Strategies of Explanatory Abstraction in Molecular Systems Biology

Nicholaos Jones*†

I consider three explanatory strategies from recent systems biology that are driven by mathematics as much as mechanistic detail. Analysis of differential equations drives the first strategy; topological analysis of network motifs drives the second; mathematical theorems from control engineering drive the third. I also distinguish three abstraction types: aggregations, which simplify by condensing details; generalizations, which simplify by generalizing details; and structurations, which simplify by contextualizing details. Using a common explanandum as a reference point—namely, the robust perfect adaptation of chemotaxis in *Escherichia coli*—I argue that each strategy targets various abstraction types to different mechanistic details.

1. Introductory Remarks. The currently dominant paradigm for understanding explanation in cellular and molecular biology prioritizes mechanism (Craver and Darden 2013; Levy 2013). Leading accounts of mechanistic explanation, while differing in the particulars of their analysis of *mechanism*, agree that mechanistic explanations explain by alluding to mechanisms or models thereof (Machamer, Darden, Craver 2000; Bechtel and Abrahamsen 2005).

There is a growing literature devoted to discerning the scope of mechanistic explanation in biological practice. Some claim to identify explanations that do not allude to mechanisms (Wouters 2007; Huneman 2010; Rice 2015). Others tend to resist making scope concessions, preferring instead to accom-

*To contact the author, please write to: Department of Philosophy, 332 Morton Hall, 301 Sparkman Drive, University of Alabama, Huntsville, AL 35899; e-mail: nick.jones@uah.edu.

†I thank Philippe Huneman for organizing the symposium on explanatory strategies across the life sciences and our cosymposiast Anya Plutynski. For helpful comments, I also thank Lindley Darden, Maureen O'Malley, anonymous reviewers for this journal, and audiences from the symposium, the Alabama Philosophical Society, Mississippi State University, and the Society for Philosophy of Science in Practice.

Philosophy of Science, 85 (December 2018) pp. 955–968. 0031-8248/2018/8505-0018$10.00

modate the putative explanations as mechanistic despite initial appearances, to broaden the scope of mechanistic explanation or the analysis of *mechanism* or else to reinterpret the putative explanations as mere descriptions (Craver 2006; Bechtel and Abrahamsen 2010; Brigandt 2013; Levy and Bechtel 2013).

I set aside questions about what qualifies as an explanation as well as questions about whether only mechanisms—or representations thereof—carry explanatory power. I focus, instead, on *explanatory strategies*, understood as patterns of reasoning directed toward providing explanations (regardless of whether they succeed in doing so). I consider three explanatory strategies from recent systems biology that are driven by mathematics as much as, if not more than, mechanistic detail. Analysis of differential equations drives the first; topological analysis of network motifs, the second; mathematical theorems from control engineering, the third.

Systems biologists interpret these strategies as supplements to the explanatory power of mechanism. My aim is to identify how the strategies differ from each other, rather than how they differ from standard mechanistic explanations or what might unify them in those differences (for which, see Green and Jones 2016). Doing so helps with understanding relations among the strategies, their tactics for integrating or ignoring mechanistic detail, and the explanatory affordances of their mathematical elements.

Central to my analysis is a distinction among three abstraction types: aggregations, which simplify by condensing details; generalizations, which simplify by generalizing details; and structurations, which simplify by contextualizing details. Using a common explanandum as reference point—namely, the robust perfect adaptation of chemotaxis in *Escherichia coli*—I argue that each strategy targets various abstraction types to different mechanistic details. These differences thereby mark how the strategies differ from each other. I begin with the typology of abstraction.

2. Abstraction Typology. I am interested in abstractions as representational rather than metaphysical. Abstractions, as I understand them, are ontologically innocent, so that characterizing features of representations as abstractions over some part of reality carries no implication that abstract features correspond to abstract objects (see also Godfrey-Smith 2009, 47–48; French 2010; Levy and Bechtel 2013, 243). So, for example, representing the relation between a person, a hotel, and a date range as a reservation does not entail that some abstract object, a *reservation*, exists, nor does representing the motions of an object's constituents as the motion of the object's center of mass entail that some abstract object, a *center of mass*, exists.

Levy and Bechtel characterize a representation as abstract insofar as a more concrete representation is possible (2013, 242; also see Craver and Darden 2013, 32). I concur. I understand abstractions as representing only some of the many elements—entities, relations, parameters, activities—as-

sociated with their targets, thereby making apparent the patterns obscured by more detailed representation. I add to these insights that biologists produce (at least) three types of abstraction.

Following Ordorica, I call the first *aggregation* (2015, 163–64). An aggregation represents some relationship among multiple elements of a representational target as a higher-level object, or multiple elements of the target as a single composite object (see fig. 1*a*). Paradigm cases of aggregations include representations of person-hotel-date relations as *reservations*, of costs of services and costs of goods as *costs*, and of the motions of an object's parts as the *motion of a center of mass* (from Ordorica 2015, 164). Aggregations abstract from plurality to individual, ignoring differences among many in order to make salient some integrated unity among the elements of a representational target. They thereby simplify representations by condensing details about representational targets.

Following Pincock, I call the second abstraction type *generalization* (2015, 864; see also Darden and Cook 1994). A generalization represents some element of a representational target as a class of elements, where potential instances of the class might include elements not present in the target (see fig. 1*b*). For example, because the class of solution measures includes all soap-bubble-like surfaces, such as the cellular froth surrounding radiolarian protozoa, representing a soap-bubble surface as a "solution measure" is a generalization (Pincock 2015, 864). Generalizations abstract from an instance to a class thereof, ignoring differences between instances of the class

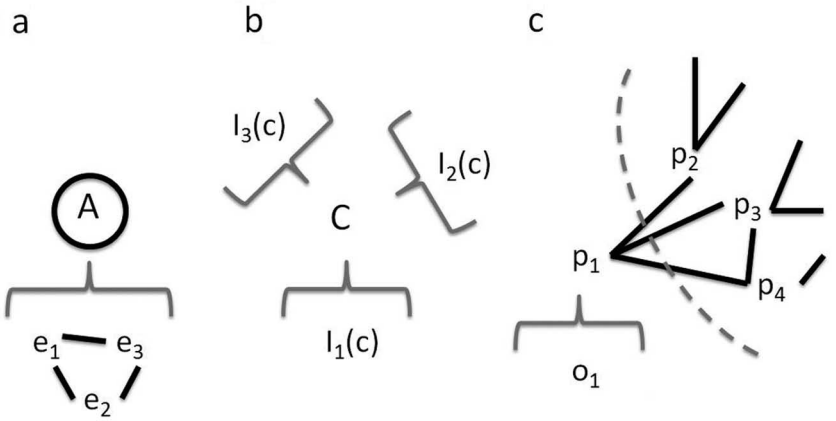

Figure 1. Visualizing abstraction types. *a*, Aggregation A represents elements e_1, e_2, and e_3 (and relations therein) as a single object. *b*, Generalization C represents $I_1(c)$ as a class, instances of which also include $I_2(c)$ and $I_3(c)$. *c*, Structuration p_1 represents element o_1 as a position relating to a larger structure that also includes p_2, p_3, and p_4.

in order to make salient some more general unity. They thereby simplify representations by generalizing from details about representational targets.

I call the third abstraction type *structuration*. A structuration represents some element of a representational target as a position in a structure, such that potential occupants of the position might include elements not present in the target (see fig. 1*c*). I follow Haslanger in understanding structures as "complex entities with parts whose behavior is constrained by their relation to other parts" (2016, 118). Paradigm cases of structurations include representing Barack Obama as president of the United States of America or representing Aineias as son of Anchises and Aphrodite. Structurations abstract to a position in a structure from an occupant of the position, ignoring intrinsic features of the occupant unrelated to its position in order to make salient the occupant's role relative to occupants of other positions in the same structure. They thereby simplify by contextualizing details about representational targets.

I understand aggregations as distinct from both generalizations and structurations, by virtue of being many-to-one, rather than one-to-one, simplifications. I also understand generalizations as distinct from structurations, by virtue of preserving relationship arity: abstractions of unary relationships (wherein a participant relates only to itself) remain unary, of binary relationships (wherein two participants relate to each other) remain binary, and so on. When a generalization is a unary relationship, its instances are unary as well; when it is binary, its instances are binary. By contrast, structurations are always at least binary relations, relating an occupant to another in a common structure. So when a structuration is binary, its occupants are unary. For example, the class of men is a generalization of the man Aineias, and it is arity preserving: "being Aineias" and "being a man" are unary relations. The position of father, by contrast, is a structuration of the man Aineias, because it is not arity preserving: "being a father" is a binary relation, because fathers are always fathers of someone else.

3. Robust Perfect Adaptation of *E. coli* Chemotaxis. My central claim is that various explanatory strategies from recent systems biology differ from each other, at least in part, by virtue of appealing to different abstraction types. I support this claim by considering a case in which multiple strategies target the same explanandum. Doing so minimizes confounds that confuse differences due to the nature of each explanatory strategy with differences due to the nature of each explanatory target (see Andersen 2018). I focus on a particular explanandum known as robust perfect adaptation of bacterial chemotaxis, following others who consider this a paradigmatic target for nonmechanistic explanation (Braillard 2010; Brigandt, Green, and O'Malley 2017; Matthiessen 2017).

3.1. Explanandum Context. *Escherichia coli* (*E. coli*) is popular model organism in biological research. It is very sensitive to small chemical changes over a very large range of background concentrations. It also has a simple and well-understood signal transduction network (Wadhams and Armitage 2004).

E. coli manages two kinds of motion (Berg 2004). It *runs* by rotating its flagellar motor counterclockwise. This aligns all of its flagella into a synchronized bundle, resulting in movement in a straight line for about 1 second. *E. coli* also *tumbles* by rotating its flagellar motor clockwise. This breaks flagellar alignment, and the asynchronized flagella produce stationary changes of direction lasting for about 0.1 seconds. *E. coli* randomly reorient after each tumble. Moreover, while these tumbles occur with regular frequency, *E. coli* with higher concentrations of CheR protein tumble more frequently (Spudich and Koshland 1976).

E. coli's motion in a uniform external environment resembles a random walk. *E. coli* has no ability to control or select its direction of motion, and its straight runs are subject to Brownian motion because of eddies. However, in the presence of a chemical attractant—amino acids such as serine or sugars such as maltose—*E. coli taxis* toward the attractant. This taxi behavior involves less frequent tumbles, leading to longer runs and so gradual motion toward the attractant. (There is an opposite behavior for repellents such as metal ions and leucine.)

The biomolecular mechanism for *E. coli* chemotaxis is well understood. When an environmental attractant (ligand) attaches to a receptor, the receptor lowers the activity of the CheW-CheA protein complex. Less activity from this complex reduces the rate of CheY phosphorylation, which results in less phosphorylated CheY diffusing to the flagella. Because CheY induces clockwise rotation of the flagellar motor, the outcome is less frequent tumbling.

3.2. Explanandum Question. Alon and colleagues have experimental verification that, in the presence of a chemical attractant mixed uniformly into the environment at a constant concentration, *E. coli* chemotaxis is *perfectly adaptive*: after a brief period of decreased tumbling frequency, the frequency of *E. coli* tumbling increases toward and returns to the exact frequency before the introduction of the attractant (Alon et al. 2009). The effect of the attractant, accordingly, becomes entirely forgotten despite its continuing presence.

The biomolecular mechanism for the adaptiveness of chemotaxis for *E. coli* is also well understood. Sometime after a new attractant has been detected by receptors, the lower activity of the CheW-CheA complex induces less CheB activity. This reduces the rate for removing methyl groups from the CheW-CheA complex. Because there is continual methylation of the CheR receptor, CheW-CheA methylation increases. More methylation

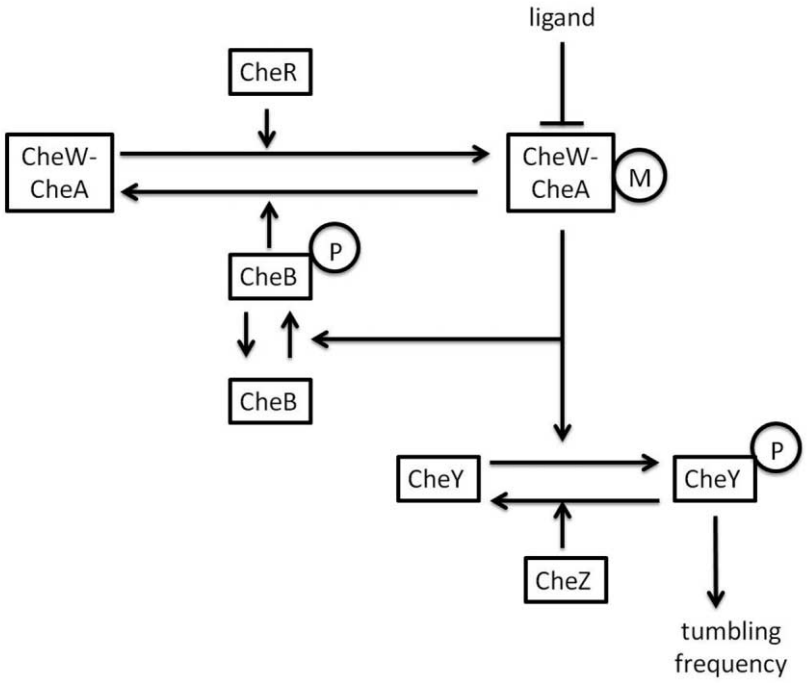

Figure 2. Mechanism sketch for chemotaxic adaptation in *E. coli*. Arrows represent "activation," whereby more activity at the tail produces more activity at the tip. "Flat-tipped" lines represent "inhibition," whereby more activity at the tail produces less activity at the tip.

means more CheW-CheA activity, which in turn induces more CheY phosphorylation. This eventually results in more phosphorylated CheY diffusing to the flagellar motor, which increases clockwise motor rotation and thereby increases the tumbling frequency. Figure 2 visualizes this mechanism (along with some omitted details).[1]

Alon and colleagues (2009) have further experimental verification that this perfectly adaptive chemotaxis of *E. coli* is robust across ranges of CheR concentrations 0.5 to 50 times higher than concentration levels in "wild

1. The CheW-CheA complex induces CheB phosphorylation, and phosphorylated CheB activates demethylation of the CheW-CheA complex. There is spontaneous dephosphorylation of CheB, but this is not part of the feedback for adaptation. There is a further protein, CheZ, that activates dephosphorylation of CheY. But, while important for chemotactic behavior, there is good evidence that CheZ is irrelevant to adaptation (Alon et al. 2009).

type" *E. coli*.[2] This is the explanandum of interest: why is the perfect adaptation of *E. coli* chemotaxis, in the presence of a well-distributed chemical attractant, robust to CheR protein concentrations?

4. Distinguishing Explanatory Strategies through Abstraction Types. Three strategies for answering this question appear in recent systems biology literature. Each targets the same explanandum. Yi and colleagues (2000) cite Barkai and Leibler (1997), while Ma and colleagues (2009) cite both of the preceding. So there is a bibliographical sense in which proponents of each strategy understand themselves as offering alternative, albeit not necessarily incompatible, explanations of the same phenomenon. I consider each in turn, first sketching the general strategy and then identifying its associated abstractions.

4.1. Dynamical Modeling. I call the first strategy *dynamical modeling*. This strategy begins by constructing a chemotaxis network for *E. coli*. This network represents more than the mechanism for *E. coli* chemotaxis, by also including specific biochemical details about when and how relevant proteins affect each other. For example, Barkai and Leibler (1997) construct a model according to which, among many other specifics, CheB demethylates only the active form of the CheW-CheA complex, and CheR works only at saturation. Here is part of the network, simplified through omission of kinetic constants associated with each reaction and followed by translations into English:

$$\bar{E}_m^* E_m^* + B \leftrightarrows \left\{E_m^* B\right\} \rightarrow E_{m-1}^*. \tag{1}$$

$$\bar{E}_m^* E_m^* + R \leftrightarrows \left\{E_m^* R\right\} \rightarrow E_{m+1}^*. \tag{2}$$

$$E_m^u + L \leftrightarrows E_m^o. \tag{3}$$

(Equation 1) Active and inactive forms of the CheW-CheA complex, each with m methyl groups, reversibly bind with phosphorylated CheB to form a bound receptor complex, and the bound complex removes one methyl group from the inactive form of the CheW-CheA complex.

(Equation 2) Active and inactive forms of the CheW-CheA complex reversibly bind with CheR to form a bound receptor complex, and the bound complex adds one methyl group to the inactive form of the CheW-CheA complex.

(Equation 3) A ligand reversibly binds with an inactive form of an unmethylated CheW-CheA complex to form an inactive CheW-CheA complex with a methylated receptor.

2. By contrast, *E. coli*'s adaptation time—the time to return to 50% of its prestimulus tumbling frequency—is not robust to different CheR concentrations because more CheR entails longer adaptation times.

The dynamical modeling strategy proceeds by constructing a dynamical model—typically a set of differential equations—from the reaction network (see Jones and Wolkenhauer 2012). One then demonstrates, via mathematical proof or simulation, that this model predicts perfect adaptation in the presence of a well-distributed chemoattractant for CheR concentration values varying over several orders of magnitude. The demonstration supports the inference that *E. coli* chemotaxis exhibits robust perfect adaptation because of its biochemical specifics.

Bechtel and Abrahamsen (2010) call the product of this strategy a *dynamical mechanistic explanation*. I set aside the issues of whether the dynamical modeling strategy produces explanations and of whether the dynamical models are mechanistic. I also do not attempt to identify the various ways in which Barkai and Leibler's dynamical model—and its associated reaction network—is abstract. Instead, I focus on the ways in which models associated with other explanatory strategies are more abstract and, in particular, on the abstractions present in these other models that are absent from Barkai and Leibler's model.

4.2. Topological Analysis. I call the second explanatory strategy *topological analysis*. This strategy begins by identifying all possible network motifs capable of predicting robust perfect adaptation. These motifs, like the reaction networks for dynamical modeling, are meant to explain the robust perfect adaptation of *E. coli* chemotaxis. Yet, unlike the reaction networks, these motifs are minimal: they contain the fewest possible nodes and links that suffice for robustly perfectly adaptive chemotaxis. The procedure for identifying all possible motifs of this sort is brute computational search. It turns out that there are exactly three, each of which has exactly three nodes and no more than three links (Ma et al. 2009).

The topological analysis strategy proceeds by identifying a reaction network known to predict robust perfect adaptation for *E. coli* chemotaxis. This strategy thereby relies on the dynamical modeling strategy, but only for mathematical results. The biochemical details of the chosen network turn out to be largely irrelevant because the topological analysis strategy proceeds by demonstrating that a *reduced form* of the chosen network is topologically equivalent to one of the network motifs. These reduced forms group mechanistic details together into functionally equivalent modules and consider interactions among modules rather than among entities. They are thereby less sensitive than reaction networks to variations of detail.

Consider, for example, one of the motifs Ma and colleagues (2009) discover for *E. coli* chemotaxis (see fig. 3*a*). The motif represents some unspecified input (*dashed arrow*) activating an "input-receiving" node A; a negative feedback loop in which A inhibits a "buffering control" node B that, in

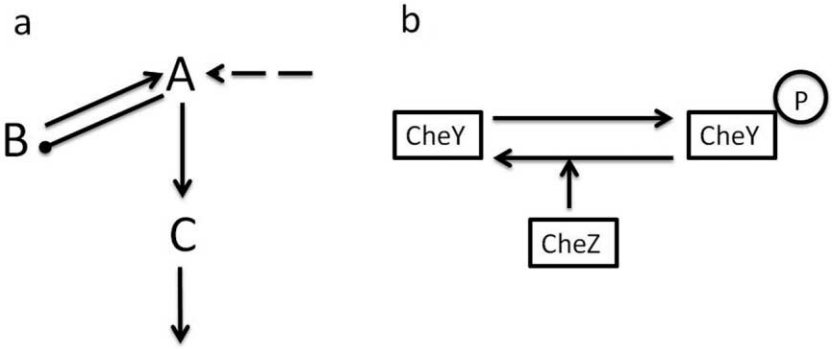

Figure 3. *a*, Motif for *E. coli* chemotaxis. *b*, Subnetwork relative to which the module controlling motor rotation is an aggregation.

turn, activates A; and an "output-transmitting" node C being activated by A and, in turn, producing some unspecified output.

Ma and colleagues demonstrate that Barkai and Leibler's (1997) reaction network for *E. coli* chemotaxis reduces to, and is thereby topologically equivalent to, this network motif. Barkai and Leibler's reaction network involves a ligand activating the methylated CheW-CheA complex, continual methylation of this complex by CheR along with a negative feedback loop in which the complex inhibits its own demethylation by inhibiting CheB phosphorylation, and the methylated complex affecting tumbling frequency by affecting CheY phosphorylation.

The topological analysis strategy infers, from the topological equivalence between a network motif and a reaction network known to predict robust perfect adaptation for chemotaxis, that *E. coli* chemotaxis exhibits robust perfect adaptation because of the topology for its chemotaxis network. Huneman (2010) calls the product of this strategy a *topological explanation*. Regardless of whether such analyses are explanatory, they are topological by virtue of demonstrating some consequence about the topological properties of a network. So even if the mechanistic details of *E. coli*'s chemotaxis network were different, and even if the biochemical specifics of the original reaction network were different, the product of the topological analysis strategy would remain the same, provided that the alternative networks preserve topological equivalence with the originals (see Jones 2014).

The topological model driving this second explanatory strategy is more abstract than the dynamical model driving Barkai and Leibler's explanatory strategy. The topological model contains two aggregations: from a subnetwork of interactions among unmethylated CheW-CheA, CheR, and CheB to a module that controls methylation and from another subnetwork of interactions among CheY and CheZ to a module that controls motor rotation (see

fig. 3*b*). I interpret these abstractions as aggregations because they are many to one: they group together entities and activities working together for some common function.

The model also constructs two structurations from these aggregations: from the methylation-controlling module to buffering control node B and from the motor-rotation module to output-transmitting node C. I interpret these as structurations rather than generalizations. The modules are well defined without reference to their broader role. But Ma and colleagues speak of the nodes relationally, providing functional characterizations for each: C is the "output-transmitting" node, while B is a "dedicated regulation node . . . that functions as a 'buffer' [for A]" (2009, 764). So conceptualized, the abstraction from module to node increases relationship arity, from module to module for an output or for an input receiver. For similar reasons, I interpret three further abstractions as structurations: from ligand to input (for a receiver node), from tumbling frequency to output (of a transmitting node), and from methylated CheW-CheA complex to input-receiving node A.

4.3. Organizational Design. I call the third explanatory strategy *organizational design*. This strategy begins with a formal proof to the effect that systems exhibit robust perfect adaptation if and only if they satisfy the characteristic equation for integral feedback control (IFC). The proof is purely mathematical, well known from control engineering theory in contexts involving mechanical systems that exhibit IFC such as thermostats. There is debate about whether the proof holds only for deterministic systems (Briat, Gupta, and Khammash 2016). I set this aside, focusing on the explanatory strategy rather than its effectiveness.

The organizational design strategy infers, from the equivalence between robust perfect adaptation and IFC, that *E. coli* chemotaxis exhibits robust perfect adaptation if and only if it satisfies the characteristic equation for IFC. It also adds that *E. coli* chemotaxis exhibits robust perfect adaptation because it satisfies the characteristic equation for IFC: the organizing principle, IFC, is supposed to explain the phenomenon of interest.

This strategy invokes neither mechanistic specifics about the chemotaxis network for *E. coli* nor topological details about the structure of that network. The strategy takes the explanandum phenomenon as given, uses a mathematical equivalence result to identify a principle both necessary and sufficient for the phenomenon, and infers that the phenomenon obtains by virtue of the principle. The strategy thereby has affinities with explanatory strategies that appeal to organizing principles (Green and Wolkenhauer 2013).

The workings of this strategy are difficult to follow without understanding the sense in which *E. coli* exhibits IFC. So consider figure 4, a standard block diagram for visualizing IFC. An input acts on a system (near the top right), generating an output and a "measure signal." A comparison between

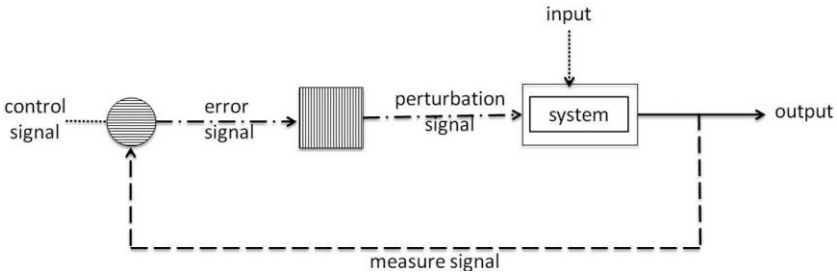

Figure 4. Block diagram for integral feedback control.

this signal and a control signal generates an "error signal." A controller (*square with vertical stripes*) then transforms the error signal into a perturbation signal designed to modulate the system's activity toward some desired output state.

E. coli exhibits IFC organization as follows. The system is the methylated CheW-CheA complex. The input is the ligand, which interacts with this complex; the output is the activity of the complex on CheY and everything downstream of that activity. The "measure signal" is the activity of the complex on CheB, which determines the rate of CheB phosphorylation. The "control signal" is the maximum rate at which CheR methylates the CheW-CheA complex, determined by the total concentration of CheR. There is a comparison between this rate and the rate of CheB phosphorylation. Nothing performs this comparison. But there is an "ideal" balance among the rates determined by the tumbling frequency before ligand exposure, and the "error signal" is the difference between this "ideal" balance and the actual balance. The controller is phosphorylated CheB, and the perturbation signal is its activity in demethylating the CheW-CheA complex.

Yi and colleagues' organizational design strategy for explaining the robust perfect adaptation of *E. coli* involves a model that is abstract in all the ways its cousin topological model is abstract and more. Some abstractions are the same: the motor rotation module aggregation and the structuration output from this module,the structuration input from the ligand, and the structuration system from the methylated CheW-CheA complex. Others differ. There is an aggregation of the module controlling phosphorylation and dephosphorylation of CheB and the structuration controller from this module. There are several generalizations as well: a control signal, abstracted from total concentration of CheR; a measure signal, from the activity of methylated CheW-CheA on phosphorylated CheB; and a perturbation signal, from the activity of phosphorylated CheB on methylated CheW-CheA. Hence, disregarding judgments about abstraction types, Yi and colleagues' model targets its abstractions to mechanistic details different from those of Ma and colleagues' topological model.

5. Confirming the Analysis. I consider the foregoing to establish that each explanatory strategy targets its various abstractions to different mechanistic details. The topological analysis strategy constructs aggregations by grouping mechanistic details into modules, and structurations by contextualizing the modules in motifs. The organizational design strategy goes further, explicitly abstracting over types of activities and, in some cases, over protein concentration levels.

Whether this result generalizes to other cases awaits future research. There is some reason to expect affirmative results. For example, Ueda and colleagues invoke a network motif to establish that *Cryptochrome 1* is necessary for properly functioning circadian clocks in mammals, and they seem to construct that motif from aggregations and structurations of more detailed mechanistic networks (Ukai-Tadenuma et al. 2011, 278).[3] Rather than amass anecdotal support, I confirm the result by issuing a prediction. If dynamical, topological, and design explanatory strategies differ as I claim, we should expect the more abstract strategies to allow for wider scope of application (see also Craver and Darden 2013, 36). More generalized models likely allow more instances, and more structurated models likely allow more position occupants.

We find confirmation of this prediction for the case of robust perfect adaptation of *Bacillus subtilis* (*B. subtilis*) chemotaxis. Details of the organization design strategy for explaining why *E. coli* chemotaxis exhibits robust perfect adaptation also apply for explaining why *B. subtilis* chemotaxis exhibits robust perfect adaptation. But details of the corresponding dynamical mechanistic strategy do not. I provide only brief details.

Rao and Ordal (2009) pursue the dynamical modeling strategy, using the same techniques as Barkai and Leibler in the case of *E. coli*. But details differ. For example, according to Barkai and Leibler's model, CheB in *E. coli* demethylates only active receptor complexes; according to Rao and Ordal, CheB in *B. subtilis* demethylates inactive ones too. Again, according to Barkai and Leibler's model, without CheY *E. coli* runs but does not tumble; according to Rao and Ordal, without CheY *B. subtilis* tumbles but does not run. One more: according to Barkai and Leibler's model, *E. coli* without CheB cannot run; according to Rao and Ordal, *B. subtilis* without CheB can run.

Details of Barkai and Leibler's dynamical modeling strategy do not apply for the case of *B. subtilis*. Details of Yi and colleague's organizational design strategy, by contrast, apply to this further case. For if they are correct, *B. subtilis*, like *E. coli*, exhibits robust perfect adaptation for chemotaxis if and only if it satisfies the characteristic equation for IFC. Because the organization design strategy involves more aggregation and structuration than the dynamical modeling strategy, the prediction is confirmed.

3. I thank Daniel Burnston for bringing this example to my attention.

6. Concluding Remarks. Dynamical, topological, and organizational design strategies apply different mathematical techniques in efforts to explain the same phenomenon—the robust perfect adaptation of bacterial chemotaxis. Each strategy applies its techniques to network models and differs with respect to the abstractions from which it constructs these models. Because each strategy targets its various abstractions to different mechanistic details, differences among abstractions help to show how the explanatory strategies differ from each other. These differences also help to explain why topological and organizational design strategies provide explanatory affordances unavailable through standard mechanistic explanations: the use of structuration allows their models to support more general conclusions, with wider scope, than the kind of differential equation models available for dynamical mechanistic explanations.

REFERENCES

Alon, Uri, Mike G. Surette, Naama Barkai, and Stan Leibler. 2009. "Robustness in Bacterial Chemotaxis." *Nature* 397:168–71.

Andersen, Holly. 2018. "Complements, Not Competitors: Causal and Mathematical Explanations." *British Journal for the Philosophy of Science* 69:485–508.

Barkai, Naama, and Stan Leibler. 1997. "Robustness in Simple Biochemical Networks." *Nature* 387:913–17.

Bechtel, William, and Adele Abrahamsen. 2005. "Explanation: A Mechanistic Alternative." *Studies in History and Philosophy of Biological and Biomedical Science* 36:421–41.

———. 2010. "Dynamic Mechanistic Explanation: Computational Modeling of Circadian Rhythms as an Exemplar for Cognitive Science." *Studies in History and Philosophy of Science* 41:321–33.

Berg, Howard C. 2004. *E. coli in Motion*. New York: Springer.

Braillard, Philippe A. 2010. "Systems Biology and the Mechanistic Framework." *History and Philosophy of the Life Sciences* 32:43–62.

Briat, Corentin, Ankit Gupta, and Mustafa Khammash. 2016. "Antithetical Integral Feedback Ensures Robust Perfect Adaptation in Noisy Biomolecular Networks." *Cell Systems* 2:15–26.

Brigandt, Ingo. 2013. "Systems Biology and the Integration of Mechanistic Explanation and Mathematical Explanation." *Studies in History and Philosophy of Biological and Biomedical Sciences* 44:477–92.

Brigandt, Ingo, Sara Green, and Maureen A. O'Malley. 2017. "Systems Biology and Mechanistic Explanation." In *The Routledge Handbook of Mechanisms and Mechanical Philosophy*, ed. S. Glennan and P. Illari, 362–73. New York: Routledge.

Craver, Carl F. 2006. "When Mechanistic Models Explain." *Synthese* 153:355–76.

Craver, Carl F., and Lindley Darden. 2013. *In Search of Mechanisms: Discoveries across the Life Sciences*. Chicago: University of Chicago Press.

Darden, Lindley, and Michael Cook. 1994. "Reasoning Strategies in Molecular Biology: Abstractions, Scans and Anomalies." In *PSA 1994: Proceedings of the 1994 Biennial Meeting of the Philosophy of Science Association*, vol. 2, ed. David Hull, Micky Forbes, and Richard M. Burian, 179–91. East Lansing, MI: Philosophy of Science Association.

French, Steven. 2010. "Keeping Quiet on the Ontology of Models." *Synthese* 172:231–49.

Godfrey-Smith, Peter. 2009. "Abstractions, Idealizations, and Evolutionary Biology." In *Mapping the Future of Biology: Evolving Concepts and Theories*, ed. A. Barberousse, M. Morange, and T. Pradeu, 47–56. Dordrecht: Springer.

Green, Sara, and Nicholaos Jones. 2016. "Constraint-Based Reasoning for Search and Explanation: Strategies for Understanding Variation and Patterns in Biology." *Dialectica* 70:343–74.

Green, Sara, and Olaf Wolkenhauer. 2013. "Tracing Organizing Principles: Learning from the History of Systems Biology." *History and Philosophy of the Life Sciences* 35:553–76.
Haslanger, Sally. 2016. "What Is a (Social) Structural Explanation?" *Philosophical Studies* 173:113–30.
Huneman, Philippe. 2010. "Topological Explanation and Robustness in Biological Sciences." *Synthese* 177:213–45.
Jones, Nicholaos. 2014. "Bowtie Structures, Pathway Diagrams, and Topological Explanation." *Erkenntnis* 79:1135–55.
Jones, Nicholaos, and Olaf Wolkenhauer. 2012. "Diagrams as Locality Aids for Search and Explanation in Molecular Cell Biology." *Biology and Philosophy* 27 (5): 705–21.
Levy, Arnon. 2013. "Three New Kinds of Mechanism." *Biology and Philosophy* 28:99–114.
Levy, Arnon, and William Bechtel. 2013. "Abstraction and the Organization of Mechanisms." *Philosophy of Science* 80:241–61.
Ma, Wenzhe, Ala Trusina, Hana El-Samad, Wendell A. Lim, and Chao Tang. 2009. "Defining Network Topologies That Can Achieve Biochemical Adaptation." *Cell* 138:760–73.
Machamer, Peter, Lindley Darden, and Carl F. Craver. 2000. "Thinking about Mechanisms." *Philosophy of Science* 67:1–25.
Matthiessen, Dana. 2017. "Mechanistic Explanation in Systems Biology: Cellular Networks." *British Journal for Philosophy of Science* 68:1–25.
Ordorica, Sergio Armando Gallegos. 2015. "The Explanatory Role of Abstraction Processes in Models: The Case of Aggregations." *Studies in History and Philosophy of Science* 56:161–67.
Pincock, Christopher. 2015. "Abstract Explanations in Science." *British Journal for the Philosophy of Science* 66:857–78.
Rao, Christopher V., and George W. Ordal. 2009. "The Molecular Basis of Excitation and Adaptation during Chemotaxic Sensory Transduction in Bacteria." In *Bacterial Sensing and Signaling*, ed. M. Collin and R. Schuch, 33–64. Basil: Karger.
Rice, Collin. 2015. "Moving beyond Causes: Optimality Models and Scientific Explanation." *Nous* 49:589–615.
Spudich, John L., and D. E. Koshland Jr. 1976. "Non-genetic Individuality: Chance in the Single Cell." *Nature* 262:467–71.
Ukai-Tadenuma, Maki, Rikuhiro G. Yamada, Haiyan Xu, Jurgen A. Ripperger, Andrew C. Liu, and Hiroki R. Ueda. 2011. "Delay in Feedback Repression by *Cryptochrome 1* Is Required for Circadian Clock Function." *Cell* 144:268–81.
Wadhams, George H., and Judith P. Armitage. 2004. "Making Sense of It All: Bacterial Chemotaxis." *Nature Reviews: Molecular Cell Biology* 5:1024–37.
Wouters, Arno G. 2007. "Design Explanation: Determining the Constraints on What Can Be Alive." *Erkenntnis* 67:65–80.
Yi, Tau-Mu, Yun Huang, Melvin I. Simon, and John Doyle. 2000. "Robust Perfect Adaptation in Bacterial Chemotaxis through Integral Feedback Control." *Proceedings of the National Academy of Sciences of the USA* 97:4649–53.