

How to Model Mechanistic Hierarchies*

Lorenzo Casini[†]

November 25, 2014

please do not quote or cite without permission

Abstract

Mechanisms are usually viewed as inherently 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*. The proposal was recently criticized by Gebharter (2014) and Gebharter and Kaiser (2014), who instead propose to decompose *arrows*. In this paper, I defend the RBN account from the criticism and argue that it offers a better representation of mechanistic hierarchies than the rival account.

Introduction

Mechanisms are usually viewed as inherently 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 is an extension of 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.

This proposal was recently 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 these alleged limitations, Gebharter (2014) and Gebharter

*Paper presented at *PSA 2014*, Chicago, 6–8 Nov 2014, in the symposium “How Adequate Are Causal Graphs and Bayesian Networks for Modeling Biological Mechanisms?”


[†]Address: Department of Philosophy, University of Geneva, 5, Rue de Candolle, CH-1211 Genève 4, Switzerland. Email: lorenzo.casini@unige.ch

and Kaiser (2014) have 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). Instead, Gebharter and Kaiser (2014) make an informal proposal, which as we shall see, does not coincide with MLCMs.

Decomposing variables and decomposing arrows are two very natural options for representing mechanistic hierarchies, if one’s starting point is already a probabilistic interpretation of causality. In this paper, I argue that the former option is superior to the latter. I proceed as follows. In §1 I present and illustrate RBNs and MLCMs. In §2 I argue against decomposing arrows. MLCMs lead to counterintuitive notions of mechanistic decomposition and mechanistic explanation; and Gebharter and Kaiser (2014)’s informal proposal goes only halfway towards a solution. Finally, in §3 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.

1 The two formalisms

Both RBNs and MLCMs are extensions of the BN formalism. A BN consists of a finite set $V = \{V_1, \dots, V_n\}$ 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 conditional on its parents Par_i in the DAG. Here is an example:



DAG and probability function are linked by the *Markov Condition* (MC):

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

In words, each variable is probabilistically independent of its non-descendants, conditional on its parents. The above figure implies for instance that V_4 is independent of V_1 and V_5 conditional on V_2 and V_3 . In the 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 are interpreted as direct causal relations and the network can be used to infer the effects of interventions as well as to make probabilistic predictions (Pearl, 2000). In this case, MC is called the *Causal Markov Condition* (CMC).

1.1 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—which was 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, thereby adding an additional justification to it.

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*.² 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.

¹ Famously, Craver (2007) has proposed a criterion for identifying constitutive relations, namely the ‘mutual manipulability’ of higher- and lower-level properties that stand in the relation. Casini et al. (2011) refer to Craver’s intuition to further motivate RBNs. Arguments against the compatibility between Craver (2007)’s account of constitution and interventionism (Woodward, 2003), on which Craver’s account is based, have been offered by Leuridan (2012) and Baumgartner and Gebharter (2014). Two remarks are in order. First: in the light of Gebharter and Kaiser (2014, 3.5.3)’s own endorsement of Craver (2007)’s interpretation of constitution, these arguments may be negatively relevant to *both* RBNs and Gebharter and Kaiser (2014)’s proposal. Although this issue is certainly worth considering, I do not discuss it further here. I should however point out that RBNs do not *define* constitution. They only *characterize* it, probabilistically—and *not* even in interventionist terms (cf. fn. 4). Interventions are only used to reason about interlevel causation. Second: Gebharter (2014)’s MLCM formalism does *not* interpret hierarchy in terms of constitution—let alone constitution in one specific sense. It is thus immune to this critique. However, instead of being an advantage, this threatens to undermine MLCMs’ ability to represent *mechanistic* hierarchies (see §2).

²A variable can have several superiors. If a variable appears more than once in an RBN, the network should not imply incompatible things about it. Consistency is discussed in detail in (Williamson, 2005, §§10.4–10.5).

Consider a toy RBN on $V = \{C, S\}$, where C represents whether some tissue in an organism is cancerous, taking the possible values 1 and 0, while S is survival after 5 years, taking the possible values *yes* and *no*. The corresponding BN is:




Figure 1.1.1

Suppose S is a simple variable but C is a network variable, with each 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 that results from mutations in factors that control cell division, usually labelled ‘growth factor’, in short GF. When C is assigned value 1 we have a network c_1 representing a functioning control mechanism, with a probabilistic dependence (and a causal connection) between growth factor G and cell division D .




Figure 1.1.2

On the other hand, when C is assigned value 0 we have a network c_0 representing a malfunction of the growth mechanism, with no dependence (and no causal connection) between G and D .




Figure 1.1.3

Since these two lower-level networks are standard BNs, the RBN is well-founded and fully described by the three networks.³

³Note that, as this example shows, an RBN may be used to represent several states of one and the same mechanism—in this case, the RBN represents a functioning state as well as a malfunctioning state. However, it need not be so used—it is also possible to build an RBN that represents just one mechanism state by having the network variable take a unique possible value.

If an RBN is to be used to model a mechanism, it is natural to interpret the arrows at the various levels of the RBN as signifying causal connections. 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* (RMC). 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 $V_i \in \mathcal{V} = \{V_1, \dots, V_n\}$, $V_i \perp\!\!\!\perp NID_i \mid DSup_i \cup Par_i$.

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 the condition that variables at different levels also stand in relations that fulfil a MC, namely variables at any level are probabilistically independent of non-inferiors or peers given their direct superiors. Intuitively, if one knows the value of C , knowledge of 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 to the effects of C (here, S). 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.

Notice that, while some authors treat CMC as a necessary truth, others argue against its universal validity (see, e.g., Williamson, 2005). Here a similar stance is adopted with respect to RCMC. RCMC is a *modelling assumption* in need of testing or justification, rather than as a necessary truth. From this, it follows that whether or not the formalism allows one to adequately represent a mechanism is an empirical matter, rather than a matter of stipulation. For instance, whether or not C adequately screens off S from G and D depends, among other things, on the assumption that G and D affect S only via C . If this is not true, because S or G participate in other mechanisms for S , RCMC is violated. Recovering RCMC would then require including other network variables that cause S , and that decompose into, among other variables, G and/or D .

Inference in RBNs proceeds via a formal device called a *flattening*. Let $\mathcal{V} = \{V_1, \dots, V_m\}$ ($m \geq n$) be the set of variables of an RBN closed under the inferiority relation: i.e., \mathcal{V} contains the variables in V , their direct inferiors, their direct inferiors, and so on. Let $\mathcal{N} = \{V_{j_1}, \dots, V_{j_k}\} \subseteq \mathcal{V}$ be the network variables in \mathcal{V} . For each assignment $n = v_{j_1}, \dots, v_{j_k}$ 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}, \dots, 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.⁴

The flattenings suffice to determine a joint distribution over the variables in \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 $v_1 \cdots v_m$.⁵ In the cancer example, for assignment c_1 of network variable C we have the flattening c_1^\downarrow :




Figure 1.1.4

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 1.1.5

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, in just the same way as can a standard causal BN, be used to draw quantitative inferences for explanation and intervention, inferences that may involve variables at the same level as well as—so we claimed in (Casini et al., 2011, §2)—*across* levels.

⁴ It should now be clear that the role of RCMC—and of RBNs more generally (see fn. 1)—is not to define constitutional relations. With respect to the flattening, the choice of calling some arrows ‘causal’ and other arrows ‘constitutional’ is not dictated by MC. Any use of RCMC to find out what does (not) constitute what *presupposes* a prior distinction between the variables at the different levels. Yet, given the distinction between the levels, RCMC does characterize constitutional relations in terms of certain probabilistic dependencies and independencies.

⁵ $P_{v_j}(V_i | Par_i)$ may be obtained from observed frequencies in a dataset. Instead, $P(V_i | Par_i DSup_i)$ can be obtained in either of two main ways. Either one determines the corresponding observed frequencies from the original dataset, or one selects from all functions that satisfy the probabilistic constraints imposed by the RBN the function Q with maximum entropy (Williamson, 2010), and sets $P(V_i | Par_i DSup_i) = Q(V_i | Par_i DSup_i)$.

1.2 Multilevel causal models

According to Gebharter (2014), RBNs fail to allow interlevel causal inferences, due to the lack of an explicit representation of interlevel causal arrows, over which causal influence propagates. These objections, I maintain, are based on the (mis)interpretation of RBNs. I postpone this discussion to §3.

Gebharter’s proposed formalism purports to remedy these alleged deficiencies by decomposing causal *arrows* rather than variables. More precisely, mechanistic hierarchy has for him to do with ‘marginalizing out’ variables when moving from a lower-level graph to a higher-level graph.

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. 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 . Also, let us indicate with $P^* \uparrow V$ the ‘restriction’ of the probability distribution P^* to 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

- 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 mechanism for 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 all and only the directed paths and common cause paths in $\langle V^*, E^* \rangle$ are 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 then so defined (2014, 148):

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

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

The information on the hierarchical relations among the nested mechanisms in the MLCM is contained in a “level graph”, which is so defined (2014, 149):

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

- a $V = \{M_1, \dots, M_n\}$, 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 restriction of M_j . Here is an example of level graph from (Gebharter, 2014, 150):

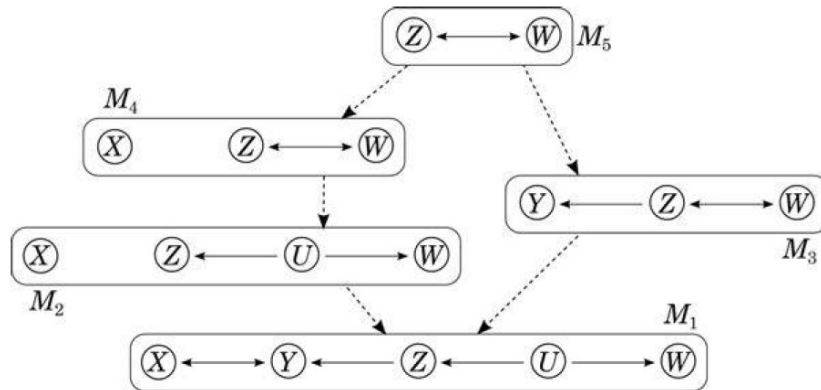


Figure 1.2.1

Notice that the ordering among graphs is not strict, so there may be pairs of graphs (e.g.: M_2 and M_3 ; M_4 and M_3) that do not stand in a restriction relation.

Below is a more concrete illustration from (Gebharter, 2014, 151), the representation of a water dispenser mechanism, on two levels,




Figure 1.2.2

such that M_1 contains the following direct causal relations: the room temperature T activates (and is measured by) a sensor S ; S , together with the status of a tempering button, B , cause the heater to be on or off, H ; H in turn causes the temperature of the water dispensed, W .⁶

2 Criticism of MLCMs

It is debatable whether hierarchies, as represented by the level graphs in figures 1.2.1 and 1.2.2, are *mechanistic*—whether they represent mechanistic decompositions, and grant mechanistic explanations.

First, it is not clear if MLCMs adequately represent mechanistic decompositions. High-level causal models in a MLCM, for instance models M_2 in figure 1.2.1, are just more coarse-grain representations of one and the same mechanism, viz. M_1 , such that some of the information in M_1 is *missing* at the higher level, as the term ‘restriction’ suggests. Is, for instance, $T \rightarrow S \rightarrow H \rightarrow W$ a mechanistic decomposition of $T \rightarrow W$, although *entities* and *properties* involved are the same at both levels, and only some *activities* (or relations) are different? Perhaps this counts as a different, equally legitimate, notion of decomposition, call it *decomposition**. The question is: How intuitive is *decomposition**?

Second, it is not clear if MLCMs adequately represent mechanistic explanations. One may concede that there is a legitimate sense in which one explains

⁶Gebharter contrasts the virtues of this MLCM with an RBN of the ‘same’ mechanism (2014, 142-3). However, this is somewhat misleading. Gebharter’s RBN is defined over a larger variable set, which includes a network binary variable D , superior to S and H , caused by T and B , and causing W . It is obvious that his RBN cannot represent the *same* mechanism as his MLCM. On the assumption that the RBN is faithful, it should be possible to order the RBN’s flattening (Gebharter, 2014, 144), call it M_0 , as *prior* with respect to M_1 —since M_1 ’s variable set V_1 is $\{V_0 \setminus D\}$. However, M_1 is incompatible with the restriction of M_0 obtained by marginalizing out D , call this M_{1^*} . (M_{1^*} would contain $S \leftrightarrow H, S \leftrightarrow W, H \leftrightarrow W$ and $B \rightarrow S$. Instead, M_1 contains $S \rightarrow H, H \rightarrow W$ and $B \rightarrow H$.) Thus, rather than one model being a correct representation and the other being a wrong representation of one and the same mechanism, the two models represent different mechanisms, and are thus not directly comparable. In the following, I shall defend RBNs with reference to the toy model introduced in §1.1.

the relation between, say, the room temperature T and the water temperature W by blowing up the process from the former to the latter and uncovering the mediating role of the sensor S and the heater H . However, this sort of explanation is different from the equally legitimate explanation whereby one redescribes the cancer mechanism C in figure 1.1.1 into more fine-grain terms, and uncovers the role of damage G and response D . G and D have an obvious mechanistic role. Instead, S and H seem to have an *etiological* role. Perhaps S and H still explain mechanistically, according to some different notion of mechanistic explanation, call it *explanation**. But just how intuitive is *explanation**?

The counterintuitive nature of *decomposition** and *explanation** is made more explicit by a careful scrutiny of the level graph in figure 1.2.1. To begin with, consider the ‘decompositions’ that 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. Let us consider (i). Here 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 alleged 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? Next, consider cases (ii) and (iii). Here 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 and more obvious issue, which is clearly non-intentional, is that the presence of a bidirected arrow in M_1 violates condition **c** of a MLCM, namely that M_1 satisfies CMC.⁷ Still, even if condition **c** is satisfied, the more general problem remains 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 just that, *hidden*, and thus non-explanatory. They add a mystery rather than remove it. A—drastic—solution that immediately 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

⁷Gebharder himself emphasizes that “the graph of a causal model that contains bidirected arrows no longer determines the Markov factorization [...]” (2014, 146, fn. 8).

undesirable because—if one buys into the MLCM framework—the corresponding decompositions would seem (more) explanatory. One may of course impose further conditions that distinguish good from bad restrictions. However, it is not obvious how one should proceed in a non *ad hoc* way, in the absence of clear intuitions on the explanatory value of bidirected causal arrows.

The above reasons lead to scepticism about the formalism’s capacity to represent mechanistic decompositions and explanations. Such worries are in part, but not fully, mitigated by the (orthogonal) suggestion in (Gebharter and Kaiser, 2014) that levels be ontologically distinct and the requirement that hierarchical relations are (partly) defined by constitutional part-whole relations.

In our approach one can generate a hierarchic causal model by replacing such a causal arrow [between two variables X and Y] by another causal structure. This causal structure should be on a lower ontological level than X and Y , it should contain at least one constitutively relevant part of X and at least one of Y , and there should be at least one causal path going from the former to the latter at the micro- level. (Gebharter and Kaiser, 2014, §3.6)

In the paper, Gebharter and Kaiser focus on modelling this sort of hierarchical relation with reference to the inhibitory feedback mechanism for the regulation of the biosynthesis of fatty acids in *Brassica napus*. The mechanism is represented as follows (Gebharter and Kaiser, 2014, figure 3.11).




Figure 2

The product of a reaction pathway, in this case the 18:1-acyl carrier protein (P) acts as a feedback signal, which inhibits an enzyme earlier in the pathway, in this case the plastidic acetyl-CoA carboxylase (ACCase), whose operation promotes the production of P itself via the transformation of the substrate acetyl-CoA (S). ACCase has two relevant properties: it is a (positive) cause of the concentration of P (E_{active}); and it is (in its P -bound state) an effect of the concentration of P ($E_{P-bound}$). $E_{P-bound}$ is in turn a (negative) cause of E_{active} (because P -bound ACCase becomes inactive) and so on and so forth, in a cycle. In addition, E_{active}

is also a negative cause on $E_{P-bound}$, which is represented by a negative influence on S . Between the binding of P to E and the inactivation of E a lower-level mechanism takes place, namely the conformational change of the substrate binding site. The binding B between functional groups of 18:1-acyl and the effector interaction site of the enzyme causes an allosteric switch X , which in turn brings about changes at sites $S2$ and $S4$ of the enzyme ACCase. This, then, prevents the substrate from being able to bind to the enzyme.⁸

It is now demanded that the levels be ontologically distinct, partly by way of decomposing properties, rather than just the relation $E_{P-bound} \rightarrow E_{active}$, as follows: B is a property of a part contained in the whole that has the property $E_{P-bound}$; and $S4$ and $S2$ are properties of parts contained in the whole that has the property E_{active} . Between parts and wholes there are relations of constitutive rather than causal relevance, in the sense of Craver (2007): a change in a part results in a change in the whole, and vice versa. More precisely, constitutive relations are represented by dashed two-headed arrows that stand at either side of the decomposition relation. As a result, decomposing arrows should apparently explain both causally *and* constitutionally.

Gebharter and Kaiser require that a causal arrow $X \rightarrow Y$ is decomposed by a lower-level causal structure only if it contains at least one constitutively relevant part of X and at least one of Y , and there is at least one causal path going from the former to the latter at the microlevel (2014). This eliminates two counterintuitive features of MLCMs, namely that mechanistic decompositions may ‘fill blanks’ (there must be a higher-level relation to begin with) and ‘add stuff’ (there must be at least one lower-level causal path). Still, two questions arise, related to the interpretation of the dashed bidirected arrows.

First, is this interpretation of mechanistic hierarchy *compatible with* MLCMs? As Gebharter and Kaiser notice, “since the two-headed dashed arrows in our hierarchic dynamic CM transport the influences of interventions in both directions, CMC does not hold in such models”. Since M_1 would contain bidirected arrows, too, it would not satisfy CMC. This entails that the Brassica napus mechanism cannot be represented by the MLCM formalism, as it currently stands.⁹

⁸ To get a causal model, Gebharter and Kaiser propose that the causal graph in figure 1.2 be associated with a probability distribution over a variable set that unrolls the cycle, so as to get a *dynamic* causal graph. This way of treating cycles is similar to the one adopted in the RBN approach (Clarke et al., 2014), with the notable difference that MC is not satisfied here (see below).

⁹Gebharter and Schurz (2014) are now extending MLCMs to account for causal cycles—in analogy with how RBNs have been so extended (cf. fn. 8). In the modified MLCM formalism, CMC need not be satisfied at the bottom level, and bidirected *causal* arrows that stand on cycles can figure in it. There is an apparent problem with this move: since bidirected arrows may in principle represent *constitutional* relations, too (see Brassica Napus), there seems to be nothing that formally distinguishes cycles from hierarchies. This has to do with constitution coming with no distinctive formal properties in MLCMs (see below).

Second, does the causal model in (Gebharter and Kaiser, 2014, §3.5) offer an adequate formal representation of a mechanistic hierarchy, *alternative to* MCLMs? I think that a positive answer would require that constitutional relations be ascribed distinctive *formal* properties. Although constitutional relations are characterized informally by part-whole relations, they don't come with distinctive probabilistic features, as one would expect from a probabilistic representation of mechanistic hierarchies. In contrast, RBNs do offer a probabilistic characterization of constitution: properties at different levels that stand in a constitutional relation relate to other properties as described by RCMC.

3 Defense of RBNs

Still, the shortcomings of MLCMs would be a small consolation for the RBN advocate, if RBNs did not survive the objections raised by Gebharter (2014) and Gebharter and Kaiser (2014). In this section I will consider, and try to rebut, such objections one by one. RBNs interpret mechanistic hierarchy via the operation of 'recursive decomposition', which in turn depends on RCMC. Two kinds of objections are 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. But Casini et al. (2011) don't 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 C screening off G and D from S .

Gebharter and Kaiser then argue that the RBN approach would be unable to adequately model the $E_{P-bound} \rightarrow E_{active}$ mechanistic decomposition:

it is not clear how the submechanism represented by $E_{P-bound} \rightarrow E_{active}$ could be analyzed in Casini et al.'s (2011) approach. They would need to add a network variable N between $E_{P-bound} \rightarrow E_{active}$ ($E_{P-bound} \rightarrow N \rightarrow E_{active}$). But then and because there is no intermediate (macro-level) cause N between $E_{P-bound}$ and E_{active} , it is unclear what this network variable N should represent at the mechanism's macro-level. (Gebharter and Kaiser, 2014, §3.5.3)

I do not dispute that there may be cases where it is hard or implausible to find network variables that stand for lower-level causal structures. However, this is an empirical problem, and not necessarily a problem with the formalism. RBNs are meant to represent a natural decomposition strategy of functional properties into structural properties. The structural properties may be then regarded as functional with respect to other structural properties, and so on and so forth. When does a network variable N exist? This depends on identifying properties at different levels, which in turn depends on a meaningful distinction between the levels.

I propose a few conditions for distinguishing between variables in a constitutional relation.¹⁰ First, between the whole and its parts are mereological relations, such that properties of the whole can be explained by their probabilistic dependence on the structure of (causal relations among) its parts' properties. Second, properties at the different levels have different explanatory roles, such that they typically enter into causal explanations involving different sets of properties. Third, there is a difference in epistemic conditions, such that the way one observes, or intervenes on, a variable at some higher level does not coincide with the the way one observes, and intervenes on, one of its constituting variables at the lower level.¹¹ When a distinction between variables informed by the above conditions is possible, the distinction between the levels seems legitimate.¹²

A network variable N exists insofar as the lower-level BN is the decomposition of one functional property, which, according to the aforementioned criteria, corresponds to a *whole's* property that has its own *explanatory role* and *epistemic autonomy*. These conditions seem satisfied by many descriptions of mechanisms in science. For instance, tissues are made of cells. Scientists talk of the cancerous state of a tissue as having an explanatory role with respect to survival. One may observe the state of a tissue or change it, for instance by replacing the whole tissue. One may use this knowledge to then infer to the probability of survival. This does not require knowing, or (surgically) intervening on, the state of GF.¹³

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

¹⁰I don't claim that the list is exhaustive or that each of the listed conditions is necessary.

¹¹Baumgartner and Gebharder (2014) develop this intuition into a 'fat-handedness' criterion for constitution. (Ironically, there an argument is proposed to defend an interpretation of mechanistic hierarchy based on decomposing variables rather than arrows.)

¹²The conditions only provide a useful heuristics. They do not belong to the RBN formalism. Still, RBNs give a probabilistic characterization of constitution, thanks to RCMC (cf. fn. 4).

¹³For more realistic examples, see (Casini et al., 2011), (Clarke et al., 2014) and (Casini, 2014).

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 higher- and lower-level variables (cf. 2014, 143-4). I reply that it is 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. So, 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 1.1.4 and 1.1.5, 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 3.1

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 . There is either one of the following problems with this objection.

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. Still, 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, there may be a more basic interpretive problem regarding how interventions are represented in RBNs. True, 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 notice that RCMC is assumed to hold true of variables in $\mathcal{V} = \{M, S, G, D\}$, and *not* of such an expanded $\mathcal{V}^+ = \{M, S, G, D, I_D\}$. 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. As a result, one *cannot* 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 should not be seen as a limitation, but as a means to achieve some end. In the RBN formalism one cannot represent interventions *as variables*—unless the variables describe properties of either the top level mechanism or of *submechanisms* at some lower level, obtained by way of (recursive) decompositions. But this would mean that the intervention is *not external* to the mechanism, contrary to the original intention. One can, instead, straightforwardly represent interventions *as (new) values* of either top-level variables or lower-level variables into which network variables (recursively) decompose. The two ways correspond to two well-known ways for representing interventions. Woodward (2003)’interventionist semantics, which represents interventions as variables, is an example of the former. Pearl (2000)’s *do*-calculus, which represents interventions as values of variables, is an example of the latter. Although both representations are legitimate, only the latter is suitable to the task for which RBNs were developed, namely to represent mechanistic decompositions.

Conclusion

Decomposing variables and decomposing arrows are two very natural options for representing mechanistic hierarchies by means of BNs. These two options have been made precise by two formalisms, RBNs and MLCMs. I argued that RBNs are better than MLCMs at analysing mechanistic hierarchies and interpreting the interlevel reasoning that depends on them. Still, one might think that

the two formalisms are not in competition against one another. Perhaps RBNs and MLCMs represent two different ways in which mechanistic decompositions can obtain? Since ‘marginalizing out’ and ‘recursively decomposing’ are very different notions, I want to caution against interpreting the two formalisms as two species of the same genus. Having said this, I do not exclude that there is a sound way to formalize the intuition in (Gebharter and Kaiser, 2014), and thus develop an alternative analysis of mechanistic hierarchy with respect to RBNs. In that case, it would be interesting to see how this alternative relates to RBNs.

Acknowledgments I thank the participants to the *Lake Geneva Biological Interest Group* and the *Philosophy of Science Association*, Chicago 6–8 November 2014, where this paper was presented. I am particularly grateful to Michael Baumgartner, Alexander Gebharter, Guillaume Schlaepfer and Jon Williamson for very helpful discussions.

References

- Baumgartner, M. and Gebharter, A. (2014). Constitutive Relevance, Mutual Manipulability, and Fat-handedness. *British Journal for the Philosophy of Science*. forthcoming.
- Casini, L. (2014). Failures of Modularity and Recursive Bayesian Networks. Unpublished.
- 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.
- Gebharter, A. and Schurz, G. (2014). A Modeling Approach for Mechanisms Featuring Causal Cycles. Paper presented at *PSA 2014*, Chicago 6–8 November 2014.
- Leuridan, B. (2012). Three problems for the mutual manipulability account of constitutive relevance in mechanisms. *British Journal for the Philosophy of Science*, 63:399–427.
- Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge.
- 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.