Chapter 16
When Is a Mechanistic Explanation Satisfactory? Reductionism and Antireductionism in the Context of Mechanistic Explanations

Tudor M. Baetu

16.1 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 (Bechtel and Abrahamsen 2005; Glennan 2002; Machamer et al. 2000; McKay Illari and Williamson 2012). There is no mystery, however, that the entities of a biological mechanism can be further decomposed into subparts, activities into sub-activities, and mechanisms into more fine grained sub-mechanisms (Bechtel and Richardson 2010; Craver 2007). Nor there is any doubt that biological mechanisms are parts of progressively more comprehensive systems of mechanisms, ranging from molecular networks to planetary ecosystems, where more systemic mechanisms can both depend on the functioning of the sub-mechanisms of which they are composed and impose constraints on their mode of operation (Hooker 2011). Thus, in the realm of mechanistic explanations, the issue of reductionism in biology[1] can be

[1] In biology, the reductionism debate is primarily about the relationship between molecular biology and other branches of biology, such as classical genetics [e.g., (Waters 1990) vs. (Kitcher 1984)] and developmental biology [e.g., (Rosenberg 2006) vs. (Oyama 1985)]. In the contemporary literature, reductionists agree that a mechanistic explanation does not need to bottom down at the most fundamental building blocks of physical reality, and that a satisfactory explanation can be articulated at the level of molecular interactions. Likewise, even the most fervent proponents of antireductionism agree that some, but not all contexts, and certainly not the totality of the universe, are important for understanding biological phenomena. If there is a resistance to molecular or genetic reductionism, the concern is that certain features of the cell, organism or the direct

T.M. Baetu (✉)
Programa de Filosofia, Universidade do Vale do Rio dos Sinos, São Leopoldo, Brazil
e-mail: tudormb@unisinos.br

G. Sandu et al. (eds.), *Romanian Studies in Philosophy of Science*, Boston Studies in the Philosophy and History of Science 313, DOI 10.1007/978-3-319-16655-1_16

reformulated as a combo of questions, one about the level of composition at which mechanistic descriptions bottom out, and the second, about whether mechanisms act as independent modules that can continue to function when separated from the systems in which they are embedded.

The goal of this paper is to provide an answer to these questions. I argue that the solution lies in the elaboration of norms for evaluating the completeness of mechanistic explanations. According to current accounts (Craver 2006, 2007; Machamer et al. 2000), a satisfactory mechanistic explanation should describe the mechanism actually producing the phenomenon of interest and include all of the relevant features of the mechanism, its component entities and activities, their properties and their organization, as well as exhibit productive continuity. It is not specified, however, how this kind of mechanistic completeness can be demonstrated in scientific practice. Current accounts emphasize the role of experimental interventions demonstrating that various components of a mechanism are actually involved in the production of the phenomenon (Baetu 2012; Craver 2006, 2007). However, a strictly interventionist approach is not enough. I argue that an increasingly popular strategy for determining whether all the relevant mechanistic components and information about these components have been taken into consideration relies on mathematical modeling. Once it is possible to demonstrate that a given mechanism is actually involved in the production of a phenomenon and that it can produce that phenomenon solely in virtue of its identified components, their known properties, organization and activities, then there is no need to further elaborate the description of the mechanism by bottoming out at deeper levels of composition or to expand it in order to include a more systemic perspective, thus providing a principled way of determining when a mechanistic explanation is satisfactorily complete for the purposes of accounting for the phenomenon of interest.

The paper is organized as follows. In Sect. 16.2, I discuss currently elaborated guidelines for developing norms of mechanistic explanation. In Sect. 16.3, I discuss the role of experimental interventions in demonstrating that a mechanism and its components are necessary, and actually involved in the production of phenomena, as well as the limitations of a strictly interventionist approach. In Sect. 16.4, I elaborate the notions of quantitative and parameter sufficiency inferences from mathematical models and show how they can provide a principled way of determining where a mechanistic explanation can safely bottom out and what is the cutoff point beyond which external factors can be ignored. Finally, some broader-interest implications are discussed in Sect. 16.5.

environment of an organism have been neglected. It is within these boundaries that the issue of reductionism is considered here.

16.2 Guidelines for Developing Norms of Mechanistic Explanation

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 deemed satisfactory when (1) it is known by means of which particular 'mechanistic recipe' the phenomenon of interest is actually produced in the biological system of interest, and (2) there are no missing ingredients and no missing lines in the description of the 'recipe' for producing the phenomenon. In more technical terms, it is important "(1) to distinguish how-possibly explanations from how-actually explanations, and (2) to distinguish mechanism sketches from mechanism schemata" (Craver 2007, 111). Aim (1) refers to the distinction between conjectures about possible mechanisms that might be able to produce the phenomenon and descriptions of the actual components, activities, and organizational features of the mechanism that in fact produce the phenomenon (Craver 2007, 112). Aim (2) alludes to the completeness of the description of a mechanism. A mechanism schema is a "truncated abstract description of a mechanism that can be filled with descriptions of known component parts and activities. [. . .] When instantiated, mechanism schemata yield mechanistic explanations of the phenomenon that the mechanism produces" (Machamer et al. 2000, 15, 17). 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 incomplete explanation "for which bottom out entities and activities cannot (yet) be supplied or which contains gaps in its stages" (2000, 18).

If it were possible to demonstrate that the 'mechanistic recipe' for producing a phenomenon is actual and complete, the explanation of the phenomenon could safely be reduced to this 'recipe' in the sense that adding further ingredients or lines to the 'recipe' would either not cause any changes in the phenomenon, meaning that such additions are causally and explanatorily irrelevant, neutral or redundant, or interfere with the functioning of the mechanism causing a failure to produce the phenomenon as it is measured in the biological system of interest, in which case the explanation would fail. The task, therefore, is to determine what kind of evidence is necessary to support the claim that the 'mechanistic recipe' is actual and complete.

16.3 The Role of Experimental Interventions in the Elucidation of Biological Mechanisms

In the life sciences, mechanisms are usually elucidated experimentally, by carefully circumscribing a putative mechanism within the boundaries of a well characterized experimental setup (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 (Baetu 2012); by performing specific interventions aimed at demonstrating the causal relevance of the entities, activities, and organizational features of a hypothesized mechanism (Craver 2007; Woodward 2002, 2003) and elucidating their causal roles relative to the operation of the mechanism (Craver 2001); and by tracking causal pathways (Craver 2007; Darden 2006).

Currently elaborated norms for evaluating mechanistic explanations are inspired from the experimental practice of the life sciences. By intervening on the components of a mechanism, it is possible to demonstrate that a mechanism is necessary to produce the phenomenon, as well as that the mechanism in question is 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 (Baetu 2012). 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 (Fig. 16.1).

Experimental interventions are used to demonstrate that mechanistic 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 don't tell us if and when all the explanatorily relevant details have been filled in or whether there are gaps in the productive continuity of a mechanism. One way of framing the problem is in terms of the ability to physically construct biological mechanisms: if the mechanism described in the proposed explanation 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 as it was originally measured? To clarify, it is not question here of further explaining why the components have the properties they have, why they are organized the way they are, or why they are doing whatever they are doing. Nor is there any doubt about the fact the identified mechanistic components, along with their experimentally demonstrated properties, organization and activities are necessary for and actually involved in the production of the phenomenon. Rather, the issue under scrutiny is 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.

Consider, for instance, the following example. When exposed to certain stimulants, such as pathogens, white blood cells, and T-cells in particular express a

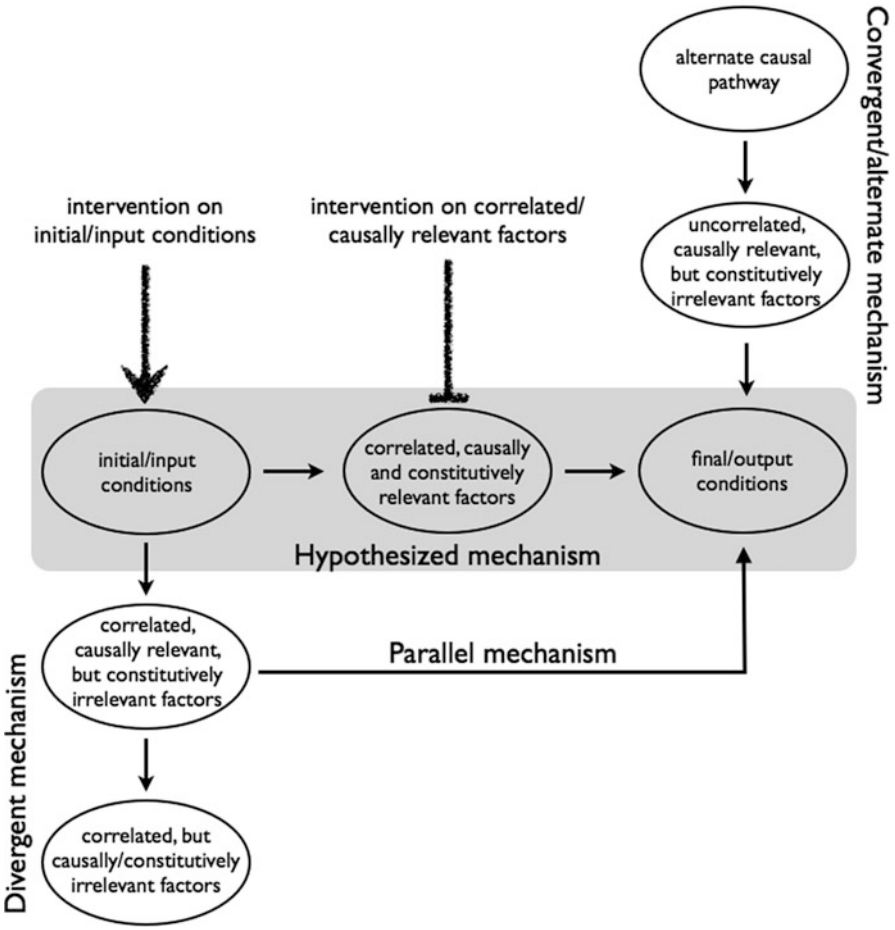

Fig. 16.1 Sorting putative mechanistic components by a two-variable knockout experiment

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, then subsequently inactivated by an inhibitory protein (inhibitor of κB, or IκB) coded by a gene under its transcriptional control (Fig. 16.2).

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 depicted in Fig. 16.2, most notably by including additional information about the tridimensional conformations of the proteins involved and their role vis-à-vis molecular function, such as structural motifs involved in specific binding

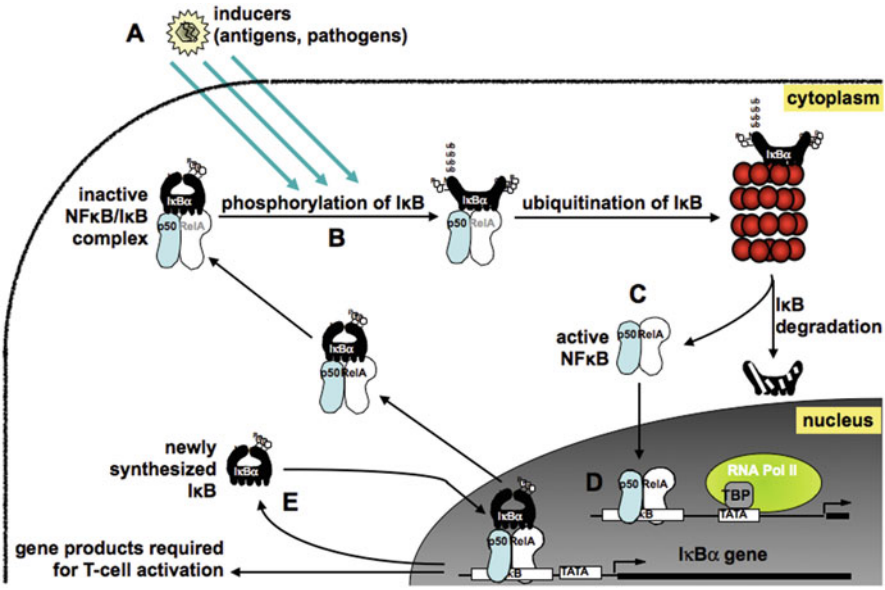

Fig. 16.2 The NF-κB negative feedback loop regulatory mechanism. In resting cells, the NF-κB transcriptional activator is held in the cytoplasm by the IκB inhibitor. When cells are stimulated (**a**), a chain of protein-protein interactions leads to the degradation of IκB (**b**); NF-κB is freed (**c**), translocates to the nucleus (**d**) where it binds specific sequences in the promoter regions of target genes drastically enhancing their transcription. NF-κB also binds the promoter of the IκB gene (**e**), and the newly synthesized IκB binds NF-κB, trapping it back in the cytoplasm

(Fig. 16.3, panel C). 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 discover new ways in which mechanistic components can be manipulated for experimental and medical purposes. This kind of knowledge and the interventions it renders possible play a crucial role in elucidating mechanisms. At the same time, the mechanistic description can also be expanded by taking into account other molecular mechanisms, most notably upstream signaling pathways and downstream mechanisms triggered via the expression of new genes (Fig. 16.3, Panel A). In this particular case, the negative feedback loop mechanism is known to be involved in a number of rather diverse biological phenomena, ranging from development and cell differentiation to immune responses and cell death. By adopting a more 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 deeper levels of composition and taking into consideration of a more systemic perspective amount to a net gain of knowledge, it is not obvious how this additional information can support the conclusion that the mechanism described in Fig. 16.2 generates the

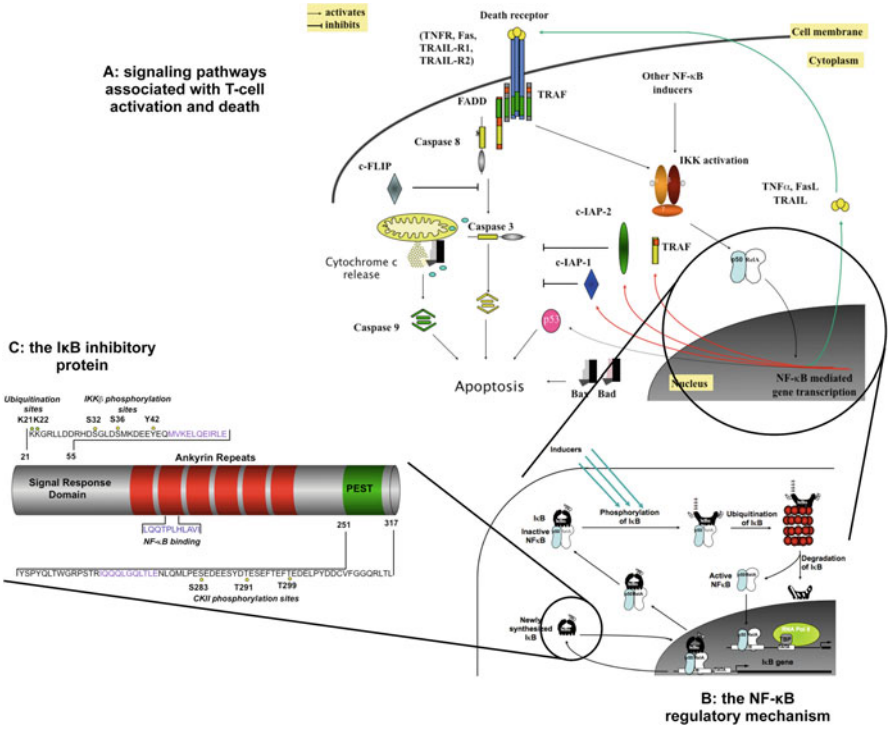

Fig. 16.3 Molecular levels [Panels adapted from (Baetu and Hiscott 2002; Baetu et al. 2001)]

phenomenon of interest in virtue of its identified components, 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 ability. 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 Fig. 16.2. There is, therefore, a worry that the mere fact that the various components of a mechanism can be analyzed at progressively lower levels of composition creates the reductionist illusion that biological phenomena are ultimately explainable by and reducible to the theories of particle physics, while in truth this analysis does not necessarily contribute to the mechanistic explanation, which should be a story about how an organized system of parts succeeds in producing a phenomenon. From the standpoint of mechanistic thinking, the goal is to figure out that precise level of composition at which parts, organized and acting as described in the proposed explanation, can generate the

phenomenon in need of an explanation, and not to explain how or why the parts have the properties they have, which is a different research question.

Likewise, if a more systemic understanding of how this regulatory mechanism contributes to a variety of biological activities is crucial for grasping the physiological and evolutionary relevance of the mechanism, as well as evaluating the therapeutic potential of interventions, this knowledge does not tell us if the regulatory mechanism's contribution to the regulation of immune responses is mediated solely by means of the feedback loop regulation of gene expression of the genes required for mounting an immune response, and independently of the mechanism's involvement in other biological activities.[2] The worry here is that the fact that biological mechanisms are embedded in or connected to other mechanisms creates the antireductionist illusion that everything is inextricably interconnected and ultimately irreducible to parts or sums of parts, while in truth the mere fact of connectedness does not allow us to determine whether or not a given mechanism can be treated as an independent module.

16.4 Inferences from Mathematical Models of Experimentally Elucidated Mechanisms

Specifying where an explanation can safely bottom out and when the mechanism can be considered an independent module requires a different kind of evidence, which is not likely to emerge from the accumulation of information bought about the further decomposition of mechanistic components or by taking into account progressively more systemic contexts. What is needed, is a reconstruction of the mechanism starting from a set of parts having the properties and organization specified by a proposed mechanistic explanation, in order to determine if, thus reconstructed, the mechanism can indeed produce the phenomenon it is supposed to produce. While the physical reconstruction of mechanisms is something documented in contemporary biology (Morange 2009; Weber 2005, Chap. 5), a much more common and accessible alternative relies on the mathematical modeling of experimentally elucidated mechanisms.[3]

[2]The organicist debate that raged in the nineteenth century biology centered on the claim that living things are organic wholes that cannot be decomposed into a set of independent mechanisms. Critics of molecular biology and its methods often appeal organicist arguments to defend more holistic approaches, and part of the manifesto of systems biology is precisely to provide a more holistic understanding of life. Contemporary echoes of this debate can be found in Nicholson (2013).

[3]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 or rules governing the behavior of certain systems. These same models played an important role in guiding the subsequent elucidation of the molecular mechanisms. More recently, mathematical models have

Commenting on a study by Hoffmann et al. (2002), where the authors constructed and tested a mathematical model of the NF-κB negative feedback regulatory mechanism described earlier (Fig. 16.2 above), 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 of computational modeling to which they allude is not one due to abstraction, idealizations 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 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 modeled mechanism, then this can be an indication that something is missing from the mechanistic explanation. That is, the mechanistic explanation might be incomplete because not all the components of the mechanism have been identified, or other mechanisms are needed to produce the phenomenon of interest.

It should be immediately noted that the kind of explanatory completeness evaluated by mathematical models has nothing to do with a 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.

If the output of the mathematical model of the proposed mechanism matches experimental measurements of the phenomenon, 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, by contemplating Fig. 16.2, one can intuitively understand how a negative feedback loop switching

---

been used to account for quantitative-dynamic features of phenomena meant to complement traditional qualitative descriptions of mechanisms (Baetu 2014a; Bechtel 2012; Bechtel and Abrahamsen 2010; Brigandt 2013). In such cases, mathematical models act as *in silico* surrogates for investigating the properties of systems they model (Baetu 2014b).

gene expression 'on' and 'off' in response to persistent exposure to triggering conditions can generate oscillating peaks of gene expression. Nevertheless, the question remains whether the feedback loop mechanism described in Fig. 16.2 can generate oscillations matching the amplitude, frequency, dampening and other minute quantitative-dynamic quirks of experimentally measured NF-κB mediated peaks of gene expression. Mathematical modeling provides the means to address this question.

When quantitative sufficiency is demonstrated by means of a detailed and realistic model, *parameter sufficiency* is further inferred. If the model simulations and predictions match experimental data, it can be argued that a more complex model, including additional parameters, is not needed. In as much as all the parameters have a clear physical interpretation, meaning that they describe known physical properties of the components of the mechanism, and at least some values of these parameters are based on independent empirical measurements, a close much between simulation and experimental measurements of the phenomenon of interest is taken as evidence in support of the claim that a more complex mechanism, including additional components, or additional mechanisms are not likely to be needed to produce the phenomenon.[4]

Parameter sufficiency plays an important role in guiding the design of artificial molecular mechanisms aimed at producing a desired phenomenon. Most famously, the repressilator (Elowitz and Leibler 2000), 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.). Even though this first attempt to construct a synthetic mechanism turned out to be only a partial success – the mechanism did produce oscillations, but lacked the desired degree of robustness – , it did demonstrate that mathematical models can be in principle 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 also provides the means to figure out whether it is safe to bottom out at the level of composition at which the mechanism is described, in the sense that a more

[4]Klipp (2005, 8–9) makes a clear distinction between 'black-box' input-output correlations and realistic models in which known mechanistic details are taken into account: "It must be noted that different system structures may produce similar system behavior (output). The structure determines the behavior, not the other way around. Therefore the system output is often not sufficient to predict the internal organization [. . .] The intention of modeling is to answer particular questions. Modeling is, therefore, a subjective and selective procedure. It may, for example, aim at predicting the system output. In this case it might be sufficient to obtain precise input-output relation, while the system internals can be regarded as black box. However, if the function of an object is to be elucidated, then its structure and the relations between its parts must be described realistically".

detailed description is not required for the immediate purpose of explaining how the components of the mechanism produce the phenomenon in 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 (Fig. 16.2 above), 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.[5] 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 and how these molecular components interact the way they do, how these components were produced in the cell or how they evolved. The expectation here is that certain changes would not influence in any way the ability of the mechanism to produce its target phenomenon. Most notably, the NF-κB activator, its DNA binding motifs and the IκB inhibitor could tolerate significant changes in sequence and structure, yet the mechanism would continue to function undisturbed on condition that some key features are preserved, such as the dual activation pathway and the affinity and kinetics of chemical interactions.[6] There is therefore a clear sense in which certain lower-level details can be ignored and the phenomenon of interest can be satisfactorily explained in terms of higher-level description of mechanistic components, their properties, organization and activities.

Likewise, a tight quantitative match between the predictions of the model and experimental measurements support the claim that, at least relative to the timeframe 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 the ability to produce it. It is expected therefore that an in vitro reconstituted NF-κB regulatory mechanism should produce peaks of gene activation closely resembling those produced in vivo, thus acting as an independent module. Again, this specifies a sense in which a more systemic context can be ignored such that a satisfactory explanation can be focused on a local mechanism.

[5]For a more detailed discussion, see (Baetu 2014a).

[6]This occurs, for example, when complementary mutations in several components rescue the wild-type phenotype.

16.5 Conclusion

Mathematical modeling provides an accessible substitute for something which is missing in biology: a rich theoretical apparatus from which one could derive detailed hypotheses capable of guiding experimental research from an initial description of the phenomenon of interest to the final explanation. In the absence of such theories, 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 it cannot rival with the all encompassing theories of physics, mathematical models can nevertheless provide a useful workaround by providing a principled way of evaluating the completeness of the information included in a mechanistic explanation, thus specifying where a mechanistic explanation can safely bottom out and what is the cutoff point beyond which external factors can be ignored.

Beyond the philosophical interest relative to the problem of reductionism, there are practical implications to be considered as well. During the discovery process, evidence that an explanation is satisfactory is an indication 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 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, and how the mechanism as a whole integrates the greater whole which the living organism, in the absence of evidence that the mechanism described in the proposed explanation can produce the phenomenon to be explained. At various points in project, researchers can stop, recompose the many bits and pieces of experimental results into 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, at the level of composition at which they are described, can indeed produce the phenomena which they are supposed to explain. Furthermore, since mechanistic explanations often provide the rationale for developing technologies for gaining control over phenomena and medical treatments, evidence that the explanation is satisfactory is key for making an enlightened decision about how much trust to put on the probability of a successful outcome, especially when there is a little room for trial and error.

Acknowledgments This work was supported 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.

Bibliography

Baetu TM (2012) Filling in the mechanistic details: two-variable experiments as tests for constitutive relevance. Eur J Philos Sci 2(3):337–353

Baetu TM (2013) Chance, experimental reproducibility, and mechanistic regularity. Int Stud Hist Philos Sci 27(3):255–273

Baetu TM (2014a) From mechanisms to mathematical models and back to mechanisms: quantitative mechanistic explanations. In: Braillard P-A, Malaterre C (eds) Explanation in biology. an enquiry into the diversity of explanatory patterns in the life sciences. Springer, Dordrecht

Baetu TM (2014b) Models and the mosaic of scientific knowledge. The case of immunology. Stud Hist Philos Biol Biomed Sci 45:49–56

Baetu TM, Hiscott J (2002) On the TRAIL to apoptosis. Cytokine Growth Factor Rev 13:199–207

Baetu TM, Kwon H, Sharma S, Grandveaux N, Hiscott J (2001) Disruption of NF-kB signalling reveals a novel role for NF-kB in the regulation of TNF-related apoptosis-inducing ligand expression. J Immunol 167:3164–3173

Bechtel W (2012) Understanding endogenously active mechanisms: a scientific and philosophical challenge. Eur J Philos Sci 2:233–248

Bechtel W, Abrahamsen A (2005) Explanation: a mechanist alternative. Stud Hist Philos Biol Biomed Sci 36:421–441

Bechtel W, Abrahamsen A (2010) Dynamic mechanistic explanation: computational modeling of circadian rhythms as an exemplar for cognitive science. Stud Hist Philos Sci Part A 41:321–333

Bechtel W, Richardson R (2010) Discovering complexity: decomposition and localization as strategies in scientific research. MIT Press, Cambridge, MA

Brigandt I (2013) Systems biology and the integration of mechanistic explanation and mathematical explanation. Stud Hist Philos Biol Biomed Sci. doi:10.1016/j.shpsc.2013.06.002

Craver C (2001) Role functions, mechanisms, and hierarchy. Philos Sci 68:53–74

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

Craver C (2007) Explaining the brain: mechanisms and the mosaic unity of neuroscience. Clarendon, Oxford

Darden L (2006) Reasoning in biological discoveries: essays on mechanisms, interfield relations, and anomaly resolution. Cambridge University Press, Cambridge

Elowitz M, Leibler S (2000) Synthetic Gene oscillatory network of transcriptional regulators. Nature 403:335–338

Glennan S (2002) Rethinking mechanistic explanation. Philos Sci 69:S342–S353

Hoffmann A, Levchenko A, Scott M, Baltimore D (2002) The I κB – NF-κB signaling module: temporal control and selective gene activation. Science 298:1241–1245

Hooker C (2011) Philosophy of complex systems. Elsevier, New York

Kitcher P (1984) 1953 and All that: a tale of two sciences. Philos Rev 93:335–373

Klipp E (2005) Systems biology in practice: concepts, implementation and application. Wiley, Weinheim

Machamer P, Darden L, Craver C (2000) Thinking about mechanisms. Philos Sci 67:1–25

McKay Illari P, Williamson J (2012) What is a mechanism? thinking about mechanisms *across* the sciences. Eur J Philos Sci 2(1):119–135

Morange M (2009) Synthetic biology: a bridge between functional and evolutionary biology. Biol Theory 4(4):368–377

Nicholson D (2013) Organisms ≠ machines. Stud Hist Philos Biol Biomed Sci. doi: 10.1016/j.shpsc.2013.05.014

Oyama S (1985) The Ontogeny of information: developmental systems and evolution. Cambridge University Press, New York

Rosenberg A (2006) Darwinian reductionism, or, How to stop worrying and love molecular biology. University of Chicago Press, Chicago

Ting AY, Endy D (2002) Decoding NF-κB signaling. Science 298:1189–1190

Waters CK (1990) Why the Anti-reductionist consensus won't survive: the case of classical Mendelian genetics. In: Proceedings to the biennial meeting of the philosophy of science association, pp 125–139
Weber M (2005) Philosophy of experimental biology. Cambridge University Press, Cambridge
Woodward J (2002) What is a mechanism? a counterfactual account. Philos Sci 69:S366–S377
Woodward J (2003) Making things happen: a theory of causal explanation. Oxford University Press, Oxford