

THE IDEALIZATION OF CAUSATION IN MECHANISTIC EXPLANATION*

Alan C. Love[†] and Marco J. Nathan[‡]

Philosophy of Science (forthcoming)

Abstract

Causal relations among components and activities are intentionally misrepresented in mechanistic explanations found routinely across the life sciences. Since several mechanists explicitly advocate accurately representing factors that make a difference to the outcome, these idealizations conflict with the stated rationale for mechanistic explanation. We argue that these idealizations signal an overlooked feature of reasoning in molecular and cell biology—mechanistic explanations do not occur in isolation—and suggest that explanatory practices within the mechanistic tradition share commonalities with model-based approaches prevalent in population biology.

1 More Thoughts About Mechanisms

The concepts of *mechanism* and *mechanistic explanation* have recently received much attention in philosophy of science. This increased scrutiny has had a polarizing effect. On the one hand, supporters suggest that thinking about mechanisms sheds light on many central issues, such as causation, explanation, reduction, and emergence. For instance, it has been claimed that the “open-endedness” of mechanistic explanations, which are not limited to linguistic representations and may involve diagrams or simulations, constitutes a substantial advantage over deductive-nomological inferences (Bechtel and Abrahamsen 2005). On the other hand, critics have argued that these concepts are insufficiently characterized or suffer from distinctive problems. For example, systems

*Acknowledgements: Both authors contributed equally to this work. Earlier versions were presented at the 2014 *Philosophy of Biology at Madison* workshop and the 2014 *PSA Biennial Meeting* in Chicago. We are grateful to both audiences and to an anonymous reviewer for helpful feedback.

[†]Department of Philosophy, Minnesota Center for Philosophy of Science, University of Minnesota, 831 Heller Hall, 271 19th Ave S, Minneapolis, MN, 55455. Email: aclove@umn.edu

[‡]Department of Philosophy, University of Denver, 264 Sturm Hall, 2000 E. Asbury Ave, Denver, CO, 80208. Email: marco.nathan@du.edu

biology and neuroscience allegedly violate two central features of mechanistic explanation: *localization* and *decomposition* (Silberstein and Chemero 2013).

Analyses of mechanisms have entered the mainstream of philosophy of science in no small part because of their prominence in scientific literature. One of the purported advantages of focusing on mechanisms is the ability to flesh out a notion of causal explanation that corresponds to actual scientific practice.¹ In this article, we argue that the intentional misrepresentation of causal relations, which are the source of explanatory power in a description of a mechanism's components and activities, generates a significant—albeit neglected—problem for the mechanistic framework. We begin by rehearsing how causation is typically depicted in molecular explanations (§2) and then argue that this practice has puzzling implications for extant mechanistic accounts (§3). We conclude by addressing two objections (§4) and sketching an alternative solution that reveals a methodological continuity between molecular and population biology (§5).

2 Causal Relations in Mechanistic Explanation

Providing a concise yet informative definition of “mechanistic explanation” is no easy task, as different proponents of a “new mechanistic philosophy” provide distinct and incompatible views of the relevant notions (Woodward 2013). For example, whereas Craver (2007) embraces a more restricted conception of mechanism, which emphasizes the importance of providing as much detail as possible in its description, Bechtel (2011) prefers a more ecumenical interpretation. In an attempt to keep our analysis general and to avoid technical disputes, we treat mechanistic explanation as the claim that many areas of science explain by decomposing systems into their constituent parts, localizing their characteristic activities, and articulating how they are organized to produce a particular effect. Thus, instead of providing a systematic account of the structure of mechanistic explanation, our emphasis will be on a core conception intended to capture common ground among various approaches. *Mechanistic explanations illustrate and display the generation of specific phenomena by describing the organization of a systems constituent components and activities.*

Mechanistic explanations in molecular and cell biology typically involve both a verbal description and a pictorial depiction (fig. 1). The specific details of this complex molecular process can be ignored here; the important point is that this kind of mechanistic representation is ubiquitous in biology. Causal relations are regularly represented by arrows, sometimes with “+” signs denoting initiation or activation and “-” signs denoting inhibition or repression. These depicted causal relations are simplified and reflect only a small subset of those occurring in the cellular context. For example, figure 1 does not distinguish between different kinds of biochemical interactions and ignores background conditions, such as features of the cytological environment or the exact duration of the

¹There is a longstanding debate about whether *all* explanation is causal. Here we only assume that causal explanations are *an* important subclass of scientific explanations and especially salient in discussions of mechanisms.

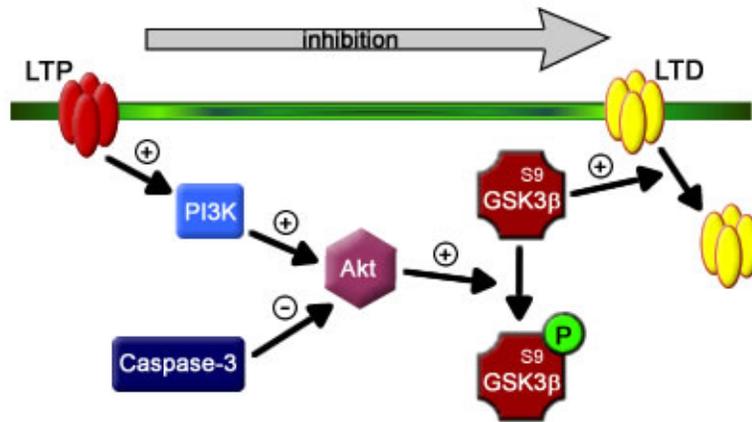


Figure 1: A pictorial depiction of the mechanism for the inhibition of a lasting reduction in synaptic transmission—Long Term Depression (LTD)—by intracellular interactions initiated by the lasting enhancement of synaptic transmission—Long Term Potentiation (LTP). (Source: <http://www.bris.ac.uk/synaptic/research/projects/mechanisms>; reproduced with permission.)

process. The arrows simply stand in for causal relations, regardless of how they are instantiated. In short, the typical representation of biochemical components as distinguishable geometrical shapes and the exclusion of known components involves *abstraction*: the intentional omission of detail.

Abstraction must be distinguished from *idealization*, the deliberate misrepresentation of detail in a model. Paraphrasing Godfrey-Smith (2009), abstract descriptions (such as a vector representation of forces in physics) “leave out a lot”; in contrast, idealized descriptions (such as the billiard ball model of a gas) “fictionalize in the service of simplification.” Although the significance of both abstraction and idealization in model construction is well known (Weisberg 2013), the literature on mechanistic explanation has stressed the former (Levy and Bechtel 2013). This is striking given that some of these idealizations are localized to the exact place in the mechanism description where the explanatory force obtains—the causal relations.

3 Causal Relations Idealized

The practice of abstracting and idealizing in mechanistic explanations is unsurprising. The cell and its myriad constituents compose an extremely sophisticated apparatus; a realistic representation of this plethora of entities and interactions—assuming that such a “complete” depiction is even feasible—would make the description impractical and the explanation unilluminating. Abstraction and idealization are thus necessary practices. Features that do not play a

central role in explanations can (and should) be abstracted away or distorted to make models more perspicuous. Mechanistic models are no exception; they should contain all and only the core explanatory components. Although determining exactly which elements should be included or depicted accurately in a mechanistic explanation constitutes a substantial question (Strevens 2008), the basic criterion of inclusion is straightforward: it involves deciding whether an accurate description of the element contributes to the explanation. A corollary of this principle is that we would not expect features that play a central explanatory role to be abstracted away or distorted in a mechanistic description. Yet, in molecular biology, the causal relations responsible for the explanandum are deliberately misrepresented on a regular basis.

As an illustration, consider *gene expression*, the process by which a portion of DNA is transcribed into RNA. This RNA is then translated into a protein product that can interact with distinct biochemical moieties or can be active in its own right within specific cellular processes (as in the case of enzymes or regulatory molecules). Descriptions of gene expression are paradigmatic examples of mechanistic explanations (Robins and Craver 2009). First, the explanandum is well understood and formulated precisely. Second, the component entities and their activities have been thoroughly investigated and described in detail. Finally, the structure of the system is well defined, including its spatial and temporal organization. Despite the impressive achievements of the last few decades, our current knowledge of gene expression remains incomplete and additional details are being uncovered constantly. Thus, current accounts of gene expression can be viewed as *mechanism sketches*, temporary depictions awaiting further detail, or *mechanism schemas*, general representations that abstract away from specific detail (Darden 2006). Although these incomplete mechanistic models fall short of *ideally complete mechanistic descriptions*, they are explanatory nonetheless. These models facilitate better predictions and can be verified or controlled through surgical experimental manipulations. Despite inevitable empirical shortcomings, gene expression is one of the most extensively studied and best-known mechanisms in molecular biology; if it cannot be understood and explained mechanistically (in a robust sense), it is hard to see what else would fit the bill.

How is causation treated in these mechanistic models of gene expression? Standard diagrammatical depictions (fig. 2) share both the abstractions and idealizations exemplified in the mechanistic description for LTD inhibition (fig. 1). In the gene expression model, the binding of a transcription factor molecule to a DNA binding site upstream of transcription initiation regulates the gene by triggering the transcription of DNA into RNA and, subsequently, the translation of RNA into protein. While these kinds of diagrammatic representations are common in textbooks, one finds increasingly detailed representations of gene expression and more precise narrative descriptions of the mechanism in more advanced discussions (fig. 3; Ptashne and Gann 2002). This more specific description of the apparatus for the regulation of eukaryotic gene expression exposes a variety of abstractions that were present in figure 2. For instance, figure 3 shows that transcription factors operate in conjunction when binding

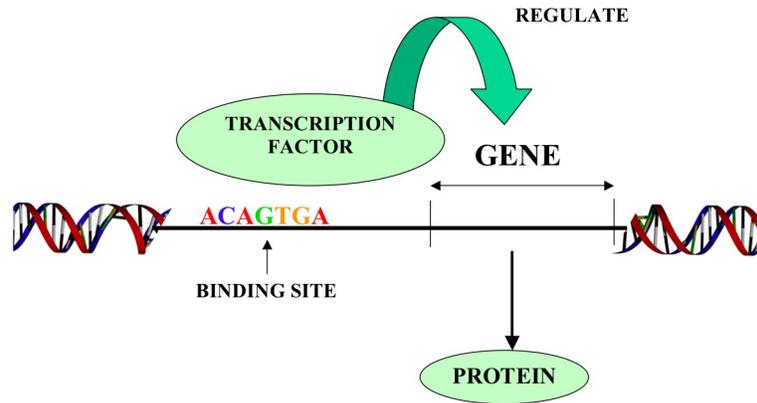


Figure 2: A diagrammatic representation of gene expression. The transcription factor molecule binds to the DNA at its binding site, and thereby regulates the production of a protein from a gene. (Source: <https://howardhughes.trinity.duke.edu/blogs/2011/06/24/zinc-fingers-plasmids-and-my-battle-with-e-coli>; reproduced by permission of Saurabh Sinha.)

to the upstream promoter region (represented by blue-colored ovals, such as TFIID) and also require the operation of cofactors (represented by peach-colored ovals: CRPS/ARP), both of which were omitted in figure 2. Importantly, figure 3 is neither a complete description of gene expression nor the most complete description presently available. Many other necessary intermediary steps and components are known, such as enzymes that catalyze biochemical reactions. But—and this is the crucial point—the lack of further detail does not undermine the explanatory force of these diagrams, which identify core features for the mechanism of gene expression. Differences in the level of incorporated detail depend on the specific explanatory goals of concrete investigative contexts.

These features of mechanistic models should be neither controversial nor problematic. Trouble begins when we combine the above discussion with the assumption, commonplace among philosophers of science, that an adequate causal explanation (whether or not it is mechanistic) should include all and only the difference makers that ensure the production of the explanandum (Waters 2007). This generates a tension because the ways in which the various components purportedly make a difference to gene transcription are misrepresented in the above diagrams. Consider the green-colored oval labeled “transcription factor” in figure 2, which represents the binding of a transcription factor to an operator site. This feature marks the first of three important misrepresentations in mechanistic explanations.

- (i) The operator site does not bind a *single* molecule; as indicated in figure

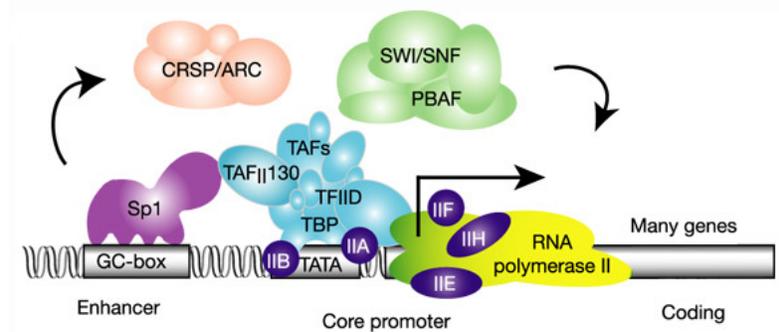


Figure 3: A more detailed diagram of gene expression, which emphasizes the many different components that combine to initiate transcription. Reprinted by permission from Macmillan Publishers Ltd: Nature (Vol. 424, 148), copyright (2003).

3, the site binds a *complex* of molecules.

While this might seem to be a pedantic observation, it is critical for understanding gene expression. Under normal circumstances, individual molecules do not act as difference makers, but complex functional units do. Thus, the diagram does not “merely leave things out” (abstraction), but “fictionalizes in the service of simplification” (idealization). Highlighting the significance of molecular complexes brings us to the second misrepresentation.

- (ii) Figures 2 and 3 depict gene expression as triggered by a single transcription factor or, more accurately, a single complex of molecules—call this functional unit p_1 . While p_1 unquestionably plays a role in the process, it is not a difference maker by itself; its presence (or absence) makes virtually no difference to the outcome. This is because even if p_1 was not there, another molecular complex of the same type ($p_2, p_3, \dots, p_{546}, \dots$) would take its place.

Some readers might protest that the entities in the diagrams are meant as *types*, not *tokens*. Specifically, the transcription factor oval represents not an individual molecular complex (p_1), but any functional unit that attaches to the binding site. What this schema does is capture, in simplified form, the necessary and sufficient conditions for gene expression. The problem with this response is that the conditions depicted in the diagram are not sufficient. Transcription factors and their associated molecular complexes do not remain indefinitely attached to the operator site once they bind. These functional units are constantly dislodged and other molecules or functional unit tokens of the same type take their place. Hence, the causal initiation of gene expression cannot be represented accurately as an individual binding event; a diachronic sequence of binding events involving

many different molecules is required. Individual molecular complexes are not difference makers but represent, at best, *necessary* conditions (if interpreted at the type level). These mechanism depictions contain idealized components that do not make a difference to the outcome of the process.

A third misrepresentation of causal relations amplifies the difficulty because it is not just that individual molecular complexes are insufficient as difference makers in the mechanism; in addition, known difference makers are intentionally omitted from the representation.

- (iii) A critical part of what makes a difference as to whether or not the gene is transcribed is the concentration of transcription factor present in the entire system (Nathan 2014).² Hence, the problem is not merely that the diagrams introduce components that do not make a difference, but that they fail to include entities that do play an actual difference-making role.

In order to circumvent this difficulty, one might reply that the oval is really shorthand for the concentration of transcription factor and, therefore, it implicitly represents the entire pattern of actual bindings. This response is inappropriate for two reasons. First, it is contradicted by the explicit descriptions associated with the diagrams, which do not invoke concentrations.³ Second, the ovals that represent transcription factors and other proteins in figures 2 and 3 fail to capture the main feature of concentrations. Suppose that a 3:1 relative concentration of transcription factor to repressor is necessary to activate transcription. The diagrams do not contain this information about relative concentration, either explicitly or implicitly, and are consistent with concentrations that do not make a difference (e.g., 2:1). Biologists are perfectly aware of the difference-making role of concentrations in cellular contexts (Ptashne and Gann 2002), but they deliberately misrepresent them. The causal relations that produce the explanandum are idealized in mechanistic diagrams; their representation intentionally ignores known variations in properties and other components that make an actual difference.

This feature takes on added significance because it clashes with a widespread and explicit criterion of adequacy for mechanistic explanations.

How-possibly models are often heuristically useful in constructing and exploring the space of possible mechanisms, but they are not adequate explanations. How-actually models, in contrast, describe real components, activities, and organizational features of the mechanism that in fact produces the phenomenon. They show how a mechanism works, not merely how it might work (Craver 2007, 112).

²More precisely, the difference maker is a *relative concentration*—the concentration of transcription factor relative to the concentration of repressor that would inhibit transcription of the gene were it to bind to the operator site.

³For figure 2: “a transcription factor molecule binds to DNA at its binding site, *and thereby regulates* the production of protein from a gene” (emphasis added). For figure 3: “The regulation of gene expression usually depends on DNA sequences located immediately 5’ of the transcription start site . . . Most core promoters contain a TATA element, which serves as a binding site for TBP (TATA-binding protein)” (Levine and Tijan 2003, 147-48).

If the actual difference-making causes are idealized, they do not show how the mechanism actually works. The dilemma should now be apparent. A practice widely utilized in describing mechanisms—the deliberate misrepresentation of the productive continuity between difference-makers—conflicts with the explicit goal of accurately representing causal relations, which is often taken as the hallmark of mechanistic explanation. The idealization of causal relations demonstrates that these models do not depict how the mechanism actually works. If actual difference makers are represented in such a way that they are not difference makers, according to what is already known about the mechanism, mechanistic explanations appear to fail according to their own criteria.

4 Objections and Replies

There are various strategies that proponents of mechanistic explanation can adopt to address the dilemma discussed above. As several commentators have noted, the new mechanistic philosophy tends to be divided on one fundamental point. Whereas one set of authors emphasizes the importance of providing mechanistic models that are as complete and specific as possible (Darden 2006; Craver 2007), others have recognized, more or less explicitly, the importance of abstracting away from unnecessary details (Bechtel 2011; Levy and Bechtel 2013). Depending where one falls on this spectrum, there are (at least) two different strategies that a mechanist could employ to address the problem.

Philosophers who value completeness and richness of detail in mechanistic explanations might respond by appealing to a distinction between *complete mechanistic descriptions*, *mechanism sketches*, and *mechanism schemas*. Diagrams such as figures 2 and 3 were never intended as complete descriptions; they are preliminary sketches that need to be progressively filled out or general schemata that capture salient features in detail but abstract away from others. On this view, the idealization of causation only appears as a dilemma because of our current lack of knowledge. Once all the details have been figured out, the various constituents and activities connecting them will be specified, no difference maker will be left idealized, and the description will become complete and fully explanatory.

This reply has the merit of emphasizing an important aspect of scientific explanation. It is possible to explain phenomena, when knowledge is lacking, by using terms that function as placeholders until more detailed descriptions become available. Indeed, the history of science is replete with such episodes, from Darwin's black-boxing of the mechanisms of ontogeny to the attempts of early psychologists to explain mental processes while ignoring the underlying neural mechanisms. Nevertheless, this argument ultimately fails to address the dilemma because the gradual elimination of *idealized* diagrams is rarely—if ever—witnessed in scientific practice. Even when the relevant details are known, researchers do not replace idealized causal relations with more accurate or realistic representations. Figure 2 deliberately abstracts away from and idealizes known details for the sake of simplicity and perspicuity. Furthermore, even

when additional details are provided, as in figure 3, many idealizations remain unacknowledged and uncorrected. While mechanists committed to the explanatory virtue of completeness might view this as a flaw in biological practice, we argue that it signals an overlooked feature of reasoning in molecular and cell biology: mechanistic explanations do not occur in isolation (see below, §5).

Not all mechanistic philosophers are committed to the ideal of completeness. For example, Levy and Bechtel (2013) argue that it is necessary to abstract away from the structural specifics of a mechanism and represent it in a skeletal, coarse-grained manner in order to understand its organization: “[I]t is often the connectivity, treated abstractly, that explains why a mechanism exhibits the particular behavior it does” (245). While the recognition of the importance of abstraction in schematic representations of mechanisms is welcome, it fails to resolve the problems raised by the idealization of causation. Although philosophers of science tend to address abstraction and idealization in similar ways, these two features pose different and independent issues for explanation. The widespread use of irreducible abstractions challenges the ideal of descriptive completeness, but it is compatible with the goal of describing how mechanisms *actually* work; abstractions make the model more perspicuous. Idealizations, in contrast, provide a further layer of complexity as they overtly violate the actuality requirement. The introduction of deliberate misrepresentations in a model directly clashes with the claim that mechanistic representations should represent how systems (or their subcomponents) actually work. A mechanist could respond by relaxing the actuality requirement and denying that realistic descriptions are necessary for mechanistic explanation. But if we give up the criterion of realism, it becomes unclear what exactly is doing the explanatory work in mechanistic descriptions.

5 Multiple Modeling in Molecular Biology

In §3, we argued that extant accounts of mechanistic explanation face a problem in accommodating the deliberate misrepresentation of causal relations among components and activities that play a difference-making role in producing the explanandum. In §4, we considered two responses to the dilemma that ultimately failed. How then should we interpret the idealization of causation? One promising approach is to recognize that mechanistic explanations do not occur in isolation. This strategy has been neglected, we surmise, for two reasons. First, the solution requires distinguishing different—but equally important—forms of idealizations that are often lumped together. Second, it weakens ontological commitments to mechanisms that several philosophers regard as important.

Weisberg (2013, Ch. 6), distinguishes between three kinds of idealization. The first is what he dubs *Galilean idealization*: the practice of introducing distortions for the sake of simplifying theories. Despite their prominence in science—especially in fields, such as computational chemistry, which require simplifying assumptions for the sake of computational tractability—Galilean idealizations are not germane to pictorial depictions of mechanisms. This is be-

cause the practice is largely pragmatic and non-permanent. Theorists idealize for reasons of computational tractability that are a function of our cognitive inabilities and with the expectation of future de-idealization in more accurate representations. Neither of these comports with the mechanism depictions discussed above (figs. 1-3). There is no expectation or practice of de-idealization and the misrepresentations are not required for pragmatic reasons like computational tractability. Philosophers tend to assume, more or less explicitly, that mechanistic descriptions involve a different kind of practice that Weisberg calls *minimalist idealization*.

Minimalist idealizations introduce only the causal factors that make a difference in producing a phenomenon, a desideratum that has been refined and endorsed by recent accounts of causal explanation (Strevens 2008). While Galilean and minimalist idealizations could, in principle, lead to identical models that idealize the same features, there would be important differences in the justification of these strategies (Weisberg 2013, 105ff). Whereas a Galilean idealizer would claim that the variables are introduced because of pragmatic usefulness, a minimalist would maintain that the model captures the relevant causal factors and, therefore, one should not expect the idealizations to be gradually eliminated as science progresses. Although minimalist idealization fits this feature of scientific practice better than the Galilean strategy, it fails to address the idealization of causation in mechanistic explanations. This is because the general goal of minimalist idealization (i.e., include all and only the relevant causal factors) is hard to reconcile with the deliberate misrepresentation of precisely those factors that account for the explanandum. This violates the assumption that the explanatory power of a model is related to its representation of the relevant difference makers.

Weisberg's third type of idealization, *multiple-model idealization*, involves constructing multiple, related-but-incompatible models that capture distinct aspects of the causal structure of a complex system. Multiple-model idealization has many parallels with the above strategy of relaxing the actuality criterion for mechanistic descriptions. The key difference is that multiple-model idealization does not aim to produce a single best model. Understanding how the explanandum is produced derives from comparing and contrasting different models. Although new knowledge can be added to individual models, there is no expectation that the idealizations will be progressively removed or that the need for multiple models will fade over time.

The strategy of multiple-model idealization explains each of the three misrepresentations identified in §3, suggesting a way to interpret the idealization of causation in mechanistic explanation. First, it accounts for the simultaneous use of coarser and finer models of the same phenomenon. Figures 2 and 3 provide different representations of gene expression, but they are not competing descriptions. These models coexist because they are not intended to be exhaustive. Various models are required to account for the occurrence of difference-making relations. In this respect, multiple-model mechanistic explanations in molecular and cell biology are closely associated with the "model-based science" of population phenomena in ecology and evolutionary biology (Godfrey-Smith 2006). In

both cases, it is a mistake to assume that more accurate models should always replace less accurate ones. *Descriptive accuracy* is an explanatory virtue, but one that can only be obtained at the expense of another virtue: *generality*.

Second, the multiple-models approach shows how it is possible to offer mechanistic models that are *dynamic* with respect to transition events. While concentration-based stochastic models would represent the dynamics of gene expression more accurately, eliminating the deterministic model would entail a significant loss of perspicuity and generality. This explains why few, if any, textbooks discuss concentrations. The reason is not, as is often assumed, that these models are too complex for beginning scholars; rather, they are *too specific*. Simpler deterministic models are more general and, therefore, applicable to a greater variety of phenomena. The mechanism of gene expression is explained with multiple models; stochastic models can (and should) be used to complement deterministic models where specific descriptions are required.

Third, the multiple-modeling approach allows scientists to offer models that can treat idealized difference makers as separate mechanisms. This permits exploring the models quasi-independently. The lesson here is that concentration-based, quantitative stochastic models are not merely adding details left out or misrepresented by single-molecule, qualitative models. The two models need not be considered alternative descriptions of the same mechanism. Rather, biologists can interpret these models as representations of different mechanisms that capture and explain different features of a single process. Given that these models can be used in different ways for different aims, asking which representation is better or more perspicuous is an ill-posed question. Criteria for the scope and detail of mechanistic models always depend on the explanandum in view.

In sum, the misrepresentation of causation in mechanism descriptions can be addressed by interpreting it as part of a multiple-modeling strategy. The goal of mechanistic explanation is not an all-inclusive single model, but a series of many complementary diagrams and descriptions comprising different idealizations, similar to what is observed in population biology. We are not suggesting that multiple-modeling is the *only* significant form of idealization, or even the most important one. Our claim is that if the idealization of causation in mechanistic explanation is understood in this way, the stated dilemma can be resolved.

6 Concluding Remarks

In conclusion, we offer two general remarks. First, abstraction and idealization are essential and irreducible features of scientific representation. This is because *describing* and *explaining* involve a necessary tradeoff between explanatory power and descriptive accuracy (Cartwright 1983). Whereas accuracy encourages focusing on individual tokens or events, explanation presupposes that they can be classified under general types. This requires abstracting away or idealizing the differences between the tokens or events. The broader and more diverse the class of items to be grouped together, the more extensive and radical

the abstractions and idealizations will become. This is precisely the situation encountered by biologists dissecting molecular mechanisms: generating a more precise explanation of entities and causal processes that have been misrepresented, black-boxed, or treated as placeholders in the original mechanism description requires new models that will contain different abstractions and idealizations. Together, these models increase explanatory power without sacrificing generality.

Second, as noted above, a major reason why mechanists tend to assume a minimalist approach to idealization is to preserve the ontological import of explanations: “explanations [mechanisms] are objective features of the world” (Craver 2007, 27). The multiple-models approach is harder to reconcile with this kind of objectivity, since inconsistent models cannot be simultaneously true. The multiple model strategy implies that we should moderate the ontological implications drawn from our mechanistic models. This is a meaningful philosophical moral: distilling metaphysical implications from scientific explanations requires close attention to explanatory practice. Regardless of this broader methodological lesson, our analysis of a widespread but neglected feature of mechanistic explanation—the idealization of causation—offers a new perspective on the ubiquitous appeals to mechanisms in molecular and cellular biology.

References

- Bechtel, William. 2011. “Mechanism and Biological Explanation.” *Philosophy of Science* 78:533–557.
- Bechtel, William and Adele Abrahamsen. 2005. “Explanation: A Mechanist Alternative.” *Studies in the History and Philosophy of Biology and Biomedical Sciences* 36:421–441.
- Cartwright, Nancy. 1983. *How the Laws of Physics Lie*. Oxford: Clarendon.
- Craver, Carl F. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. New York: Oxford University Press.
- Darden, Lindley. 2006. *Reasoning in Biological Discoveries: Essays on Mechanisms, Interfield Relations, and Anomaly Resolution*. New York: Cambridge University Press.
- Godfrey-Smith, Peter. 2006. The Strategy of Model-Based Science. *Biology and Philosophy* 21:725–740.
- 2009. “Abstractions, Idealizations, and Evolutionary Biology.” In *Mapping the Future of Biology: Evolving Concepts and Theories*, ed. Anouk Barberousse, Michel Morange, and Thomas Pradeu, 47–55. Berlin: Springer.
- Levine, Michael. and Robert Tijan. 2003. “Transcription Regulation and Animal Diversity.” *Nature* 424:147–151.

- Levy, Arnon and William Bechtel. 2013. "Abstraction and the Organization of Mechanisms." *Philosophy of Science* 80:241–261.
- Nathan, Marco J. 2014. "Causation by Concentration." *British Journal for the Philosophy of Science* 65: 191–212.
- Ptashne, Mark. and Alexander Gann. 2002. *Genes and Signals*. New York: Cold Spring Harbor Laboratory Press.
- Robins, Sarah K. and Carl F. Craver. 2009. "Biological Clocks: Explaining with Models of Mechanisms." In, *The Oxford Handbook of Philosophy and Neuroscience*, ed. John Bickle, 41–67. New York: Oxford University Press.
- Silberstein, Michael and Anthony Chemero. 2013. "Constraints on Localization and Decomposition as Explanatory Strategies in the Biological Sciences." *Philosophy of Science* 80:958–970.
- Strevens, Michael. 2008. *Depth. An Account of Scientific Explanation*. Cambridge: Harvard University Press.
- Waters, C. Kenneth. 2007. "Causes that Make a Difference." *The Journal of Philosophy* 104: 551–579.
- Weisberg, Michael. 2013. *Simulation and Similarity: Using Models to Understand the World*. New York: Oxford University Press.
- Woodward, James. 2013. "Mechanistic Explanation: Its Scope and Limits." *Proceedings of the Aristotelian Society* 87: 39–65.