The Completeness of Mechanistic Explanations

Tudor M. Baetu*†

The paper discusses methodological guidelines for evaluating mechanistic explanations. According to current accounts, a satisfactory mechanistic explanation should include all of the relevant features of the mechanism, its component entities and activities, and their properties and organization, as well as exhibit productive continuity. It is not specified, however, how this kind of mechanistic completeness can be demonstrated. I argue that parameter sufficiency inferences based on mathematical model simulations provide a way of determining whether a mechanism capable of producing the phenomenon of interest can be constructed from mechanistic components organized, acting, and having the properties described in the mechanistic explanation.

1. Introduction. Some of the most successful and influential explanations in the life sciences amount to descriptions of mechanisms, where mechanisms are characterized as organized systems of parts that operate in such a way as to produce phenomena. The most systematic attempt to develop norms for evaluating the completeness of mechanistic explanations can be found in Carl Craver's *Explaining the Brain*, although guidelines and suggestions are present in the works of other proponents of the new mechanistic philosophy (Machamer, Darden, and Craver 2000; Darden 2006; Bechtel and Richardson 2010; Baetu 2012; Craver and Darden 2013). Craver (2007, 111) argues that criteria for evaluating mechanistic explanations should be able to address two main aims: "(1) to distinguish how-possibly explanations from how-actually explanations, and (2) to distinguish mechanism sketches from mechanism schemata." Aim 1 refers to the distinction between con-

*To contact the author, please write to: Department of Philosophy, University of Bristol, Cotham House, Bristol, BS6 6JL, UK; email: tudor.baetu@bristol.ac.uk.

†This work was supported in part by a generous fellowship from the KLI Institute. I thank Stuart Glennan, Mathieu Charbonneau, Dan Nicholson, Maarten Boudry, Argyris Arnellos, Laura Nuño de la Rosa, and Michael Rammerstorfer for their much-appreciated input.

Philosophy of Science, 82 (December 2015) pp. 775–786. 0031-8248/2015/8205-0004$10.00

jectures about 'how-possibly' mechanisms that might produce a phenomenon of interest and 'how-actually' descriptions of the real components, activities, and organizational features of the mechanism that in fact produces the phenomenon (2007, 112). Aim 2 refers to the distinction between complete and incomplete explanations. A satisfactory mechanistic explanation should "include all of the relevant features of the mechanism, its component entities and activities, their properties, and their organization" (Craver 2006, 367) and "exhibit productive continuity without gaps from the set up to termination conditions" (Machamer et al. 2000, 3). By contrast, a mechanism sketch "is an abstraction for which bottom out entities and activities cannot (yet) be supplied or which contains gaps in its stages" (Machamer et al. 2000, 18).

According to the above account, a mechanistic explanation is analogous to a recipe for producing a phenomenon starting from a list of ingredients, where the ingredients are mechanistic entities and their properties, and the recipe amounts to the organization and sequence of activities these entities perform. The mechanistic explanation is satisfactory when it is known by means of which particular 'mechanistic recipe' the phenomenon of interest is actually produced (aim 1), and when there are no missing ingredients and no missing lines in the description of the 'mechanistic recipe' for producing that phenomenon (aim 2). Note that it is not a question here of explaining why the mechanistic ingredients have the properties they have, how they came to be the way they are, or why they are organized the way they are. Rather, mechanistic completeness is understood in terms of whether entities, properties of entities, activities, or organizational features have been omitted, such that the mechanistic explanation amounts to an incomplete recipe missing some ingredient or step in the sequence of events necessary for the production of the phenomenon of interest.

I argue that the norms for evaluating mechanistic explanations elaborated by Craver (2007) and others are inspired by the experimental practice of the life sciences and that while these norms provide an adequate assessment of aim 1, they are of limited use for addressing aim 2. The purpose of the paper is to provide additional guidelines for an account of mechanistic completeness meant to complement previously elaborated interventionist norms of mechanistic explanation. To this end, I argue that an increasingly popular strategy to evaluate aim 2 relies on mathematical modeling. In certain cases, the information generated by testing mathematical models of known mechanisms can be used to infer whether all the relevant mechanistic components and information about these components, including their properties, organization, and activities they perform, have been taken into consideration.

2. The Explanatory Leakage Problem. In the life sciences, mechanisms are typically elucidated experimentally, by carefully circumscribing puta-

tive mechanisms within the boundaries of well-characterized experimental setups (Baetu 2013); by means of decomposition strategies (Bechtel and Richardson 2010); by conducting exploratory interventions aimed at identifying correlating factors providing an initial pool of putative mechanistic components (Woodward 2002); by performing specific interventions aimed at demonstrating the causal relevance of the entities, activities, and organizational features of a hypothesized mechanism (Craver 2007) and elucidating their causal roles relative to the operation of the mechanism (Craver 2001); and by tracking causal pathways (Darden 2006; Craver 2007).

By intervening on the components of a mechanism, it is possible to demonstrate that the mechanism is necessary and actually involved in the production of the phenomenon (Craver 2007). Given a suitable experimental design (e.g., standardized quantitative measurements, multivariable intervention experiments), experimental interventions can provide further evidence that no parallel or convergent causal pathways are actually involved in the production of a phenomenon in a particular experimental setup. For example, in a typical knockout experiment, two factors, the initial conditions and a mechanistic component, are simultaneously manipulated on an independent basis, and the effects on the output conditions are observed. If the knocking out of the component results in a complete inhibition of the output, one can infer that the mechanism is necessary and sufficient for producing the phenomenon of interest, in the sense that there are no other mechanisms that produce the phenomenon via alternate causal pathways that do not involve the knocked-out component (Baetu 2012).

Experimental interventions can demonstrate that mechanisms involving specific components are necessary and actually involved in the production of phenomena, thus providing methodological criteria for distinguishing how-possibly explanations from how-actually explanations. However, interventions cannot tell us whether all the components have been filled in or whether there are gaps in the productive continuity of a mechanism. This raises what William Bechtel calls the 'challenge of recomposing mechanisms' (Bechtel and Abrahamsen 2010). Perhaps the most striking way of framing the challenge is in terms of the ability to physically construct biological mechanisms (Morange 2009; Craver and Darden 2013, 92–94): if the mechanism described in the proposed explanation of a phenomenon were to be artificially synthesized from components organized, acting, and having the properties described in the mechanistic explanation, would it succeed in producing the phenomenon of interest?

The guiding idea behind the distinction between mechanism schemas and mechanism sketches is that the latter are missing details required in order to demonstrate that the proposed mechanism can produce the phenomenon for which it is responsible by virtue of the components of the mechanism, along with the properties, organization, and activities of these components. Unfor-

tunately, there are no methodological criteria for determining which details matter, thus creating an 'explanatory leakage' problem. In principle, there is always more to be said relative to the causal basis of any given biological phenomenon. The entities of a biological mechanism can be further decomposed into subparts, activities into subactivities, and mechanisms into more fine-grained submechanisms, thus raising questions about where a mechanistic explanation is expected to bottom out. At the same time, biological mechanisms are also parts of progressively more comprehensive systems of mechanisms, ranging from molecular networks to planetary ecosystems, where these systems both depend on the functioning of the mechanisms of which they are composed and impose constraints on these very same mechanisms. This raises questions about whether there are such things as local mechanisms operating independently of ever more encompassing systems. It seems therefore that a mechanistic explanation can be indefinitely detailed and expanded by bottoming out at lower levels of composition and by taking into consideration integration into higher-level systems. Meanwhile, it is not in the least clear how this relentless accumulation of information can ever demonstrate that the mechanism described in the explanation can produce the phenomenon for which it is responsible by virtue of its identified components and their properties, organization, and activities.

To illustrate the problem, consider the following example. When exposed to certain stimulants, T lymphocytes express a variety of genes required for mounting an immune response, after which they automatically return to their initial resting state. This spike of gene expression following stimulation is explained by a negative feedback regulatory mechanism whereby a transcriptional factor (nuclear factor κB, or NF-κB) is initially activated and then subsequently inactivated by an inhibitory protein (inhibitor of κB, or IκB) coded by a gene under its transcriptional control. There are many details missing from the above mechanistic description. The mechanistic description can be further elaborated by bottoming down at the deeper level of biochemical details rather than the lower resolution level of molecular interactions, most notably by including additional information about the tridimensional configurations of the proteins involved and their role vis-à-vis molecular function (e.g., structural motifs involved in specific binding). By digging deeper, researchers typically hope to gain a better understanding of why and how mechanistic components are able to do what they are doing, as well as to discover new ways in which mechanistic components can be manipulated for experimental and technological purposes. The mechanistic description can also be expanded by taking into account how this particular regulatory mechanism connects with other molecular mechanisms, most notably upstream signaling pathways and downstream mechanisms triggered by the expression of new genes. The mechanism is known to be involved in a number of rather diverse biological phenomena, ranging from development to immunity. By

adopting a systemic viewpoint, one may hope to gain a better understanding of how immunity relates to other biological activities. This is particularly important for understanding possible side effects of therapies designed to enhance desirable immune responses or inhibit deleterious ones.

While both a more fine-grained description bottoming out at lower levels of composition and a more systemic perspective amount to a net gain of knowledge, it is not obvious which, if any, of these many additional details are needed in order to conclude that the negative feedback loop mechanism can generate the phenomenon of interest by virtue of identified mechanistic components and their properties, organization, and activities. Higher-resolution structural details of the NF-κB transcriptional activator and the IκB inhibitor are crucial for understanding how these two proteins bind each other and which alterations (e.g., mutations) result in a loss in binding. Nevertheless, given experimentally gained knowledge that the two bind, further knowing how and why they bind does not tell us whether it is possible to artificially synthesize the feedback regulatory mechanism starting from a pool of NF-κB transcriptional activators, IκB inhibitor proteins, and other molecular components organized as described in the explanation. Likewise, if a more systemic understanding of how this regulatory mechanism contributes to a variety of biological activities is crucial for assessing the physiological and evolutionary relevance of the mechanism, this knowledge does not tell us whether the regulatory mechanism's contribution to the regulation of transcription is mediated solely by means of the feedback loop regulation of gene expression and independently of the mechanism's involvement in other biological activities. In this respect, it is not clear whether the mechanism can be detached from the system in which it is embedded and treated as an independent module capable of producing the phenomenon (spikes of gene expression) on its own.

3. Pragmatic and Philosophical Motivations for Elaborating Norms of Mechanistic Completeness. The pragmatic import of developing norms of mechanistic completeness is twofold. During the discovery process, evidence supporting the completeness of the explanation indicates that the research project is on the right path. Before worrying about the countless ways in which a mechanistic explanation could be further detailed and expanded, it is crucial to gather at least some evidence that the proposed mechanism, at the level of composition at which it is currently described, can and does produce the phenomenon of interest. It would be misguided to try to understand how and why the components of a mechanism do what they are doing, how the mechanism and its organizational features came into being, or how the mechanism integrates the greater whole that is the living organism in the absence of evidence that the mechanism described in the explanation can produce the phenomenon to be explained. The second point

of pragmatic relevance links to the fact that mechanistic explanations often provide the rationale for developing technologies for gaining control over phenomena, such as experimental techniques and medical treatments. In this case, evidence supporting the completeness of the explanation is needed for making an enlightened decision about the probability of a successful outcome, especially in situations where a trial-and-error approach is not an ethically viable option.

From a conceptual point of view, the interest in developing norms for evaluating mechanistic explanations links primarily to the problem of reduction. Given the possibility of an indefinite descent to lower levels of composition, how deep does one need to go in order to claim that the explanation is satisfactory for the purposes of accounting for the phenomenon of interest? A related problem is figuring out whether mechanisms act as independent modules that can continue to function when separated from the systems in which they are embedded, especially when the system in question is the physiological context of a living thing. As discussed earlier, failure to specify levels of compositional detail where mechanistic explanations can safely bottom out, as well as criteria for determining when and to what extent mechanisms behave like independent modules, generates an explanatory leakage problem, whereby mechanistic explanations can be continuously elaborated without necessarily demonstrating that the mechanism described in the explanation can produce the phenomenon for which it is responsible by virtue of its components and their properties, organization, and activities.

4. Quantitative and Parameter Sufficiency Inferences. In this section I discuss how mathematical modeling provides a means to draw a principled distinction between mechanism schemas and mechanism sketches, thus specifying where an explanation can safely bottom out and at what point the mechanism can be considered an independent module.

Mathematical modeling is by no means a novel practice in biology. The Hodgkin–Huxley model of the action potential, the Michaelis–Menten model of enzyme kinetics, and Knudson's two-hit model of cancer development made use of theoretical tools in order to demonstrate that biological and biochemical phenomena can be accounted for as consequences of laws and rules governing the behavior of certain systems. These same models played an important role in guiding the subsequent elucidation of molecular mechanisms. Nevertheless, the claim that models may also be useful tools for evaluating the completeness of mechanistic explanations is relatively recent. Thirty years ago, Harold Morowitz (1984) argued that experiments that can be carried out in the laboratory can also be carried out on the computer, and that the extent to which the two match measures the completeness of the paradigm of molecular biology. The suggestion here is that the abil-

ity to accurately simulate the behavior of a mechanism—or, more ambitiously, that of a whole cell or organism—by means of empirically informed mathematical models provides the evidence necessary to demonstrate that life is indeed explainable in strictly molecular-mechanistic terms. A decade later, this strategy was effectively put into practice with the explicit aim of evaluating the completeness of mechanistic explanations (Hartwell et al. 1999).

Today, there are many studies illustrating this practice in biology, some of which involve extensive systems of mechanisms or even whole cells. For simplicity, I will use as an example a study conducted by Hoffmann et al. (2002) involving a mathematical model of the NF-κB negative feedback regulatory mechanism briefly described earlier. Commenting on this study, Alice Ting and Drew Endy make the following point:

> A limitation of computational modeling is that, in the absence of complete information about cell parts and interconnections, it is easy to omit critical parameters that might influence the state of a cell or signaling pathway. This is illustrated in the Hoffmann et al. work. . . . When they used this model to predict the behavior of wild-type cells, the outcome was very different from what was actually measured, even though many of the parameters were empirically obtained. Such discrepancies could be due to compensatory changes in expression and signaling state from one cell line to the next, or to additional pathway components and regulatory mechanisms beyond the current model. (2002, 1190)

The limitation to which they allude is not one due to abstraction, idealization, or the instrumental nature of the models used, but rather the concern that, even when constructing detailed and highly realistic mathematical models of previously elucidated molecular mechanisms, and even when the values of the parameters of the model are based on empirical measurements, these models can only be as complete as our knowledge of the modeled mechanisms is. However, as the authors quickly point out, there is a bright side to this limitation. If the output of the model fails to closely match the phenomenon known to be produced by the mechanism, this can be an indication that something is missing from the mechanistic explanation. That is, the mechanistic explanation might be incomplete in the sense that not all the components of the mechanism have been identified, or that a more complex system including other mechanisms is needed to produce the phenomenon.

Conversely, if the output of the mathematical model matches experimental measurements of the phenomenon of interest, this is taken as evidence supporting the claim that the proposed mechanism is quantitatively sufficient for generating that phenomenon. This is an important piece of information. Qualitative descriptions associated with traditional mechanistic explanations usually suffice to provide an intuitive understanding of how a

mechanism may produce something roughly resembling the phenomenon to be explained. For instance, one can intuitively understand how a negative feedback loop switching gene expression 'on' and 'off' in response to persistent exposure to triggering conditions can generate oscillating peaks of gene expression. Nevertheless, a qualitative description cannot account for quantitative-dynamic details such as values of the amplitude, frequency, and dampening of the oscillations. Numerical computation, rather than qualitative description, is required in order to account for such details.

When quantitative sufficiency is demonstrated by means of a detailed and realistic model, parameter sufficiency is further inferred. If model simulations match experimental data, it can be argued that a more complex model including additional parameters is not needed. Inasmuch as all the parameters have a clear physical interpretation—which is to say that they describe physical properties of the components of the modeled mechanism—and at least some values of these parameters are based on independent empirical measurements, a close match between simulation and experimental measurements of the phenomenon is taken as evidence supporting the claim that a more complex mechanism including additional components and the physiological context of other mechanisms is not needed to produce the phenomenon.

The notion of parameter sufficiency plays an important role in guiding the design of artificial molecular mechanisms aimed at producing a desired phenomenon. The repressilator, an artificial molecular oscillator, was designed on the basis of mathematical models predicting that sustained oscillations (the desired outcome) are favored by transcriptional regulation mechanisms constructed from molecular components organized in a certain way (in this case, negative feedback loops) and having a particular set of properties (strong promoters, low leakiness, etc.; Elowitz and Leibler 2000); for a philosophical discussion see Morange (2009). Even though this first attempt to construct a synthetic mechanism turned out to be only a partial success—the mechanism did produce oscillations, but it lacked the desired degree of robustness—it demonstrated that, in principle, mathematical models can be used to evaluate and predict whether a mechanism synthesized from the components described in the designed mechanism can generate the phenomenon of interest down to minute quantitative-dynamic aspects.

Beyond the specific needs of synthetic biology, parameter sufficiency inferences can provide more rigorous and objective norms for distinguishing between mechanism schemas and mechanism sketches: the former satisfy the requirement for parameter sufficiency, while the latter do not. It is precisely this criterion that Hoffmann et al. used to evaluate whether the NF-κB regulatory mechanism can generate the peaks of gene expression it was supposed to explain. A mathematical model of the mechanism revealed that there is a mismatch between the oscillations simulated by the

model and the oscillatory responses measured experimentally. The nature of the mismatch further suggested that there is a missing component of the mechanism responsible for stabilizing the activation of gene expression following persistent stimulation. After identifying the missing component, a revised mechanism was proposed and models of the revised mechanism were able to generate oscillations closely matching experimentally measured values. In turn, this fit provided evidence that all the relevant mechanistic components and all the relevant information about these components have been included in the explanation.

Again, it should be noted that the kind of explanatory completeness evaluated by means of parameter sufficiency has little to do with an ultimate understanding of how everything works at the level of systemic interactions between the most fundamental building blocks of physical reality. Rather, it is an engineer's understanding of completeness framed in terms of information required to reconstruct in silico a mechanism capable of producing the phenomenon of interest starting from components organized, acting, and having the properties described in the mechanistic explanation. This notion of mechanistic completeness addresses the explanatory leakage problem described in section 2. Parameter sufficiency provides the means to determine whether it is safe to bottom out at the level of composition at which the mechanism is described, in the sense that a more detailed description is not required for the purpose of explaining how the components of the mechanism produce the phenomenon by virtue of their properties, organization, and activities; and whether it is safe to treat the mechanism as an independent module that can be separated from the system in which it is embedded and yet continue to produce the phenomenon for which it is responsible.

In the NF-κB regulatory mechanism example, the key finding amounted to the realization that the initial negative feedback loop mechanism needs to be augmented to include a parallel pathway of activation not subjected to negative feedback, and that it takes the combined activity of both pathways in order to produce peaks of gene expression matching experimental observations. The bottoming-out argument here is that in order to produce the phenomenon of interest, the key requirement is that of a double activation pathway involving experimentally identified molecular components shown to be necessary for the production of the phenomenon and shown to interact in such a way as to make possible the double activation pathway. For the immediate purpose of explaining the phenomenon of interest, it is not essential to further understand why these molecular components interact the way they do, how these components were produced in the cell, or how they evolved. Furthermore, it is expected that certain changes would not impact the ability of the mechanism to produce its target phenomenon. For instance, the NF-κB activator, its DNA-binding motifs, and the IκB

inhibitor could tolerate certain changes in sequence and structure, yet the mechanism would continue to function on condition that some key features are preserved, such as the dual activation pathway and the affinity and kinetics of chemical interactions (as might occur, for instance, when complementary mutations in several components rescue the wild-type phenotype). There is therefore a clear sense in which certain lower-level structural biochemical details can be ignored and the phenomenon of interest can be satisfactorily explained in terms of higher-level molecular description of mechanistic components, and their properties, organization, and activities.

Likewise, the tight quantitative match between the predictions of the model and experimental measurements supports the claim that, at least within the time frame in which the phenomenon is characterized, other mechanisms at work in the cell, as well as effects triggered downstream as a result of the functioning of the mechanism, are not required to produce the phenomenon of interest or interfere with its ability to produce it. It is therefore expected that an in vitro reconstituted NF-κB regulatory mechanism should produce spikes of gene activation closely resembling those produced in vivo. Or again, a genetically modified organism in which the coding sequences of the genes activated by the mechanism are replaced with those of other genes would nevertheless display the same pattern of gene expression. This specifies a sense in which a more systemic context can be ignored such that a satisfactory explanation can amount to the description of a local mechanism operating as an independent module.

5. Some Final Thoughts. It is perhaps wise to end on a cautionary note. Despite its many promises, mathematical modeling is by no means a miracle solution providing ultimate, foolproof answers. Inferring that a mechanism operates as an independent module and that all the relevant mechanistic components have been taken into consideration is dependent on how finely grained is the description of the phenomenon of interest, something that is contingent on the resolution of the measurement techniques of the day. Since the description of the phenomenon of interest is susceptible to revisions, so are the claims that the proposed explanation is complete. Furthermore, pragmatic interests often dictate that an explanation is 'good enough' as long as the manipulation of some key mechanistic components suffices to achieve the desired results (e.g., the emphasis on the genetic basis of medical conditions given the potential for gene therapy).

More importantly, the inference that a mechanistic explanation is incomplete is likely to be trustworthy only inasmuch as the assumptions and data on the basis of which the model is constructed reflect with a sufficient degree of adequacy features of the actual mechanism responsible for the phenomenon. Depending on the model, a failure to match actual measure-

ments can also be attributed to distortions inherent to the modeling process, such as lack of experimentally measured values, idealizing assumptions about the dynamic behavior of the mechanism, and simplifications needed in order to enhance computability. Disentangling these various sources of error is difficult, although evidence accumulating over time can eventually favor an interpretation over its rivals. In this respect, the study by Hoffmann et al. is particularly interesting because the failure to simulate the phenomenon of interest was attributed to an incomplete description of the molecular mechanism rather than an inadequacy of the model, thus prompting a revision of the mechanism. Evidence supporting this interpretation came from experimental evidence for the revised mechanism, coupled with the fact that the revised mechanistic explanation was able to answer additional questions about some seemingly unrelated phenomena; for discussion, see Baetu (2015).

Despite its caveats, there is a clear sense of excitement about mathematical modeling in all branches of biology. Mathematical modeling provides an accessible substitute for something missing in biology: a theoretical apparatus formulated in mathematical language allowing for the elaboration of explanations and hypotheses capable of precise quantitative predictions. In the absence of such an apparatus, experimental research is bound to remain largely exploratory, and exploration implies a fundamental incertitude about how much is known and how much remains to be investigated. While mathematical modeling cannot rival the all-encompassing theories of physics, it can nevertheless provide a useful work-around by providing a principled way of evaluating the completeness of the information included in mechanistic explanations. At any given point during a project, researchers can stop, put together the many bits and pieces of experimental data into putative mechanistic descriptions, and then model these descriptions in order to gain at least a rough estimate of whether, thus far, they 'got things right' and the proposed mechanisms can indeed produce the phenomena they are supposed to explain.

REFERENCES

Baetu, T. M. 2012. "Filling in the Mechanistic Details: Two-Variable Experiments as Tests for Constitutive Relevance." *European Journal for Philosophy of Science* 2 (3): 337–53.

———. 2013. "Chance, Experimental Reproducibility, and Mechanistic Regularity." *International Studies in History and Philosophy of Science* 27 (3): 255–73.

———. 2015. "From Mechanisms to Mathematical Models and Back to Mechanisms: Quantitative Mechanistic Explanations." In *Explanation in Biology: An Enquiry into the Diversity of Explanatory Patterns in the Life Sciences*, ed. P.-A. Braillard and C. Malaterre. Dordrecht: Springer.

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

Bechtel, W., and R. Richardson. 2010. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Cambridge, MA: MIT Press.
Craver, C. 2001. "Role Functions, Mechanisms, and Hierarchy." *Philosophy of Science* 68:53–74.
———. 2006. "When Mechanistic Models Explain." *Synthese* 153:355–76.
———. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Clarendon.
Craver, C., and L. Darden. 2013. *In Search of Biological Mechanisms: Discoveries across the Life Sciences*. Chicago: University of Chicago Press.
Darden, L. 2006. *Reasoning in Biological Discoveries: Essays on Mechanisms, Interfield Relations, and Anomaly Resolution*. Cambridge: Cambridge University Press.
Elowitz, M., and S. Leibler. 2000. "Synthetic Gene Oscillatory Network of Transcriptional Regulators." *Nature* 403:335–38.
Hartwell, L. H., J. J. Hopfield, S. Leibler, and A. W. Murray. 1999. "From Molecular to Modular Cell Biology." *Nature* 402 (6761 Supp.): C47–52.
Hoffmann, A., A. Levchenko, M. Scott, and D. Baltimore. 2002. "The I κB–NF-κB Signaling Module: Temporal Control and Selective Gene Activation." *Science* 298:1241–45.
Machamer, P., L. Darden, and C. Craver. 2000. "Thinking about Mechanisms." *Philosophy of Science* 67:1–25.
Morange, M. 2009. "Synthetic Biology: A Bridge between Functional and Evolutionary Biology." *Biological Theory* 4 (4): 368–77.
Morowitz, H. 1984. "The Completeness of Molecular Biology." *Israel Journal of Medical Sciences* 20:750–53.
Ting, A. Y., and D. Endy. 2002. "Decoding NF-κB Signaling." *Science* 298:1189–90.
Woodward, J. 2002. "What Is a Mechanism? A Counterfactual Account." *Philosophy of Science* 69 (Proceedings): S366–S377.