufficiently explain any particular phenomenon, even in systems biology.

within a given level from the constitutive relations holding between upper and lower levels. Their explanations may thus be upward- as well as downward-looking, aiming to characterize the activity of a mechanism in terms of its functional role as a constituent of a higher-level mechanism, especially when ‘the behavior of this system determines the causal factors impinging on it in a systematic manner’ (Bechtel [2008], p. 156).

Biologists may often work at a high degree of abstraction when considering system-level dynamics, but what they are doing is not different in kind from other mechanistic strategies—they draw on data to develop schemas that are then evaluated internally and through experimental interventions. While systems biologists studying cellular networks clearly acknowledge the constitutive role of the detailed processes involved in individual instances of protein synthesis, RNA transcription, and the like for the maintenance of a cell, they can also abstract away from these details when considering the regulative roles that these entities play in a higher-level system. Moreover, it is often recognized that current techniques are insufficient for the ultimate explanatory task at hand:

To understand the complexity of living cells future research will need to build models including all these layers [genomic, transcriptomic, proteomic, and so forth]. Statistical inference on parts of the system will not provide the mechanistic insights functional genomics is seeking for. (Markowitz and Spang [2007]. p. 13)

A degree of functional localization is possible in the case of cellular networks. By means of network analysis, researchers are able to identify distinct functional roles with sub-network clusters and interaction motifs. Contrary to the claim that localization can be wholly disregarded, authors argue for the use of location data to assist in selecting an accurate model (Herrgard *et al.* [2004], p. 72).¹² It is this bi-directional aspect of schema construction that provides the grounds for systems biological explanations.

4 Design Explanations and Bacterial Chemotaxis

In point (3) of the introduction, I noted that the modelling methods of systems biologists have inspired authors to posit distinct, non-mechanistic explanatory strategies in the field. In opposition to the compatibility with mechanistic

¹² As an addendum to this argument, I’d like to point out that there are several articles that try to show that emergence is compatible with mechanistic philosophy: ‘The cases where macromolecules are not separable from their environment would lead to strong emergence. We would here suggest that it will be possible to make all essential properties of living organisms emerge from silicon-cell-type models. This then implies that all functional properties of living systems come from weak emergence. We base this conjecture on the experience that free-energy transduction, gene expression, cell cycling and developmental biology can be generated by such models (cf. www.siliconcell.net)’ (Westerhoff and Kell [2007], p. 58, see also Booger *et al.* [2005]).

science that I have been claiming, critics argue for an alternate notion: structural or design explanation. An important feature of these explanations is that they are not causal. According to Braillard ([2010], p. 56), design explanations reveal functional constraints for the systems to which they apply: ‘What is at stake is that some specific design principles might be necessary for a biological function to be performed or for the entire system to be able to exist in a changing environment independently from the evolutionary path taken’. Since functional constraints are not the product of evolution, they ‘have their origin in the functioning of the whole system in the context of some environment’.

As an example, Braillard refers to a design explanation given for an aspect of bacterial chemotaxis. Chemotaxis is the process by which motile bacteria sense changes in their chemical environment and move to more favourable conditions (Bren and Eisenbach [2000]). Motion is directed towards areas with a greater concentration of particular chemoattractants or a lesser concentration of repellents. For organisms like *Escherichia coli*, motion occurs by means of alternating clockwise and counter-clockwise rotation of a flagellum, cycling between directed ‘runs’ and redirecting ‘tumbles’. The network of molecules that compose the chemotactic signalling pathway mediating chemical environment and flagellar motion in *E. coli* has been studied since the work of Julius Adler in the 1960s (Baker *et al.* [2005]). It is now one of the best-understood networks of its kind, and homologues of the genes known to encode its components are found in most motile bacteria (Wadhams and Armitage [2004]; Porter *et al.* [2008]). Add in its relative simplicity with regard to the number and kind of interacting elements and we have excellent reason to use this as a case study for how cellular networks are explained in systems biology.

Figure 2 is one of many similar diagrammatic representations of the chemotaxis network in *E. coli* that are found in scientific texts. It is composed of transmembrane chemoreceptors (methyl-accepting chemotaxis proteins, or MCPs) spanning the cell membrane, several intra-cellular chemotaxis proteins (Che’s), and a flagellar motor. The interactions are as follows: CheA is able to draw phosphate molecules off of ambient ATP in the bacterial cell, a process called autophosphorylation. These can then phosphorylate CheY or CheB, passing the phosphate to one of these proteins. Phosphorylated CheY, CheY_P, is able to bind to the flagellar motor, causing it to reverse its motion from the default counter-clockwise to clockwise rotation. This induces tumbling in the bacterium, altering the direction of its next run.

CheW links the CheAs to MCPs to form a molecular complex such that, when an environmental attractant binds to an MCP,¹³ this brings about

¹³ These are mainly amino acids, sugars, and oxygen (Porter *et al.* [2008]).




Figure 2. The phosphorelay signalling network governing flagellum activity in *E. coli*.

a conformational shift in the receptor that is transmitted through CheW, ultimately suppressing the CheAs' capacity for autophosphorylation. As the concentration of environmental chemoattractants increases, more bind to the receptors, reducing the frequency of reorientations.¹⁴ Thus the concentration of CheY_P serves as an internal measure of the bacterium's environment, inducing it to engage in longer runs according to the chemical gradient of its environment, thereby promoting movement towards areas of higher attractant concentration. Several molecular mechanisms are in place to mitigate these longer runs so that the bacterium may adapt to new environmental concentrations. CheRs act on a cytoplasmic signalling domain of the receptors, adding a methyl group that alters the protein's conformation. This increases CheA activity, thereby counteracting the effect of attractant binding. This is balanced by a negative feedback loop: CheBs demethylate receptors, and do so at a much higher rate when phosphorylated by CheA, counteracting increases in CheA activity caused by CheR. Lastly, while CheR and CheB intervene on receptors at the beginning of the signalling chain to effect signal adaptation, CheZ promotes signal termination by intervening near the end. It acts on CheY_{Ps}, dephosphorylating them so that they can no longer bind to the motor.

While this schema allows for a thorough characterization of the individual components and interactions of the network, it fails to provide a rich sense of how multiple interactions take place and affect one another as a dynamic system. In particular, it cannot fully account for a noteworthy characteristic observed in bacterial chemotaxis: the behaviour of the flagella can adapt to changes in environmental conditions with near perfect precision. That is, flagellar activity remains extremely consistent over large variations in

¹⁴ As should be expected, environmental repellents have the reverse effect.

concentration of a surrounding chemoattractant or repellent as long as this concentration is homogeneous. This behaviour, termed ‘robust perfect adaptation’, allows *E. coli* to remain sensitive to small gradations in its chemical environment regardless of the overall concentration.

Differential equations allow for a more exact understanding of synchronic network interactions. Using these, Barkai and Leibler ([1997]) presented a model in which robust adaptation is achieved in such a way that the action of internal proteins that modify system activity is determined by the net activity of the system itself:

A mechanism for robust adaptation [...] can be obtained when the rates of the modification and the reverse-modification reactions depend solely on the system activity, A [the total amount of active proteins], and not explicitly on the concentrations E_m and E [proteins whose presence depends on environmental factors]. This system can be viewed as a feedback system, in which the output A determines the rates of modification reactions, which in turn determine the slow changes in A . (Barkai and Leibler [1997], pp. 915–6)

According to the Barkai–Leibler model, then, robust perfect adaptation does not depend on specific values of internal parameters, but is built into the network structure itself.

That this is the case in *E. coli* was subsequently confirmed experimentally (Alon *et al.* [1999]). Authors reported that varying the concentration of receptors CheB or CheY altered adaptation time and steady-state tumbling frequency, but the bacteria would still reliably return to this steady-state behaviour over a range of attractant concentrations. The Barkai–Leibler model was later analysed by Yi *et al.* ([2000]). They determined, by means of the analytical techniques of control theory, that the model exhibited a precise level of control over system output (flagellum activity) due to its inclusion of a specific control structure, an integral feedback loop (Figure 3).

Braillard takes this to mean that the particular organization of the chemotaxis pathway governing flagellar behaviour in *E. coli* is explained by the identification of a feedback loop in the Barkai–Leibler model because only this could produce robust perfect adaptation: ‘Furthermore, and this is crucial, [the researchers] claim that integral feedback control is not only sufficient but also necessary for robust perfect adaptation’ (Braillard [2010], p. 49). Yi *et al.* ([2000], p. 4652) claim not only that the structure is necessary for robust perfect adaptation, but that ‘if [Barkai and Leibler’s] specific model is later found to be contradicted by experimental data, another mechanism implementing integral feedback is likely to be present’. But this is not an independent explanation. While the discovery of integral feedback helps to explain how the flagellum behaviour could alter so smoothly, it is only the beginning of the




Figure 3. A block diagram of integral feedback control. u is the input for a process with gain k . The difference between the actual output, y_1 , and the steady-state output, y_0 , represents the error, y . Integral control arises through the feedback loop in which the time integral of y , x , is fed back into the system (Yi *et al.* [2000]).

explanation. Knowledge of the specific molecular components, their network of interactions, and how those interactions are mediated—in other words, mechanistic details—is all crucial to precisely characterizing how exactly this behaviour can be produced in this organism.

Take, for instance, a thermostat that can precisely adjust the activity of a heater based on a continuous temperature reading. From the perspective of control theory, there is no difference between this behaviour and that of the *E. coli*'s flagella insofar as they both involve integral feedback control.¹⁵ If identifying the dependence of robust adaptation on integral control counts as explanation, then we must say that bacterial chemotaxis and thermostats are here explained in the same way. Such is the generality of design explanation that Braillard ([2010]) celebrates, but surely we want our sciences to employ a form of explanation that enables us to distinguish between these systems. This can only occur if models are understood to be based on and correspond to the actual organization of such systems. Only then can we understand how underlying mechanistic interactions give rise to certain abstract patterns of behaviour and are capable of maintaining them. Otherwise, we are simply characterizing *abstracta* and their properties, not explaining particular phenomena.

Braillard ([2010]) holds that there are general design principles that can be observed in the functional relations between a whole system and its environment from which one can explain why the system is structured as it is. As his commentary on Yi *et al.*'s ([2000]) work shows, Braillard ([2010]) holds this explanation to reside in the necessity (or high preferability) of a particular design for a particular biological function to occur, a necessity that he claims can hold without reference to evolutionary theory. The idea is that, regardless

¹⁵ This happens to work in thermostats because temperature is proportional to the integral of heat.

of the exact process by which an organism comes to have the functions that it does, these functions can only be realized if they constrain the structure of the organism in ways that accord with general design principles.

Braillard's ([2010], p. 59) appeal to the 'transcendence of design principles' over their material instantiation thus appears to be the claim that the need for organism structure to conform to design principles comes from the necessity of design principles themselves, just as natural laws have sometimes been described as 'bestowing' necessity. Given an environmental function, design explanations are taken to show 'why a certain structure is present' (Braillard [2010], p. 50). They achieve this by 'point[ing] to synchronic and non-causal functional dependences between the system's structure and its environment' (Braillard [2010], p. 51). But this turns out to be trivial, simply explaining structure with structure. It is hardly clear that noting the contribution of organismal structure to environmental functions can sufficiently explain why that structure is present to begin with.

Consider our case study: the apparent explanation for why the chemotactic network is organized such that it must contain an integral feedback loop is that, by design constraints, this organization is necessary to achieve the observed function of robust perfect adaptation. Such an explanation need not specify the actual system of interacting molecules in the organism under consideration, but is satisfied by showing that a general principle is hereby instantiated. Yet, foregoing evolutionary reasoning, the only available explanation for why this particular organism possesses integral feedback—why the principle is instantiated in this case—is because it exhibits robust perfect adaptation. One cannot help but feel that this ricochet from the phenomenon of the organism's environmental function to general design and back is vacuous, failing to fully explain either.

I grant that the work of Yi *et al.* ([2000]) does not look like mechanistic explanation. They analyse a model for a phenomenon, identify a necessary structural feature, and predict that models for similar phenomena will require this structural feature. But it is not clear that they intend to give a full explanation of why the chemotactic network is structured in the way that it is. In fact, they appear to be working at one particular stage of mechanistic science, namely, schema construction: 'The recognition that integral control is responsible for the robustness of perfect adaptation in the Barkai–Leibler model allows one to evaluate the importance of the various assumptions of the model' (Yi *et al.* [2000], p. 4650). From this they determine constraints for the construction of future mechanism schemas: 'When combined with biological realizability, this may greatly constrain, on the basis of external behavior, the *possible internal mechanisms* that can be used to achieve the observed behavior' (Yi *et al.* [2000], p. 4651, emphasis added). Analysing models allows researchers to prune down the space of how-possibly mechanisms,

contributing to a more exact determination of the causal and functional organization of the relevant entities and interactions.

It is important to retain this sense that explanations in biology are constantly orientated towards particular phenomena. Models serve as representations, more or less abstract, of the real systems in which these phenomena occur. Observation of model properties alone does not guarantee that they accurately represent the properties of the system. While experimentation on *E. coli* did show that adaptation is robust under internal parameter change, Barkai and Leibler's model does not settle all of the questions regarding the organization and functionality of the chemotactic network. First, it should be noted that, while highly abstract, their model is clearly developed from a prior mechanistic understanding. This much is acknowledged by the authors in their paper with Alon: 'Once the components of a biochemical network are isolated and their interactions characterized, the mechanisms of the network's functioning can be addressed' (Alon *et al.* [1999]). Certain mechanistic details are also assumed in the model itself. For instance, it depends on the assumption that CheB only demethylates actively bound receptors and that receptor-attractant binding happens on a faster timescale than methylation/demethylation (Tindall *et al.* [2008]).

Recently, there has been reason to doubt the notion that Barkai and Leibler's model sufficiently characterizes the active organization of the *E. coli* network. Modellers aim to account not only for a system's capacities for adaptation and robustness, but also its sensitivity and gain (Tindall *et al.* [2008]). It cannot be assumed that these capacities function independently of one another, and so a failure to correctly model one may produce inaccuracies in another. For example, Barkai and Leibler's model left out phosphorylation interactions between CheA and CheY or CheB, which are thought to play a role in the system's sensitivity and gain (Tindall *et al.* [2008]). Attempts to expand on their model by including these interactions revealed that if, as Barkai and Leibler's model requires, CheR is active at saturation levels and acts on both bound and unbound receptors, then perfect adaptation cannot take place (Morton-Firth *et al.* [1999]). Other authors, citing all of the work that Braillard takes to support his claims, contest the importance of perfect adaptation:

Each individual in a population of bacteria exhibits its own characteristic flagella switching frequency. This variability is often not addressed in discussions of 'perfect adaptation' of the chemotaxis system. Even in a constant environment, the behavior of an individual fluctuates too much over time to indicate a perfectly adapted state. (Baker *et al.* [2006])

Braillard claims generality as a distinct virtue of design explanations contra mechanistic ones, yet there is reason for scepticism here as well. He fails to

explicitly point out that the network in Barkai, Leibler, and Yi *et al.*'s work is that of a single species, *E. coli*, and not a model of bacterial chemotaxis *tout court*. Further, the article (Rao *et al.* [2004])—cited in Braillard's claim that 'different species have mechanisms that are different, but the design principles are the same'—uses a model of the bacteria, *Bacillus subtilis*, that is openly based on the Barkai–Leibler model of *E. coli*. Thus, that both adapt to the environment robustly is deemed a 'somewhat unsurprising result', and the model fails to shed light on why *B. subtilis* has been observed to be the far more robust of the two (Tindall *et al.* [2008], p. 1540). One possible explanation for this is the added complexity of its signalling pathway.

Studies of the *B. subtilis* network (Figure 4) show that it employs three different adaptation systems (Rao *et al.* [2008]). Interestingly, removal of any two of these has been found to severely impact bacterial motility, suggesting that it is their integrated, coordinated activity that results in robust adaptive behaviour. It appears that the control system, observationally comparable to that found in *E. coli*, requires a much more involved, distributed organization in this case.¹⁶ The authors note that 'An open question is why does *B. subtilis* need three adaptation systems when organisms such as *E. coli* require only one' (Rao *et al.* [2008], p. 484). For these biologists, noting similar design principles—as Rao and co-authors did ([2004])—does not exhaust the question of why organisms have the particular organization that they do. Yet this is precisely the question that Braillard takes design principles to answer.

More generally, the literature since Barkai and Leibler's work speaks to a burgeoning awareness of the diversity in chemotactic systems. With respect to *E. coli*, 'it is becoming increasingly apparent that chemotaxis in other bacteria, although based on similar principles, might be far more complex' (Wadhams and Armitage [2004]). Others write, 'Although the chemistry of the signaling reactions is conserved across all known chemotaxis pathways, the way in which these reactions are assembled to produce functional pathways differs between species' (Porter *et al.* [2011]). Divergences like those between the *E. coli* and *B. subtilis* networks become even more pronounced when one considers organisms that are, unlike these two, non-enteric. While *B. subtilis* differs in having multiple adaptation systems, many bacteria have multiple chemotaxis pathways. The non-enteric bacteria *Rhodobacter sphaeroides*, for instance, is estimated to have three, including one triggered by wholly internal, cytoplasmic chemoreceptors (Porter *et al.* [2011]). Cross-talk between these pathways and other sensory systems allows the bacterium to integrate complex information regarding its external environment and internal metabolic

¹⁶ Where *E. coli* realizes adaptation through the CheA-activating methylation of receptors by CheB_P and deactivating demethylation by CheR, methylation both activates and deactivates CheA in *B. subtilis* depending on the receptor binding sites that methyl groups are shuffled between—a process that cannot yield robust adaptation on its own.




Figure 4. The phosphorelay signalling network governing flagellum activity in *B. subtilis*.

state in determining motile behaviour. I take such diversity to raise doubts about the claim that highly general design principles alone will sufficiently answer the question that Braillard takes them to, namely: ‘why it is this particular organization and not another one?’ (Braillard [2010], p. 58).

In sum, Braillard sees systems biology as engaging in explanatory projects that deviate from mechanistic concerns. Instead of seeking explanations of how a phenomenon arises through an account of constituent entities, their activities, and their organization, design explanations use mathematical abstractions to locate general principles determining the structural organization of living creatures. I have questioned the explanatory adequacy of this approach and have given some reasons to doubt its ability to give an independent unifying account of diverse species’ organizations. On the contrary, I find the strategies and orientation of research among systems biologists to largely support the claim that they are engaged in mechanistic inquiry rather than treating model analysis as a separate practice. As one group of authors puts it:

What is the future of mathematical modeling in helping to understand bacterial chemotaxis? In order to answer this question, we reflect upon the exact goal of understanding bacterial chemotaxis systems. Is it not to elucidate the *mechanisms* of sensing and moving in order that we may be able to predict the behavior of bacterial chemotactic systems in the environment? (Tindall *et al.* [2008], emphasis added)

The analysis of mathematical models is compatible with mechanistic explanation not because it is complementary to it, but because it is part of the general

strategy of schema construction and evaluation. This is echoed by the same authors who gave the control theoretic analysis of the Barkai–Leibler model: ‘A promising aspect of this broader theory is in providing further necessity results to help biologists greatly narrow their search for specific mechanisms’ (Yi *et al.* [2000], p. 4652).

Rather than affirm the explanatorily vacuous ‘transcendence of design principles’, we may note their positive role in enabling certain inferences about the active organization of a mechanism’s components. Assuming that design principles help us systematically predict functional constraints in a system’s organization, one might consider an alternate strategy of using them as evidence in constructing a plausible mechanism schema. Since constraint-based reasoning is regularly employed in schema construction, specification of necessary limitations on the structural possibilities of a mechanism would seem to work well here. Kuhlmann ([2011], p. 21), although he also ascribes to the functionalist language of design explanations, can help me make this point. He states that ‘the reason why the same explanatory [design-based] strategies can successfully be applied [in diverse fields] is the observable fact that there are structural similarities in the dynamics of compound systems with completely different kinds of subunits’. Again, I am inclined to think that Kuhlmann is not really describing a form of explanation here, but a stage in the logic of discovery.

Reasoning using abstract structural analogies is rightly identified as a useful and sometimes powerful heuristic for schema construction, but explanation cannot stop there. Fascination with the apparent autonomy of design should be tempered by a realization that abstract principles bar us from explaining the actual workings of any given system of sufficient complexity. Knowing that diverse biological processes employ integral feedback without knowing the functional role of the entities involved prevents researchers from being able to understand the most basic effects of intervention into the system. Systems biologists are concerned not only with the possible ways a structure may be realized, but how it is actually realized in a living system. As a consequence, they are aware of how the field is still afflicted with a poverty of data and complications involving the construction of accurate models. They are aware, that is, of how much more work is needed before they possess full-fledged explanations of the phenomena of life.

5 Conclusion

I have used cellular regulatory networks as an example to show how the type of model-based reasoning involved in systems biology corresponds to mechanistic strategies. In response to point (1) of the introduction, I claimed that cellular networks are best understood as complex mechanisms that underlie

and maintain phenomena. In response to point (2), I argued that this mechanistic view of cellular networks is compatible with the forms of abstraction and holism used to analyse them. And in response to point (3), I argued that it is more plausible to think that this programme is guided by mechanistic explanatory strategies than by an entirely separate kind of explanation. Due to the sheer complexity of the subject matter and the current impossibility of collecting sufficiently precise data, much of the work in the field centres on techniques of adequate modelling. What critics refer to as an alternate form of explanation by design is better understood as a contribution to the analogical and constraint-based reasoning involved in mechanistic schema construction.

Acknowledgements

My sincere thanks to Lindley Darden for offering her time and patience to read, discuss, and re-read this article. I also want to thank Carl Craver for answering questions via video conference, both authors for making available the manuscript of their collaborative effort, and three anonymous reviewers for their informative and insightful comments.

*Department of Philosophy
Georgetown University
Box 571133, New North 215
Washington, DC 20057-1133, USA
dem67@georgetown.edu*

References

- Albert, R. [2007]: ‘Network Inference, Analysis, and Modeling in Systems Biology’, *The Plant Cell*, **19**, pp. 3327–38.
- Alon, U. [2006]: *An Introduction to Systems Biology: Design Principles of Biological Circuits*, London: Chapman and Hall.
- Alon, U., Surette, M. G., Barkai, N. and Leibler, S. [1999]: ‘Robustness in Bacterial Chemotaxis’, *Nature*, **397**, pp. 168–71.
- Andersen, H. [2012]: ‘Mechanisms: What Are They Evidence for in Evidence-Based Medicine?’, *Journal of Evaluation in Clinical Practice*, **18**, pp. 992–9.
- Baker, M. D., Wolanin, P. M. and Stock, J. B. [2005]: ‘Signal Transduction in Bacterial Chemotaxis’, *Bioessays*, **28**, pp. 9–22.
- Baker, M. D., Wolanin, P. M. and Stock, J. B. [2006]: ‘Systems Biology of Bacterial Chemotaxis’, *Current Opinion in Microbiology*, **9**, pp. 187–92.
- Barkai, N. and Leibler, S. [1997]: ‘Robustness in Simple Biomedical Networks’, *Nature*, **387**, pp. 913–7.
- Beatty, J. [1995]: ‘The Evolutionary Contingency Thesis’, in G. Wolters and J. G. Lennox (eds), *Concepts, Theories, and Rationality in the Biological Sciences*, Pittsburgh: University of Pittsburgh Press, pp. 45–81.

- Bechtel, W. [2007]: 'Biological Mechanisms: Organized to Maintain Autonomy', in F. Boogerd, F. Bruggeman, J. Hofmeyr and H. Westerhoff (eds), *Systems Biology: Philosophical Foundations*, New York: Elsevier, pp. 269–302.
- Bechtel, W. [2008]: *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*, New York: Routledge.
- Bechtel, W. [2011]: 'Mechanism and Biological Explanation', *Philosophy of Science*, **78**, pp. 533–57.
- Bechtel, W. and Abrahamsen, A. [2005]: 'Explanation: A Mechanist Alternative', *Studies in History and Philosophy of Biological and Biomedical Sciences*, **36**, pp. 421–41.
- Bechtel, W. and Abrahamsen, A. [2013]: 'Thinking Dynamically about Biological Mechanisms: Networks of Coupled Oscillators', *Foundations of Science*, **18**, pp. 707–23.
- Bechtel, W. and Richardson, R. C. [1993]: *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*, Princeton: Princeton University Press.
- Boogerd, F. C., Bruggeman, F. J., Hofmeyr, J. S. A. and Westerhoff, H. V. [2007a]: 'Towards Philosophical Foundations of Systems Biology: Introduction', in F. Boogerd, F. Bruggeman, J. Hofmeyr and H. Westerhoff (eds), *Systems Biology: Philosophical Foundations*, New York: Elsevier, pp. 3–23.
- Boogerd, F. C., Bruggeman, F. J., Hofmeyr, J. S. A. and Westerhoff, H. V. [2007b]: 'Afterthoughts as Foundations for Systems Biology', in F. Boogerd, F. Bruggeman, J. Hofmeyr and H. Westerhoff (eds), *Systems Biology: Philosophical Foundations*, New York: Elsevier, pp. 321–36.
- Boogerd, F. C., Bruggeman, F. J., Richardson, R. C., Stephan, A. and Westerhoff, H. V. [2005]: 'Emergence and Its Place in Nature: A Case Study of Biochemical Networks', *Synthese*, **145**, pp. 131–64.
- Braillard, P. [2010]: 'Systems Biology and the Mechanistic Framework', *History and Philosophy of Life Sciences*, **32**, pp. 43–62.
- Bren, A. and Eisenbach, M. [2000]: 'How Signals Are Heard during Bacterial Chemotaxis: Protein–Protein Interactions in Sensory Signal Propagation', *Journal of Bacteriology*, **182**, pp. 6865–73.
- Brigandt, I. [forthcoming]: 'Evolutionary Developmental Biology and the Limits of Philosophical Accounts of Mechanistic Explanation', in P.-A. Braillard and C. Malaterre (eds), *Explanation in Biology: An Enquiry into the Diversity of Explanatory Patterns in the Life Sciences*, New York: Springer.
- Cartwright, N. [1989]: *Nature's Capacities and Their Measurement*, Oxford: Oxford University Press.
- Craver, C. [2007]: *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*, Oxford: Oxford University Press.
- Craver, C. and Darden, L. [2013]: *In Search of Mechanisms: Discoveries across the Life Sciences*, Chicago: Chicago University Press.
- De Backer, P., De Waele, D. and Van Speybroeck, L. [2010]: 'Ins and Outs of Systems Biology vis-à-vis Molecular Biology: Continuation or Clear Cut?', *Acta Biotheoretica*, **58**, pp. 15–49.

- Dennett, D. [1991]: 'Real Patterns', *The Journal of Philosophy*, **88**, pp. 27–51.
- Frigg, R. and Hartmann, S. [2012]: 'Models in Science', in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, <plato.stanford.edu/archives/fall2012/entries/models-science/>.
- Gardner, T. S., di Bernard, D., Lorenz, D. and Collins, J. J. [2003]: 'Inferring Genetic Networks and Identifying Compound Mode of Action via Expression Profiling', *Science*, **301**, pp. 102–5.
- Glennan, S. [1996]: 'Mechanisms and the Nature of Causation', *Erkenntnis*, **44**, pp. 49–71.
- Hanson, N. R. [1970]: 'Is There a Logic of Scientific Discovery?', in B. Brody (ed.), *Readings in the Philosophy of Science*, Englewood Cliffs, NJ: Prentice Hall, pp. 620–33.
- Hartemink, A. J., Gifford, D. K., Jaakkola, T. S. and Young, R. A. [2002]: 'Combining Location and Expression Data for Principled Discovery of Genetic Regulatory Network Models', in R. B. Altman, A. K. Dunker, L. Hunter, T. E. D. Klein and K. Lauderdale (eds), *Pacific Symposium on Biocomputing*, Singapore: World Scientific, pp. 437–49.
- Herrgard, M. J., Covert, M. W. and Palsson, B. [2004]: 'Reconstruction of Microbial Transcriptional Regulatory Networks', *Current Opinion in Biotechnology*, **15**, pp. 70–7.
- Kuhlmann, M. [2011]: 'Mechanisms in Dynamically Complex Systems', in P. McKay-Illari, F. Russo and J. Williamson (eds), *Causality in the Sciences*, Oxford: Oxford University Press, pp. 880–907.
- Levy, A. and Bechtel, W. [2013]: 'Abstraction and the Organization of Mechanisms', *Philosophy of Science*, **80**, pp. 241–61.
- Ljung, L. [1999]: *System Identification: Theory for the User*, New Jersey: Prentice Hall.
- Machamer, P., Darden, L. and Craver, C. [2000]: 'Thinking about Mechanisms', *Philosophy of Sciences*, **67**, pp. 1–25.
- Markowitz, F. and Spang, R. [2007]: 'Inferring Cellular Networks: A Review', *BMC Bioinformatics*, **8**, p. S5, doi:10.1186/1471-2105-8-S6-S5.
- Morton-Firth, C. J., Shimizu, T. S. and Bray, D. [1999]: 'A Free-Energy-Based Stochastic Simulation of the Tar Receptor Complex', *Journal of Molecular Biology*, **286**, pp. 1059–74.
- Porter, S. L., Wadhams, G. H. and Armitage, J. P. [2008]: 'Rhodobacter Sphaeroides: Complexity in Chemotactic Signaling', *Trends in Microbiology*, **16**, pp. 251–60.
- Porter, S. L., Wadhams, G. H. and Armitage, J. P. [2011]: 'Signal Processing in Complex Chemotaxis Pathways', *Nature Reviews Microbiology*, **9**, pp. 153–65.
- Ramsey, J. L. [2008]: 'Mechanisms and Their Explanatory Challenges in Organic Chemistry', *Philosophy of Science*, **75**, pp. 970–82.
- Rao, C. V., Glekas, G. D. and Ordal, G. W. [2008]: 'The Three Adaptation Systems of *Bacillus subtilis* Chemotaxis', *Trends in Microbiology*, **16**, pp. 480–7.
- Rao, C. V., Kirby, J. R. and Arkin, A. P. [2004]: 'Design and Diversity in Bacterial Chemotaxis: A Comparative Study in *Escherichia coli* and *Bacillus subtilis*', *PLoS Biology*, **2**, pp. 0239–52.

- Richardson, R. C. and Stephan, A. [2007]: 'Mechanism and Mechanical Explanation in Systems Biology', in F. Boogerd, F. Bruggeman, J. Hofmeyr and H. Westerhoff (eds), *Systems Biology: Philosophical Foundations*, New York: Elsevier, pp. 123–44.
- Silberstein, M. and Chemero, T. [2013]: 'Constraints on Localization and Decomposition as Explanatory Strategies in the Biological Sciences', *Philosophy of Science*, **80**, pp. 958–70.
- Skipper, R. A. and Millstein, R. L. [2005]: 'Thinking about Evolutionary Mechanisms: Natural Selection', *Studies in History and Philosophy of Biological and Biomedical Sciences*, **36**, pp. 327–47.
- Strogatz, S. [1994]: *Non-linear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*, Cambridge, MA: Westview Press.
- Teller, P. [2001]: 'Twilight of the Perfect Model Model', *Erkenntnis*, **55**, pp. 393–415.
- Tindall, M. J., Porter, S. L., Maini, P. K., Gaglia, G. and Armitage, J. P. [2008]: 'Overview of Mathematical Approaches Used to Model Bacterial Chemotaxis I: The Single Cell', *Bulletin of Mathematical Biology*, **70**, pp. 1525–69.
- Wadhams, G. H. and Armitage, J. P. [2004]: 'Making Sense of It All: Bacterial Chemotaxis', *Nature Reviews Molecular Cell Biology*, **5**, pp. 1024–37.
- Weisberg, M. [2007]: 'Three Kinds of Idealization', *The Journal of Philosophy*, **104**, pp. 639–59.
- Weiskopf, D. A. [2011]: 'Models and Mechanisms in Psychological Explanation', *Synthese*, **183**, pp. 313–38.
- Westerhoff, H. V. and Kell, D. B. [2007]: 'The Methodologies of Systems Biology', in F. Boogerd, F. Bruggeman, J. Hofmeyr and H. Westerhoff (eds), *Systems Biology: Philosophical Foundations*, New York: Elsevier, pp. 23–70.
- Woodward, J. [2008]: 'Explanation, Invariance, and Intervention', *Philosophy of Science*, **64**, pp. S26–41.
- Yi, T., Huang, Y., Simon, M. and Doyle, J. [2000]: 'Robust Perfect Adaptation in Bacterial Chemotaxis through Integral Feedback Control', *Proceedings of the National Academy of Sciences*, **97**, pp. 4649–53.