How to Model Mechanistic Hierarchies

Lorenzo Casini*†

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

1. Introduction. Mechanisms are usually viewed as hierarchical, with lower levels of a mechanism influencing, and decomposing, its higher-level behavior. To draw quantitative predictions from a model of a mechanism, the model must 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 higher- and lower-level variables 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

*To contact the author, please write to: 5, Rue de Candolle, Geneva CH-1211, Switzerland; e-mail: lorenzo.casini@unige.ch.

†I thank the Lake Geneva Biological Interest Group and the audience of the Philosophy of Science Association symposium in Chicago, November 6–8, 2014, where this article was presented. I am especially grateful to Michael Baumgartner, Alexander Gebharter, Guillaume Schlaepfer, and Jon Williamson. This work was supported by the Swiss National Science Foundation (grant CRSII 1_147685/1).

Philosophy of Science, 83 (December 2016) pp. 946–958. 0031-8248/2016/8305-0026$10.00
Copyright 2016 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/687877 Published online by Cambridge University Press

intervention). To overcome such limitations, Gebharter (2014) has made the alternative proposal that decomposition involves arrows rather than variables. In particular, he proposes an alternative formalism, also extending the BN formalism, namely, multilevel causal models (MLCMs).

Decomposing variables and decomposing arrows are two alternative ways of modeling mechanistic hierarchies probabilistically. Here, I argue that the former option is superior to the latter. I proceed as follows. In section 2, I present and illustrate RBNs and MLCMs. In section 3, I argue against decomposing arrows. MLCMs lead to counterintuitive notions of mechanistic decomposition and mechanistic explanation. In section 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 constitutional arrows or demanding that RBNs 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 directed acyclic graph (DAG), whose nodes are the variables in a finite set $V = \{V_1, \ldots, V_n\}$ (each variable taking finitely many possible values), and the probability distribution $P(V_i|Par_i)$ of each variable V_i conditional on its parents Par_i. DAG and the probability function are linked by the Markov Condition:

(**MC**) For any $V_i \in V = \{V_1, \ldots, V_n\}$, $V_i \perp\!\!\!\perp ND_i|Par_i$.

In words, each variable is probabilistically independent of its non-descendants, conditional on its parents. For instance, in the DAG in figure 1, V_4 is independent of V_1, and V_5 is 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 $P(v_1 \ldots v_n) = \prod_{i=1}^{n} P(v_i|par_i)$, 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 $v = v_1 \ldots v_n$. In a causally interpreted BN, the arrows in the DAG stand for direct causal relations, and the network can be used to infer

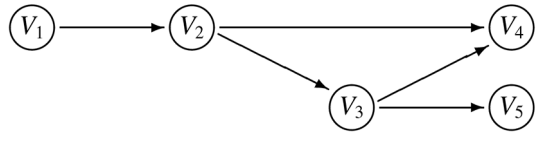


Figure 1. Example of a Bayesian network.

https://doi.org/10.1086/687877 Published online by Cambridge University Press

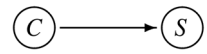


Figure 2. Top level of an RBN representing the relation between a network variable, cancer (C), and a simple variable, survival (S).

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] over $V = \{C, S\}$ with joint distribution P, where $C = \{1, 0\}$ represents whether an organism's tissue is cancerous, and $S = \{\text{yes, no}\}$ is survival after 5 years (fig. 2). Suppose S is a simple variable, but C is a network variable, with each of its two values denoting a

1. Craver (2007) proposed that constitution is established by the 'mutual manipulability' of higher- and lower-level properties. Casini et al. (2011) referred to Craver's intuition to further motivate RBNs. The compatibility between Craver's (2007) account and interventionism (Woodward 2003), on which Craver's account rests, was later questioned (see, e.g., Baumgartner and Gebharter 2016; Baumgartner and Casini 2017). While I agree that more work is required to use RBNs for constitutional discovery, I should emphasize that RBNs were proposed as tools for (among other things) representing constitutional relations based on background knowledge and not for establishing them by interventions.

2. For a more realistic example, see Casini et al. (2011).

https://doi.org/10.1086/687877 Published online by Cambridge University Press



Figure 3. Lower-level networks decomposing the binary network variable C in figure 2 into, respectively, a relation and the absence of a relation between growth factor (G) and division (D).

lower-level (standard) BN that represents a state of the mechanism for cancer. I 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 $C = 1$ corresponds a lower-level network c_1 (see fig. 3, left) with joint distribution P_{c_1} representing a functioning control mechanism, with a probabilistic dependence (and a causal connection) between G and D. And to $C = 0$ corresponds a lower-level network c_0 (fig. 3, right) with joint distribution P_{c_0} representing a malfunctioning growth mechanism, with no dependence (and no causal connection) between G and D. Since these two lower-level networks are standard BNs, the RBN is well founded and fully described by the above 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 $\mathcal{V} = \{V_1, \ldots, V_m\}$ ($m \geq n$) be the variable set of an RBN closed under the inferiority relation (i.e., $\mathcal{V}$ contains the variables in V, their direct inferiors, their direct inferiors, etc.). Let NID_i indicate the set of non-inferiors-or-descendants of V_i, and $DSup_i$ the set of direct superiors of V_i. Then,

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

In words, each variable 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. Applied to the toy example, given the value of C, the value of its constituents G and D is redundant for inferring to C's effect S.

RCMC adds to CMC a recursive MC (RMC), to the effect that variables at any level are probabilistically independent of non-inferiors or peers given their direct superiors. 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

https://doi.org/10.1086/687877 Published online by Cambridge University Press



Figure 4. Flattenings of the RBN represented in figures 2 and 3.

with respect to RCMC. RCMC is a modeling assumption in need of testing or justification, not a necessary truth. Thus, whether 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 $\mathcal{N} = \{V_{j_1}, \ldots, V_{j_k}\} \subseteq \mathcal{V}$ 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, such that in the RBN where V_{j_i} is the direct superior of V_i, $P(V_i|Par_i \cup DSup_i) = P_{v_{j_i}}(V_i|Par_i)$. 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 $v_1 \ldots v_m$.[3] Notice that MC holds in the flattening because RCMC holds in the RBN. Only, since the arrows that link variables to their direct inferiors are constitutional, CMC is not satisfied.[4]

In the cancer example, the flattening with respect to c_1 is $c_1^{\downarrow}$ (see fig. 4, left), where $P(c_1) = 1$ and $P(S|c_1)$ are determined by the top-level distribution P and where $P(G|c_1) = P_{c_1}(G)$ and $P(D|Gc_1) = P_{c_1}(D|G)$ are determined by the lower-level distribution P_{c_1}. Analogously, the flattening with respect to c_0 is $c_0^{\downarrow}$ (see fig. 4, right), where $P(c_0) = 1$ and $P(S|c_0)$ are determined by the top-level distribution P and where $P(G|c_0) = P_{c_0}(G)$ and $P(D|Gc_0) = P_{c_0}(D|G)$ are determined by the lower-level distribution P_{c_0}.

In each case, the required probabilities are determined by the original RBN. Given the joint distribution, the causally interpreted RBN may be

3. The probabilities $P_{v_{j_i}}(V_i|Par_i)$ may be obtained from observed frequencies in a data set, whereas $P(V_i|Par_i DSup_i)$ can be obtained by either determining the corresponding observed frequencies from the original data set or first selecting from all functions that satisfy the probabilistic constraints imposed by the RBN the function Q with maximum entropy (Williamson 2010) and then setting $P(V_i|Par_i DSup_i) = Q(V_i|Par_i DSup_i)$.

4. Whether an arrow in the flattening is causal or constitutional cannot be read off the flattening. Rather, it depends on background knowledge; cf. footnote 1.

https://doi.org/10.1086/687877 Published online by Cambridge University Press

used to draw quantitative inferences for explanation and intervention, both within and across levels.

2.2. Multilevel Causal Models. Differently from RBNs, MLCMs decompose arrows rather than variables.[5] 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 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, $\Sigma_{i=1}^{n} P(Y|Par(Y)_{X=x_i})$, thereby marginalizing X out. As a result, one gets a truncated distribution over $V \setminus \{X\}$, consistent with the original one over V.

Let us indicate a causal model as $\langle V, E, P\rangle$, where V and E define a DAG over a variable set V and a set of edges E, and P is the probability distribution associated to the DAG. 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 (Gebharter 2014, 147):

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

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

5. Gebharter and Kaiser (2014, sec. 3.6) propose that levels be ontologically distinct (partly) on the basis of constitutional relations between whole's and parts' properties. Since it is unclear how, if at all, this proposal relates to MLCMs, I will not discuss it here.

https://doi.org/10.1086/687877 Published online by Cambridge University Press

$P^*\rangle$ 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 'multilevel causal model' is so defined (Gebharter 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.

That is, an 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' (Gebharter 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 an 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 restriction of M_j.

Figure 5 represents a level graph. 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 6 depicts a more concrete example, that is, a two-level water dispenser mechanism.[6] The room temperature T influences a sensor S; S and the status of a tempering button, B, cause the heater H to be on/off; H causes the temperature of the water dispensed, W.

6. Gebharter contrasts the virtues of the MLCM in fig. 6 with an RBN of the 'same' mechanism (2014, 142–43). This is misleading. In a nutshell, since the RBN contains constitutional arrows, the two models cannot possibly represent the same mechanism. This motivates my choice of defending RBNs by reference to the toy model in sec. 2.1.

https://doi.org/10.1086/687877 Published online by Cambridge University Press

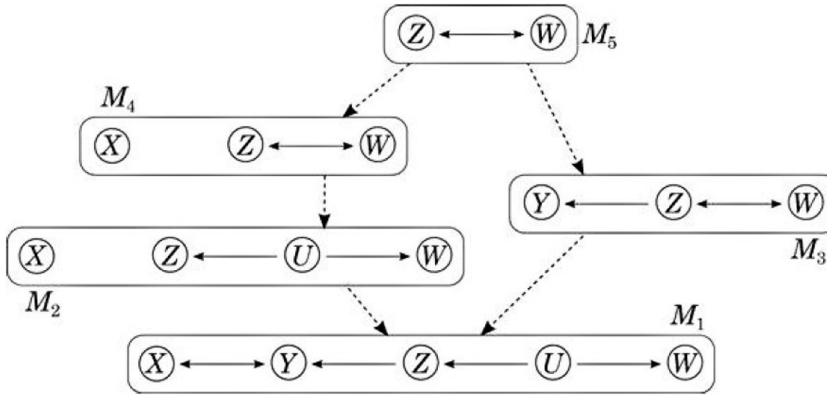


Figure 5. A level graph (reprinted from Gebharter 2014, 150).

3. Criticism of MLCMs. It is unclear whether hierarchies, as analyzed in terms of the notion of 'marginalizing out', are mechanistic—that is, whether they represent mechanistic decompositions and grant mechanistic explanations. First, it is unclear whether MLCMs represent mechanistic decompositions. High-level causal models in an MLCM, for instance, M_2 and M_3 in figure 5, are just more coarse-grain representations of one and the same structure, that is, 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

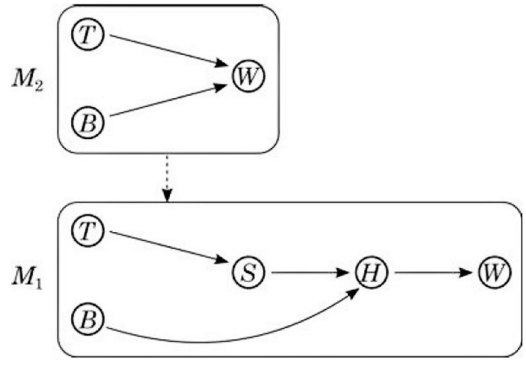


Figure 6. Dispenser mechanism (reprinted from Gebharter 2014, 151).

https://doi.org/10.1086/687877 Published online by Cambridge University Press

of explanation is different from the explanation whereby one decomposes the cancer mechanism C and uncovers the role of damage G and response D. Variables 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 5, 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 dependencies, what mechanism is $X \leftrightarrow Y \leftarrow Z$ in M_1 a decomposition of? In (ii) and (iii), in contrast, 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, 'explanations' do not seem to correspond to some of the represented restrictions 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 unintentional (cf. Gebharter 2014, 146 n. 8)—is that the bidirected arrow in M_1 violates condition **c** of an 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 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.

https://doi.org/10.1086/687877 Published online by Cambridge University Press

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's (2014) and Gebharter and Kaiser's (2014) 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 and thus whether 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, sec. 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? Intuitively, RCMC holds when higher-level differences in some functional property, or phenomenon, depend on differences in its underlying structure, or mechanism, such that the state of the phenomenon makes the states of its constituents in the underlying mechanism redundant with respect to (among other things) the phenomenon's causes or effects. Not all higher-level phenomena are so dependent on structures and thus representable by network variables. Thus, RBNs may incur a problem of too limited applicability, which is an empirical matter. On the face of it, many biological phenomena seem representable by means of network variables. For instance, it seems appropriate to represent the different effects of a tissue on survival as dependent on differences in the tissue's underlying cellular structure. In contrast, if my argument in section 3 is correct, MLCMs appear conceptually inadequate—marginalizations may satisfy the restriction condition and yet not correspond to mechanistic decompositions.

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 which causal paths probabilistic influence propagates between such

https://doi.org/10.1086/687877 Published online by Cambridge University Press

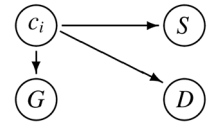


Figure 7. Flattening representing the structure assumed by the flattenings in figure 4 after an intervention on D.

higher- and lower-level variables (2014, 143–44). True, there are no such arrows. But this is because, by assumption, screened-off variables influence each other, if at all, only via network variables. When RCMC is satisfied, probabilistic influence propagates constitutionally (rather than causally) across the flattening's dashed arrows and causally across same-level solid arrows.

Let us now show how the second objection is ill founded, with reference to the difference in the toy example in section 2.1 between the unconditional probability of $S = s_1$ and the probability of $S = s_1$ conditional on a 'do' intervention (Pearl 2000) that sets $D = d_1$. The former equals to $P(c_0)\,P(s_1 \mid c_0) + P(c_1)\,P(s_1 \mid c_1)$. The latter is obtained by first removing the arrow $G \rightarrow D$ from c_1, so that both flattenings have the same structure (see fig. 7) and then calculating $P(s_1 \mid \mathrm{do}(D = d_1)) = P(s_1 d_1) / P(d_1)$, where

$$P(s_1 d_1) = P(c_0 s_1 d_1) + P(c_1 s_1 d_1) = P(c_0)P(s_1 \mid c_0)P_{c_0}(d_1) + P(c_1)P(s_1 \mid c_1)P_{c_1}(d_1)$$
$$P(d_1) = P(c_0)P_{c_0}(d_1) + P(c_1)P_{c_1}(d_1).$$

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" (145). 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, interventions on D can still make a difference to S, as 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 downward in the flattening, this is because of 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.

https://doi.org/10.1086/687877 Published online by Cambridge University Press

Second, it is true that RCMC says that S is independent of any noninferior or descendent (here, none), conditional on its direct causes (here, C) and direct superiors (here, none). But RCMC is assumed to hold in $\mathcal{V} = \{M, S, G, D\}$ and not in the expanded set $\mathcal{V}^+ = \{M, S, G, D, I_D\}$. The reason is that 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, such as I_D. The graph topology cannot represent such parts. Thus, one cannot read off the graph topology that such intervention 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 a limitation of RBNs but a means to an end. One cannot represent interventions as variables.[7] Yet, one can represent interventions as operators, which change the 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, exemplified by respectively Woodward's (2003) interventionist semantics and Pearl's (2000) do-calculus. Although both strategies are in principle legitimate, only the latter is relevant to the task for which RBNs were developed, that is, to represent mechanistic decompositions.

5. Conclusion. Decomposing variables and decomposing arrows are alternative ways of modeling 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 analyzing 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.

REFERENCES

Baumgartner, M., and L. Casini. 2017. "An Abductive Theory of Constitution." *Philosophy of Science* 84 (2), forthcoming.

Baumgartner, M., and A. Gebharter. 2016. "Constitutive Relevance, Mutual Manipulability, and Fat-Handedness." *British Journal for the Philosophy of Science* 67 (3): 731–56.

Casini, L., P. M. Illari, F. Russo, and J. Williamson. 2011. "Models for Prediction, Explanation and Control: Recursive Bayesian Networks." *Theoria* 26 (70): 5–33.

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

7. Unless, of course, the variables describe properties of either top-level mechanisms or lower-level submechanisms obtained by (recursive) decomposition—in which case, however, the intervention is not external to the mechanism, contrary to the original intention.

https://doi.org/10.1086/687877 Published online by Cambridge University Press

Craver, C. F., and W. Bechtel. 2007. "Top-Down Causation without Top-Down Causes." *Biology and Philosophy* 22 (4): 547–63.
Gebharter, A. 2014. "A Formal Framework for Representing Mechanisms?" *Philosophy of Science* 81 (1): 138–53.
Gebharter, A., and M. I. Kaiser. 2014. "Causal Graphs and Biological Mechanisms." In *Explanation in the Special Sciences: The Case of Biology and History*, ed. M. I. Kaiser, O. Scholz, D. Plenge, and A. Hüttemann, 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.
———. 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.

https://doi.org/10.1086/687877 Published online by Cambridge University Press