

# A Causal Bayes Net Analysis of Glennan's Mechanistic Account of Higher-Level Causation (and Some Consequences)\*

Alexander Gebharter

## Abstract

One of Stuart Glennan's most prominent contributions to the new mechanist debate consists in his reductive analysis of higher-level causation in terms of mechanisms (Glennan, [1996]). In this paper I employ causal Bayes nets (CBNs) to reconstruct his analysis. This allows for identifying general assumptions which have to be satisfied to get the analysis working. I show that once these assumptions are in place, they imply (against the background of the CBN machinery) that higher-level causation indeed reduces to interactions between component parts of mechanisms. I also briefly discuss the plausibility of these assumptions and some consequences for the mechanism debate.

---

\*This is a draft paper. The final version of this paper is published under the following bibliographical data: Gebharter, A. (2019). A causal Bayes net analysis of Glennan's mechanistic account of higher-level causation (and some consequences). *British Journal for the Philosophy of Science*. Advance online publication. doi:10.1093/bjps/axz034

# Contents

<b>1 Introduction</b>	<b>2</b>
<b>2 Glennan’s Analysis of Higher-Level Causation</b>	<b>6</b>
<b>3 Causal Bayes Nets</b>	<b>9</b>
<b>4 Reconstructing Glennan’s Analysis</b>	<b>14</b>
<b>5 Consequences for the Mechanism Debate</b>	<b>23</b>
5.1 Do mechanisms add anything to lower-level causal structure? . . .	23
5.2 Mechanisms and inter-level causation . . . . .	26
5.3 Interventions into mechanistic hierarchies . . . . .	31
<b>6 Conclusion</b>	<b>36</b>
<b>References</b>	<b>38</b>

## 1 Introduction

Mechanisms play an important role for many sciences such as biology, medicine, and neuroscience. Mechanisms are systems that can be described at different levels. They consist of several lower-level component parts whose causal interactions give rise to the overall behaviour of the system as a whole (Glennan, [1996]; Illari and Williamson, [2000]; Machamer, Darden, and Craver, [2000]). Most modern mechanists are mainly interested in how mechanisms and models of mechanisms can be used for purposes of explanation and prediction (Craver, [2007]; Bechtel and Abrahamsen, [2005]). Mechanistic explanations come with the following key advantage over purely causal explanations: While purely causal knowledge only allows for an horizontal (or etiological) explanation of a

phenomenon, the understanding of the mechanism underlying this phenomenon provides a basis for a vertical (or synchronic) explanation across levels as well. So mechanisms provide a basis for causal and constitutional reasoning across different levels of organisation.<sup>1</sup>

Since the concept of mechanism is intimately connected to that of causation, it seems promising to try to reduce (or explain) one of these two concepts to (or in terms of) the other. Attempts to analyse mechanisms in terms of causation (Gebharter, [2014], [2017b]; Woodward, [2002]) seem quite natural, simply because mechanisms are understood as systems consisting of causally interacting parts anyway. The other direction, namely to analyse causation in terms of mechanisms, seems more challenging. The probably most prominent mechanistic account of causation has been put forward by Glennan ([1996]). In a nutshell, Glennan proposes that causal relations (or laws) connecting higher-level phenomena occur (or hold) due to mechanisms underlying these phenomena.

Several philosophers such as Craver ([2007]) and Psillos ([2004]) have attacked Glennan's ([1996]) mechanistic analysis of higher-level causation, and Glennan has defended his approach against these and other criticisms (see, for example, Glennan, [2011]). The debate is still ongoing (Casini, [2016]) and the question of whether a mechanistic analysis of higher-level causation can be successful is still controversial. This paper is intended as a contribution to the debate about how mechanisms relate to higher-level causation. I use the causal Bayes net (CBN) machinery (Pearl, [2000]; Spirtes, Glymour, and Scheines, [1993]) as a tool for investigating Glennan's mechanistic analysis of higher-level causation. The framework allows for testing causal hypotheses, it provides the basis for a multitude of sophisticated procedures for uncovering causal structure

---

<sup>1</sup>How to define and identify levels is an important but still controversial question (Craver and Bechtel, [2007]; Eronen, [2015]; Kästner, [2018]). In this paper I will bracket this problem and assume that it can be solved in some way.

(Spirtes *et al.*, [1993]), and the approach satisfies standards successful empirical theories satisfy as well. In particular, it provides the best explanation for why sets of variables describing different empirical systems show different probabilistic dependence and independence patterns and the theory as a whole can be tested on empirical grounds (Gebharter, [2017b]; Schurz and Gebharter, [2016]). The hope is that the framework can, thus, also be used to get an empirically informed grasp on philosophical issues such as the question of how higher-level causation relates to mechanisms.<sup>2</sup>

Note that CBNs have already been used for modelling mechanisms (see, for example, Casini, Illari, Russo, and Williamson, [2011]; Clarke, Leuridan, and Williamson, [2014]; Gebharter, [2014]; Koch, Eisinger, and Gebharter, [2017]; Murray-Watters and Glymour, [2015]).<sup>3</sup> I am also not the first one using causal modelling techniques to shed new light on the question of whether higher-level causation can be reduced to mechanisms. Casini ([2016]), for example, but also Glennan ([2011]) himself borrows from the causal modelling literature. The main difference between these approaches and my own approach is that I combine different levels as well as causal and constitutional relationships in a single causal model.<sup>4</sup> The reconstruction of Glennan's ([1996]) mechanistic analysis to be developed in this paper will show that a supporter of such an analysis needs to make three interrelated basic assumptions in order to reduce higher-

---

<sup>2</sup>Note that alternative frameworks might be used for investigating how higher-level causation relates to mechanisms as well. However, I think that CBNs are especially nice for the reasons mentioned.

<sup>3</sup>For possible problems with such approaches, see (Kaiser, [2016]; Weber, [2016]). For a possible solution to some of these problems, see (Gebharter and Koch, [2018]).

<sup>4</sup>Casini *et al.* ([2011]) and Clarke *et al.* ([2014]) combine causal and constitutional relationships as well. They are, however, rather interested in modelling mechanistic hierarchies than in analysing higher-level causation in terms of mechanisms.

level causation to mechanisms. Once one subscribes to these assumptions, it follows from the CBN apparatus that higher-level causation reduces to interactions between component parts of mechanisms. Though an in-depth evaluation of these assumptions would go beyond the scope of this paper, I briefly discuss them in order to get a first impression of their putative plausibility.

The paper is structured as follows: In [section 2](#) I present Glennan's ([1996]) mechanistic account of higher-level causation. In [section 3](#) I introduce the basics of the CBN framework required for subsequent sections. In [section 4](#) I use the CBN framework as a basis for reconstructing Glennan's analysis. I also formulate and briefly discuss the three basic assumptions mentioned. In [section 5](#) I discuss several interesting consequences of my analysis for the mechanism debate. In particular, I discuss (i) the worry that a mechanistic analysis of higher-level causation might just reduce higher-level causation to lower-level causation, (ii) several consequences for the debate about inter-level causation in the presence of mechanisms, and (iii) how interventions on mechanisms might work according to the analysis provided. I conclude in [section 6](#).

Summarising, the paper aims at the following three main goals:

- Fleshing out Glennan's ([1996]) basic idea by providing a general formal framework for reducing higher-level causation in terms of mechanisms.
- Using the framework as a tool for identifying several assumptions one has to make (in addition to the framework) in order to get the reduction done.
- Sketching how the framework might shed new light on several other issues discussed in the mechanism literature.

## 2 Glennan's Analysis of Higher-Level Causation

Glennan starts his ([1996]) seminal paper with Hume's problem. In a nutshell, Hume's problem consists in the question of how we can acquire causal knowledge though we seem to be unable to qualitatively distinguish between regularities due to causal laws and spurious correlations. From a phenomenological point of view, the two kinds of regularities are indistinguishable simply because the 'secret connexion' which binds cause and effect together does not produce any specific impression or experience in the observer. While Hume's own skeptical solution to the problem consists in redefining causation as regular conjunction, Glennan proposes that one might be able to distinguish between regularities due to causal laws and spurious correlations by checking whether there exists a mechanism which connects the phenomenon of interest with its purported cause. According to Glennan ([1996]), Hume's claim that 'no number of observations can yield an impression of a connection' (p. 49) between cause and effect is correct, but his claim that this connection is 'secret' is false. Actually, one can establish a causal relation by discovering a mechanism that connects the cause to the effect phenomenon, which will often require sophisticated scientific methods.

To further spell out the idea of reducing higher-level causation to mechanisms, Glennan ([1996]) provides the following definition of the concept of mechanism:

A mechanism underlying a behaviour is a complex system which produces that behaviour by of the interaction of a number of parts according to direct causal laws. (Glennan, [1996], p. 52)

According to this definition, a mechanism is always a mechanism for a specific behaviour. This nicely fits the idea that one and the same system might

behave in quite different ways, where interactions between different parts might be responsible for these different behaviours. This observation leads to the distinction between the parts of a system which are constitutively relevant for the behaviour of interest and those which are not. Let us call the former, in accordance with Craver ([2007]), the mechanism's components. As the definition says, these components are supposed to bring about the behaviour of interest by causally interacting with each other. The part of the definition which requires them to interact 'according to direct causal laws' is intended to guarantee genuine causal interactions between component parts as well as that the overall 'behaviour of the aggregate stems from a series of local interactions' (Glennan, [1996], p. 56). Note that later on Glennan ([2002], p. S344) replaced the phrase 'according to direct causal laws' in the original definition by 'where the interactions between parts can be characterised by direct, invariant, change-relating generalisations', where 'direct, invariant, change-relating generalisations' should be understood in the sense of Woodward ([2003]). However, since whether one prefers the one or the other phrase to characterise mechanisms will play no role for the reconstruction of Glennan's analysis of higher-level causation in section 4, I will stick with the simpler definition from the 1996 paper.

Now Glennan's ([1996]) mechanistic analysis of higher-level causation states that '[t]wo events are causally connected when and only when there is a mechanism connecting them' (p. 64). The analysis comes with a number of philosophical merits: Firstly, it renders questions whether higher-level phenomena are causally connected scientifically tangible. It suggests that questions like these can be answered on the basis of empirical investigations. This is, for example, also supported by the Russo-Williamson thesis (Russo and Williamson, [2007]) which states that establishing causal claims requires that difference-making evidence is supported by an understanding of the mechanism responsible for the

phenomenon of interest. Secondly, it provides a more sophisticated solution to Hume's problem than regularity theories of causation have to offer. While regularity theories redefine causation as constant conjunction and, thus, run the risk of rendering too many regularities causal, a mechanistic analysis allows for phenomena which are not causally related though they regularly occur together. Finally, a mechanistic analysis of higher-level causation allows for multiple realisability. Smoking, for example, is a cause of premature death, but there are many different mechanisms for explaining how exactly smoking can lead to premature death.

There are, of course, also problems with Glennan's ([1996]) mechanistic analysis of higher-level causation. One of the most frequently discussed problems is that a mechanistic account does not provide any possibility for analysing causal interactions among component parts at the fundamental physical level. While the account seems to be capable of solving Hume's problem for higher-level causation, the problem remains a threat at the fundamental physical level. Glennan ([2002], [2011]) argues that counterfactuals stating what would happen under interventions can be used to achieve knowledge of the causal relations among component parts at the fundamental physical level. He emphasises, however, that though interventions can be used for such epistemic purposes, they do not provide a basis for a thorough metaphysical analysis of the concept of causation.<sup>5</sup> Another problem with Glennan's mechanistic analysis might be that it seems to simply reduce higher-level causation to causation at the fundamental physical level. Glennan ([1996]) counters this worry by pointing out that 'it is not [the fundamental causal] laws which make the causal claim true; rather, it is the structure of the higher level mechanism and the properties of its parts'

---

<sup>5</sup>For a discussion of the limitations of interventions for testing causal relations underpinned by complex mechanisms, see (Casini, [2016]).

(p. 66). The idea here is that the presence of the mechanism as a structured whole makes—in addition to the causal structure at the fundamental physical level—a substantial contribution in bringing about cause-effect relationships at higher levels. I will come back to this and other issues in [section 5](#). But first, I will briefly introduce the basics of the CBN formalism relevant for subsequent sections.

### 3 Causal Bayes Nets

Bayes nets (BNs) have originally been developed to graphically store independence information and to simplify reasoning under uncertainty. The formalism can, however, be used to represent all kinds of relations and dependencies that have just the right formal properties. In particular, they must conform to the Markov condition ([Spirtes \*et al.\*, \[1993\]](#), p. 33), which establishes a connection between probabilistic dependence and graphical structure. Before I can introduce the Markov condition properly, a few preliminaries are required. In the following,  $\mathbf{V}$  shall be a set of random variables  $X_1, \dots, X_n$  which might represent events or properties,  $\mathbf{E}$  a set of arrows ( $\rightarrow$ ) connecting pairs of variables in  $\mathbf{V}$ , and  $P$  a probability distribution over  $\mathbf{V}$ .  $\mathbf{G} = \langle \mathbf{V}, \mathbf{E} \rangle$  is called the BN's graph and  $\mathbf{Par}(X_j)$  shall stand for the set of a variable  $X_j$ 's parents, which is the set of all  $X_i \in \mathbf{V}$  with  $X_i \rightarrow X_j$ . Variables  $X_i$  that are connected to  $X_j$  via a path of the form  $X_i \rightarrow \dots \rightarrow X_j$  are called  $X_j$ 's ancestors, and variables  $X_j$  connected to a variable  $X_i$  over such a path are called  $X_i$ 's descendants. The Markov condition can be formulated as follows:

**Definition 1** (Markov condition).  $\langle \mathbf{V}, \mathbf{E}, P \rangle$  satisfies the Markov condition if and only if  $\mathbf{G} = \langle \mathbf{V}, \mathbf{E} \rangle$  and probability distribution  $P$  conform to

$$P(X_1, \dots, X_n) = \prod_{i=1}^n P(X_i | \mathbf{Par}(X_i)). \quad (1)$$

The conditional probabilities  $P(X_i|\mathbf{Par}(X_i))$  appearing in Equation 1 are called  $X_i$ 's parameters. The Markov condition is basically a screening off condition. If satisfied, it guarantees that every variable  $X_i \in \mathbf{V}$  is independent of its non-descendants (which are the variables  $X_j$  not connected via a path  $X_i \rightarrow \dots \rightarrow X_j$  to  $X_i$ ) conditional on its direct ancestors (which are the variables in  $\mathbf{Par}(X_i)$ ).<sup>6</sup> Now one kind of relation that conforms to Equation 1 and, thus, could be represented by a BN's arrows, is direct causal dependence (Pearl, [2000]; Spirtes *et al.*, [1993]). We shall call a BN in which some (or all) arrows are causally interpreted a causal Bayes net (CBN). For a model in which all arrows are causally interpreted, the Markov condition implies that every variable becomes probabilistically independent of all its non-effects given its direct causes. It also implies Reichenbach's ([1956]) insights that common causes screen off their effects and that an effect's direct causes screen it off from its indirect causes.

Other relations that seem to conform to Equation 1 are, for example, supervenience and constitution (Gebharder, [2017a], [2017c]). Recent work by Schaffer ([2016]) suggests that the grounding relation might conform to the Markov condition as well. In the following, however, I will only be interested in BNs that represent causal and/or constitutional relationships. I will, from now on, represent direct causal connections by continuous arrows ( $\longrightarrow$ ) and direct constitutional relationships by dashed arrows ( $\dashrightarrow$ ).<sup>7</sup> I will refer to a variable

---

<sup>6</sup>Probabilistic independence between  $X_i$  and  $X_j$  conditional on  $X_k$  is defined as  $P(x_i|x_j, x_k) = P(x_i|x_k) \vee P(x_j, x_k) = 0$  for all  $X_i$ -,  $X_j$ -, and  $X_k$ - values  $x_i$ ,  $x_j$ , and  $x_k$ , respectively. Probabilistic dependence is defined as the negation of probabilistic independence.

<sup>7</sup>Note that representing these different kinds of relations by different arrows is just a cosmetic procedure. Both kinds of arrows are assumed to technically work in exactly the same way.

$X_i$ 's direct causes as its causal parents and to  $X_i$ 's direct constituents as its constitutional parents. Likewise, I will refer to the variables  $X_i$  connected to a variable  $X_j$  via a path of the form  $X_i \longrightarrow \dots \longrightarrow X_j$  as  $X_j$ 's causal ancestors, and to the variables  $X_j$  connected to a variable  $X_i$  over such a path as  $X_i$ 's causal descendants. If  $X_i$  and  $X_j$  are connected over a path of the form  $X_i \dashrightarrow \dots \dashrightarrow X_j$ , I will refer to  $X_i$  as one of  $X_j$ 's constitutional ancestors and to  $X_j$  as one of  $X_i$ 's constitutional descendants.

Given the terminology introduced we can state the Markov condition (Definition 1) for systems possibly featuring both causal and constitutional relations as follows:

**Definition 2** (causal/constitutional Markov condition).  $\langle \mathbf{V}, \mathbf{E}, P \rangle$  satisfies the causal/constitutional Markov condition if and only if every  $X_i \in \mathbf{V}$  is probabilistically independent of its non-descendants conditional on its causal and constitutional parents.

The CBN framework comes with a simple test for whether a particular arrow  $X_i \longrightarrow X_j$  (or  $X_i \dashrightarrow X_j$ ) can mediate probabilistic influence between  $X_i$  and  $X_j$ . To determine whether  $X_j$ -variations are correlated with  $X_i$ -variations in some circumstances due to the arrow  $X_i \longrightarrow X_j$  (or  $X_i \dashrightarrow X_j$ ), one has to check whether  $X_j$  probabilistically depends on  $X_i$  conditional on  $X_j$ 's other parents (Schurz and Gebharter, [2016], p. 1087), meaning that one has to check whether  $P(x_j|x_i, \mathbf{r}) \neq P(x_j|\mathbf{r}) \wedge P(x_i, \mathbf{r}) > 0$  holds for some combination of  $X_i$ -values  $x_i$ ,  $X_j$ -values  $x_j$ , and values  $\mathbf{r}$  of  $\mathbf{Par}(X_j) \setminus \{X_i\}$ . If the answer to this question is positive, we say that the arrow  $X_i \longrightarrow X_j$  (or  $X_i \dashrightarrow X_j$ ) is productive. Arrows that do not pass the productivity test are redundant and can be eliminated without making the model less informative.<sup>8</sup> This fact will play a crucial role

---

<sup>8</sup>Assuming productivity for all of a model's arrows amounts to assuming minimality (Schurz and Gebharter, [2016], Theorem 2). Note that applying the productivity test to single

for evaluating Glennan’s ([1996]) mechanistic analysis of higher-level causation in section 4.

Since the productivity test introduced will do a lot of work in subsequent sections, further elaboration and illustration seems appropriate. I will use the test as a kind of Occam’s razor: If two (or more) parents of a variable compete for relevance, the one (or ones) not making any difference must go. Assume, for example, we are interested in whether Suzy throws a stone ( $S = 1/0$ ) and in whether Billy throws a feather ( $B = 1/0$ ) at a vase which might, as a result, shatter ( $V = 1/0$ ). Suppose we are interested in the following two competing causal hypotheses:

$h_1$ : Whether Suzy throws ( $S = 1/0$ ) is causally relevant for whether the vase breaks ( $V = 1/0$ ).

$h_2$ : Whether Billy throws ( $B = 1/0$ ) is causally relevant for whether the vase breaks ( $V = 1/0$ ).

$S$  passes the productivity test. If Billy throws the feather, Suzy throwing the stone still makes a difference for whether the vase breaks ( $P(V = 1|S = 1, B = 1) \neq P(V = 1|B = 1)$ ) and we can conclude that  $h_1$  is true. But since whether Billy throws does neither make a difference if Suzy throws ( $P(V = 1|B = 1, S = 1) = P(V = 1|B = 0, S = 1) = P(V = 1|S = 1)$ ) nor if she does not throw ( $P(V = 1|B = 1, S = 0) = P(V = 1|B = 0, S = 0) = P(V = 1|S = 0)$ ),  $B$  fails the test and we conclude that  $h_2$  is false. Note that this does not yet mean that  $h_2$  can be reduced to (or explained in terms of)  $h_1$ . However, depending on the context, one can use the productivity test for answering questions like these

---

causal or constitutional arrows does only require that the model satisfies the Markov condition; neither minimality (Spirtes *et al.*, [1993], p. 34) nor the stronger faithfulness assumption (Spirtes *et al.*, [1993], p. 56) is required. Consequently, I will neither assume minimality nor faithfulness in this paper.

as well. While  $h_1$  and  $h_2$  describe same-level phenomena, we can construct a similar case involving phenomena located at different levels. Assume, for example, I expect it to rain later on and we are interested in the causal impact of whether I decide to take an umbrella with me. Suppose we are interested in the following two competing causal hypotheses:

$h_3$ : My decision ( $D = 1/0$ ) is causally relevant for whether I take an umbrella with me ( $U = 1/0$ ).

$h_4$ : The corresponding brain processes ( $B = 1/\dots/n$ ) are causally relevant for whether I take an umbrella with me ( $U = 1/0$ ).

Depending on how exactly my decisions are related to what is going on in my brain one will get different results from applying the productivity test to this scenario. If my decisions are multiply realisable and supervene on brain processes, then  $h_3$  will fail and  $h_4$  will pass the test. In the context of different levels this could be interpreted as  $h_3$  being reducible to  $h_4$ . While it makes perfect sense to say that my decisions are causally relevant for whether I take an umbrella with me in everyday live situations and higher-level causal claims might also be useful for scientific explanation they can, in the end, be reduced: Higher-level phenomena are only causally relevant because they relate to lower-level phenomena in a certain way, but the lower-level phenomena do all the actual causal work. In subsequent sections I will use the productivity test in this latter sense: as an indicator for when certain causal relations can be reduced in terms of more fundamental causal relations (and interlevel relationships of constitution).

## 4 Reconstructing Glennan’s Analysis

In this section, I use the CBN framework as a basis for developing a reconstruction of Glennan’s ([1996]) mechanistic account of higher-level causation.<sup>9</sup> For illustrative purposes, I introduce an abstract toy example that shall stand proxy for all possible cases involving higher-level causation. Now, for Glennan, the analysis of higher-level causation starts with two higher-level phenomena that stand in a causal relationship. Let us model the possible behaviours of the systems exhibiting these phenomena with the two random variables  $H_1$  and  $H_2$ . In addition, we assume that  $H_1$  is causally relevant for  $H_2$ . Hence,  $H_1 \longrightarrow H_2$  has to be a part of our model. Further details about how the possible behaviours of the systems described by  $H_1$  and  $H_2$  are related to each other is modelled by a probability distribution  $P$  over  $\mathbf{V}$ , where  $\mathbf{V}$  is a set of variables containing  $H_1$ ,  $H_2$ , and several other variables to be introduced soon. Since  $H_1$  is assumed to be a cause of  $H_2$ , we assume that  $H_2$  depends on  $H_1$ , meaning that  $P(h_2|h_1) \neq P(h_2) \wedge P(h_1) > 0$  holds for some  $H_1$ - and  $H_2$ -values  $h_1$  and  $h_2$ , respectively.

The next step in Glennan’s ([1996]) analysis would be to decompose the two systems whose higher-level behaviours are described by  $H_1$  and  $H_2$  and to identify the component parts of the mechanisms responsible for these behaviours. There are currently several proposals for how to identify a mechanism’s component parts on the market (see, for example, Baumgartner and Casini, [2017]; Baumgartner and Gebharder, [2016]; Craver, [2007]; Gebharder, [2017c]; Harbecke, [2015]; Krickel, [2018]). However, since we are first and foremost inter-

---

<sup>9</sup>The basic idea how to model causal and inter-level relations in one and the same model is the same as in (Gebharder, [2017a]). While the productivity test introduced in section 3 is used to answer questions about the causal efficacy of mental properties in (Gebharder, [2017a]), I use it as a device for exploring issues concerning reduction in terms of mechanisms in this paper.

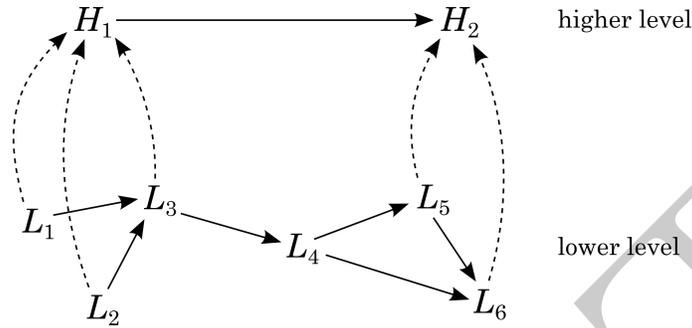


Figure 1

ested in reconstructing Glennan’s analysis rather than in the question of how to distinguish the parts constitutively relevant for certain behaviours from those that are not, we can ignore the latter question for the moment. For now it is only important to stress that some parts of our systems stand in constitutive relevance relations to the higher-level behaviours of interest, while others do not. Let us assume that the mechanism underlying  $H_1$  consists of three constitutively relevant parts, while the mechanism responsible for  $H_2$  consists of two. We model  $H_1$ ’s component parts with the variables  $L_1$ ,  $L_2$ , and  $L_3$  and  $H_2$ ’s components with  $L_5$  and  $L_6$ . Let us assume that we also found a non-component part which we will model with  $L_4$ . (There might be additional non-component parts which we will ignore for now.) Next, let us assume that the lower-level variables  $L_1, \dots, L_6$  are causally connected to each other like in Figure 1.<sup>10</sup> By adding a dashed arrow for every component part we finally arrive at the graph depicted in Figure 1.<sup>11</sup>

<sup>10</sup>Note that nothing hinges on the particular causal structure among lower-level variables as long as the assumptions to be introduced are satisfied.

<sup>11</sup>Constitutive relevance arrows are drawn from lower- to higher-level variables because higher-level phenomena behave w.r.t. their component parts similarly to how effects behave w.r.t. to their causes. For a more detailed argumentation, see (Gebharder, [2017c]).

Now the proxy model in Figure 1 and the CBN framework can be used as a tool for investigating the commitments of Glennan’s ([1996]) mechanistic account of higher-level causation. If higher-level causation can actually be reduced in terms of mechanisms, this would amount to the claim that certain constraints that allow for eliminating higher-level causal arrows such as  $H_1 \longrightarrow H_2$  have to hold for all models like the one in Figure 1. Now the crucial question is whether we can formulate such constraints. They would have to come in the form of general assumptions one has to make in order to render higher-level arrows such as  $H_1 \longrightarrow H_2$  unproductive in the sense explained in section 3. The reduction of the higher-level causal relationship between  $H_1$  and  $H_2$  would, according to this idea, consist in demonstrating that the arrow  $H_1 \longrightarrow H_2$  is redundant, meaning that it does no causal work that is not already done by some lower-level causal structure.

Here comes the first assumption required to get the reduction of higher-level causation in terms of the model: All kinds of higher-level phenomena occur because of underlying mechanisms. The assumption guarantees that for every higher-level phenomenon described by some variable  $H$  there will be lower-level constitutively relevant parts that can be described by variables  $L_1, \dots, L_n$ . It is typically also assumed that the overall behaviour of a mechanism  $H$  supervenes on the behaviours of the mechanism’s component parts  $L_1, \dots, L_n$  (see, for example, Baumgartner and Gebharder, [2016]; Craver, [2007]; Glennan, [1996]).<sup>12</sup>

---

<sup>12</sup>Note that in mechanistic models actually used for explanation the supervenience assumption will not always hold. However, the project of reducing higher-level causation in terms of mechanisms is not epistemic, but metaphysical in nature, and almost everyone seems to agree that, given the full metaphysical picture, a mechanism’s macro behaviour supervenes on the behaviours of its component parts. Throughout the paper I assume that lower-level variables  $L_1, \dots, L_n$  describing the component parts of a mechanism can in principle be defined, even if actually unknown.

We can summarise and formulate these requirements in terms of CBN models as follows:

**Assumption 1** (supervenience/constitution). For each higher-level phenomenon described by a variable  $H$  there are component parts described by variables  $L_1, \dots, L_n$  such that  $L_1 \dashrightarrow H, \dots, L_n \dashrightarrow H$  and  $P$  is constrained by

$$\forall h \forall h' \exists \mathbf{l} : \text{If } h \neq h', \text{ then } P(\mathbf{l}|h) \neq P(\mathbf{l}|h'), \text{ and} \quad (2)$$

$$\forall \mathbf{l} \exists h : P(h|\mathbf{l}) = 1, \quad (3)$$

where  $\mathbf{l}$  (with values  $\mathbf{l}$ ) is an individual variable ranging over all possible combinations of instantiations of lower-level variables  $L_1, \dots, L_n$ .

Equation 2 and Equation 3 are constraints on the probability distributions that come with models of mechanisms. They reflect probabilistic implications of supervenience and constitutive relevance: If  $H$  supervenes on  $\{L_1, \dots, L_n\}$ , then every change of  $H$ 's value will be accompanied by a change in the probability distribution over  $\{L_1, \dots, L_n\}$  (Equation 2). And vice versa: Whenever the values of all variables  $L_1, \dots, L_n$  representing constitutively relevant parts are fixed, then also  $H$ 's value will be determined (Equation 3).

Now Assumption 1 renders higher-level causal relations redundant. This can be illustrated on the basis of our toy model as follows:<sup>13</sup> Let  $\mathbf{l}$  be an arbitrarily chosen combination of instantiations of  $H_2$ 's constitutively relevant variables  $L_5$  and  $L_6$ . It follows with Equation 3 that there is exactly one  $H_2$ -value  $h_2$  for every  $\mathbf{l}$  such that  $P(h_2|\mathbf{l}) = 1$ , while  $P(h'_2|\mathbf{l}) = 0$  for all  $H_2$ -values  $h'_2$  different from  $h_2$ . Now there are two possible cases for every  $H_1$ -value  $h_1$ : Either (i)  $h_1$  and  $\mathbf{l}$  are compatible, meaning that  $P(h_1, \mathbf{l}) > 0$ , or (ii) they are not, meaning

---

<sup>13</sup>An analogous argument can be formulated against the efficacy of mental causes in the context of the causal exclusion problem (Gebharder, [2017a]).

that  $P(h_1, \mathbf{1}) = 0$ . If (i) is the case, then  $P(h_2|h_1, \mathbf{1}) = P(h_2|\mathbf{1})$  and  $P(h'_2|h_1, \mathbf{1}) = P(h'_2|\mathbf{1})$  hold. Thus, no  $H_2$ -value depends on  $h_1$  conditional on  $\mathbf{1}$ . If (ii) is the case, then no  $H_2$ -value depends on  $h_1$  conditional on  $\mathbf{1}$  by definition. So conditionalising on  $\mathbf{1}$  renders  $H_2$  independent from  $H_1$  in both cases (i) and (ii). Because  $\mathbf{1}$  was arbitrarily chosen, this result generalises:  $H_2$  and  $H_1$  are independent conditional on any possible combination of instantiations of the variables  $L_5$  and  $L_6$ , meaning that  $H_2$  and  $H_1$  are independent conditional on  $\text{Par}(H_2) \setminus \{H_1\} = \{L_5, L_6\}$ . Thus, the arrow  $H_1 \longrightarrow H_2$  is unproductive. Note that this reasoning does not hinge on the particular structure of our exemplary model. It can be generalised for all kinds of CBN models involving mechanisms that satisfy [Assumption 1](#).

Note, however, that [Assumption 1](#) is still not sufficient for a general reduction of higher-level causation to mechanisms. One needs to assume, in addition, that at least some of the components of the cause mechanism are causally relevant for at least some of the components of the effect mechanism. In other words, one needs to assume that for every higher-level causal relation there is a corresponding lower-level causal connection. In [Figure 1](#), this assumption is already satisfied due to each one of the directed causal chains going through  $L_4$ . We can state this second assumption as follows:

**Assumption 2** (lower-level causal relevance #1). If a higher-level phenomenon described by a variable  $H_i$  is causally relevant for another higher-level phenomenon described by a variable  $H_j$ , then at least some of  $H_i$ 's component parts are causally relevant for  $H_j$ 's component parts as well.

But why exactly do we need [Assumption 2](#) in addition to [Assumption 1](#)? The need for [Assumption 2](#) can be illustrated by means of our exemplary model in [Figure 1](#) as follows: Assume that the component variables of the two mechanisms represented would not be causally connected. This could, for example, be the

case if there were no arrow between  $L_3$  and  $L_4$ . However, we still assume that  $H_1$  is a higher-level cause of  $H_2$ , which suggests that some changes in the cause mechanism (consisting of  $H_1$ ,  $L_1$ ,  $L_2$ , and  $L_3$ ) can bring about changes in the target mechanism (consisting of  $H_2$ ,  $L_5$ , and  $L_6$ ). If there is no causal connection between the supervenience bases of  $H_1$  and  $H_2$ , then the only causal relation that can account for this fact is  $H_1 \rightarrow H_2$ . In order to establish such a connection between  $H_1$  and  $H_2$ , however,  $H_1 \rightarrow H_2$  must pass the productivity test introduced in [section 3](#). Hence, the arrow  $H_1 \rightarrow H_2$  can be rendered redundant only if  $H_1$ 's supervenience base and  $H_2$ 's supervenience base are causally connected. There are two possibilities for such a causal connection:  $H_1$ 's and  $H_2$ 's supervenience bases have a common cause or  $H_1$ 's supervenience base is a cause of  $H_2$ 's supervenience base. In order to account for the fact that  $H_2$ -changes can be brought about by producing changes in  $H_1$ , there must be a directed causal chain from  $H_1$ 's supervenience base to  $H_2$ 's supervenience base.

There is one final assumption one needs to make for a successful reduction of higher-level causation to interactions among component parts of the mechanisms involved. Let me briefly illustrate why one needs to make it. Assume that  $H_1$  has four possible values (0, 1, 2, 3),  $H_2$  has three possible values (0, 1, 2), and  $L_5$  and  $L_6$  each have two possible values (0, 1). (Since details about the other variables of the model are not required for the point I want to make here, I will ignore those variables for the moment.) Now assume that  $H_1$ ,  $L_5$ , and  $L_6$  are connected in such a way that the following conditionals hold:

If  $H_1 = 0$ , then  $L_5 = 0$  and  $L_6 = 0$ .

If  $H_1 = 1$ , then  $L_5 = 0$  and  $L_6 = 1$ .

If  $H_1 = 2$ , then  $L_5 = 1$  and  $L_6 = 0$ .

If  $H_1 = 3$ , then  $L_5 = 1$  and  $L_6 = 1$ .

In addition, we assume that  $H_2$ 's parameters are defined as follows, where

$i \in \{0, 1, 2, 3\}$  and  $h_{2,0}$  stands short for  $H_2 = 0$ ,  $h_{2,1}$  stands short for  $H_2 = 1$ , and

so on:

$$\begin{aligned}
 P(h_{2,0}|\mathbf{r}) = 1 \quad P(h_{2,1}|\mathbf{r}) = 0 \quad P(h_{2,2}|\mathbf{r}) = 0 & \quad \text{if } \mathbf{r} = h_{1,i}, l_{5,0}, l_{6,0} \\
 P(h_{2,0}|\mathbf{r}) = 0 \quad P(h_{2,1}|\mathbf{r}) = 1 \quad P(h_{2,2}|\mathbf{r}) = 0 & \quad \text{if } \mathbf{r} = h_{1,i}, l_{5,1}, l_{6,0} \\
 P(h_{2,0}|\mathbf{r}) = 0 \quad P(h_{2,1}|\mathbf{r}) = 1 \quad P(h_{2,2}|\mathbf{r}) = 0 & \quad \text{if } \mathbf{r} = h_{1,i}, l_{5,0}, l_{6,1} \\
 P(h_{2,0}|\mathbf{r}) = 0 \quad P(h_{2,1}|\mathbf{r}) = 0 \quad P(h_{2,2}|\mathbf{r}) = 1 & \quad \text{if } \mathbf{r} = h_{1,i}, l_{5,1}, l_{6,1}
 \end{aligned}$$

These parameters are chosen in such a way that whenever two of the three parents of  $H_2$  take any of their values, also the third parent's value is determined. This means that no change in  $H_2$ 's value can lead to a change in the probability distribution over one of  $H_2$ 's parents conditional on  $H_2$ 's remaining parents. Given these rather extreme constraints, the productivity test implies that the higher-level causal arrow  $H_1 \rightarrow H_2$  is not productive. However, it also follows that the constitutional arrows  $L_5 \dashrightarrow H_2$  and  $L_6 \dashrightarrow H_2$  are not productive. In addition, deleting  $H_1 \rightarrow H_2$  renders  $L_5 \dashrightarrow H_2$  and  $L_6 \dashrightarrow H_2$  and deleting  $L_5 \dashrightarrow H_2$  or  $L_6 \dashrightarrow H_2$  renders  $H_1 \rightarrow H_2$  productive again. In cases such as the one constructed, the productivity test proposed in [section 3](#) cannot identify the culprit; it only tells us that at least one of the competing arrows cannot represent an efficacious relation.

How should we interpret this result? First of all, note that such cases require a very specific fine-tuning of the model's parameters and can, thus, be expected to be extremely rare.<sup>14</sup> However, the logical possibility of such scenarios demands for an answer to the question raised. Since a supporter of a mechanistic analysis of higher-level causation aims at reducing higher-level causation in terms of mechanisms, it is advisable for her to favour the constitutional

---

<sup>14</sup>The problem should, for example, not arise in cases of non-deterministic causal dependencies. That the causal influences among variables to be analysed are indeterministic or at least pseudo-indeterministic is a common assumption in the causal modelling literature (see, for example, [Spirtes et al., \[1993\]](#)).

arrows over the causal ones.<sup>15</sup> In other words, she needs to assume that the higher-level behaviour’s component parts rather than the higher-level causes are efficacious:

**Assumption 3** (constitutional priority). If constitutional as well as causal arrows pointing at a variable  $H$  describing some higher-level phenomenon do not pass the productivity test (described in [section 3](#)), then  $H$ ’s causal parents rather than  $H$ ’s constitutional parents are inefficacious w.r.t.  $H$ .

Summarising, the three assumptions introduced imply (to the background of the CBN framework) that higher-level causation reduces to causal interactions among component parts of the mechanisms involved. [Assumption 1](#) warrants that for every higher-level phenomenon there exist several component parts which, together, render any causal relation between that phenomenon and its higher-level causes unproductive. [Assumption 3](#) guarantees that those higher-level causes are actually inefficacious and that the causal relations do not fail the productivity test simply because one or more constitutively relevant parts are inefficacious. Finally, [Assumption 2](#) is required in order to account for the fact that the effect mechanism can be influenced by manipulating the cause mechanism. Let me stress again that the possibility of reducing higher-level causation in terms of mechanisms does not hinge on the specific structure of the model in [Figure 1](#). The results of this section apply to all CBN models involving mechanisms as long as one subscribes to Assumptions 1-3.

Before pointing at some interesting consequences of the results obtained in this section, let me say a few words about the plausibility of Assumptions 1-3.

---

<sup>15</sup>Another reason to go for the constitutional arrows is that a mechanism’s components are typically assumed to be difference makers. Recall from [section 3](#) that the productivity test is basically a test for whether a parent of a variable is a difference maker for that variable. So if  $L_5$  and  $L_6$  are components of  $H_2$ , the arrows  $L_5 \dashrightarrow H_2$  and  $L_6 \dashrightarrow H_2$  should stay and  $H_1 \rightarrow H_2$  must go.

[Assumption 1](#) nicely fits the new mechanist movement and probably every mechanist would subscribe to it: Higher-level phenomena are brought about by causal interactions among their constitutively relevant parts and the behaviour of the system as a whole is connected to the behaviours of the component parts via supervenience. Note that the supervenience assumption is not only supported by mechanists, but also by almost everyone who believes in a layered view of the world (see, for example, [Kim, \[2007\]](#)). [Assumption 2](#) is quite plausible as well. If my higher-level decision is causally relevant for raising my left arm, then it seems that also the processes in my brain over which this decision supervenes are causally relevant for that phenomenon. My arm going up without any neurological causal basis would be more than miraculous. Finally, also [Assumption 3](#) seems plausible. It grants priority to constitutional relationships over causal relationships. This nicely fits the idea that constitution is a metaphysical relation that is more essential for an object or phenomenon than causation. While the presence of a cause might bring about quite different effects in different possible worlds due to different causal laws, the presence and behaviour of the constituents of a higher-level object or phenomenon determine the behaviour of this object or phenomenon in all possible worlds. It is, for example, hard to see how certain spatial changes in the molecules constituting my left arm could not result in a corresponding arm movement at the higher-level. Summarising, it seems that all three assumptions—at least at first sight—are quite plausible and, thus, that Glennan's ([\[1996\]](#)) mechanistic account of higher-level causation might be supported by a CBN analysis such as the one developed in this section. But note that for a final verdict about the three assumptions' plausibility a much more thorough investigation would be required. Anyway, the main aim of this paper consists in spelling out which assumptions one would have to subscribe to in order to reduce higher-level causation in terms of mechanisms.

## 5 Consequences for the Mechanism Debate

In this section I discuss three interesting consequences and connections of the results obtained in [section 4](#) to other issues discussed in the literature on mechanisms. I discuss (i) the worry that a mechanistic analysis of higher-level causation might just reduce higher-level causation to lower-level causation, (ii) several consequences for the debate about inter-level causation in the presence of mechanisms, and (iii) how interventions on mechanisms might work according to the analysis provided in [section 4](#).

### 5.1 Do mechanisms add anything to lower-level causal structure?

If [Glennan \(\[1996\]\)](#) is right and higher-level causation can actually be reduced to interactions between component parts of mechanisms, then one might wonder whether the analysis could not also be done without mechanisms. This is basically [Psillos' \(\[2004\]\)](#) worry briefly mentioned in [section 1](#): At some point one runs out of mechanisms when analysing higher-level causation mechanistically and it seems that, in the end, the causal structure at the lowest level is doing all the work in reducing higher-level causation. [Glennan \(\[2011\]\)](#) countered this objection by arguing that his mechanistic account and an interventionist treatment of causation<sup>16</sup> depend on each other and, thus, stand and fall together: While the mechanistic approach depends on interventionism when it comes to accounting for causal interactions among component parts, interventionism, in turn, '[...] relies on the mechanical approach because the truth-conditions for counterfactuals depend upon the structure of mechanisms' ([Glennan, \[2011\]](#), p. 806). Glennan goes on by arguing that interventionism faces a regress problem

---

<sup>16</sup>[Psillos \(\[2004\]\)](#) and [Glennan \(\[2011\]\)](#) both presume an interventionist understanding of causation.

similar to the one of the mechanistic approach. I agree with [Casini \(\[2016\]\)](#) that Glennan's argumentation does not go through because both accounts run into different kinds of regresses: The mechanist runs into a vicious ontological regress, while the interventionist only faces a probably rather harmless conceptual regress.

In this section I am not particularly interested in how the mechanistic account and interventionism relate to each other. I still take Psillos' ([\[2004\]](#)) worry as a starting point, but approach it from a different angle. Interestingly, Glennan already thought about the possible objection that his analysis might just reduce higher-level causation to causal relations at the fundamental physical level back then in the 1990s. He addressed it by pointing out that '[a]lthough the mechanism responsible for connecting two events may supervene upon lower-level mechanisms, and ultimately on mechanistically inexplicable [causal] laws of fundamental physics, it is not these laws which make the causal claim true; rather, it is the structure of the higher level mechanism and the properties of its parts' ([Glennan, \[1996\]](#), p. 66). Which role mechanisms play exactly in making the higher-level causal claim true is, however, not explained in more detail and there are several possibilities how one could interpret this quote. Here is one such possibility: The structure of the higher-level mechanism that is required in addition to the lower-level causal structure basically consists of the particular component parts whose causal interactions give rise to the higher-level phenomenon of interest. Information about this structure consists primarily in information about which parts are constitutively relevant for that phenomenon and about their causal interactions.<sup>17</sup> This kind of information is required to

---

<sup>17</sup>I am indebted to anonymous referee for pointing out that the structure of a mechanism meant here might also include the spatiotemporal ordering of the component parts. Since the point I want to make in this subsection does not require this kind of structural information, I will abstract from it for now.

link the higher-level phenomenon to what is causally going on at the lower level and, thus, is essential for getting the reduction done. However, just making such a claim does not yet establish its truth. The reconstruction and the CBN model developed in [section 4](#) can be used to further support it.

The reconstruction comes with clear criteria for when a higher-level causal relation is reducible. The model shows that more is needed than just the lower-level causal structure in order to reduce higher-level causation. In particular, one also needs to connect the higher-level causal structure with the lower-level causal structure by adding several dashed arrows representing constitutive relevance relations ([Assumption 1](#)). Without these additional arrows, the productivity test would yield that  $H_1 \rightarrow H_2$  is productive (and, thus, the reduction would fail) and the Markov condition ([Definition 1](#)) would be violated because there would be probabilistic dependencies between variables at different levels (for example, between  $H_2$  and  $L_5$  and between  $H_2$  and  $L_6$ ) not accounted for by any connection in the graph. This shows that a successful and consistent reduction requires mechanisms as mediators between the behaviours of higher- and lower-level variables. Note that acquiring information about the structure of mechanisms, information about which parts of a system are constitutively relevant for the overall behaviour of that system, is not a trivial task. Constitutive relevance relations cannot be read off the variables and causal relations represented in the model. Decomposing mechanisms and finding out which parts are actually components are quite demanding processes which cannot be done on a priori grounds. This is further supported by the fact that there are currently several quite sophisticated theories of constitutive relevance and approaches for constitutive relevance discovery on the market (see, for example, [Baumgartner and Gebharder, \[2016\]](#); [Baumgartner and Casini, \[2017\]](#); [Craver, \[2007\]](#); [Gebharder, \[2017c\]](#); [Harbecke, \[2010\]](#), [\[2015\]](#)).

## 5.2 Mechanisms and inter-level causation

Several authors have recently argued for or against inter-level causation in the presence of mechanisms. [Leuridan \(\[2012\]\)](#), for example, is sympathetic towards interpreting constitutive relevance as a bidirectional causal relation between variables at different levels, and [Krickel \(\[2017\]\)](#) argues for the view that mechanistic hierarchies do not exclude inter-level causation if one takes it seriously that ‘the relata of the mechanistic level relation are acting entities’ (p. 543). Authors such as [Craver and Bechtel \(\[2007\]\)](#), [Kistler \(\[2009\]\)](#), and [Romero \(\[2015\]\)](#), on the other hand, claim that inter-level causation can be reduced to truths about constitutive relevance relations and lower-level causal structure. In the following, I will explore how the analysis provided in [section 4](#) may be used in order to support the reductionist camp.

Let me illustrate this by means of the toy example introduced in [section 4](#). Assume, in addition to the assumptions already in place, that the model features the two inter-level causal relations  $L_1 \rightarrow H_2$  and  $H_1 \rightarrow L_5$  ([Figure 2](#)).<sup>18</sup> A reduction in the sense of [Craver and Bechtel \(\[2007\]\)](#) would, again, consist in showing that these two arrows are redundant and can, in principle, be deleted from the model and that all the work these arrows seem to do is actually done by constitutive relevance relations and the lower-level causal structure. Let us start with bottom-up causation and with the arrow  $L_1 \rightarrow H_2$ . With [Assumption 1](#) it follows that  $L_1 \rightarrow H_2$  is unproductive for exactly the same reasons the higher-level causal arrow  $H_1 \rightarrow H_2$  turned out as unproductive in [section 4](#) ( $L_1$  and  $H_2$  are independent conditional on  $\mathbf{Par}(H_2) \setminus \{L_1\} = \{L_5, L_6\}$ ); and for the few cases in which the productivity test would output  $L_1 \rightarrow H_2$  as well as one or both of the constitutional arrows  $L_5 \dashrightarrow H_2$  and  $L_6 \dashrightarrow H_2$

---

<sup>18</sup>Since the focus is on inter-level causation now, I have deleted the higher-level causal arrow  $H_1 \rightarrow H_2$ .

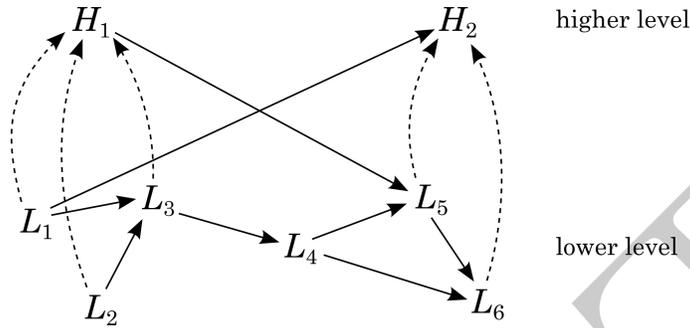


Figure 2

as unproductive, [Assumption 3](#) tells us that the bottom-up cause should be considered as inefficacious. Finally, one needs to guarantee that manipulating  $L_1$  can bring about changes in  $H_2$ . If the bottom-up arrow  $L_1 \rightarrow H_2$  should turn out as unproductive, one needs to account for this fact by assuming that one or more directed chains from  $L_1$  to  $H_2$  go through  $H_2$ 's supervenience base  $\{L_5, L_6\}$ . This can be guaranteed by subscribing to the following assumption:

**Assumption 4** (lower-level causal relevance #2). If a lower-level phenomenon described by a variable  $L_i$  is causally relevant for a higher-level phenomenon described by a variable  $H_j$ , then  $L_i$  is causally relevant for at least some of  $H_j$ 's component parts as well.

[Assumption 4](#) has the same function for reducing bottom-up causation as [Assumption 2](#) has for reducing higher-level causation. It seems plausible for almost the same reasons, too:<sup>19</sup> It would be miraculous if there were bottom-up causes without a corresponding lower-level causal story to be told. Assume, for example, that certain processes in my brain caused the raising of my left arm. Then, surely, there is a causal story to be told about how these brain processes caused the muscles in my left arm to contract in specific ways which

<sup>19</sup>Like in the case of the other assumptions, a more thorough evaluation of [Assumption 4](#) would go beyond the scope of this paper.

constitute the raising of my left arm. A common cause path connecting these brain processes and the muscle contractions in my left arm would not be enough to explain the dependence of the raising of my left arm on these brain processes.

Reducing the top-down causal arrow  $H_1 \longrightarrow L_5$  is more cumbersome. Assumptions 1 and 3 which, together, render causes inert in the presence of constitutive relevance relations have no bearing on lower-level variables such as  $L_5$ . To render  $H_1$  inefficacious w.r.t.  $L_5$ , one needs to subscribe to additional assumptions. Here is a possible candidate:

**Assumption 5** (extended productivity).  $X_i \longrightarrow X_j$  (or  $X_i \dashrightarrow X_j$ ) is productive if and only if  $X_j$  depends on  $X_i$  conditional on  $\mathbf{Par}(X_j) \setminus \{X_i\}$ . If  $X_i$  is a higher-level direct cause of  $X_j$ , then  $X_j$  must also depend on  $X_i$  conditional on the set  $\mathbf{V}'$  of  $X_i$ 's constitutive ancestors.

Assumption 5 is basically an extended version of the productivity test introduced in section 3. It states that for a causal arrow to be productive, it does not always suffice that the two variables the arrow connects cannot be screened off by the effect variable's other parents. If the cause variable is a higher-level variable, then the variables constitutively relevant for the cause variable have to be considered as well. Assumption 5 seems plausible on first sight. The question on which everything seems to hinge is whether the behaviour of the system at the higher level ( $H_1$ ) can contribute anything to the causal work done by the mechanism's component parts ( $L_1$ ,  $L_2$ , and  $L_3$ ).

Unfortunately, it turns out that Assumption 5 is not sufficient to reduce top-down causation in all cases. The problem is that though the set of constituents  $\{L_1, L_2, L_3\}$  clearly screens  $H_1$  off from  $L_5$ , Assumption 5 does not exclude probability flow along the arrow  $H_1 \longrightarrow L_5$  in general. In particular, the causal arrow  $H_1 \longrightarrow L_5$  might be responsible for a probabilistic dependence between  $H_1$  and  $L_5$  conditional on  $L_4$ . As a consequence,  $H_1 \longrightarrow L_5$  would not be

redundant since  $H_1 \rightarrow L_5$  would pass the productivity test and, thus, deleting  $H_1 \rightarrow L_5$  would violate the Markov condition. Since constitution is assumed to formally work exactly like causation, the situation is equivalent to a case in which a variable  $X_1$  is a deterministic direct cause of another variable  $X_2$  and, at the same time, a probabilistic direct cause of another variable  $X_3$ , where both  $X_2$  and  $X_3$  are probabilistic direct causes of a fourth variable  $X_4$ . In such a case, nothing logically excludes the possibility that  $X_2$  is dependent on  $X_4$  conditional on its other parent  $X_3$ , meaning that nothing excludes that  $X_2 \rightarrow X_4$  passes the original productivity test and is actually required to avoid a violation of the Markov condition.

Another—and, as far as I can see, the only—possible way to go consists in subscribing to a variant of the closure of the physical domain, which says that if there are physical causes for a physical event, then citing all of these physical causes gives us all the information we might get for whether the physical event of interest actually occurs (see, for example, Kim, [2007]). In the light of this assumption, higher-level causes cannot provide any additional information. This assumption can be adopted for levels of mechanisms:

**Assumption 6** (causal closure of levels). If there exist lower-level direct causes  $L_1, \dots, L_n$  as well as higher-level direct causes  $H_1, \dots, H_m$  for a lower-level variable  $L_{n+1}$ , then  $H_1, \dots, H_m$  do not pass the productivity test introduced in section 3.

Assumption 6 amounts to the claim that higher-level causes always fail the original productivity test in the presence of same-level causes. So while the reduction of higher-level as well as bottom-up causation came quite naturally due to assumptions which have—at least at first glance—some plausibility, the reducibility of top-down causation needs to be assumed by brute force. Anyway, it seems clear that to reduce top-down causation in terms of mechanisms, a variant of the causal closure of levels is required. If this is so and subscribing to

[Assumption 6](#) amounts to nothing over and above begging the question, then the question of whether there is top-down causation in the presence of mechanisms becomes, in the end, an empirical question; it entirely depends on whether there are overall behaviours of mechanistic systems that pass the productivity test.

Note that even if one is ready to subscribe to the causal closure of levels or is able to establish it on empirical grounds, one still has to guarantee that lower-level causal arrows have priority over competing top-down causal arrows if both kinds of arrows turn out to be unproductive and the deletion of one of these arrows would render the others productive again. One could do this analogously to how constitutional arrows were prioritised by [Assumption 3](#):

**Assumption 7** (lower-level causal priority). If lower-level as well as higher-level causal arrows pointing at a variable  $X$  describing some phenomenon do not pass the productivity test (described in [section 3](#)), then  $X$ 's higher-level causal parents rather than  $X$ 's lower-level causal parents are inefficacious w.r.t.  $X$ .

[Assumption 7](#) has some plausibility. If one has to give up causal efficacy either for the higher- or for the lower-level cause, almost everyone would decide in favour of the lower-level cause. This is, for example, an assumption commonly accepted in the debate surrounding causal exclusion (see [Kim, \[2007\]](#)). It can also be supported by the arm raising example mentioned in support of [Assumption 2](#) in [section 4](#).

Summarising, the analysis provided in [section 4](#) reduces bottom-up causation to lower-level causation plus constitutive relevance if one is ready to accept [Assumption 4](#). In order to reduce top-down causation as well, one has to make additional assumptions such as the more problematic [Assumption 6](#) and [Assumption 7](#). Supporters of reductive accounts of inter-level causation, such as [Craver and Bechtel \(\[2007\]\)](#), might be able to strengthen their position by com-

ing up with plausible arguments for Assumptions 4 and 7 and by supporting Assumption 6 on empirical grounds.

### 5.3 Interventions into mechanistic hierarchies

How interventions work in the presence of mechanistic hierarchies is still not well understood. Craver ([2007]), for example, suggests that it might be possible to surgically intervene on the whole as well as on any part of a mechanism, where interventions are to be understood in the sense of Woodward ([2003], p. 98).<sup>20</sup> Baumgartner and Gebharder ([2016]) and Romero ([2015]) show that there is no way to surgically intervene on a mechanism’s overall behaviour as long as one accepts that the higher-level behaviour of a mechanism supervenes over the behaviours of its component parts. They argue that all interventions on the mechanism as a whole are actually fat-handed, meaning that they are common causes of the higher-level behaviour as well as of the behaviours of some component parts. Let us see whether we can add something to that debate.

For illustrating how interventions might work according to the analysis provided in section 4, let us have another look at the exemplary model introduced in that section. Let us introduce a higher-level intervention variable  $I_{H_1}$  for  $H_1$ .<sup>21</sup> According to the mechanistic account of higher-level causation, there must be some lower-level component parts for  $I_{H_1}$ . Let us assume that  $I_{L_1}$  and  $I'_{L_1}$  model two of them, that  $I'_{L_1}$  is a direct cause of  $I_{L_1}$ , and that both lower-level variables are direct causes of  $L_1$ . In addition, we add the causal arrows

---

<sup>20</sup>An intervention on a variable is ‘surgical’ if it causes no other variable directly.

<sup>21</sup>I use a very weak notion of ‘intervention variable’ here. For  $I_X$  to count as an intervention variable for another variable  $X$  I will only assume that  $I_X$  is a direct cause of  $X$  or of one or several of  $X$ ’s constitutional ancestors. Under an intervention on a mechanism I will understand an intervention on a variable of that mechanism.

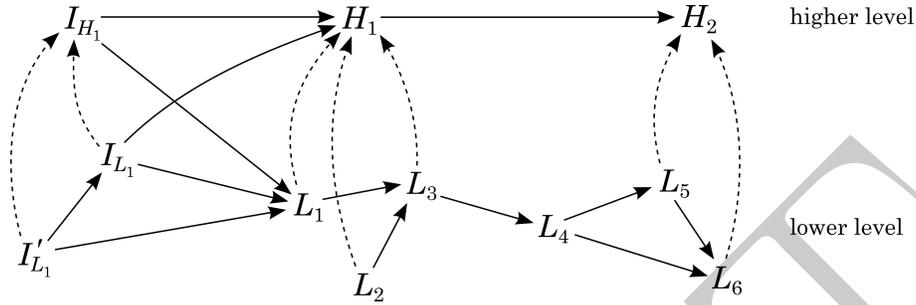


Figure 3

$I_{H_1} \rightarrow L_1$  and  $I_{L_1} \rightarrow H_1$  and arrive at the structure depicted in Figure 3.<sup>22</sup>

As we have seen, subscribing to Assumptions 1-4 renders possible causes of higher-level variables such as  $H_1$  causally inert. (Note that the problematic Assumption 6 and also Assumption 7 are not required for this result.) Under these conditions, the result of Baumgartner and Gebharder ([2016]) and Romero ([2015]) that the behaviour of a mechanism at the higher level can only be manipulated while, at the same time, intervening on some of that mechanism's component parts, turns out as correct. But we can also add something to this result: Actually, all the work is done by the lower-level variables. So interventions on mechanisms are in some sense fat-handed. They are, however, not fat-handed in the sense of Baumgartner and Gebharder and Romero. By intervening on a mechanism one does not intervene on higher- and lower-level variables via different effective causal paths. One rather influences higher-level variables in virtue of exerting causal influence on the mechanism's constitutively relevant parts.

Our analysis can be used to shed new light on the following two related ques-

<sup>22</sup>The  $I$ -variables should represent all the possible ways of how a mechanism consisting of  $H_1$ ,  $L_1$ ,  $L_2$ ,  $L_3$  could be manipulated. They will stand proxy for all kinds of possible intra- and inter-level interventions on mechanisms and serve as a theoretical tool for investigating how such interventions on mechanistic systems might actually work.

tions: (i) How can post intervention distributions be computed in CBN models capturing mechanistic hierarchies, and (ii) how can constitutive relevance relations be inferred on the basis of interventions? Let us start with question (i). The effect of setting a variable  $X_i$  to value  $x_i$  by means of an intervention in a purely causal BN can be computed as follows (see Pearl, [2000]): As a first step, one deletes (or ‘breaks’) all the arrows pointing at  $X_i$ . Next, one applies Equation 1 to the resulting truncated graph in order to compute the conditional probabilities  $P(\mathbf{V}'|X_i = x_i)$  for any set  $\mathbf{V}' \subseteq \mathbf{V}$  of variables of interest. These conditional probabilities are then identified with the post intervention distribution over  $\mathbf{V}'$ . This procedure guarantees that intervening on a variable  $X_i$  can have an effect on  $X_i$ ’s descendants (its effects), but not on its ancestors (its causes) or causally independent variables.

Here is a first sketch of how this technique could also be used in order to compute the effects of interventions on mechanisms in systems spanning no more than two levels. According to our analysis, there are at least two different kinds of background knowledge one might have when intervening on mechanisms that nicely correspond to the different kinds of experiments biologists and neuroscientists carry out when they study how manipulating multi-level systems might influence behaviours at different levels (see, for example, Craver, [2007]; Kästner, [2018]). For each kind of knowledge/experiment there is a different way to compute the corresponding post intervention distribution.

**Bottom-up experiment.** This kind of experiment corresponds to a situation in which the experimenter knows exactly which lower-level variable(s) she actively changes, but not the effect of the intervention on the system’s behaviour at the higher level. In this case, one just applies the arrow breaking technique described: One first deletes the arrows into those lower-level variables which are the targets of the intervention and then computes the post intervention distri-

bution on the basis of the resulting graph and Equation 1. By doing this, one can learn about how changing the values of these lower-level variables probabilistically influences the overall behaviour of the mechanism or the higher- or lower-level behaviours of other mechanisms causally connected to it.

**Top-down experiment.** One might, the other way round, know that an experimental intervention produced a specific higher-level behaviour  $Y = y$ , but not which parts of the system the intervention influenced directly and which of these parts were actually responsible for that higher-level behaviour. According to our analysis, this corresponds to a top-down experiment. The effect of such an intervention on any variable  $X$ 's value  $x$  (describing, for example, a specific behaviour of a lower-level part) can be computed as follows: First, consider all the possibilities how the mechanism's lower-level variables could have been influenced by a (possibly fat-handed) intervention. For a set  $\mathbf{L}$  of  $m$  lower-level variables there will be  $|\mathcal{P}(\mathbf{L})| - 1$  such possibilities. Let us label them  $S_1, \dots, S_n$ . Next, produce the truncated graph  $\mathbf{G}_i$  for each possible scenario  $S_i$  by deleting the arrows into the variables intervened on in that scenario. For each pair  $S_i$  and  $\mathbf{G}_i$  list all those possible value combinations  $\mathbf{l}$  of the variables intervened on that are compatible with  $X = x$  and  $Y = y$ . For each of them, compute the conditional probability  $P(x|y, \mathbf{l})$  on the basis of Equation 1 and the truncated graph  $\mathbf{G}_i$ . Record all of these probabilities in a set  $\mathbf{P}$ . After having done this for every pair  $S_i$  and  $\mathbf{G}_i$ , compute the average of all the probabilities recorded in  $\mathbf{P}$  and identify the post intervention distribution with this average.<sup>23</sup> The

---

<sup>23</sup>This reflects the situation in which the experimenter considers all possible intervention scenarios  $S_i$  as equally likely. If one has a more concrete idea of how likely it is that  $Y = y$  has been produced by intervening on specific parts of the system, one could use the weighted average instead. The weights would then correspond to the experimenters confidence in particular intervention hypotheses.

suggested procedure corresponds to observing  $Y = y$  and going through all the possibilities how this observation could have been produced by interventions on parts of the mechanism. On its basis one can learn about how an experiment might have influenced a mechanism's parts as well as how the intervention and the observation  $Y = y$  might influence the probability distributions over variables of other mechanisms causally connected to the first one.

Our analysis shows that whether an experiment counts as a bottom-up or a top-down experiment depends on the experimenter's background knowledge and explanatory goals. The result that these different kinds of experiments correspond to shifts in epistemic perspectives also nicely fits Kästner's ([2018]) observation that '[b]y fleshing out the perspectival aspect inherent in the mechanistic view we gain a better understanding of the explanatory practices in science [...]' (p. 77). There is clearly also space for several kinds of mixed experiments that combine the strategies for bottom-up and top-down experiments sketched.

Let me now come to the related point (ii) mentioned: constitutive relevance discovery. While it is not possible to surgically intervene on a mechanism's higher-level behaviour as a whole—which is exactly what Craver's ([2007]) prominent mutual manipulability criterion for constitutive relevance relations demands—, especially bottom-up experiments as characterised turn out to be useful for identifying constitutive relevance relations. According to the analysis developed, the following procedure would, for example, be sufficient to do the job: First, decompose the system (described by  $H$ ) into several parts (described by  $L_1, \dots, L_n$ ). Next, pick one of the lower-level variables  $L_i$ . Now check whether setting  $L_i$  to one of its values by intervention leads to any change in the probability distribution of  $H$  while the values of all other lower-level variables  $L_j$  are fixed by interventions. If  $H$  wiggles in such a bottom-up experiment, then  $L_i$  must be a component part of  $H$  or a cause of a component part not included

in  $L_1, \dots, L_n$ . This procedure requires, of course, that the lower-level variables can be controlled independently. More sophisticated discovery procedures for constitutive relevance relations might be developed on the basis of the analysis presented in this paper. The analysis might also be compatible with some of the approaches already on the market. The algorithmic procedure discussed in (Gebharter, [2017c]), for example, should be compatible.

## 6 Conclusion

In this paper I provided a reconstruction of Glennan's ([1996]) mechanistic account of higher-level causation in terms of CBNs. Reducing higher-level causation in a CBN setting means to specify general conditions under which higher-level causal arrows can be rendered redundant in all kinds of models involving mechanisms. In section 4 I specified three such assumptions: Higher-level phenomena can always be decomposed in component parts in such a way that they supervene on the behaviours of these component parts (Assumption 1). If a higher-level phenomenon is causally relevant for another one, then at least some component parts of the former must be causes of some component parts of the latter (Assumption 2). And, finally, constitutional arrows are always prioritised over causal arrows competing for productivity (Assumption 3). Since all three assumptions seem—at least at first glance—quite plausible from a mechanist's perspective, the CBN reconstruction might be used to provide additional support for Glennan's mechanistic analysis of higher-level causation from the perspective of an empirically informed theory of causation.

In section 5 I pointed to some interesting consequences of the findings of section 4 for other issues discussed in the mechanism literature. The analysis provided (i) shows that and also why lower-level causal structure as well as mechanisms are required for reducing higher-level causation. It also (ii) shows

that a reductive analysis of inter-level causation in terms of causal and constitutive relevance relations such as Craver and Bechtel's ([2007]) follows from the CBN machinery if one is ready to subscribe to three additional assumptions (Assumptions 4, 6, and 7). Finally, the analysis sheds new light on (iii) how interventions on mechanisms might actually work, how post intervention distributions could be computed for systems involving mechanisms, and how different kinds of experiments could be used to study mechanisms and to identify a mechanism's constitutively relevant parts. This is, of course, only the tip of the iceberg and much more work needs to be done in order to develop a full-fledged account of interventions on mechanisms and reliable methods for the discovery of causal and constitutive relevance relations in mechanisms.

**Acknowledgements:** I would like to thank Michael Baumgartner, Lorenzo Casini, Christian J. Feldbacher-Escamilla, Stuart Glennan, Mario Günther, Stephan Hartmann, Lena Kästner, Beate Krickel, Gerhard Schurz, Reuben Stern, and Jim Woodward for their input and important discussions. Thanks also to three anonymous referees for many helpful comments.

*Alexander Gebharter*  
*Department of Theoretical Philosophy*  
*University of Groningen*  
*Groningen, Netherlands*  
*alexander.gebharter@gmail.com*

## References

- Baumgartner, M. and Casini, L. [2017]: ‘An abductive theory of constitution’, *Philosophy of Science*, **84**, pp. 214–33.
- Baumgartner, M. and Gebharder, A. [2016]: ‘Constitutive relevance, mutual manipulability, and fat-handedness’, *British Journal for the Philosophy of Science*, **67**, pp. 731–56.
- Bechtel, W. and Abrahamsen, A. [2005]: ‘Explanation: A mechanist alternative’, *Studies in the History and Philosophy of the Biological and Biomedical Sciences*, **36**, pp. 421–41.
- Casini, L. [2016]: ‘Can interventions rescue Glennan’s mechanistic account of causality?’, *British Journal for the Philosophy of Science*, **67**, pp. 1155–83.
- Casini, L., Illari, P. M., Russo, F. and Williamson, J. [2011]: ‘Models for prediction, explanation and control: Recursive Bayesian networks’, *Theoria – An International Journal for Theory, History and Foundations of Science*, **26**, pp. 5–33.
- Clarke, B., Leuridan, B. and Williamson, J. [2014]: ‘Modelling mechanisms with causal cycles’, *Synthese*, **191**, pp. 1651–81.
- Craver, C. [2007]: *Explaining the brain*, Oxford: Clarendon Press.
- Craver, C. and Bechtel, W. [2007]: ‘Top-down causation without top-down causes’, *Biology and Philosophy*, **22**, pp. 547–63.
- Eronen, M. I. [2015]: ‘Levels of organization: A deflationary account’, *Biology & Philosophy*, **30**, pp. 39–58.
- Gebharder, A. [2014]: ‘A formal framework for representing mechanisms?’, *Philosophy of Science*, **81**, pp. 138–53.
- Gebharder, A. [2017a]: ‘Causal exclusion and causal Bayes nets’, *Philosophy and Phenomenological Research*, **95**, pp. 353–75.
- Gebharder, A. [2017b]: *Causal nets, interventionism, and mechanisms: Philo-*

*sophical foundations and applications*, Cham: Springer.

- Gebharder, A. [2017c]: ‘Uncovering constitutive relevance relations in mechanisms’, *Philosophical Studies*, **174**, pp. 2645–66.
- Gebharder, A. and Koch, D. [2018]: ‘Combining causal Bayes nets and cellular automata: A hybrid modeling approach to mechanisms’, *British Journal for the Philosophy of Science*.
- Glennan, S. [1996]: ‘Mechanisms and the nature of causation’, *Erkenntnis*, **44**, pp. 49–71.
- Glennan, S. [2002]: ‘Rethinking mechanistic explanation’, *Philosophy of Science*, **69**, pp. S342–53.
- Glennan, S. [2011]: ‘Singular and general causal relations: A mechanist perspective’, in P. M. Illari, F. Russo and J. Williamson (eds), *Causality in the sciences*, Oxford: Oxford University Press, pp. 789–817.
- Harbecke, J. [2010]: ‘Mechanistic constitution in neurobiological explanations’, *International Studies in the Philosophy of Science*, **24**, pp. 267–85.
- Harbecke, J. [2015]: ‘The regularity theory of mechanistic constitution and a methodology for constitutive inference’, *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, **54**, pp. 10–9.
- Illari, P. M. and Williamson, J. [2012]: ‘What is a mechanism? Thinking about mechanisms across the sciences’, *European Journal for Philosophy of Science*, **2**, pp. 119–35.
- Kaiser, M. I. [2016]: ‘On the limits of causal modeling: Spatially-structurally complex biological phenomena’, *Philosophy of Science*, **83**, pp. 921–33.
- Kästner, L. [2018]: ‘Integrating mechanistic explanations through epistemic perspectives’, *Studies in History and Philosophy of Science Part A*, **68**, pp. 68–79.

- Kim, J. [2007]: *Physicalism, or something near enough*, Princeton: Princeton University Press.
- Kistler, M. [2009]: ‘Mechanisms and downward causation’, *Philosophical Psychology*, **22**, pp. 595–609.
- Koch, D., Eisinger, R. S. and Gebharder, A. [2017]: ‘A causal Bayesian network model of disease progression mechanisms in chronic myeloid leukemia’, *Journal of Theoretical Biology*, **433**, pp. 94–105.
- Krickel, B. [2017]: ‘Making sense of interlevel causation in mechanisms from a metaphysical perspective’, *Journal for General Philosophy of Science*, **48**, pp. 453–68.
- Krickel, B. [2018]: ‘Saving the mutual manipulability account of constitutive relevance’, *Studies in History and Philosophy of Science Part A*, **68**, pp. 58–67.
- Leuridan, B. [2012]: ‘Three problems for the mutual manipulability account of constitutive relevance in mechanisms’, *British Journal for the Philosophy of Science*, **63**, pp. 399–427.
- Machamer, P., Darden, L. and Craver, C. [2000]: ‘Thinking about mechanisms’, *Philosophy of Science*, **67**, pp. 1–25.
- Murray-Watters, A. and Glymour, C. [2015]: ‘What’s going on inside the arrows? Discovering the hidden springs in causal models’, *Philosophy of Science*, **82**, pp. 556–86.
- Pearl, J. [2000]: *Causality*, Cambridge: Cambridge University Press.
- Psillos, A. [2004]: ‘A glimpse of the secret connexion: Harmonizing mechanisms with counterfactuals’, *Philosophy of Science*, **12**, pp. 288–319.
- Reichenbach, H. [1956]: *The direction of time*, Berkeley: University of California Press.
- Romero, F. [2015]: ‘Why there isn’t inter-level causation in mechanisms’, *Syn-*

*these*, **192**, pp. 3731–55.

Russo, F. and Williamson, J. [2007]: ‘Interpreting causality in the health sciences’, *International Studies in the Philosophy of Science*, **21**, pp. 157–70.

Schaffer, J. [2016]: ‘Grounding in the image of causation’, *Philosophical Studies*, **173**, pp. 49–100.

Schurz, G. and Gebharter, A. [2016]: ‘Causality as a theoretical concept: Explanatory warrant and empirical content of the theory of causal nets’, *Synthese*, **193**, pp. 1073–103.

Spirtes, P., Glymour, C. and Scheines, R. [1993]: *Causation, prediction, and search*, Dordrecht: Springer.

Weber, M. [2016]: ‘On the incompatibility of dynamical biological mechanisms and causal graphs’, *Philosophy of Science*, **83**, pp. 959–71.

Woodward, J. [2002]: ‘What is a mechanism? A counterfactual account’, *Philosophy of Science*, **69**, pp. S366–77.

Woodward, J. [2003]: *Making things happen*, Oxford: Oxford University Press.