ARTICLE

Six Theses on Mechanisms and Mechanistic Science

Stuart Glennan[1] · Phyllis Illari[2] · Erik Weber[3]

Abstract

In this paper we identify six theses that constitute core results of philosophical investigation into the nature of mechanisms, and of the role that the search for and identification of mechanisms play in the sciences. These theses represent the fruits of the body of research that is now often called New Mechanism. We concisely present the main arguments for these theses. In the literature, these arguments are scattered and often implicit. Our analysis can guide future research in many ways: it provides critics of New Mechanism with clear targets, it can reduce misunderstandings, it can clarify differences of opinion among New Mechanists and it helps to define a research agenda for New Mechanists.

Keywords Explanatory pluralism · Mechanism discovery · Mechanistic explanation · Minimal mechanism · Models of mechanisms

1 Introduction

In this paper we identify six theses that constitute core results of philosophical investigation into the nature of mechanisms and of the role that the search for and identification of mechanisms play in the sciences. The term "New Mechanism" (Craver and Tabery 2019; Craver and Krickel 2019; Glennan and Illari 2018b) has proved a convenient label to refer to this research, which has emerged over the last 30 + years, but, like any philosophical ism, it can be hard to pin down. New Mechanism is often associated with the authors of a few works from the 1990s who drew attention to the

✉ Phyllis Illari
phyllis.illari@ucl.ac.uk

Stuart Glennan
sglennan@butler.edu

Erik Weber
Erik.Weber@UGent.be

[1] Department of Philosophy, Religion and Classics, Butler University, 4600 Sunset Avenue, Indianapolis, IN 46208, USA

[2] Department of Science and Technology Studies, University College London, Gower Street, London WC1E 6BT, UK

[3] Centre for Logic and Philosophy of Science, Ghent University (UGent), Blandijnberg 2, 9000 Ghent, Belgium

centrality of the search for mechanisms in the scientific enterprise (Bechtel and Richardson 1993; Glennan 1996; Machamer et al. 2000), but many other contemporary philosophers have thought and written about mechanisms and mechanistic science – sometimes, but not always, supportive of, informed by, or in response to, arguments made in these early works.

Given the large literature, we think there is value in distilling what can be learned from this research into a small number of succinct theses. The insights expressed in the theses have taken some time to emerge, and terminological disputes, as well as some shifts from earlier positions among New Mechanists, have made these points easy to miss. By succinctly formulating these overarching theses, we are able to identify key insights of the New Mechanist program. Additionally, we concisely present arguments for the theses that are scattered and sometimes only implicit. By making the arguments explicit and presenting them in a systematic way, we consolidate the theses.

Let us make clear what we mean by calling these six theses "core results." To say that they are *results* is to say that they are theses we have come to know as a consequence of thirty years of research into mechanisms. They are achievements that required considerable effort. To say that they are *core* results is to say that they summarize what we take to be the most important insights to be gleaned from this dialog. While we have endeavored to identify a set of theses that span breadth of recent philosophical research on mechanisms and mechanistic science, we take it as entirely possible that others might argue for additional core results we have missed. However, our theses cover crucial topics such as the definition of mechanisms (thesis 1), the role of models (thesis 2), mechanism discovery (thesis 5), and the use of mechanistic knowledge in scientific explanations (theses 3 and 4) and in other contexts (thesis 6).

All of these theses have been defended in one form or another by New Mechanist philosophers, and we believe that there is now a significant consensus among mechanistically inclined philosophers on each of these points. But it is not our goal to review the literature with an eye towards describing points of agreement or disagreement. It is instead to provide a succinct formulation of insights that we draw from this literature, and to provide arguments that make clear why each of the theses should be accepted.

While our project looks backward for arguments to support the theses, the motivation behind it is forward-looking, in the sense of providing resources that will guide future research. First, our exposition of core theses and arguments will provide critics of mechanistic approaches with a clear target. Second, our clarifications of these theses may help alleviate certain misunderstandings of the claims and motivations of some prominent advocates of New Mechanism that, in our view, have distracted researchers from less explored and more promising questions about mechanisms and mechanistic science. In particular, our work will show that New Mechanists are more amenable to pluralistic approaches to explanation and discovery than is sometimes acknowledged. We hope this insight can encourage more fruitful exchanges between New Mechanists and their critics in the future. Third, our project provides New Mechanists with a set of largely shared and safe points of departure. It establishes common ground but also helps to clarify differences in opinion. For instance, scholars can identify themselves as "standard" New Mechanists (if they accept all six theses) or pinpoint their views exactly by rejecting one of the theses. Finally, as we shall elaborate on in the conclusion, our project has heuristic potential: the theses generate a research agenda for mechanistic philosophers.

2 Thesis 1: The Definition of Mechanisms

Our first thesis is:

The most fruitful way to define mechanisms is that a mechanism for a phenomenon consists of entities (or parts) whose activities and interactions are organized so as to be responsible for the phenomenon.

We call this the *Minimal Mechanism Thesis* because the definition that occurs in it is the definition of minimal mechanism as put forward in Glennan (2017, 19) and Glennan and Illari (2018b, 2).

A number of the early papers in the New Mechanism literature were concerned with offering a suitable definition or characterization of mechanism. Three early and frequently cited proposals are (Glennan 1996, modified in Glennan 2002; Machamer et al. 2000; Bechtel and Abrahamsen 2005), and a number of subsequent papers have discussed these proposals, arguing for one over the other, or identifying limitations and possible counterexamples to some or all of them. Illari and Williamson (2012) summarize these debates, and argue on their basis for a consensus definition of mechanism. Minimal mechanism (a designation coined in Glennan 2017) is a close cousin of the proposal of Illari & Williamson.

Our overarching rationale for adopting minimal mechanism as an "industry standard" is that doing so will be fruitful. In arguing for this standard we do not assert that we have discovered the true nature of mechanisms or suggest that the definition completely clarifies what do and do not count as mechanisms. Neither do we suggest that we have discovered the most felicitous possible expression of the concept. What we do believe is that the two decades of discussion of how to define mechanisms has led to some points of consensus, embodied in minimal mechanism, and that using this consensus as a starting point will advance rather than stifle future research.

The first virtue of minimal mechanism is that it has been tested and has survived these tests. The inevitable consequence of critiques of the earlier proffered accounts of mechanisms is to have identified problematic cases—chiefly cases where structures intuitively described as mechanisms fail to meet a stipulation in the definition. For instance, Machamer et al. (2000) stipulated that mechanisms be regular, but subsequent discussions have identified a variety of things intuitively identified as mechanisms which fail to meet this requirement. Similarly, Glennan's (1996) attempt to characterize interactions between parts of mechanisms in terms of laws has been judged, in light of current philosophical discussions of laws, misleading. And Bechtel and Abrahamsen's definition of mechanisms appeals to the concept of "function," which might problematically be taken to exclude mechanisms whose phenomena are not the product of design or selection. At the same time, minimal mechanism seems to identify a series of points of consensus about what goes into a mechanism. Mechanisms are identified and individuated by what they do (the "phenomenon"), what they are made of ("components"/"entities"/"parts"), what their parts do ("activities"/"interactions"/"operations"), and how they are structured ("organization").

A second virtue of minimal mechanism is that it allows mechanisms to have both a causal and a part-whole dimension. It is a point of consensus among New Mechanists that mechanisms generally can have two dimensions (Craver and Tabery 2019; Glennan 2017). In their horizontal (causal) dimension, mechanisms move from start points to end points, or maintain states over periods of time; for instance, a protein synthesis mechanism might be activated by the binding of RNA polymerase to a promoter region. In their vertical or part-whole dimension, the activity of some composite entity is made up of the activities of its component parts. For instance, when a muscle tissue contracts,

the muscles' contracting is made up of or constituted by the contractions of the individual muscle cells of which the tissue is composed. The definition of minimal mechanism characterizes the relation between a mechanism and its phenomenon as one of responsibility, allowing us to consider both causal/processual cases, where the output of a mechanism depends causally on its inputs, and constitutive cases where the activity of a mechanism as a whole depends upon the activities of its parts. The proper understanding of the relationship between these two dimensions is far from settled—but minimal mechanism provides a framework for organizing these debates.

A third virtue of minimal mechanism is its generality. When Illari & Williamson first proposed a synthesized version of earlier definitions, a major motivation was to provide an account of sufficient generality to account for mechanisms that are objects of investigation across the sciences—encompassing not just the early areas of focus among New Mechanists, like molecular biology and neuroscience, but also mechanisms in physics and astrophysics, and social mechanisms. Some of the stipulations culled from earlier accounts of mechanisms, like Machamer, Darden & Craver's regularity requirement, were eliminated precisely because they are features of some, but not all, mechanisms—in the minimal sense. Some critics (e.g., Dupré 2013) worry that the generality of the mechanistic account makes it vacuous, but this is not the case. Mechanisms are a distinctive metaphysical category—a composite of entities and activities organized in a particular way. Mechanisms are thus not objects or substances, nor are they properties, nor are they laws. For many, they are particulars rather than universals, located at different places in space and time (Glennan 2017; Krickel 2018). To identify these structures as central features of the natural and social world and as objects of investigation across the sciences is to say something non-trivial about both nature and science.

A fourth virtue of minimal mechanism is that it clarifies what is required for the construction of mechanistic models and explanations—namely the identification (with varying degrees of completeness and specificity) of a target phenomenon, and of the ways in which entities, activities and interactions are organized so as to account for that phenomenon.

A final virtue of minimal mechanism is that the definition serves to organize philosophical research. It embodies broad agreement that mechanisms are for phenomenon, and that they are appropriately identified by their entities, activities, interactions and organization; but it is consistent with a number of different metaphysical interpretations of the nature of these constituents. For instance, some mechanists (e.g., Glennan 2017; Krickel 2018) have adopted a view of activities and interactions that fits within a broadly neo-Aristotelian framework; others may be content with an account that is neutral on the nature of causation or which reduces activities and interactions to patterns of counterfactual dependence (Craver 2007; Menzies 2012; Woodward 2000). To date, relatively little has been written as to what should and should not count as an entity—and we think this will be an increasingly pressing matter as New Mechanistic analyses are extended to novel domains. Organization is important, but under-discussed. Both entities and activities are typically composites—made of other entities and activities—but it is far from clear under which conditions a set of components make up a composite, and it is also not clear the degree to which conditions we might adopt reflect ontological as opposed to pragmatic considerations. Similarly, there is much question about the ways in which a phenomenon depends upon its mechanism. It is evident that mechanisms involve both compositional and causal relations, but how best to sort out these dependencies is still far from settled. The definition and concepts employed in minimal mechanism can help organize these and other debates.

3 Thesis 2: The Necessity and Multiplicity of Models

Our second thesis is:

Scientists can only discover, describe, and explain mechanisms through the construction of models, and these models are inevitably partial, abstract, idealized and plural.

We call this the *Representation Thesis*.

Mechanisms, as described in the Minimal Mechanism Thesis, are bits of worldly stuff, but to get a hold on these—to find them, describe them, explain how they work, and to make them do things—requires creating representations of mechanisms. These representations we call models.

We use the term "model" inclusively to refer to all sorts of representations—mathematical models, simulations, diagrams, material models, and so forth. In using this term, our chief contrast is with theory, at least where theories are understood to be collections of law statements taken to hold truly and without exception over their domain. Discussion of modeling, and, with it, approximation and idealization, is a longstanding part of the mechanisms literature (Bechtel and Richardson 1993; Craver 2007; Glennan 2005; Leuridan 2010; Machamer et al. 2000; Wimsatt 1972).

Our chief object in offering this thesis is to head off two misunderstandings of the mechanistic program. The first has to do with the debate between ontic and epistemic conceptions of explanation, and the second has to do with what has come to be called the "more details are better" thesis.

The distinction between the ontic and epistemic conception of explanation dates back to Salmon (1984, 15–20), but resurfaces in the mechanisms literature as a debate chiefly between Craver, on the one hand (Craver 2014), and Bechtel and Wright on the other (Wright 2012; Wright and Bechtel 2007). While Salmon originally equated the epistemic conception narrowly with Hempel's covering law model and the view that explanations are arguments, in the recent literature it has sometimes been portrayed as a debate between those who think that explanations are epistemic entities (models, representations) and those who think that explanations are "in the world." We think, however, that New Mechanist philosophers on both sides of this argument should grant that models are required to provide mechanistic explanations.

Let us first point out, as others have observed (e.g., Craver 2014), that as a matter of linguistic practice, we use the word "explain" with reference both to models and theories, and to facts about or features of the world. We can equally well say that interactions with the lac operon explain how lactose metabolism is regulated, or that the lac operon model explains how lactose metabolism is regulated. We do not see how any charitable reading of either the ontic or epistemic conceptions denies the reality or the legitimacy of these locutions. Moreover, the use of ontic locutions in which some facts are said to explain other facts is perfectly consistent with the view that all acts of explanation require models to represent those facts. We thus take belief in the ubiquity and necessity of models to be common ground among mechanistic philosophers, regardless of whether they support the epistemic or ontic conception, or some alternate or intermediate position. We take the still-live debate between advocates of ontic and epistemic conceptions to be not about whether models are involved in explanation at all, but rather on how they explain. What kinds of features must models have, and what relations to the world must they bear, in order to explain at all?

The second part of thesis 2 is that models are abstract, idealized and plural. In asserting this, we hope to head off a common misunderstanding of the claims of some mechanists

(especially Craver, Darden and Kaplan), namely the idea that mechanists believe that more detailed models are always better models.

The idea that some mechanists believe that more details are always better stems, we think, from a reading of the account of mechanism discovery and explanatory norms found in Machamer et al. (2000), Darden and Craver (2002), Craver (2007) and Craver and Darden (2013). The kernel of this account is a picture of discovery whereby scientists begin with a "mechanism sketch," which they gradually fill in to provide a complete "mechanism schema" (see especially Craver 2007, 112–113). The metaphor of filling-in is highly suggestive of a conception of mechanism description where there is, ideally at least, one ultimate best description of a mechanism, and this metaphor is no doubt the source of the belief that Craver, Darden, and like-minded mechanists believe that more details are always better (Batterman and Rice 2014; Chirimuuta 2014; Levy and Bechtel 2013).

Whatever may have been said or believed by Craver, Darden and Machamer in some of their early works, we think there is nothing inconsistent between the account of mechanism discovery there given and the view that models are abstract, idealized and plural. Recent New Mechanist literature has emphasized models-pluralism (Glennan 2017; Hochstein 2015), and Craver and Kaplan, the two mechanists most commonly tarred with accusations of abhorring abstraction and idealization, have made clear in a recent (2020) paper that they do not. In that paper, they defend a norm of explanatory completeness that suggests that, all things being equal, the more *relevant* details a scientists knows, the better, but they are at pains to emphasize that many details are irrelevant and make explanations worse. They also argue that all things are seldom equal; that under many conditions abstraction and idealization are necessary; that models can be plural, and that partial, plural, and idealized models can explain.

4 Thesis 3: The Ubiquity of Mechanistic Explanations

Our third thesis is:

Mechanistic explanations are ubiquitous across the empirical sciences.

We call this the *Ubiquity Thesis*.

Mechanistic philosophers in the first decade of this century sought to get a grip on explanations in the life sciences, and there was a widely held conviction that the covering law model was not suited for this purpose because there are hardly any laws in biology. While we certainly agree with this assessment of explanation in the life sciences, the point of the Ubiquity Thesis is to remind us that mechanistic explanations are abundant outside the life sciences, including in many areas of both the physical and social sciences.

As a result of the efforts of mechanistically minded philosophers of science working in biology, psychology and neuroscience (e.g., Cummins 2000; Thagard 2006; Craver 2007; Bechtel 2008) there is now little debate about the centrality of mechanistic explanation in the life sciences. Even those who argue that there are non-mechanistic explanations in these domains (e.g., Chemero and Silberstein 2008; Shapiro 2019) do not disagree that mechanistic explanation is a crucial part of explanatory practice in the life sciences.

The triumph of mechanistic explanation in the life sciences is sometimes taken to imply that mechanistic explanation is peculiar to the life sciences, but this, we think, is a mistake. Reflecting on the lack of attention to mechanistic explanation, Bechtel & Abrahamsen offered this diagnosis:

Given the ubiquity of references to mechanism in biology, and sparseness of reference to laws, it is a curious fact that mechanistic explanation was mostly neglected in the literature of 20th century philosophy of science. This was due both to the emphasis placed on physics and to the way in which explanation in physics was construed. (2005, 423)

We agree with their diagnosis, but we think their most significant point has to do with the way philosophers of science construed explanation in physics. We must add that, given the ubiquity of references to mechanisms in the physical sciences, it is a curious fact that, before the rise of mechanistic philosophy, explanation in physics was construed as being exclusively of the covering law type.

To substantiate our point we shall give an example of a mechanistic explanation in physics. Boyle's law states that for gases held at a fixed temperature and mass, the product of pressure (P) and volume (V) is constant (P x V = k, where k depends on the temperature and amount of gas). This law (and many other gas laws such as the more general ideal gas law) can be explained by means of the kinetic model of gases (which fits into the general kinetic theory of matter; the overall theory also deals with properties of fluids and solids). As to entities, the most important tenets of the kinetic gas model are:

The system consists of gas molecules and a solid container.
The number of molecules is large.

With respect to activities the crucial tenets are:

The molecules are in perpetual linear motion at constant speed, except during collisions
The walls of the container are rigid (so they do not move).
The molecules collide with each other and with the walls of the container.
These collisions are perfectly elastic (no loss of kinetic energy).

The crucial organizational tenets are:

The molecules are located within the container.
The number of molecules moving in any one direction is, on average, the same as that moving in any other direction (i.e. the motion is on average random).

The explanation of Boyle's law by means of the kinetic gas model is mechanistic; it describes entities, activities and organization.

We count this explanation as mechanistic in part because we treat random motion as a kind of organization. In keeping with a minimal definition of mechanism (Thesis 1) our understanding of organization is permissive:

The purport of the term "organized" is simply to indicate that a mechanism's phenomenon depends not just on what the parts of the mechanism are, or on what activities those parts engage in, but on how the parts and their activities are arranged. (Glennan 2017, 23)

Random motion counts as a kind of organization, because it is an arrangement on which the phenomenon depends. A corollary of this permissive notion of organization is that some mechanisms are "orderly" while others are not. Arnon Levy (2014) characterizes orderly mechanisms as those in which (a) different components play a different role in bringing about the phenomenon (he calls this "internal division of labor") and (b) these

components play their role in virtue of local relations to other components. In this terminology, gases in a container are not orderly mechanisms, but they are mechanisms.

There are many other textbook examples of mechanistic explanations in the physical sciences: for instance the explanation of the behavior of rotating rigid bodies (see e.g. Feynman et al. 2011, vol. I, ch. 18) and of the refractive index (ibid, chapter 31). Philosophers also have recently begun to pay increased attention to mechanistic explanations across the physical sciences (van Eck 2015; Falkenburg and Schiemann 2019; Illari 2019; Kuhlmann 2015; Ramsey 2008; Teller 2010).

The Ubiquity Thesis implies that mechanistic explanations also occur in the social sciences. We shall offer one example from political science that is based on Bueno de Mesquita et al. 1999 and Maoz and Russett 1993. The scientific study of international relations theory has produced a few widely accepted generalizations. One of them is:

Democracies do not fight wars with one another.

Some other regularities related to war-proneness and democracy are also widely accepted, for instance:

Overall, democracies are not less war-prone than nondemocracies.
Democracies are more likely to initiate wars against autocracies than are autocracies against democracies.

Maoz & Russett aim at a unified explanation for the two first regularities:

Specifically, a better understanding is required of the causal mechanism explaining simultaneously both the democratic-peace phenomenon and the lack of difference between democracies and nondemocracies in terms of their overall conflict proneness. (1993, 624)

The way they formulate their explanatory task already reveals that they are after mechanistic explanations. The explanation they develop for the democratic peace regularity can be summarized as follows:

[D]ue to the complexity of the democratic process and the requirement of securing a broad base of support for risky policies, democratic leaders are reluctant to wage wars, except in cases wherein war seems a necessity or when the war aims are seen as justifying the mobilization costs. The time required for a democratic state to prepare for war is far longer than for nondemocracies. Thus, in a conflict between democracies, by the time the two states are militarily ready for war, diplomats have the opportunity to find a nonmilitary solution to the conflict. (1993, 626)

This explanation refers to features of social agents, actions and structures that exemplify what social scientists and philosophers call "social mechanisms," which are clearly a species of minimal mechanism. In social mechanisms, the entities are human individuals (in the example: political leaders and diplomats, but also citizens that have to be mobilized). The activities are mainly intentional actions. The phenomenon that is produced is a causal or other regular relationship between aggregate-level social variables (in the example: a causal relation between democratic/autocratic political structure and war-proneness).

The attention to social mechanisms and mechanism-based explanations in the social sciences originates in the work of Elster (1989) and Little (1991). As Hedström and Ylikoski (2010) note, analytical sociologists and other social scientists developed the idea largely independently of the mechanistic tradition in the philosophy of science, but the two traditions embrace similar conceptions of core features of mechanisms. While a pair of

examples is not enough to establish ubiquity, they suffice to show that mechanistic explanations are not essentially tied to any scientific discipline or group of disciplines.

5 Thesis 4: Mechanisms and Explanatory Pluralism

Our fourth thesis is:

Emphasizing that mechanistic explanations are ubiquitous in all scientific disciplines does not entail that all scientific explanations are mechanistic.

We call this the *Explanatory Pluralism Thesis*.

Mechanistic philosophers spend a lot of time studying mechanistic explanations, and there are lots of those (thesis 3). Some have also argued that certain dynamical models which do not map to any mechanistic details are ipso facto non-explanatory (Kaplan and Craver 2011). Together these facts might create the impression that mechanist philosophers are convinced that all scientific explanations are mechanistic, and some critics (e.g. Psillos and Ioannidis 2017) have claimed that New Mechanists are committed to such a view.

Nonetheless, we know of no mechanist who would endorse such a radical position, and there appears to us to be increasing agreement that a pluralist approach is required. A case for pluralism can be found in a series of papers that accept the legitimacy of mechanistic explanations, but which identify distinctive forms of explanations that complement them. These include design explanations (Wouters 2007), covering law explanations (Gervais & Weber 2013) and topological explanations (Huneman 2018). Glennan (2017, ch. 8) has developed a general account of explanation that embraces complementary roles for mechanistic explanation, causal but non-mechanistic explanation, and even non-causal explanation.

To develop our argument for thesis 4, we consider explanations of pigeons' capacity to navigate homeward (an example discussed in Gervais and Weber 2013, taken from Keeton and Gould 1986, 575–585). A trained pigeon can be taken from home, transported over very long distances (hundreds of miles are not uncommon) and still find the way back to its home after being released. Moreover, they are able to exercise this capacity both in sunny weather and on cloudy days. This behavior gives rise to two explananda:

Pigeons usually have the capacity to find their way back home on sunny days.
Pigeons usually have the capacity to find their way back home on cloudy days.

These explananda are generalizations about pigeons as a species. The term "usually" expresses the fact that the generalizations have exceptions: they describe a strong trend in the set of pigeons, but they are not universal laws. For instance, if a breeder feeds alcohol to their birds before releasing them, they probably will not return home (at least not immediately). As we will see below, the explanation for sunny days differs from the explanation for cloudy days. That is why we distinguish the two conditions and formulate two explananda.

With respect to the first explanandum, it has been shown that pigeon navigation depends on the position of the sun as a reference point. Thus, the following covering law explanation can be formulated:

Pigeons usually have a solar compass.

Animals with a solar compass usually have the capacity to find their way back home on sunny days.

Therefore, pigeons usually have the capacity to find their way back home on sunny days.

Like the explanandum, the two generalizations in the explanans describe strong but not exceptionless regularities in the set of pigeons. For the first regularity, exceptions are pigeons that did not develop normally or have certain kinds of brain damage. For the second regularity, the exceptions are temporary disturbances, such as the administration of alcohol. The explanation for the capacity of pigeons to find their way back on cloudy days is similar, but involves a magnetic compass. Experimental research shows that the magnetic compass is a backup, not used when days are sunny.

In our view, these explanations of pigeons' navigational capacities support thesis 4 because they are legitimate but non-mechanistic explanations. We first argue that they are legitimate, and then that they are non-mechanistic.

Our argument for the legitimacy of these explanations starts from Woodward's account of explanatory power. Its main tenet is that derivations in explanations should mirror physical dependency relations:

> The idea is that these derivations trace or mirror the relations of physical dependency that hold between the explanans conditions and the explananda phenomena – relations that would be revealed if, for example, we were to physically intervene to alter the explanans conditions. (2003, 201)

Woodward uses this idea to distinguish explanatory derivations from non-explanatory ones. In the famous flagpole example, the explanation of the length of the shadow answers questions such as "what would have happened if the flagpole was higher or lower?" or "what would have happened if the sun was at a different angle above the horizon?". Derivations that do not mirror a physical dependency (and hence are not explanatory) do not allow us to answer "what-if-things-had-been-different" questions. In the flagpole example, the non-explanatory reversed derivation does not provide a correct answer to the question "what would have happened to the height of the flagpole if the shadow had been made shorter?" (e.g. by means of an additional light beam). The reversed derivation suggests that the flagpole would have become shorter (given that the position of the sun is unchanged), which is not the case. We cannot produce a shorter flagpole by intervening on the length of the shadow.

Let us apply these ideas to the pigeon example. The presence of the explanandum capacities in the pigeons physically depends on the proper functioning of the solar compass, the proper functioning of the magnetic compass, and the weather condition (sunny vs. cloudy). The physical dependency can be revealed by certain interventions. To demonstrate the dependence of sunny-day navigation on a solar compass, biologists alter light cues: If you put pigeons in a windowless room in which the light goes on at midnight and goes off at noon, and then release them on a sunny day, these pigeons fly in the wrong direction (90° to the left of the optimal direction). On cloudy days the "clock-shifting" has no impact. To test the dependence of cloudy-day navigation on a magnetic compass, biologists put magnets around the neck of pigeons; these pigeons become disoriented on cloudy days, though they still succeed in getting home on sunny days.

Note that our position entails the rejection of what Kaplan and Craver (2011) have called "predictivism" – the view that explanations are structurally equivalent to predictions. The explanatory pluralism that we propose incorporates physical dependence as a necessary condition for explanatory power. Hence, there is a structural difference between explanation and prediction: derivations in predictions do not have to mirror physical dependency

relations in the world. It is the fact that these explanations track physical dependency relations that allows us to answer what-if-things-had-been-different questions.

There are two reasons for characterizing these explanations as non-mechanistic. First, they do not contain a model of the mechanism (the solar compass or the magnetic compass). The generalizations in the explanans state that pigeons have such a mechanism, but there is no information about how this mechanism works; no information is given about the parts of the solar compass and their activities. This is in conflict with the general idea that mechanistic explanations explain how some phenomenon comes about via a model of the mechanism that is responsible for it (Glennan 2017, 228). Models of mechanisms characterize the entities, activities, and organization of mechanisms. There is no such model of the solar compass in the explanation, not even a sketchy one.

The second reason not to take these explanations as mechanistic is because the explanations derive their explananda from covering laws. It is true that mechanistic explanations sometimes appeal to generalizations (e.g. "Democratic leaders are reluctant to wage wars") but these describe the activities of the entities that are components of the mechanism. Mechanistic explanations do not proceed by subsumption under a covering law. As we have seen in the discussion of thesis 3, the mechanistic model of explanation was developed as a tool to analyze explanatory practices in contexts where no overarching covering laws were available. In the pigeon example, the explanation works by subsumption under a more general regularity about all animals.

6 Thesis 5: Diverse Tools for Mechanism Discovery

Our fifth thesis is:

The diversity of kinds of mechanisms requires and explains the diversity of tools, strategies and heuristics for mechanism discovery.

We call this the *Discovery Pluralism Thesis*.

Mechanism discovery has received relatively little attention in the wider New Mechanism literature, despite the fact that many prominent New Mechanists examine how empirical work gets us knowledge of mechanisms. A superficial reading might suggest that New Mechanists like Bechtel and Richardson, Craver, and Darden, offer *conflicting* accounts of "the way" to find mechanisms. But this is not true. Instead, the mechanisms literature as a whole offers a toolbox of mechanism discovery methods, an array of what Bechtel and Richardson call "heuristics" and Darden calls "strategies." In articulating this thesis we wish to point to these tools as an underutilized resource, and to encourage further investigations into new methods for mechanism discovery.

The first publication in a certain strand of New Mechanist literature, Bechtel and Richardson (1993), is firmly focused on discovery. They frame all their work in terms of "fallible heuristics." Fallible heuristics, following Polya (1945) and Simon (1957), are used when comprehensive or otherwise optimal search is impossible—which all mechanists agree happens a lot, and many take to be the typical case. Under these circumstances, heuristics can still offer fallible guidance to take you from what you know to fruitful places to search next. Once this much-missed or misunderstood idea is given its due place, it becomes clear that different New Mechanist accounts of discovery are not in competition, but are instead a catalog of ways to explore mechanisms.

Bechtel and Richardson (1993) lay out the idea of *decomposition and localization* as fallible heuristics of mechanism discovery. These seem to be targeted at the constitutive

aspect of mechanisms, like the movement of muscles contracting explaining arm movement, or activities of the brain constituting a person's memory. Broadly, the problem is that hypotheses about mechanisms responsible for a phenomenon are underdetermined by evidence. In finding a mechanism, we first try to decompose (i) the phenomenon into a number of tasks or activities that make it up, perhaps distinguishing working memory from occurrent memory, for example, and (ii) the material into parts or entities that might contribute to producing the phenomenon, such as parts of the brain (see Bechtel 2008). Then we localize, trying to identify parts that accomplish particular tasks, such as matching activities of working memory to the parts of the brain that perform it.

This entire process is fallible, in that we frequently get the decomposition into both tasks and parts wrong. Even when we start to localize a task to the right part, we begin with simplifying assumptions, like the assumption that a single part accomplishes the whole of the task. But our initial hypotheses are almost always too simple, and we have to repeatedly modify them until eventually we have a reasonably adequate mechanistic explanation. So decomposition and localization may be a classic mechanism discovery story, but the fallibility of experimental tools and assumptions at every stage is *crucial* to interpreting it correctly. Finally, Bechtel and Richardson also note that we begin this kind of discovery with a description of a phenomenon, and this is also fallible, and is often revised (which they call "reconstituting the phenomenon").

Craver's mutual manipulability (MM) account of constitutive relevance (2007) is probably now the most discussed view of the problem of identifying mechanism constituents. For Craver, an entity is a constitutively relevant part (or "working part") of a mechanism if (i) it is a part of a larger entity and (ii) the activities of the part and the whole are mutually manipulable. This second condition means, roughly, that some interventions that change the activities of the mechanism as a whole (e.g., heating a cell) can lead to detectable changes in the parts (e.g. faster reactions), *and* that some interventions on the parts (e.g. providing more energy) can lead to detectable changes in the whole (e.g. faster repair of the cell wall).

While the MM account is sometimes treated as a metaphysical story about the nature of the constitutive relevance relation (e.g., Baumgartner and Gebharter 2016), Craver designed it in response to experimental practices in neuroscience. This means we can *also* consider it in a way that may be more useful, as a discovery story, reflecting experimental techniques for finding parts in many mechanisms. In a recent paper Craver et al. (2021) have emphasized that MM, and an improved version of the criterion they call Matched Interlevel Experiments (MIE), are merely sufficient criteria for establishing that some entity or activity is part of a mechanism; they have also pointed to possible strategies for identifying parts that are not amenable to manipulation. Similarly, Aizawa and Headley (forthcoming) argue that practical limitations on experimental precision sometimes prevent certain bottom-up interventions, and that in these cases abductive inference strategies are required.

Darden (2006) and Craver and Darden (2013) have also offered an account of discovery heuristics for mechanisms. For these mechanisms, the aim is to understand the mechanism by which some end state is reached or product is produced from some initiating condition. For instance, in a protein synthesis mechanism, the initiating condition might be the binding of RNA polymerase to a promoter region, and the terminating condition might be the production of a protein. Darden and Craver discuss what they call "forward/backward chaining." Chaining is when you start from something you know about a mechanism, which can be anywhere in the mechanism, and try to move backwards or forwards along the mechanism into the unknown, but towards some other part or activity that is already

known. So what you know can guide your search. This strategy could use almost any kind of technology or study, but what you know about nearby entities, activities, or organization helps you choose a tool or study. Other strategies include "schema instantiation" and "modular subassembly." Darden has also explored how scientists draw mechanism diagrams in the discovery process, using them to indicate "gaps," sometimes with a question mark, to represent what they know, and draw attention to what is unknown. That these techniques are used successfully is something all New Mechanists can recognize. The visual, particularly mechanism diagrams, and the material basis of scientific reasoning, are now central in Bechtel's concerns, too (Abrahamsen et al. 2018).

So while different mechanists have paid attention to different discovery methods, these methods are best seen not as competitors but as different tools in the toolbox of scientific practice. Very different investigative tools, and scientific study designs may be involved in different strategies for probing the phenomena, entities, activities and organization of mechanisms. There is room for lots of further work uncovering new heuristics, as is done in Bechtel (2019), and examining how they are put together to build up evidence for a mechanism, and also how this evidence may be related to other kinds of evidence.

7 Thesis 6: New Mechanism as a Resource for Science, Technology and Evidence-Based Policy

Our sixth thesis is:

The mechanisms literature is a rich source of insights that can be used to address challenging reasoning problems in science, technology and evidence-based policy.

We call this the *Fruitfulness Thesis*.

The philosophy of science community has addressed a variety of questions about mechanisms, including the ontology of mechanisms, the constitutive relevance question, the modeling and representation of mechanisms, and the nature and extent of mechanistic explanations. This has resulted in a substantial body of literature that can collectively provide a rich resource for tackling reasoning problems outside philosophy, viz. in science and in evidence-based policy.

Our judgement about fruitfulness is largely prospective; it is a judgement of how this literature could be applied, more than a claim that it has been applied. But to justify our optimism, we will illustrate one domain in which this literature has already had an impact—the problem of evaluating evidence in medicine.

Worldwide, evidence in medicine is assessed by hundreds of public agencies. These assessments are typically done by expert panels convened for each disease or condition, but public agencies report on how the clinical literature is searched and how panels are asked to make decisions. Their assessments have implications for science, technology and policy.

Evidence panels have many reasoning tasks, but typically share the key aim of establishing efficacy of treatments, i.e. establishing whether treatments such as drugs, or health technologies like pacemakers cause better outcomes for patients. Such decisions are generally based upon a hierarchy of evidence assessment criteria that strictly prioritizes evidence gained in randomized clinical trials (RCTs) over observational studies, and observational studies over "pre-clinical," "basic" or "bench" research. But a burgeoning literature in philosophy of medicine is now examining the benefits of being more pluralistic about evidence, and many have specifically argued that there are ways

in which evidence of mechanisms can usefully complement evidence gained in RCTs or observational studies (Williamson 2019; Aronson et al. 2018; Clarke et al. 2014; Russo and Williamson 2012).

Illari (2011) examines how scientific research in medicine might be better used in evidence assessment. She argues that evidence of mechanism gained in, for example, laboratory and animal studies, and evidence of correlation such as is gained in RCTs or observational studies, are complementary in establishing efficacy, as each helps address the major weakness of the other. On the one hand, if we have the commonly recognized concern that the correlation found in a trial is due to bias or confounding by a common cause of treatment and recovery, then evidence that the mechanism of action of the treatment really works in vitro or in animal studies helps to reduce that concern. On the other hand, if we have identified some evidence of a mechanism explaining the link between treatment and recovery, we may be concerned about the existence of other undiscovered mechanisms, that may swamp, or "mask," the action of the one we have identified. The evidence needed to ameliorate this concern is evidence of a correlation between treatment and recovery in a real human population, coming from an RCT or other trial (see also Parkkinen et al. 2018).

Turning more explicitly to evidence-based policy, public health agencies and their evidence assessment panels face various other difficult inferences, and philosophers have also examined the relevance of evidence of mechanism to several of these. Alongside efficacy, they need to assess *risk* of treatments, health technologies, and environmental factors (Osimani 2014). There is also a constant problem of deciding whether a study showing treatment efficacy will apply to the population of interest, which may be different from the study population. This is known as the problem of external validity, or of extrapolation, and is very important for public health policy. Many have thought that evidence of mechanism is relevant, and philosophers debate how best to use it (Steel 2008; Howick et al. 2013a, b). Moreover, almost all studies are on populations, so that clinical practitioners treating individuals must decide whether population-level studies will apply to their patient. Known mechanisms may well be helpful (La Caze 2011). For example, ageing tends to increase your risks for all diseases via mechanisms that are a constant target of investigation (Howick et al. 2013a, b). Finally La Caze argues more generally that "the mechanistic understanding provided by basic science is central to the design, analysis and application of clinical research" (La Caze 2011, 83).

While this is only one example, the prospective fruitfulness of the mechanisms literature is also demonstrated by the fact that philosophical work is now extending out to related fields. A debate on mechanisms, processes and factors alongside mechanisms in "Evidence-Based Policy" is now well established (Cartwright and Hardie 2012; Cartwright 2012; La Caze and Colyvan 2016). More recently, PhD scholars have been investigating issues of mechanism and evidence in psychiatry (Lenssen 2019; Green-Hovda 2020).

The lesson we draw from discussions about evidence in medicine is that philosophical theory can be very helpful in addressing practical policy issues. This is unlikely to stay confined to policy work specific to medicine, given the work on mechanisms in physical sciences we have already referred to above (in thesis 3), and emerging work such as that on information security (Spring and Illari 2019). These resources are comparatively neglected by commentators on New Mechanism. This is at least partly because they are not very controversial within the literature, but it is also because they are not widely known or well understood outside it. This means their power and potential importance for understanding scientific practices is under-appreciated.

8 Conclusion

While our six theses span a wide range of topics, there is a common thread. All of them recognize and/or endorse various types of plurality. Specifically,

- There are many types of mechanisms, e.g. physical, biological, neurological and social mechanisms (thesis 1).
- There can be many models describing the same mechanism, and there is no single model that is the best for all conceivable purposes (thesis 2).
- There are many forms of explanations; mechanistic and non-mechanistic explanations co-exist (theses 3 and 4).
- There are many ways in which mechanistic knowledge is useful (theses 3, 4 and 6).
- There are many ways in which mechanisms can be investigated (thesis 5).
- There are many projects (outside of philosophy) in which philosophical theorizing about mechanisms can be fruitful (thesis 6).

The open and pluralistic character of mechanistic philosophy has not always been recognized, and we hope by attending to these theses we can avoid a number of potential misunderstandings that have sometimes been advanced in the literature. It is *not true* that:

- Mechanistic philosophers think that all or most adequate scientific explanations are mechanistic.
- Mechanistic philosophers think that explanations are always better if they add more detail to the description of the mechanism.
- Mechanistic philosophers think that the sole or principal use of knowledge of mechanisms is to provide mechanistic explanations.

If we are correct about the pluralistic spirit of mechanistic philosophy of science, then certain research problems deserve more attention than they have received to date. For example,

- Given the ubiquity of mechanisms across scientific domains, and given the permissiveness of the minimal mechanism concept, further work should be done to clarify taxonomic differences among mechanisms, and the implications of those differences for research strategies in various scientific domains.
- Given that there is, in practice at least, no one ideal mechanical model, further work should be done to understand how multiple models may be integrated for the purposes of explanation and control.
- Given that the mechanistic approach does not imply that all legitimate explanations are mechanistic, further work should be done on understanding the relationship between mechanistic and non-mechanistic explanations within various scientific domains.
- Given that mechanisms are ubiquitous, that mechanism discovery tools are diverse, and that mechanistic knowledge is a resource for intervention in problem domains, further emphasis should be given to applying mechanistic approaches to scientific, technical and policy domains outside widely discussed areas like molecular biology and neuroscience.

While we think our case for these theses makes clear the openness to pluralism among New Mechanists, we do not thereby mean to imply that the pluralistic approach embodied in the theses is beyond doubt. It may well be that some philosophers embracing a mechanistic approach will wish to argue for, e.g., a more monistic approach to explanation. If so, we hope these theses may help provide a way of organizing different species of mechanistic philosophy of science.

The mechanist approach is now an established paradigm in philosophy of science, and one that has proven fruitful already in some research and policy domains. Our hope in this paper has been to harvest insights of this paradigm so that we may put them to good use, and so we can see more clearly what remains to be done.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

References

Abrahamsen, Adele, Benjamin Sheredos and William Bechtel. 2018. Explaining visually using mechanism diagrams. In Glennan and Illari (2018a), 238–254.

Aizawa, Kenneth, and Drew Headley. forthcoming. Abduction and composition. *Philosophy of Science.*

Andersen, Holly. 2018. Complements, not competitors: Causal and mathematical explanations. *British Journal for the Philosophy of Science* 69(2): 485–508.

Aronson, Jeffrey, Adam La Caze, Michael Kelly, Veli-Pekka Parkkinen, and Jon Williamson. 2018. The use of mechanistic evidence in drug approval. *Journal of Evaluation in Clinical Practice* 24: 1166–1176.

Batterman, Robert, and Colin Rice. 2014. Minimal model explanations. *Philosophy of Science* 81(3): 349–376.

Baumgartner, Michael, and Alexander Gebharter. 2016. Constitutive relevance, mutual manipulability, and fat-handedness. *British Journal for the Philosophy of Science* 67: 731–756.

Bechtel, William. 2008. *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. New York and London: Routledge.

Bechtel, William, and Adele Abrahamsen. 2005. Explanation: A mechanist alternative. *Studies in History and Philosophy of the Biological and Biomedical Sciences* 36: 421–441.

Bechtel, William, and Robert Richardson. 1993. *Discovering complexity. Decomposition and localization as strategies in scientific research*. Princeton: Princeton University Press.

Bechtel, William 2019. From parts to mechanisms: Research heuristics for addressing heterogeneity in cancer genetics. *History and Philosophy of the Life Sciences* 41(27). https://doi.org/10.1007/s40656-019-0266-x.

Cartwright, Nancy. 2012. Will this policy work for you? Predicting effectiveness better: How philosophy helps (Presidential Address). *Philosophy of Science* 79: 973–989.

Cartwright, Nancy, and Jeremy Hardie. 2012. *Evidence-based policy: A practical guide to doing it better*. Oxford: Oxford University Press.

Chemero, Anthony, and Michael Silberstein. 2008. After the philosophy of mind: Replacing scholasticism with science. *Philosophy of Science* 75(1): 1–27.

Chirimuuta, Mazviita. 2014. Minimal models and canonical neural computations: The distinctness of computational explanation in neuroscience. *Synthese* 191(2): 127–153.

Clarke, Brendan, Donald Gillies, Phyllis Illari, Federica Russo, and Jon Williamson. 2014. Mechanisms and the evidence hierarchy. *Topoi* 33: 339–360.

Craver, Carl. 2007. *Explaining the brain*. Oxford: Clarendon Press.

Craver, Carl, and Lindley Darden. 2013. *In Search of mechanisms: Discoveries across the life sciences*. Chicago: University of Chicago Press.

Craver, Carl, and David Kaplan. 2020. Are more details better? On the norms of completeness for mechanistic explanations. *The British Journal for the Philosophy of Science* 71(1): 287–319.

Craver, Carl, and Beate Krickel. 2019. Mechanisms in science. *Oxford Bibliographies*. https://doi.org/10.1093/OBO/9780195396577-03.

Craver, Carl, Stuart Glennan, and Mark Povich. 2021. Constitutive relevance and mutual manipulability revisited. *Synthese*. https://doi.org/10.1007/s11229-021-03183-8.

Craver, Carl and James Tabery. 2019. Mechanisms in science. *The Stanford Encyclopedia of Philosophy* (Summer 2019 Edition). https://plato.stanford.edu/archives/sum2019/entries/science-mechanisms/.

Craver, Carl. 2014. The Ontic account of scientific explanation. In *Explanation in the special sciences: The case of biology and history*, ed. Marie Kaiser, Oliver Scholz, Daniel Plenge, and Andreas Hüttemann, 27–52. Synthese Library. Dordrecht: Springer.

Cummins, Robert. 2000. "How does it work?" vs. "What are the laws?": Two conceptions of psychological explanation. In *Explanation and cognition*, ed. Frank Keil and Robert Wilson, 117–144. Cambridge, MA: MIT Press.

Darden, Lindley. 2006. *Reasoning in biological discoveries*. New York: Cambridge University Press.

Darden, Lindley, and Carl Craver. 2002. Strategies in the interfield discovery of the mechanism of protein synthesis. *Studies in History and Philosophy of Biological and Biomedical Sciences* 33: 1–28.

de Mesquita, Bueno, James Morrow Bruce, Randolph Siverson, and Alistair Smith. 1999. An institutional explanation of the democratic peace. *The American Political Science Review* 93: 791–807.

Dupré, John. 2013. Mechanism and causation in biology I: Living causes. *Proceedings of the Aristotelian Society Supplementary* 87: 19–37.

Elster, Jon. 1989. *Nuts and bolts for the social sciences*. Cambridge: Cambridge University Press.

Falkenburg, Brigitte and Gregor Schiemann. 2019. *Mechanistic explanations in physics and beyond. European Studies in Philosophy of Science*, vol. 11. Cham: Springer.

Feynman, Richard, Robert Leighton, and Matthew Sands. 2011. *The Feynman lectures on physics (The new millennium)*. New York: Basic Books.

Gervais, Raoul, and Erik Weber. 2013. Inferential explanations in biology. *Studies in History and Philosophy of Biological and Biomedical Sciences* 44: 356–364.

Glennan, Stuart. 1996. Mechanisms and the nature of causation. *Erkenntnis* 44: 49–71.

Glennan, Stuart. 2002. Rethinking mechanistic explanation. *Philosophy of Science* 69(Supplement): S342–S353.

Glennan, Stuart. 2005. Modeling mechanisms. *Studies in the History of the Biological and Biomedical Sciences* 36(2): 443–464.

Glennan, Stuart. 2017. *The new mechanical philosophy*. Oxford: Oxford University Press.

Glennan, Stuart, and Phyllis Illari, eds. 2018a. *The Routledge handbook of mechanisms and mechanical philosophy*. Abingdon and New York: Routledge.

Glennan, Stuart and Phyllis Illari. 2018b. Introduction: Mechanisms and mechanical philosophies. In Glennan and Illari 2018a, 1–9.

Green-Hovda, Sydney. 2020. *Evidence for causal claims in medicine and the mental health sciences.* PhD Dissertation. University of Antwerp.

Hedström, Peter, and Petri Ylikoski. 2010. Causal mechanisms in the social sciences. *Annual Review of Sociology* 36: 49–67.

Hochstein, Eric. 2015. One mechanism, many models: A distributed theory of mechanistic explanation. *Synthese* 193(5): 1387–1407.

Howick, Jeremy, Paul Glasziou, and Jeffrey Aronson. 2013a. Can understanding mechanisms solve the problem of extrapolating from study to target populations (the problem of external validity)? *Journal of the Royal Society of Medicine* 106: 81–86.

Howick, Jeremy, Paul Glasziou, and Jeffrey Aronson. 2013b. Problems with using mechanisms to solve the problem of extrapolation. *Theoretical Medicine and Bioethics* 34: 275–291.

Huneman, Philippe. 2018. Diversifying the picture of explanations in biological sciences: Ways of combining topology with mechanisms. *Synthese* 195: 115–146.

Illari, Phyllis. 2011. Mechanistic evidence: Disambiguating the Russo-Williamson thesis. *International Studies in Philosophy of Science* 25(2): 139–157.

Illari, Phyllis. 2019. Mechanisms, models and laws in understanding supernovae. *Journal for General Philosophy of Science* 50: 63–84.

Illari, Phyllis, and Jon Williamson. 2012. What is a mechanism? Thinking about mechanisms *across* the sciences. *European Journal for Philosophy of Science* 2: 119–135.

Ioannidis, Stavros, and Stathis Psillos. 2017. In defense of methodological mechanism: The case of apoptosis. *Axiomathes* 27: 601–619.

Kaplan, David, and Carl Craver. 2011. The explanatory force of dynamical and mathematical models in neuroscience: A mechanistic perspective. *Philosophy of Science* 78(4): 601–627.

Keeton, William, and James Gould. 1986. *Biological science*, 4th ed. New York and London: W. W. Norton and Co.

Krickel, Beate. 2018. *The mechanical world: The metaphysical commitments of the new mechanistic approach.* Dordrecht: Springer.

Kuhlmann, Meinard. 2015. A mechanistic reading of quantum laser theory. In *Why is more different? Philosophical issues in condensed matter physics and complex systems*, ed. Brigitte Falkenburg and Margaret Morrison, 251–271. Dordrecht: Springer.

La Caze, Adam and Mark Colyvan. 2016. A challenge for evidence-based policy. *Axiomathes* 27(1): 1–13.

La Caze, Adam 2011. The role of basic science in evidence-based medicine. *Biology and Philosophy* 26: 81–98.

Lenssen, Josephine. 2019. *The eccentric manners of explanatory models: Towards an account of perspectival mosaic unity in psychiatry.* PhD thesis, Vrije Universiteit Amsterdam. https://research.vu.nl/en/publications/the-eccentric-manners-of-explanatory-models-towards-an-account-of.

Leuridan, Bert. 2010. Can mechanisms really replace laws of nature? *Philosophy of Science* 77(3): 317–340.

Levi, Arnon. 2014. Machine-likeness and explanation by decomposition. *Philosophers' Imprint* 14(6): 1–15.

Levi, Arnon, and William Bechtel. 2013. Abstraction and the organization of mechanisms. *Philosophy of Science* 80(2): 241–261.

Little, Daniel. 1991. *Varieties of social explanation.* Boulder: Westview Press.

Machamer, Peter, Lindley Darden, and Carl Craver. 2000. Thinking about mechanisms. *Philosophy of Science* 67: 1–25.

Maoz, Zeev, and Bruce Russett. 1993. Normative and structural causes of democratic peace, 1946–1986. *The American Political Science Review* 87: 624–638.

Menzies, Peter. 2012. The causal structure of mechanisms. *Studies in History and Philosophy of Biological and Biomedical Science* 43(4): 796–805.

Osimani, Barbara. 2014. Hunting side effects and explaining them: Should we reverse evidence hierarchies upside down? *Topoi* 33(2): 295–312.

Parkkinen, Veli-Pekka., Christian Wallmann, Michael Wilde, Brendan Clarke, Phyllis Illari, Michael Kelly, Charles Norell, Federica Russo, Beth Shaw, and Jon Williamson. 2018. *Evaluating evidence of mechanisms in medicine: Principles and procedures.* Dordrecht: Springer.

Polya, George. 1945. *How to solve it. A new aspect of mathematical method.* Princeton: Princeton University Press.

Ramsey, Jeffrey. 2008. Mechanisms and their explanatory challenges in organic chemistry. *Philosophy of Science* 77(5): 970–972.

Russo, Federica, and Jon Williamson. 2012. EnviroGenomarkers: The interplay between mechanisms and difference making in establishing causal claims. *Medicine Studies: International Journal for the History, Philosophy and Ethics of Medicine and Allied Sciences* 3: 249–262.

Salmon, Wesley. 1984. *Scientific explanation and the causal structure of the world*. Princeton: Princeton University Press.

Shapiro, Larry. 2019. A tale of two explanatory styles in cognitive psychology. *Theory and Psychology* 29(5): 719–735.

Simon, Herbert. 1957. *Models of man: social and rational. Mathematical essays on rational human behavior in a social setting*. New York: Wiley.

Spring, Jonathan, and Phyllis Illari. 2019. Building general knowledge of mechanisms in information security. *Philosophy and Technology* 32(4): 627–659.

Steel, Daniel. 2008. *Across the boundaries: Extrapolation in biology and social science*. New York: Oxford University Press.

Teller, Paul. 2010. Mechanism, reduction, and emergence in two stories of the human epistemic enterprise. *Erkenntnis* 73: 413–425.

Thagard, Paul. 2006. *Hot thought: Mechanisms and applications of emotional cognition*. Cambridge, MA: MIT Press.

Van Eck, Dingmar. 2015. Mechanistic explanation in engineering science. *European Journal for the Philosophy of Science* 5: 349–375.

Williamson, Jon. 2019. Establishing causal claims in medicine. *International Studies in the Philosophy of Science* 32: 33–61.

Wimsatt, William C. 1972. Complexity and organization. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association 1972*, 67–86.

Woodward, James. 2000. What is a mechanism? A counterfactual account. *Philosophy of Science* 69(S3): S366–S377.

Woodward, James. 2003. *Making things happen. A theory of causal explanation*. New York: Oxford University Press.

Wouters, Arno. 2007. Design explanation: Determining the constraints on what can be alive. *Erkenntnis* 67(1): 65–80.

Wright, Cory. 2012. Mechanistic explanation without the ontic conception. *European Journal for Philosophy of Science* 2(3): 375–394.

Wright, Cory, and William Bechtel. 2007. Mechanisms and psychological explanation. In *Philosophy of psychology and cognitive science*, ed. Paul Thagard, 31–79. New York: Elsevier.