

Towards Mechanism 2.1: A Dynamic Causal Approach

Abstract: I propose a dynamic causal approach to characterizing the notion of a mechanism. Levy and Bechtel, among others, have pointed out several critical limitations of the new mechanical philosophy, and pointed in a new direction to extend this philosophy. Nevertheless, they have not fully fleshed out what that extended philosophy would look like. Based on a closer look at neuroscientific practice, I propose that a mechanism is a dynamic causal system that involves various components interacting, typically nonlinearly, with one another to produce a phenomenon of interest.

1. Introduction

The last three decades have witnessed the rise of the so-called *new mechanical philosophy* (NMP) in philosophy of science. The emergence of this NMP was largely motivated by philosophers' realization that, in contrast with the physical sciences where natural laws play a central role in offering explanation, prediction and understanding, the life sciences are best characterized as a hodgepodge of subdisciplines that focus on discovering and investigating mechanisms. Another motive for the NMP's arising is related to the shift from the focus on scientific theories to on scientific practice.

Advocates of the NMP provide philosophers with a new framework for re-examining many pivotal problems in philosophy of science, e.g., scientific explanation, causation, the autonomy of the special sciences, to name just a few. However, even though the NMP has significantly reshaped the landscape of philosophy of science, there is still a long way to go. Recently, many authors have realized that the framework has serious limitations (Brigandt 2013; Levy and Bechtel 2013; Levy and Bechtel 2016). At the heart of these limitations is the fact that previous work tends to center on qualitative aspects of mechanisms and draws on examples primarily from textbooks in cell and molecular biology, while neglects quantitative/dynamic aspects of mechanisms that are reflected in real scientific practice.

Given these limitations, Levy and Bechtel (2016) call for an extended conception

of mechanisms and mechanistic explanation, the so-called ‘mechanism 2.0’.¹ Although Levy and Bechtel, among others,² point in the right direction (or so I suppose) and highlight several crucial points regarding what the extended philosophy would look like, they have not yet fully developed their proposal. So, I here, following in their footsteps, take up the mission of developing one version of such an extended philosophy and call it ‘mechanism 2.1’. My approach, largely inspired by neuroscientific practice, is capable of capturing both the qualitative and quantitative aspects of mechanisms, and dovetails well with real scientific practice.

The essay unfolds as follows. Section 2 briefly describes the NMP, followed by Section 3 where Levy and Bechtel’s proposal for ‘mechanism 2.0’ is introduced. Section 4 proposes a dynamic causal approach to characterizing mechanisms, and Section 5 discusses what philosophical implications it can deliver.

2. The New Mechanical Philosophy

The NMP represents a bundle of closely connected but slightly different ideas

¹ Notice that Levey and Bechtel (2016)’s interest is in expanding the mechanistic explanation framework rather than the conception of mechanisms. However, I think an extended conception of mechanistic explanation must be built upon an extended conception of mechanisms, since the latter is more fundamental. Yet, their project does inform me of how to develop an extended account of mechanisms.

² E.g., Kaplan and Bechtel (2011), and Brigandt (2013).

proposed by a number of philosophers concentrating primarily on practice in the life sciences (Bechtel and Richardson 1993; Machamer et al. 2000; Glennan 2002, 2005; Bechtel and Abrahamsen 2005; Bechtel 2006, 2008; Darden 2006; Craver 2007).

These philosophers all agree that we place mechanisms on center stage when examining those traditional philosophical questions (e.g., explanation, causation), even though they have not yet reached a consensus on how to philosophically specify the notion of mechanisms. According to one most commonly cited characterization:

“Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions.”

(Machamer et al. 2000, 3)

In characterizing mechanisms, different authors employ different terminologies which reflect their distinct ontological commitments.³ Setting aside these ontological disputes, nevertheless, they all seem to agree that a mechanism involves four elements: a phenomenon/behavior, components/parts/entities, interactions/activities/operations,

³ Machamer et al. (2000) take a dualistic stance towards mechanisms, holding that a mechanism is composed of two ontologically different kinds: entities and activities. Bechtel (2006, 2008) also thinks that a mechanism is composed of two different kinds: component parts and component operations. Glennan (2002), by contrast, takes a monist position, holding that a mechanism is composed of parts that interact to produce a phenomenon of interest.

and spatiotemporal organization/structure. Another element, not clearly shown, is also worth mentioning: multilevel hierarchy.

The multilevel hierarchy is manifested by the fact that the component of a mechanism may constitute a sub-mechanism by itself, and that the mechanism may constitute a component of an even bigger mechanism. This also implies that a mechanism's identification hinges on what target phenomenon/behavior is under question. In other words, there is no mechanism *simpliciter*, but only a mechanism *for* a particular phenomenon/behavior. With respect to components and interactions—in terms of Craver (2007)'s constitutively relevant criterion—only those that contribute to producing a particular phenomenon/behavior of the mechanism count as the components and interactions of the mechanism.

This NMP has significant implications for a number of philosophical issues, e.g., explanation. This philosophy advocates a new account of explanation, i.e., mechanistic explanation. According to this account, explaining a phenomenon/behavior (at least in the life sciences) lies in uncovering a mechanism, i.e., uncovering how the various components interact with one another in a spatiotemporally orchestrated manner to produce the phenomenon of interest. Obviously, there is no role for laws to play, and explanation does not proceed in a manner suggested by the covering-law model of scientific explanation.

No doubt, this philosophy's attractiveness essentially comes down to the fact that it goes in concert with the practice in the life sciences. Yet, as many philosophers have pointed out, although this framework has come very close to practice, it does not

come close enough.

3. Mechanism 2.0: Call for An Extension

Recently, many philosophers have cast doubt on the adequacy of the NMP (Bechtel and Abrahamsen 2010, 2013; Brigandt 2013; Levy and Bechtel 2013, 2016).

According to these philosophers, the NMP has the following limitations. First, the NMP treats a mechanism as if it is composed of a *linear* causal sequence. However, scientists have recognized that a mechanism can be a very complex network of interacting components that possesses feedback/feedforward loops, whose interactions are typically non-linear and non-sequential. Second, the NMP routinely concentrates on the structural, organizational, and spatial aspects of a mechanism, ignoring that a mechanism is essentially a *dynamic* system within which the parts are changing over time. Third, these two features, linear and non-dynamic thinking, are always associated with a third feature of that philosophy: *qualitative* thinking. This feature is clearly illustrated by the way the new mechanists qualitatively describe how a mechanism is brought about, and by the simple paradigmatic examples drawn from textbooks (e.g., the *lac* operon of *E. coli*). These qualitative characterizations of mechanisms may help unravel some qualitative aspects of the mechanism, but fall short of making sense of those quantitative, often more important and more complex, aspects.

Due to these limitations, an extended philosophy of mechanisms, accompanied

by an updated account of mechanistic explanation, is called for (Bechtel and Abrahamsen 2010; Brigandt 2013; Levy and Bechtel 2016). However, although Levy and Bechtel (2016), among others, have pointed out the limitations of the NMP and signposted the direction for an extension, they have not fully fleshed out what that extended philosophy would be. For the moment, let me list those key features, as singled out and agreed upon by these philosophers, that an extended conception of a mechanism must be able to capture. First, the extended framework must treat a mechanism as a non-linear, dynamic complex system that may involve feedback/feedforward loops. Second, in addition to the qualitative thinking, the extended framework must also facilitate quantitative thinking. Third, as a result, the extended philosophy must come even closer to real scientific practice. Given these ingredients, it is time to portray the full image.

4. Mechanism 2.1: A Dynamic Causal Approach

I propose that a mechanism is a *dynamic causal system* that involves various components interacting, typically non-linearly (though sometimes linearly), with one another to produce a phenomenon of interest. In agreement with the NMP, my approach also holds that a mechanism involves four elements: a phenomenon/behavior to be explained, components/parts/entities, interactions/activities/operations, and spatiotemporal organization/structure. Besides, it also considers the multilevel character of mechanisms. However, my approach

differs from the NMP in two important aspects. First, it treats a mechanism as a dynamic system that may involve non-linear interactions and feedback/feedforward loops, and second, it *explicitly* views a mechanism as a *causal structure* composed of components and their causal connections (Here I am not denying that many advocates of the NMP also treat a mechanism as a causal structure. The point is that they only do so implicitly or qualitatively. So, by ‘explicitly’ I mean a mechanism is *formally* represented as a causal structure using certain quantitative tools, e.g., causal graphs (Spirtes et al. 2000; Pearl 2009).

This approach does not come out of the blue. Rather, it reflects how scientists—especially those neuroscientists—in practice conceptualize a mechanism (Friston et al. 2003, 2009, 2017; Stephan et al. 2007; Rubenstein et al. 2016). To see how this approach can make sense of scientific practice and therefore offer us an extended conception of mechanisms, consider an example drawn from neuroscience. Neuroscientists wonder how human brains respond to stimuli, e.g., visual words. The question they are asking is what mechanism underlies the observed pattern regarding humans’ response to visual stimuli. To answer this question, they hypothesize a mechanism involving five components (i.e., areas) in the brain: visual areas V1 and V4, the inferior temporal gyrus (BA37), the angular gyrus (BA39), and the superior temporal gyrus (STG). The hypothesized mechanism is depicted below:




Figure 1. A schematic representation of a neuronal mechanism responsible for bringing about the observed stimuli-response pattern in humans. The figure is adapted from Friston et al. (2003, 1275).

Obviously, this mechanism involves feedback loops. Also, the mechanism can be interpreted as a causal structure, for all the arrows, both the one-way and two-way arrows, denote causal connections.⁴ These causal connections are termed *effective connectivity*, denoting “the influence that one neuronal system exerts over another in terms of inducing a response” (Ibid., 1277). As can be seen from the figure, there are two kinds of stimuli/inputs that influence the system: a stimulus can induce a response by either exerting direct influences over a specific region, e.g., u_1 , or exerting indirect effects by modulating the coupling (i.e., the causal connection) among regions, e.g., u_2 . Attention to a particular feature is a case of the second kind of stimulus/input, for differing degrees of attention usually can result in different strengths of the coupling between the same set of regions. In total, there are three

⁴ Notice that this approach differs from the causal graphical theory (Spirtes et al. 2000; Pearl 2009), since it allows cyclic causal structures while the latter does not.

types of interactions: (1) the direct influence of inputs on brain areas, (2) the intrinsic coupling among brain areas, and (3) the modulation of the intrinsic coupling induced by inputs.

We have not yet seen how the mechanism can be dynamic. Given Figure 1, mental simulation may help us roughly understand how the mechanism works, but it offers no help in understanding the mechanism dynamically. To do so, we must be equipped with some mathematical tools. The deterministic differential equations are often the sought-after tools by neuroscientists.⁵ Now, we assign a state variable x_i to each region of the mechanism, describing some neurophysiological properties of that region, e.g., postsynaptic potentials. These state variables can interact with one another, namely, one state variable's change relies at least upon (the change of) one other state variable. The set of interactions between the state variables then can be expressed by a set of ordinary differential equations:

$$\frac{dx}{dt} = \begin{bmatrix} f_1(x_1, \dots, x_n) \\ \cdot \\ \cdot \\ \cdot \\ f_n(x_1, \dots, x_n) \end{bmatrix} = F(x) \quad (1)$$

Yet, this set of equations is insufficient to specify the mechanism. To begin with, the set of equations does not give us any information about the specific form, or the nature, of the causal relationships, f_i . Hence, a set of parameters, denoted by θ , that

⁵ The other options are state space models, iterative maps, etc.

encodes the information about the form and strength of the causal relationships is required. The set of dependence/causal relationships, however, does define the structure/organization of the mechanism (Stephan et al. 2007, 130). Second, since the mechanism is an open system that exchanges matter, energy and/or information with its environment, the inputs into the system, denoted by the vector function $u(t)$, should also be considered. By expanding equation (1) along these two lines, we obtain a general nonlinear state equation for the system:

$$\frac{dx}{dt} = F(x, u, \theta) \quad (2)$$

This equation describes how a state variable's change is a function of some neurophysiological influences exerted by some state variables (including itself at an earlier time) and some inputs, and establishes a mapping between the system dynamics and the system structure. It offers

“A causal description of how system dynamics results from system structure, because it describes (i) when and where external inputs enter the system; and (ii) how the state changes induced by these inputs evolve in time depending on the system's structure. Given a particular temporal sequence of inputs $u(t)$ and an initial state $x(0)$, one obtains a complete description of how the dynamics of the system [...] results from its structure [...]” (Ibid., 130).

The equation is general because it provides an overarching framework for representing neural systems that can be implemented in different ways. One such an implementation, a bilinear approximation,⁶ represents the system dynamics using a bilinear differential equation:

$$\begin{aligned}
 \frac{dx}{dt} &= F(x, u, \theta) \\
 &= Ax + \sum u_j B^j x + Cu \\
 &= (A + \sum u_j B^j)x + Cu
 \end{aligned} \tag{3}$$

where A is the connectivity matrix denoting the intrinsic coupling among brain areas when no input is present, B^j are the induced connectivity matrices denoting the change of the intrinsic coupling induced by the j th input, and C is the matrix standing for the direct influences of inputs on brain areas. Together, they constitute the parameter set $\theta = \{A, B^j, C\}$ to be estimated. With the parameter set at hand, the mechanism represented in Figure 1 can be redrawn below:

⁶ A bilinear approximation is achieved in the following way: the differential equations for each state variable and for each input are linear individually, but nonlinear jointly. For details of this method, see Svoronos et al. (1980).

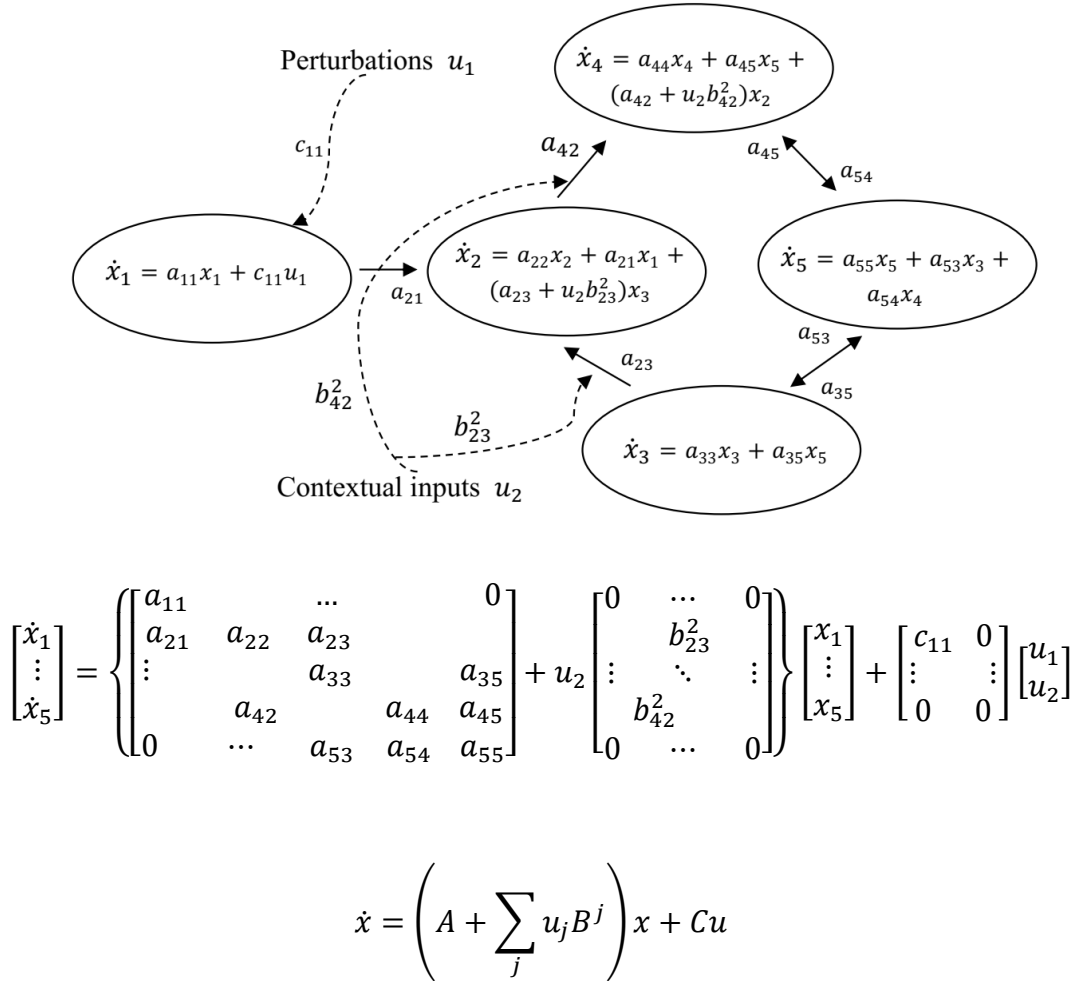
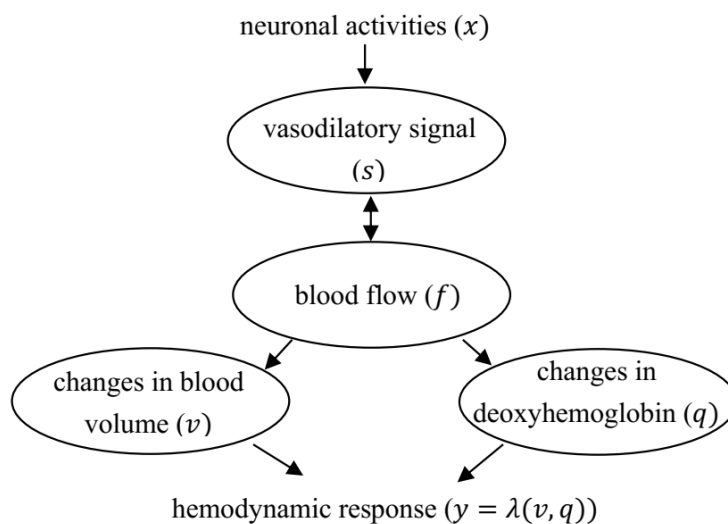


Figure 2. A schema that re-depicts the mechanism in Figure 1 using the differential equations. The lower panel presents the differential equations shown in the upper panel in a matrix form, which can be further simplified using the parameter matrices A , B^j and C . The figure is adapted from Friston et al. (2003, 1279).

In this scenario, each state variable's change, \dot{x}_i , is a function of its own state at an

earlier time, at least one other state variable, and some external inputs.

So far, we have shown in detail how a mechanism can be dynamic, and how a mechanism's dynamic character can be properly captured with the help of certain quantitative tools. However, that is not the end of the story. To fully understand a mechanism, it is standard practice that neuroscientists look deeply into each area of the mechanism and treat each as a dynamic system, i.e., a sub-mechanism.⁷ More specifically, the sub-mechanism in our example is this: changes in neuronal activity induce a vasodilatory signal which results in changes in blood flow, which in turn cause changes in blood volume and deoxyhemoglobin content. Then, blood volume and deoxyhemoglobin content nonlinearly generate measurable responses of that area. The sub-mechanism of each area is depicted below:



⁷ Doing so is partly because each state variable, as representing some neuronal activities, can induce measurable hemodynamic responses, but the causal architecture of the mechanism itself is not observable. So, this is a way to get access to the causal architecture of the mechanism.

Figure 3. A schema that depicts the sub-mechanism of each area of the mechanism. The figure is adapted from Stephan et al. (2007, 133).

This sub-mechanism involves four hemodynamic state variables (s , f , v and q), and a parameter set ϑ . To understand this sub-mechanism dynamically, we, again, need appeal to a set of differential equations that captures the (causal) relationships between these state variables employing the parameter set ϑ .⁸ Finally, we obtain a full picture of the mechanism involving two levels (the mechanism-level and the component-level):

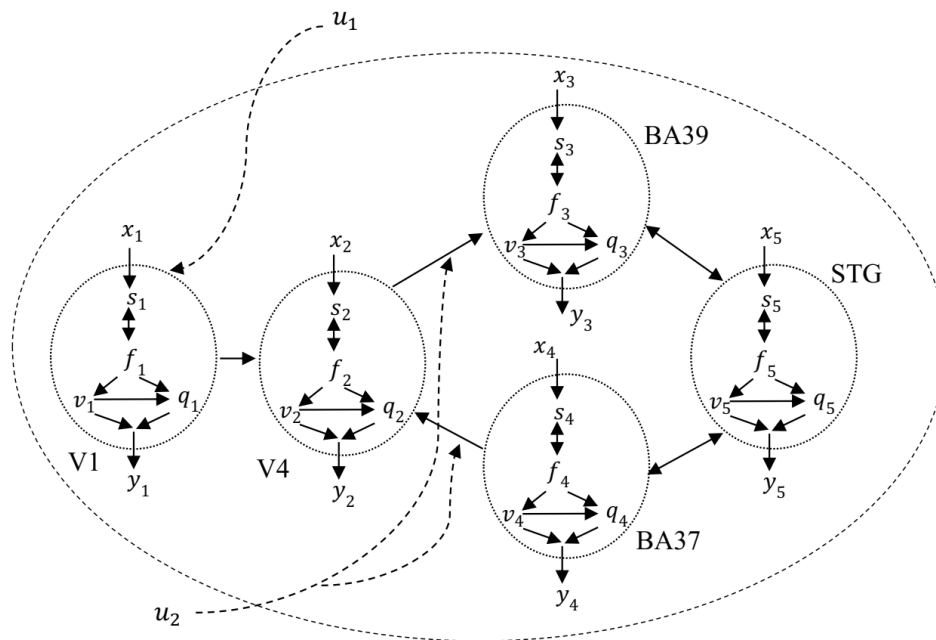


Figure 4. A schema that represents a mechanism and its sub-mechanisms.

⁸ This parameter set and the parameter set θ for the system dynamics constitute the whole parameter set $\{\theta, \vartheta\}$, which can be estimated from the measured signal data using a Bayesian estimation approach. The estimation procedure can be found in Friston et al. (2003).

This schematic graph, as depicting a causal structure, together with the quantitative tools necessary to capture the nonlinear, dynamic aspects embodied in the causal structure, constitute the basis for proposing that a mechanism is a dynamic causal system that involves various components interacting, typically non-linearly, with one another to produce a phenomenon of interest.⁹ The next section will discuss the key features of this approach, and the philosophical implication it delivers.

5. Discussion

5.1. What is a mechanism, again?

The dynamic causal approach shares with the NMP all those important insights regarding the conception of mechanisms. For example, it agrees that a mechanism consists of four basic elements: a phenomenon to be explained, various components, interactions among these components, and a spatiotemporal organization/structure. Moreover, it treats a mechanism as a multilevel system. Figure 4 in the last section

⁹ For the limitations of space, this essay does not fully show how the dynamic, quantitative aspects of the mechanism under consideration are unpacked. For those interested in these details, please see Bechtel and Abrahamsen (2010), where they demonstrate via a similar case, i.e., circadian rhythms, that the dynamic, quantitative aspects can be understood only when certain quantitative tools are employed.

unambiguously reflects this multilevel feature of a mechanism. Furthermore, this approach subscribes to the view that there is no mechanism *simpliciter*, but only a mechanism *for* a particular phenomenon/behavior. In our neuroscientific example discussed above, neuroscientists only singled out five regions of the brain plus their interactions and dismissed all the rest as irrelevant with respect to the stimulus-response pattern in question. Last but not the least, I concur that scientific practice is our best guide to understanding what a mechanism is—that is, we better look at how scientists conceptualize, hypothesize, represent, discover, and entertain mechanisms.

However, a closer look at neuroscientific practice can lead us to some key points overlooked by many new mechanists. First, as some authors have pointed out (Bechtel and Abrahamsen 2010, 2013; Brigandt 2013; Levy and Bechtel 2013, 2016), a mechanism is essentially a dynamic system. Following these authors, I further proposed that a mechanism is a dynamic causal system such that dynamic and causal aspects are a mechanism's defining features. This understanding implies that a qualitative mindset is no longer sufficient to fully understand mechanisms, so that a philosophical conception of mechanisms should be better equipped with a quantitative thinking. Second, many new mechanists emphasize the distinction between entities/parts and activities/interactions. However, an updated philosophy must be able to accommodate the fact that, being a dynamic system, the boundary between entities/parts and activities/interactions may become blurred in some cases. This is the case in our neuroscientific example, where the boundary is clear in the mechanism

involving five regions, but unclear in the sub-mechanisms since their components stand for some quantities that are not clearly entities, e.g., changes in blood flow, changes in blood volume, etc. Though many would think that these quantities are better classified as activities/interactions, the practitioners do not find this classificatory problem worrisome as long as they believe that the state variables denoting them are meaningful and well-defined.

Third, although some philosophers implicitly regard a mechanism as a causal structure, they fail to fully cash out this idea. In my approach, the organization of a mechanism now is explicitly treated as a causal structure that can be quantitatively described using some mathematical tools, e.g., differential equations. The quantitative tools facilitate understanding the nonlinear, dynamic aspects of the causal structure that a qualitative thinking usually stops short of making sense.¹⁰ Also, this dynamic causal approach largely extends the causal graphical theory in characterizing a causal structure, because it allows a causal structure to be cyclic.¹¹ The causal structure involves both spatial and temporal dimensions, as the spatial dimension is clearly represented by Figure 4 and the temporal dimension is captured by the set of differential equations (in which each region's change is a function of its own earlier

¹⁰ So, the quantitative tools also facilitate understanding the linear aspects if there are such aspects.

¹¹ Because the variables in the differential equations are somehow time-indexed, e.g., each variable's change is a function of its own state at an earlier time, the problem of circularity does not arise here.

state, at least one other state variable, and perhaps some external inputs).

Unsurprisingly, the dynamic causal approach ramifies into other issues associated with mechanisms, e.g., mechanistic explanation, the way of representing mechanisms, etc.

5.2. An updated account of mechanistic explanation

I follow those new mechanists in holding that a mechanistic explanation is one that uncovers the underlying mechanism of a phenomenon/behavior of interest. But I further add that a mechanistic explanation is a very complicated practice that often—if not at all times—involves the employment of many different epistemic means, e.g., qualitative tools such schematic drawings and verbal descriptions, and quantitative tools such as causal graphs and differential equations, to unpack the dynamic, causal aspects of a mechanism. This view does not deny the value of qualitative tools in offering mechanistic explanation, but it does insist that those qualitative tools can provide explanation only when the explanatory task does not require us to unravel the dynamic aspects of the mechanism.

So, in accordance with Levy and Bechtel (2016), this view regards mechanistic explanation as dynamic in two related senses: on the one hand, the mechanism itself is a complex, dynamic system, and on the other, the process of constructing, articulating and evaluating a mechanistic explanation based on the mechanism in question is also a dynamic matter. This dynamic nature can be reflected by, but not restricted to, the

following scenarios: some parts of a larger system regarded as irrelevant to explaining a phenomenon of interest at an earlier time may be incorporated into a new explanation that treats them as relevant, an explanation may take a different form when a new mathematical tool is invented or when a new component/interaction is identified, a mechanism may at some later stage be embedded into a larger mechanism to explain a phenomenon of interest, etc.

This view also suspects the dichotomy made between mechanistic and mathematical explanation.¹² Some authors maintain that there is a clear-cut boundary between mechanistic and mathematical explanation and that they are competitors rather than comrades (e.g., Craver 2006; Winter 2006). However, our updated account of mechanistic explanation, based on the dynamic causal approach, is able to show that mathematical elements play an indispensable role in building a mechanistic explanation. This is the case in our neuroscientific example, where the set of differential equations is the key to revealing the dynamic aspects of the mechanism. This position goes in tune with many philosophers who either show that mathematical elements are indispensable for a mechanistic explanation (e.g., Bechtel and Abrahamsen 2010, 2013; Brigandt 2013), or demonstrate that constructing mechanistic explanation in the life sciences usually takes an integrative strategy where both mechanistic and mathematical elements figure prominently and work

¹² Mathematical explanation here narrowly means those using mathematics to explain physical phenomena, rather than those *purely mathematical explanations*. See Colyvan (2012) for the distinction.

collaboratively (e.g., Fagan 2012; Boogerd et al. 2013; Green et al. 2015).¹³

5.3. A new way of representing mechanisms

A new conception of mechanisms is usually coupled with a new way of representing mechanisms, and, on the other hand, a new way of representing mechanisms typically reflects a new conception of mechanisms. This two-way dependence relationship has been instantiated in our neuroscientific example, where neuroscientists' conceptualizing mechanisms as dynamic causal systems urges them to appeal to relevant mathematical tools to capture this dynamic causal nature, and the way they represent mechanisms employing these tools also reveals that they think of the mechanisms as dynamic causal systems. Most prominently, they employ differential equations and causal graphs to capture those dynamic causal aspects of a mechanism.

We must note that there might be different ways of representing mechanisms, which may reflect distinct ways of conceptualizing mechanisms. In fact, Casini et al. (2011) and Gebharder and Kaiser (2014) have proposed two alternatives. Casini et al. (2011) attempt to represent a mechanism as a recursive Bayesian network, where each variable at a higher-level can be described as a sub-mechanism at a lower-level. However, though this approach captures the hierarchical and causal nature of

¹³ Some also argue that the mathematical elements are part of a broader practice of building mechanistic explanations (Kaplan and Craver 2011; Matthiessen 2017).

mechanisms, it seems unclear how it can treat mechanisms as dynamic systems.¹⁴

Gebharter and Kaiser (2014)'s approach comes closer to my approach, for it respects both the dynamic and causal aspects of mechanisms. But it differs from my approach since it brings the dynamics to the scene via adding time index to each variable, e.g., x_{t_1} , x_{t_2} denote *neuron_x firing at t_1* and *neuron_x firing at t_2* . This usually results in a very complicated causal structure and therefore seems unpractical.

Notice that this short section is not intended to assess the plausibility/implausibility of different representational strategies, but rather to point out that there are alternatives available and each may have its own merits and shortcomings.

6. Conclusion

Based on neuroscientific practice, I have proposed a dynamic causal approach to characterizing the notion of mechanisms. This approach shares with the NMP all those insights about mechanisms, but also offers an extended, updated conception that highlights the dynamic causal aspects of mechanisms and that comes closer to real scientific practice.

¹⁴ For a more comprehensive criticism, see Gebharter (2014).

References

- Bechtel, William. 2006. *Discovering Cell Mechanisms: The Creation of Modern Cell Biology*. Cambridge: Cambridge University Press.
- Bechtel, William, and Adele Abrahamsen. 2010. "Dynamic Mechanistic Explanation: Computational Modeling of Circadian Rhythms as an Exemplar for Cognitive Science." *Studies in History and Philosophy of Science Part A* 41 (3): 321–33.
- Bechtel, William, and Robert Richardson. 1993. *Discovering Complexity*. Princeton: Princeton University Press.
- Boogerd, Fred, Frank Bruggeman, and Robert Richardson. 2013. "Mechanistic Explanations and Models in Molecular Systems Biology." *Foundations of Science* 18 (4): 725–44.
- Brigandt, Ingo. 2013. "Systems Biology and the Integration of Mechanistic Explanation and Mathematical Explanation." *Studies in History and Philosophy of Science Part C* 44 (4): 477–92.
- Casini, Lorenzo, Phyllis McKay Illari, Federica Russo, and Jon Williamson. 2011. "Models for Prediction, Explanation and Control: Recursive Bayesian Networks." *Theoria* 26 (1): 5–33.
- Craver, Carl. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Oxford University Press.
- Darden, Lindley. 2006. *Reasoning in Biological Discoveries: Essays on Mechanisms, Interfield Relations, and Anomaly Resolution*. Cambridge: Cambridge University

Press.

Fagan, Melinda Bonnie. 2012. "Waddington Redux: Models and Explanation in Stem Cell and Systems Biology." *Biology & Philosophy* 27 (2): 179–213.

Friston, Karl, Lee Harrison, and Will Penny. 2003. "Dynamic Causal Modelling." *Neuroimage* 19 (4): 1273–302.

Friston, Karl. 2009. "Causal Modelling and Brain Connectivity in Functional Magnetic Resonance Imaging." *PLoS Biology* 7 (2): e1000033.

Gebharder, Alexander. 2014. "A Formal Framework for Representing Mechanisms?" *Philosophy of Science* 81 (1): 138–53.

Gebharder, Alexander, and Marie Kaiser. 2014. "Causal Graphs and Biological Mechanisms." In *Explanation in the Special Sciences*, edited by Marie Kaiser, Oliver Scholz, Daniel Plenge, and Andreas Huttemann, 55–85. Springer.

Glennan, Stuart. 2002. "Rethinking Mechanistic Explanation." *Philosophy of Science* 69 (S3): S342–53.

Green, Sara, Melinda Fagan, and Johannes Jaeger. 2015. "Explanatory Integration Challenges in Evolutionary Systems Biology." *Biological Theory* 10 (1): 18–35.

Kaplan, David, and William Bechtel. 2011. "Dynamical Models: An Alternative or Complement to Mechanistic Explanations." *Topics in Cognitive Science* 3 (2011) 438–44.

Kaplan, David, and Carl Craver. 2011. "The Explanatory Force of Dynamical and Mathematical Models in Neuroscience: A Mechanistic Perspective." *Philosophy of Science* 78 (4): 601–27.

- Levy, Arnon, and William Bechtel. 2013. "Abstraction and the Organization of Mechanisms." *Philosophy of Science* 80 (2): 241–61.
- Levy, Arnon, and William Bechtel. 2016. "Towards Mechanism 2.0: Expanding the Scope of Mechanistic Explanation." *PhilSci-Archive*.
- Machamer, Peter, Lindley Darden, and Carl Craver. 2000. "Thinking about Mechanisms." *Philosophy of Science* 67 (1): 1–25.
- Matthiessen, Dana. 2017. "Mechanistic Explanation in Systems Biology: Cellular Networks." *The British Journal for the Philosophy of Science* 68 (1): 1–25.
- Pearl, Judea. 2009. *Causality: Models, Reasoning, and Inference*. 2nd ed. Cambridge: Cambridge University Press.
- Rubenstein, Paul, Stephan Bongers, Bernhard Schölkopf, and Joris Mooij. 2016. "From Deterministic ODEs to Dynamic Structural Causal Models." *ArXiv Preprint ArXiv:1608.08028*.
- Spirtes, Peter, Clark Glymour, and Richard Scheines. 2000. *Causation, Prediction, and Search*. Cambridge: MIT press.
- Stephan, Klaas, Lee Harrison, Stefan Kiebel, Olivier David, Will Penny, and Karl Friston. 2007. "Dynamic Causal Models of Neural System Dynamics: Current State and Future Extensions." *Journal of Biosciences* 32 (1): 129–44.
- Svoronos, Spyros, George Stephanopoulos, and Rutherford Aris. 1980. "Bilinear Approximation of General Non-Linear Dynamic Systems with Linear Inputs." *International Journal of Control* 31 (1): 109–26.