How to Model Mechanistic Hierarchies

Abstract

Mechanisms are usually viewed as hierarchical, with lower levels of a mechanism influencing, and decomposing, its higher-level behaviour. In order to adequately draw quantitative predictions from a model of a mechanism, the model needs to capture this hierarchical aspect. The recursive Bayesian network (RBN) formalism was put forward as a means to model mechanistic hierarchies (Casini et al., 2011) by decomposing *variables* into their constituting causal networks. The proposal was criticized by Gebharter (2014). He proposes an alternative formalism, which decomposes *arrows*. Here, I defend RBNs from the criticism and argue that they offer a better representation of mechanistic hierarchies than the rival account.

wordcount: 4899

1 Introduction

Mechanisms are usually viewed as hierarchical, with lower levels of a mechanism influencing, and decomposing, its higher-level behaviour. In order to adequately draw quantitative predictions from a model of a mechanism, the model needs to capture this hierarchical aspect. The *recursive Bayesian network* (RBN) formalism was put forward as a means to model mechanistic hierarchies (Casini et al., 2011). The formalism extends the Bayesian network (BN) formalism, already used to model same-level causal relations probabilistically (Pearl, 2000). In RBNs, higher-level *variables* decompose into lower-level causal BNs. The relations between a higher-level variable and its lower-level networks are constitutive.

This proposal was criticized by Gebharter (2014) and Gebharter and Kaiser (2014), on two main grounds: descriptive adequacy—it is unclear when the formalism is applicable to real mechanisms—and conceptual adequacy—RBNs do not allow one to draw interlevel inferences for explanation and intervention. To overcome such limitations, Gebharter (2014) has made the alternative proposal that decomposition involves *arrows* rather than variables. In particular, Gebharter (2014) proposes an alternative formalism, also extending the BN formalism, namely *multilevel causal models* (MLCMs).

Decomposing variables and decomposing arrows are two alternative ways of modelling mechanistic hierarchies, by which one may extend a probabilistic interpretation of causality. In this paper, I argue that the former option is superior to the latter. I proceed as follows. In §2, I present and illustrate RBNs and MLCMs. In §3, I argue against decomposing arrows. MLCMs lead to counterintuitive notions of mechanistic decomposition and mechanistic explanation. In §4, I defend RBNs from the criticism. RBNs do allow interlevel causal explanation, via the uncoupling of interlevel causal relations into a constitutional step and a causal step. RBNs also allow reasoning about interlevel interventions; believing otherwise depends on either wrongly assuming that changes cannot transmit along the constitutional downward-directed arrows, or on demanding that the RBN formalism represent intervention variables, which the formalism is not meant to represent.

2 The two formalisms

Both RBNs and MLCMs are extensions of the BN formalism. A BN consists of a finite set $V = \{V_1, \ldots, V_n\}$ of variables, each of which takes finitely many possible values, together with a directed acyclic graph (DAG) whose nodes are the variables in V, and the probability distribution $P(V_i|Par_i)$ of each variable V_i

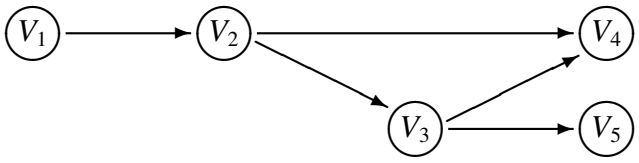

Figure 1

conditional on its parents Par_i in the DAG. DAG and probability function are linked by the *Markov Condition* (MC):

MC. For any $V_i \in \mathcal{V} = \{V_1, \ldots, V_n\}$, $V_i \perp\!\!\!\perp ND_i \mid Par_i$.

In words, each variable is probabilistically independent of its non-descendants, conditional on its parents. For instance, the DAG in figure 1 implies that V_4 is independent of V_1 and V_5 conditional on V_2 and V_3. In BN jargon, V_2 and V_3 'screen off' V_4 from V_1 and V_5. A BN determines a joint probability distribution over its nodes via $P(v_1 \cdots v_n) = \prod_{i=1}^{n} P(v_i|par_i)$ where v_i is an assignment $V_i = x$ of a value to V_i and par_i is the assignment of values to its parents induced by the assignment $v = v_1 \cdots v_n$.

In a *causally-interpreted* BN, the arrows in the DAG stand for direct causal relations and the network can be used to infer the effects of interventions and make probabilistic predictions (Pearl, 2000). In this case, MC is called the *Causal Markov Condition* (CMC).

2.1 Recursive Bayesian networks

RBNs represent hierarchies by decomposing variables (Casini et al., 2011). One of the motivations behind this choice is that scientists often talk of properties at

different levels that stand in a constitutive relation with one another.[1] Another motivation—only implicit in (Casini et al., 2011)—is that decomposing variables has the additional advantage of making 'interlevel causation' intelligible, by uncoupling (problematic) cases of interlevel downward or upward causation into two (less-problematic) steps, a constitutional, across-level step and a causal, same-level step (Craver and Bechtel, 2007). RBNs make this idea formally precise.

Mechanistic hierarchy is interpreted via the notion of 'recursive decomposition' of variables. An RBN is a BN defined over a finite set V of variables *whose values may themselves be RBNs*. A variable is called a *network variable* if one or more of its possible values is an RBN and a *simple variable* otherwise. A standard BN is an RBN whose variables are all simple. An RBN x that occurs as the value of a network variable in RBN y is said to be at a *lower level* than y; variables in y are the *direct superiors* of variables in x while variables in the same network are *peers*. If an RBN contains no infinite descending chains—i.e., if each descending chain of networks terminates in a standard BN—then it is *well-founded*. Only well-founded RBNs are considered here.

Consider a toy RBN[2] defined over $V = \{C, S\}$, where C represents whether some tissue in an organism is cancerous, taking the possible values 1 and 0, while

[1] Craver (2007) proposed that constitutive relations are established by the 'mutual manipulability' of higher- and lower-level properties that stand in the relation. Casini et al. (2011) referred to Craver's intuition to further motivate RBNs. The compatibility between Craver (2007)'s account of constitution and interventionism (Woodward, 2003), on which Craver's account rests, was questioned (see Baumgartner and Gebharter, 2015, and references therein). I will not discuss the issue here. I should only emphasize that the issue is orthogonal to the adequacy of RBNs as mechanistic models. RBNs are not tools for *establishing* constitution based on interventions, but tools for *representing* constitutional knowledge—however this may be got—and reasoning probabilistically across the levels.

[2] For more realistic examples, see Casini et al. (2011) and Clarke et al. (2014).

S is survival after 5 years, taking the possible values *yes* and *no*. The corresponding BN is: Suppose S is a simple variable but C is a network variable, with each

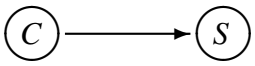

$P(C),\ P(S|C)$

Figure 2

of its two values denoting a lower-level (standard) BN that represents a state of the mechanism for cancer. I will ignore many of the factors, such as DNA damage response mechanisms, also responsible for cancer, and only focus on the unregulated cell growth and division, D, that results from mutations in the so-called 'growth factor', G. To the assignment of value 1 to C corresponds a lower-level network c_1 representing a functioning control mechanism, with a probabilistic dependence (and a causal connection) between G and D.

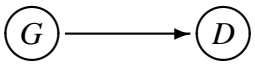

$P_{c_1}(G),\ P_{c_1}(D|G)$

Figure 3

To the assignment of value 0 to C corresponds a lower-level network c_0 representing a malfunctioning growth mechanism, with no dependence (and no causal connection) between G and D.

$P_{c_0}(G),\ P_{c_0}(D)$

Figure 4

Since these two lower-level networks are standard BNs, the RBN is well-founded and fully described by the three networks.

If an RBN is to be used to model a mechanism, the arrows at the various levels of the RBN signify causal connections. In addition, just as standard causally-interpreted BNs are subject to the CMC, a similar condition applies to causally-interpreted RBNs, called the *Recursive Causal Markov Condition* (RCMC). Let us indicate with NID_i the set of non-inferiors-or-descendants of V_i and with $DSup_i$ the set of direct superiors of V_i. Then, RCMC says that

RCMC. For any $V_i \in \mathcal{V} = \{V_1, \ldots, V_n\}$, $V_i \perp\!\!\!\perp NID_i \mid DSup_i \cup Par_i$.

In words, each variable in the RBN is independent of those variables that are neither its effects (i.e., descendants) nor its inferiors, conditional on its direct causes (i.e., parents) and its direct superiors. RCMC adds to CMC a recursive MC (RMC), viz. the condition that variables at any level are probabilistically independent of non-inferiors or peers given their direct superiors. Intuitively, if one knows the value of C, knowing the value of constituent variables G or D doesn't add anything to one's ability to infer to, say, the causes of C (here, none) or the effects of C (here, S). Notice that since the screening off that holds in virtue of RMC depends on constitutional rather than causal facts, not all dependencies identified by the RCMC can be causally interpreted.

While some authors treat CMC as a necessary truth, others argue against its universal validity (e.g., Williamson, 2005). A similar stance is adopted with respect to RCMC. RCMC is a modelling assumption in need of testing or justification, not a necessary truth. Thus, whether or not the formalism allows one to adequately represent a mechanism is an empirical rather than stipulative matter.

Inference in RBNs proceeds via a formal device called a *flattening*. Let $\mathcal{V} = \{V_1, \ldots, V_m\}$ $(m \geq n)$ be the set of variables of an RBN closed under the inferiority relation: that is, $\mathcal{V}$ contains the variables in V, their direct inferiors, their direct inferiors, and so on. Let $\mathcal{N} = \{V_{j_1}, \ldots, V_{j_k}\} \subseteq \mathcal{V}$ be the network variables in $\mathcal{V}$. For each assignment $n = v_{j_1}, \ldots, v_{j_k}$ of values to the network variables we can construct a standard BN, the *flattening* of the RBN with respect to n, denoted by $n^{\downarrow}$, by taking as nodes the simple variables in $\mathcal{V}$ plus the assignments $v_{j_1}, \ldots, v_{j_k}$ to the network variables, and including an arrow from one variable to another if the former is a parent or direct superior of the latter in the original RBN. The conditional probability distributions are constrained by those in the original RBN—in the RBN where V_{j_i} is the direct superior of V_i, $P(V_i|Par_i \cup DSup_i) = P_{v_{j_i}}(V_i|Par_i)$. Notice that MC holds in the flattening because the RCMC holds in the RBN. Only, since the arrows in the flattening that link variables to their direct inferiors are constitutional, CMC is not satisfied.[3]

The flattenings determine a joint distribution over $\mathcal{V}$ via $P(v_1 \cdots v_m) = \prod_{i=1}^{m} P(v_i|par_i dsup_i)$, where the probabilities on the right-hand side are determined by a flattening induced by $v_1 \cdots v_m$.[4]

[3] Notice that the role of RCMC—and of RBNs more generally (cf. fn. 1)—is not to establish constitution. Whether an arrow in the flattening is causal or constitutional is not dictated by MC, but depends on background knowledge.

[4] $P_{v_{j_i}}(V_i \mid Par_i)$ may be obtained from observed frequencies in a dataset. $P(V_i \mid Par_i DSup_i)$ can be obtained by either determining the corresponding observed frequencies from the original

In the cancer example, the flattening with respect to assignment c_1 is $c_1^{\downarrow}$:

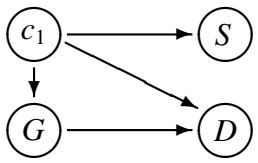

Figure 5

with probability distributions $P(c_1) = 1$ and $P(S|c_1)$ determined by the top level of the RBN, and with $P(d_1|g_1c_1) = P_{c_1}(d_1|g_1)$ determined by the lower level (similarly for g_0 and d_0). The flattening with respect to assignment c_0 is $c_0^{\downarrow}$:

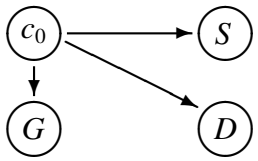

Figure 6

Again, $P(d_1|c_0) = P_{c_0}(d_1)$ etc. In each case the required conditional distributions are determined by the distributions given in the original RBN.

Having determined a joint distribution, the causally-interpreted RBN may be used to draw quantitative inferences for explanation and intervention, both within and across levels.

---

dataset, or by first selecting from all functions that satisfy the probabilistic constraints imposed by the RBN the function Q with maximum entropy (cf. Williamson, 2010) and then setting $P(V_i \mid Par_iDSup_i) = Q(V_i \mid Par_iDSup_i)$.

2.2 Multilevel causal models

According to Gebharter (2014), RBNs fail to model interlevel causal explanations and interventions, due to the lack of an explicit representation of interlevel causal arrows, over which causal influence propagates. (These objections, as I argue in §4, are based on misinterpreting RBNs.) Gebharter's proposed formalism purports to remedy these deficiencies by decomposing causal *arrows* rather than variables.[5] More precisely, a mechanistic hierarchy has to do with 'marginalizing out' variables when moving from a lower-level graph to a higher-level graph. In short, the formalism exploits the following idea: when the value of some variable X in the set of Y's parents $Par(Y)$ is unknown, $P(Y|Par(Y))$ may be calculated by summing over X's possible values, $\sum_{i=1}^{n} P(Y|Par(Y)_{X=x_i})$, thereby marginalizing X out. As a result, one gets a simpler probability distribution over $\mathcal{V}\setminus\{X\}$, consistent with the original one over $\mathcal{V}$.

Let us indicate a causal model as $\langle V, E, P\rangle$, where $\langle V, E\rangle$ is a DAG, defined over a variable set V and a set of edges E among them, and P an associated probability distribution.[6] Let $X \leftrightarrow Y$ indicate that two variables X and Y are effects of a latent common cause—i.e., a cause of X and Y not represented within the graph of some variable set V—and with $P^* \uparrow V$ the 'restriction' of the probability distribution P^* to a variable set V. The restriction of a lower-level causal model $\langle V^*, E^*, P^*\rangle$ to a higher-level causal model $\langle V, E, P\rangle$ is so defined (2014, 147):

Restriction. $\langle V, E, P\rangle$ is a restriction of $\langle V^*, E^*, P^*\rangle$ if and only if

[5] Gebharter and Kaiser (2014, §3.6) make the orthogonal suggestion that levels be *ontologically* distinct (partly) on the basis of constitutional relations between whole's and parts' properties. It is not clear how, if at all, this proposal relates to MLCMs. Thus, I will not discuss it here.

[6] To be consistent with Gebharter's notation, I henceforth denote sets with italic rather than calligraphic fonts.

a $V \subset V^*$, and

b $P^* \uparrow V = P$, and

c for all $X, Y \in V$:

c.1 if there is a directed path from X to Y in $\langle V^*, E^* \rangle$ and no vertex on this path different from X and Y is in V, then $X \rightarrow Y$ is in $\langle V, E \rangle$, and

c.2 if X and Y are connected by a common cause path π in $\langle V^*, E^* \rangle$ or by a path π free of colliders containing a bidirected edge in $\langle V^*, E^* \rangle$, and no vertex on this path π different from X and Y is in V, then $X \leftrightarrow Y$ is in $\langle V, E \rangle$, and

d no path not implied by **c** is in $\langle V, E \rangle$.

That is, the lower-level structure $\langle V^*, E^*, P^* \rangle$ represents the higher-level structure $\langle V, E, P \rangle$ iff $\langle V, E, P \rangle$ is the restriction of $\langle V^*, E^*, P^* \rangle$ uniquely determined when V^* is restricted to V. The restriction is such that information about causal relations and existence of common causes in $\langle V^*, E^* \rangle$ is preserved by $\langle V, E \rangle$, and the probabilistic information of P^* is consistent with P upon marginalizing out variables in $\{V^* \setminus V\}$.

A 'multi-level causal model' (MLCM) is so defined (2014, 148):

MLCM. $\langle M_1 = \langle V_1, E_1, P_1 \rangle, \ldots, M_n = \langle V_n, E_n, P_n \rangle \rangle$ is a multi-level causal model if and only if

a $M_1, \ldots, M_n$ are causal models, and

b every M_i with $1 < i \leq n$ is a restriction of M_1, and

c M_1 satisfies CMC.

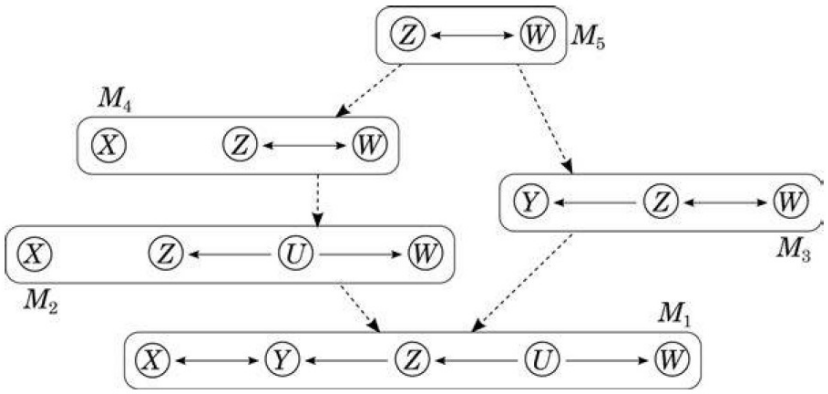

Figure 7

That is, a MLCM is an ordered set of causal models $\langle M_1 = \langle V_1, E_1, P_1\rangle, \ldots, M_n = \langle V_n, E_n, P_n\rangle\rangle$, where the bottom-level, unrestricted causal model M_1 satisfies CMC. (Higher-level models may not satisfy CMC.) Each causal model in the MLCM represents a mechanism.

The information on the hierarchical relations among the nested mechanisms in the MLCM is contained in a 'level graph' (2014, 149):

Level graph. A graph $G = \langle V, E\rangle$ is called an MLCM $\langle M_1 = \langle V_1, E_1, P_1\rangle, \ldots, M_n = \langle V_n, E_n, P_n\rangle\rangle$'s *level graph* if and only if

a $V = \{M_1, \ldots, M_n\}$, and

b for all $M_i = \langle V_i, E_i, P_i\rangle$ and $M j = \langle V_j, E_j, P_j\rangle$ in V: $M_i \rightarrow M_j$ is in G if and only if $V_i \subset V_j$ and there is no $M_k = \langle V_k, E_k, P_k\rangle$ in V such that $V_i \subset V_k \subset V_j$ holds.

A level graph $G = \langle V, E\rangle$ is constructed from a MLCM by adding dashed (non-causal) arrows between any two models M_i and M_j, $M_i \rightarrow M_j$, if and only if V_i is the largest proper subset of V_j in MLCM, so that M_i is, so to say, the smallest

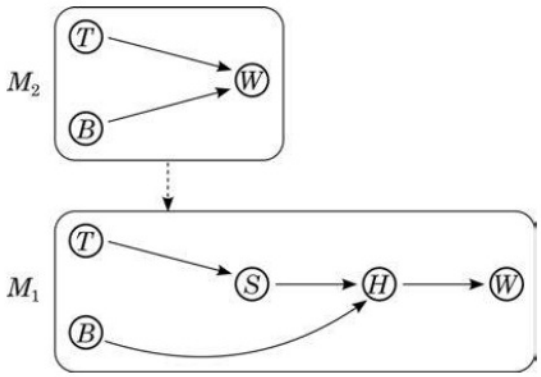

Figure 8

restriction of M_j. In figure 7 is an example of level graph (Gebharter, 2014, 150). Since the ordering among graphs is not strict, there may be graph pairs (e.g., M_2 and M_3; M_4 and M_3) that do not stand in a restriction relation.

Figure 8 depicts a more concrete example (Gebharter, 2014, 151), viz. a two-level water dispenser mechanism.[7] The room temperature T causally influences a sensor S; S, together with the status of a tempering button, B, cause the heater H to be on or off; H causes the temperature of the water dispensed, W.

3 Criticism of MLCMs

It is unclear whether hierarchies, as analysed in terms of the notion of 'marginalizing out', are *mechanistic*—that is, represent *mechanistic* decompositions and grant *mechanistic* explanations.

[7]Gebharter contrasts the virtues of this MLCM with an RBN of the 'same' mechanism (2014, 142-3). This may be misleading, since the two models cannot possibly represent the same mechanism. (In a nutshell, this is because the RBN contains constitutional arrows.) This motivates my choice of defending RBNs by reference to the toy model already introduced in §2.1.

First, it is unclear whether MLCMs represent mechanistic decompositions. High-level causal models in a MLCM, for instance M_2 and M_3 in figure 7, are just more coarse-grain representations of one and the same structure, viz. M_1, such that some of the information in M_1 is *missing* at the higher level, as the term 'restriction' suggests.

Second, it is unclear whether MLCMs represent mechanistic explanations. Admittedly, there is a sense in which one explains the relation between, say, the room temperature T and the water temperature W by uncovering the mediating role of the sensor S and the heater H. However, this sort of explanation is different from the explanation whereby one decomposes the cancer mechanism C in figure 2, and uncovers the role of damage G and response D. G and D have an obvious mechanistic role—insofar as they *constitute* C; instead, S and H seem to have a purely *causal* role.

The inadequacy of the MLCM notions of mechanistic decomposition and explanation is made more explicit by looking at the kind of hierarchical relations allowed by the formalism.

Consider the 'decompositions' in figure 7, which correspond to restricting (i) V_1 to V_2, (ii) V_1 to V_3, and (iii) V_3 to V_5. In all such cases, instead of opening a black box (as is common in mechanistic explanation), one 'creates' a box, and does not, strictly speaking, decompose anything. In (i), the decomposition is 'filling a blank': the absence of probabilistic and causal dependencies among variables is explained by direct causation, a hidden common cause structure, or combinations thereof that involve new variables, too. The absence of probabilistic and causal dependencies between X and Z in M_2 is explained by the structure $X \leftrightarrow Y \leftarrow Z$ in M_1 (more on this case of 'explanation' below). Since there is no arrow between X and Z in M_2, and since mechanisms require causal dependen-

cies, what mechanism is $X \leftrightarrow Y \leftarrow Z$ in M_1 a decomposition of? In (ii) and (iii), the decomposition is in fact 'adding stuff'. For instance, $Z \leftrightarrow W$ in M_5 is 'decomposed' into $Y \leftarrow Z \leftrightarrow W$ in M_3. But in what sense is a lower-level mechanism that includes an isolated effect not included in the higher level a decomposition of the higher level mechanism?

Relatedly, to some of the represented restrictions do not seem to correspond 'explanations' either. Consider the restriction of M_4 to M_5. Here, the common cause structure $Z \leftrightarrow W$ is 'explained' by the absence of probabilistic or causal dependence between Z and a new variable X, which is apparently disconnected from whatever mechanism is responsible for $Z \leftrightarrow W$. An even more striking case of lack of explanation is the 'decomposition' of X and Z in M_2 into $X \leftrightarrow Y \leftarrow Z$ in M_1. A first issue—arguably non-intentional (cf. Gebharter, 2014, 146, fn. 8)—is that the bidirected arrow in M_1 violates condition **c** of a MLCM, namely that M_1 satisfies CMC. Still, even if condition **c** were satisfied, the problem would remain that, if decompositions are to explain, this sort of decomposition should not be allowed at *any* level. Intuitively, hidden common cause structures such as $X \leftrightarrow Y$ are, insofar as they are *hidden*, non-explanatory. They add a mystery rather than remove it. A—drastic—solution that comes to mind is to forbid bidirected arrows at any level. This would entail, however, that restrictions that marginalize out common causes are disallowed, too, which is undesirable because—if one buys into the MLCM framework—the corresponding decompositions would seem (more) explanatory. One may of course impose further conditions to distinguish good from bad restrictions, but it is not obvious how one should proceed in a non *ad hoc* way, without clear intuitions on the explanatoriness of bidirected arrows.

In sum, the resulting account of mechanistic hierarchies is at best incomplete, and at worst inadequate. To prove RBNs' superiority, it remains to be shown

whether RBNs survive Gebharter (2014)'s and Gebharter and Kaiser (2014)'s objections. The next section endeavours to establish that they do.

4 Defense of RBNs

RBNs interpret mechanistic hierarchy via the operation of 'recursive decomposition', which in turn depends on RCMC. Two kinds of objections were raised against RCMC. First, about empirical adequacy: it is unclear when RCMC holds, so it is unclear if the formalism is applicable to real mechanisms. Second, about conceptual adequacy: RCMC prevents RBNs from being useful for interlevel reasoning for explanation and intervention. Let us begin with the first objection:

> it is neither obvious that RCMC holds in general, nor is it clear how one could distinguish cases in which it holds from cases in which it does not. (Gebharter and Kaiser, 2014, §3.5.3)

Agreed, RCMC may not hold in general. Nor did Casini et al. (2011) claim that it does. When *does* it hold, then? What RCMC adds to CMC—which is not called into question here—is RMC. RMC has to do with the (in)dependencies among variables at different levels. In the cancer example, RMC depends on the constitutional relations between C on the one hand, and G and R on the other, being such that C screens off G and D from S.

Gebharter and Kaiser (2014, §3.5.3) argue that the RBN approach would be unable to adequately model mechanistic decompositions, where there seems to exist no intermediate macro-level variable that corresponds to a micro-level structure. I do not dispute that there may be cases where it is hard or implausible to define network variables that decompose into lower-level causal structures. However, this is an empirical problem, and not necessarily a problem with the formal-

ism. It suffices to say that in many mechanisms, talk of network variables seems uncontroversial: the state of a tissue depends on the causal structure in the cells; it has a causal role with respect to survival; etc. Intuitively, network variables exist whenever there are functional states that are decomposable into structures such that RCMC holds.

In sum, I concede that RBNs are not universally applicable. Still, the conditions for their applicability are clearly spelled out. When such conditions are satisfied, RBNs provide adequate mechanistic decompositions. Notice that an analogous reply is not open to Gebharter. Whether MLCMs adequately represents hierarchies is not an empirical matter alone; it is also a conceptual matter, insofar as the restriction condition does not suffice to distinguish between legitimate and illegitimate marginalizations, that is, marginalizations that correspond to mechanistical decompositions as opposed to marginalizations that don't.[8]

Finally, let us come to the objection that RBNs do not support interlevel reasoning for explanation and for prediction of the results of interventions:

> [Casini et al.'s] approach does (i) not allow for a graphical representation of how a mechanism's macro variables are causally connected to the mechanism's causal micro structure, which is essential when it comes to causal explanation, and it (ii) leads to the fatal consequence that a mechanism's macro variables' values cannot be changed by any intervention on the mechanism's micro structure whatsoever (...) (Gebharter, 2014, 139)

Explanation first. Since there are no arrows between variable at different levels screened off by network variables, Gebharter claims that it is unclear over

[8] An analogous point applies to the explanatory power of RBNs (see below) vis-à-vis MLCMs (cf. §3).

which causal paths probabilistic influence propagates between such higher- and lower-level variables (2014, 143-4). True, there are no such arrows. But this is because, by assumption, screened off variables influence each other, if at all, only *via* the network variables. When RCMC is satisfied, the probabilistic influence propagates *constitutionally* (rather than causally) across the dashed arrows in the flattenings, and causally across the same-level solid arrows.

Let us now consider the second objection. With reference to the example in figures 5 and 6, I claimed that one may, for instance, reason about the result of a lower-level intervention on D on the probability of the higher-level variable S. Given the observed value of $P(s_1)$, calculated as

$$P(s_1) = P(c_0)P(s_1|c_0) + P(c_1)P(s_1|c_1),$$

one may ask: What is the effect of setting $D = d_1$ on the probability of observing $S = s_1$? To answer, one calculates as follows. First, one removes the arrow $G \rightarrow D$ from c_1, so that both flattenings have the same structure below.

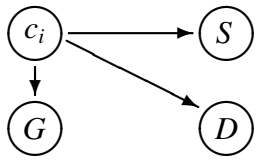

Figure 9

Then, one calculates $P(s_1||d_1) = P(s_1d_1)/P(d_1)$, where:

$$P(s_1d_1) = P(c_0s_1d_1) + P(c_1s_1d_1) = P(c_0)P(s_1|c_0)P_{c_0}(d_1) + P(c_1)P(s_1|c_1)P_{c_1}(d_1)$$

$$P(d_1) = P(c_0)P_{c_0}(d_1) + P(c_1)P_{c_1}(d_1)$$

Gebharter objects that 'according to the RBN approach, intervening on a mechanism's microvariables does not have any probabilistic influence on any one of the macrovariables whatsoever' (2014, 145) *because* if one were to use an intervention variable I to intervene on a lower-level variable, the intervention 'would—and this can directly be read off the BN's associated graph's topology (...)—not have any probabilistic influence on any macrovariable at all' (*ibid.*). In the cancer example: an intervention I_R on R would not have any effect on S. I think this objection is due to either of the following misinterpretations.

First, it is true that c_i screens off D from S, and thus there is no $D \rightarrow S$ causal arrow. However, concluding that interventions on D can make no difference to S would be wrong. The lack of *causal* connections in the flattening does not block changes along *constitutional* arrows. It is important to stress that, although the dashed arrows point downwards in the flattening, this is due to technical reasons only, having to do with the condition for MC to hold across levels. One may use the downward-pointing arrows to reason—constitutionally—in both directions. Here, changing D makes a constitutional difference to C, which makes a causal difference to S. The overall difference is calculated with the RBN.

Second, it is true that RCMC says that S is independent of any variable that is not an effect or an inferior (here, none), conditional on its direct causes (here, C) and direct superiors (here, none). But RCMC is assumed to hold true in $\mathcal{V} = \{M, S, G, D\}$, and *not* in the expanded set $\mathcal{V}^+ = \{M, S, G, D, I_D\}$. The reason for this is not *ad hoc*. RBNs are meant to represent decompositions of (properties of) wholes into (properties of) their parts. They are *not* meant to represent parts that do not belong to any whole—which is what I_D is. The graph topology cannot represent such parts. Thus, one can*not* read off the graph topology that such interventions variables have no effect.

More generally, in an RBN, everything one gets at lower levels must be the result of (recursively) decomposing the top level. This is not to be seen as a limitation of RBNs, but a means to an end. In the RBN formalism, one can*not* represent interventions as *variables*.[9] Yet, one *can* straightforwardly represent interventions as changes in *values* of either top-level variables or lower-level variables into which network variables (recursively) decompose. The two representations correspond to two well-known strategies for representing interventions, which are exemplified by respectively Woodward (2003)'interventionist semantics and Pearl (2000)'s *do*-calculus. Although both strategies are in principle legitimate, only the latter is suitable to the task for which RBNs were developed, viz. to represent mechanistic decompositions.

5 Conclusion

Decomposing variables and decomposing arrows are alternative ways of modelling mechanistic hierarchies by means of BNs. The two options have been made precise by, respectively, RBNs and MLCMs. I argued that RBNs are better than MLCMs at analysing mechanistic hierarchies and interpreting interlevel mechanistic reasoning. From a conceptual point of view, the argument establishes that the notion of mechanistic hierarchy has a tight connection to the notion of recursive decomposition, but no such connection to the notion of marginalizing out.

[9]Unless, of course, the variables describe properties of either top-level mechanisms or lower-level *sub*-mechanisms obtained by way of (recursive) decompositions—in which case, however, the intervention is not external to the mechanism, contrary to the original intention.

References

Baumgartner, M. and Gebharter, A. (2015). Constitutive Relevance, Mutual Manipulability, and Fat-handedness. *The British Journal for the Philosophy of Science*. doi: 10.1093/bjps/axv003.

Casini, L., Illari, P. M., Russo, F., and Williamson, J. (2011). Models for Prediction, Explanation and Control: Recursive Bayesian Networks. *THEORIA*, 70:5–33.

Clarke, B., Leuridan, B., and Williamson, J. (2014). Modeling Mechanisms with Causal Cycles. *Synthese*, 191:1651–1681.

Craver, C. and Bechtel, W. (2007). Top-down Causation Without Top-down Causes. *Biology and Philosophy*, 22:547–563.

Craver, C. F. (2007). *Explaining the Brain*. Oxford: Oxford University Press.

Gebharter, A. (2014). A Formal Framework for Representing Mechanisms? *Philosophy of Science*, 81(1):138–153.

Gebharter, A. and Kaiser, M. I. (2014). Causal Graphs and Biological Mechanisms. In Kaiser, M. I., Scholz, O., Plenge, D., and Hüttemann, A., editors, *Explanation in the Special Sciences: The Case of Biology and History*, pages 55–85. Dordrecht: Springer.

Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.

Williamson, J. (2005). *Bayesian Nets and Causality: Philosophical and Computational Foundations*. Oxford: Oxford University Press.

Williamson, J. (2010). *In Defence of Objective Bayesianism*. Oxford: Oxford University Press.

Woodward, J. (2003). *Making Things Happen. A Theory of Causal Explanation*. Oxford: Oxford University Press.