

Models for Prediction, Explanation and Control: Recursive Bayesian Networks

Lorenzo CASINI, Phyllis McKay ILLARI, Federica RUSSO and Jon WILLIAMSON

Received: 18.06.2010

Final Version: 17.01.2011

BIBLID [0495-4548 (2011) 26: 70; pp. 5-33]

ABSTRACT: The Recursive Bayesian Net (RBN) formalism was originally developed for modelling nested causal relationships. In this paper we argue that the formalism can also be applied to modelling the hierarchical structure of mechanisms. The resulting network contains quantitative information about probabilities, as well as qualitative information about mechanistic structure and causal relations. Since information about probabilities, mechanisms and causal relations is vital for prediction, explanation and control respectively, an RBN can be applied to all these tasks. We show in particular how a simple two-level RBN can be used to model a mechanism in cancer science. The higher level of our model contains variables at the clinical level, while the lower level maps the structure of the cell's mechanism for apoptosis.

Keywords: Bayesian network; causal model; mechanism; explanation; prediction; control.

1. Introduction

This paper seeks to integrate considerations arising from recent philosophical work on scientific explanation into the causal Bayesian network modelling formalism.

Bayesian networks were originally developed to model probabilistic and causal relationships (Pearl 1988). In the last two decades, Bayesian nets have become the model of choice for prediction and control—for making quantitative predictions and for deciding which variables to intervene on in order to control variables of interest. Thus a Bayesian net can be used to answer questions such as: *given that a patient has treatment t , what is the probability $P(r|t)$ that their cancer will recur in the next 5 years?* And: *on which variables should we intervene in order to minimise the probability of recurrence?* Causal information is important here because it is only worth intervening on the *causes* of recurrence, not on other variables which might be indicators of, or evidence of, recurrence.

The causal structure modelled by a Bayesian net can also help answer certain simple explanatory questions, such as, *what was the chain of events that led up to the recurrence of the patient's cancer?* But often we want to be able to offer explanations, not in this backward, aetiological sense, but in a downward, mechanistic sense. In order to answer *how did the patient's cancer recur?* we may need to specify the lower-level activities of the relevant cancer mechanism and the corresponding cancer response mechanisms. To answer such explanatory questions a model needs to represent the relevant mechanisms, including their hierarchical organisation.

Philosophers of science have studied this kind of mechanistic explanation in some detail in recent years. The current consensus is that a phenomenon is explained by pointing out the constitution of reality—carved up in terms of parts, what the parts do,



and the organisation of these parts—that is responsible for the phenomenon. These lower-level phenomena may themselves call for explanation, in which case yet lower-level phenomena will be invoked, and so on. This sort of explanation invokes *inter-level* constitution and responsibility relations rather than *intra-level* causal relations. It is usually downward-looking, but in some cases phenomena can be explained in upward-looking mechanistic explanations (Darden 2006, 109).

The question therefore arises as to whether Bayesian nets—which are often used to model intra-level causal relations—can be extended to model hierarchical mechanistic structure, i.e., inter-level explanatory relations. If so, then Bayesian nets could be used for mechanistic explanation as well as for prediction and control. They could also be used for inter-level prediction and control, addressing questions such as, *what components of the (low-level) DNA damage response mechanism should one intervene on in order to increase the probability of survival?*

On the other hand, the question arises as to whether models of mechanisms, which, in the biomedical science textbooks for instance, often take the form of elaborate diagrams, and which depict qualitative structure very well, can be extended to include quantitative, probabilistic information. This would allow one to use the model to answer *quantitative* inter-level explanatory questions. E.g., *why did Alfie survive 10 years rather than the 1 year that was most probable given his (higher-level) clinical symptoms? Because his (lower-level) DNA damage response mechanism had certain features that made longer survival much more probable.* That there is a need for a quantitative extension of mechanistic models is highlighted, for instance, by cancer biologist Yuri Lazebnik:

Biologists summarize their results with the help of all-too-well recognizable diagrams, in which a favorite protein is placed in the middle and connected to everything else with two-way arrows. Even if a diagram makes overall sense, it is usually useless for a quantitative analysis, which limits its predictive or investigative value to a very narrow range. (Lazebnik 2002, 181)

With the aim of addressing these two questions, this paper applies a hierarchical extension of Bayesian nets to modelling hierarchical (inter-level) mechanistic structure as well as causal (intra-level) relations and quantitative, probabilistic relations.

Bayesian nets have been extended to model hierarchy in a number of ways. For example, *recursive Bayesian multinets* model context-specific independence relationships and decisions (Peña et al. 2002), *recursive relational Bayesian networks* model relational structure and more complex dependence relationships (Jaeger 2001), *object-oriented Bayesian networks* can simplify the structure of large and complex Bayesian nets (Koller and Pfeffer 1997), *hierarchical Bayesian networks* offer a very general means of modelling arbitrary lower-level structure (Gyftodimos and Flach 2002) and *recursive Bayesian networks* were developed to model nested causal relationships (Williamson and Gabbay 2005). In this paper we shall see how recursive Bayesian networks can also be used to model mechanisms, thus providing an integrated modelling formalism for prediction, explanation and control. This is important from the philosophy of science perspective of seeking to understand modelling and its relation to goals of science such as prediction, explanation and control. It is also important from the AI perspective of needing to provide models that can be used to answer a variety of queries in decision support systems. And it is important from the bioinformatics perspective, which requires models that can integrate a variety

of data sources at different levels (e.g., clinical data and genomic data) with qualitative knowledge of the basic science involved.

In the remainder of this section we will introduce the notion of mechanistic explanation to which we appeal, the Bayesian net modelling formalism, and a cancer-science case study that we will use as our running example. Then, in §2 we explain the recursive Bayesian network formalism and show how it can be used to model mechanisms. In §3 we show how such a network can be applied to the cancer science example. In §4 we argue that the recursive Bayesian network formalism really does model mechanisms in the sense invoked by the recent philosophy of science literature. In §5 we compare the formalism advocated here with other kinds of formalisms that might be applied to prediction, explanation and control. We summarise and outline future research in §6.

It should be emphasised that in this paper we *assume* full knowledge of causal, probabilistic and mechanistic relationships with the aim of showing that recursive Bayesian networks offer a useful way of representing and reasoning with that knowledge. While the assumption of full knowledge is rather strong, it allows us to set aside certain technical questions to do with partial information; for example, causal Bayesian nets depend on the Causal Markov Condition (see below) and when common causes of measured variables are not themselves measured, this assumption becomes implausible. We leave the questions of how best to cope with partial knowledge, and of how to use recursive Bayesian networks in an exploratory way to discover causal, probabilistic and mechanistic relationships to future work; these questions will not be addressed in this paper.

Mechanistic explanation. Mechanistic explanation is an alternative to traditional approaches to explanation such as the deductive-nomological (DN) model. It is far more suited to the biomedical sciences, where there are few if any exceptionless laws, and where scientists see themselves as instead involved in discovering mechanisms. In the last decade there has been a great deal of philosophical debate about what a mechanism is, with three main contenders. Machamer, Darden and Craver have the dominant view: ‘Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions’ (Machamer et al. 2000, 3). Stuart Glennan holds: ‘A mechanism for a behavior is a complex system that produces that behavior by the interaction of a number of parts, where the interactions between parts can be characterized by direct, invariant, change-relating generalizations’ (Glennan 2002, S344). Bechtel and Abrahamsen’s view is: ‘A mechanism is a structure performing a function in virtue of its component parts, component operations, and their organization. The orchestrated functioning of the mechanism is responsible for one or more phenomena.’ (Bechtel and Abrahamsen 2005, 423). These detailed debates are interesting, but there is also clearly some consensus. The main contenders agree that much explanation in the biomedical sciences proceeds by finding the mechanism responsible for the phenomenon, and they all agree that finding a mechanism involves finding parts, what the parts do, and their organization. Mechanistic explanation is usually thought of as hierarchical rather than causal: the parts sought are those that *constitute* the phenomenon of interest.

Bayesian nets. A Bayesian net (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. Figure 1 gives an example of a directed acyclic graph; to form a Bayesian net, the probability distributions $P(V_1)$, $P(V_2|V_1)$, $P(V_3|V_2)$, $P(V_4|V_2V_3)$ and $P(V_5|V_3)$ need to be provided. The graph and the probability function are linked by the *Markov Condition* which says that each variable is probabilistically independent of its non-descendants, conditional on its parents, written $V_i \perp\!\!\!\perp ND_i \mid Par_i$. Figure 1 implies for instance that V_4 is independent of V_1 and V_5 conditional on V_2 and V_3 . A Bayesian net 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 value x 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* Bayesian net or *causal net*, the arrows in the DAG are interpreted as direct causal relationships (Williamson 2005), and the net can be used to infer the effects of interventions as well as to make probabilistic predictions (Pearl 2000); in this case the Markov Condition is called the *Causal Markov Condition*.




Figure 1: A directed acyclic graph

Cancer case study. The applicability of RBNs will be illustrated with reference to cancer science. Cancer is a complex biological and social phenomenon, initiated by exposure to DNA-damaging factors and leading, through a succession of steps, to ‘unregulated cell growth’ (King 2000, 1). At the molecular level of description, causal factors triggering cancer are commonly divided into those that are external to the individual (e.g., UV, ionising radiation, chemicals) and those that are internal (e.g., free radical formation, incomplete repair of misaligned bases).¹ These factors, in turn, are related to high-level variables such as lifestyle (e.g., dietary habits, smoking habits, exposure to solar radiation), family history, age, clinical evidence (e.g., biopsy results, X-rays), survival, etc. The bearing of these high-level variables on the molecular variables is still not well understood. What we know, however, is that internal and external factors both exercise their harmful potential by damaging DNA and that, in turn, the cell’s ability to respond to DNA damage, whether unaided (before cancer development), or aided by treatments (when cancer, after development, is being cured), influences the organism’s survival via regulating cell growth.

¹ King (2000, 24) uses the terms ‘exogenous’ and ‘endogenous’ factors, respectively. In order to avoid confusion with the technical meaning those terms have in statistical modelling, we use the more intuitive terms ‘external’ and ‘internal’.

In §3 we will present a two-level RBN comprising, at the higher level, simple variables—*age, familial factors and survival*—and a recursive variable, *DNA condition*, the latter being analysed at the lower level in terms of one mechanism for DNA damage response, *apoptosis*, i.e., cell suicide, and the role of the protein p53 within this mechanism. This can provide a valuable insight into the relationships between higher-level variables, on the one hand, and molecular indicators of correct or incorrect functioning of apoptosis, on the other.


2. Recursive Bayesian nets

Recursive Bayesian networks (RBNs) were originally developed in Williamson and Gabbay (2005) to model nested causal relationships such as *[smoking causing cancer] causes tobacco advertising restrictions which prevent smoking which is a cause of cancer*. But nested causality is not a general concern of this paper; in this section we develop the RBN formalism in the context of modelling mechanisms rather than nested causality.²

Definitions. A recursive Bayesian net is a Bayesian net defined over a finite set V of variables *whose values may themselves be RBNs*. A variable is called a *network variable* if one of its possible values is an RBN and a *simple variable* otherwise. Note that an RBN is a Bayesian net—a Bayesian net whose variables may be richly structured. On the other hand, Bayesian nets are also RBNs: a *simple Bayesian net* is an RBN whose variables are all simple.


The directed acyclic graph of an RBN A is the *top level* of A . A DAG of another RBN that is the value of a network variable V_i of A is the *next level down* in A , and so on; V_i is the *direct superior* of the variables in that DAG, and those variables are its *direct inferiors*. Variables that occur at the same level of an RBN are said to be *peers*. If an RBN contains no infinite descending chains—i.e., if each descending chain of inferiors terminates in a simple variable—then it is *well-founded*. We restrict our attention to well-founded RBNs here.

Example. To take a very simple example, consider an RBN on $V = \{M, S\}$, where M stands for some *DNA damage response mechanism* which takes two possible values, 0 and 1, while S is *survival after 5 years* which takes two possible values *yes* and *no*. The corresponding Bayesian net is:




² A mechanism can be thought of as a special case of nested causality: one in which causal relations are nested according to the levels of organisation of the mechanism. While in the general case of nested causality the same variable can appear at more than one level (e.g., smoking and cancer in the above example), this is rarely if ever plausible within the specific context of a mechanism, where different levels tend to mention different variables because each level of a mechanism is taken to *constitute* a higher-level phenomenon to be explained as well as to explain it. This fact makes the question of the consistency of a RBN (Williamson and Gabbay 2005, §4) somewhat easier to analyse in our context. (On the other hand, should mechanisms be found that do exhibit causation across levels, then it is possible to use repeated variables in the RBN formalism to represent such phenomena.)

Suppose that S is a simple variable but that M is a network variable, with each of its two values denoting a lower-level (simple) Bayesian network that represents a state of the DNA damage response mechanism. When M is assigned value 1 we have a net m_1 representing a functioning damage response mechanism, with a probabilistic dependence (and a causal connection) between damage D and response R :



On the other hand, when M is assigned value 0 we have a net m_0 representing a malfunction of the damage response mechanism, with no dependence (and no causal connection) between damage D and response R :



Since these two lower-level nets are simple Bayesian nets the RBN is well-founded and fully described by the three nets.

Modelling assumptions. Since an RBN is defined as a special kind of Bayesian net, the Markov Condition is imposed on any set of peers. In fact, since we are concerned here with modelling mechanisms under complete causal knowledge, we shall assume in this paper that each arrow is interpreted causally—though this causal interpretation is not essential to the general RBN formalism—and that the *Causal Markov Condition* holds of any set of peers.³ The Causal Markov Condition describes the causally-induced probabilistic independencies that obtain at any particular level of the RBN. But RBNs are subject to a further condition, the *Recursive Markov Condition*, which describes inter-level independencies and which says that each variable is probabilistically independent of those variables that are neither its inferiors nor peers, conditional on its direct superiors. Combining the intra-level independencies posited by the Causal Markov Condition and the inter-level independencies of the Recursive Markov Condition we have a further condition, which we take as the key condition to be satisfied by an RBN:


³ Note that the Causal Markov Condition is somewhat controversial (see, e.g., [Cartwright 2007](#), Part II). In particular, its validity depends on the way in which causality itself is analysed ([Williamson 2005](#), Chapter 4). For the purposes of this paper, we need neither pin down a specific analysis of causality nor precisely delimit the validity of the Causal Markov Condition: it is simply taken to be a modelling assumption that is open to question and to testing. See ([Spirtes et al. 1993](#), §5.5) for discussion of statistical tests for probabilistic independence.

Recursive Causal Markov Condition (RCMC). 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: $V_i \perp\!\!\!\perp NID_i \mid DSup_i \cup Par_i$ for each variable V_i , where NID_i is the set of non-inferiors-or-descendants of V_i and $DSup_i$ is the set of direct superiors of V_i .

Joint distribution. We shall now turn to the question as to how one might exploit the modelling assumption RCMC to define a joint probability distribution over the variables that occur as the various levels of the network. 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 simple Bayesian net, 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,⁴ 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: $P(V_i \mid Par_i \cup DSup_i)$ must be consistent with the $P_{v_{j_l}}(V_i \mid Par_i)$ given in the RBN for each direct superior V_{j_l} of V_i . If each variable has at most one direct superior in the RBN then this will uniquely determine the required distribution $P(V_i \mid Par_i \cup DSup_i)$; in other cases we follow [Williamson and Gabbay \(2005, §5\)](#) and take the distribution to be that, from all those that satisfy the constraints, which has maximum entropy. The Markov Condition holds in the flattening because the Recursive Causal Markov Condition holds in the RBN.


Of course in the flattening the arrows are not all interpretable causally so the *Causal* Markov Condition is not satisfied. The non-causal arrows (the arrows from the direct superiors of a variable to the variable itself) are not to be interpreted as giving the direction of explanation either—they signify the direct superiority relation. The flattening should be thought of as a formal tool for defining a joint distribution, rather than a part of the RBN model itself, so the fact that the arrows in the flattening do not admit a uniform interpretation is neither here nor there: the arrows are a formal device for representing probabilistic independencies via the Markov Condition. (For the flattening to satisfy the Markov Condition, the extra arrows have to be directed from direct superior to direct inferior; the opposite orientation will not work as it will imply different independence relationships which are not all supported by the Recursive Causal Markov Condition.)

In our example, for assignment m_0 of network variable M we have the flattening m_0^\downarrow :



⁴ These can be thought of as variables that can only take one possible value, i.e., constants. In the Bayesian net literature they are called *instantiations* of the corresponding variables V_{j_1}, \dots, V_{j_k} .

with probability distributions $P(m_0) = 1$, $P(S|m_0)$ determined by the top level of the RBN and with $P(r_1|m_0) = P_{m_0}(r_1)$ and similarly for r_0, d_0 and d_1 . The flattening with respect to assignment m_1 is:



Again $P(r_1|d_1m_1) = P_{m_1}(r_1|d_1)$ etc. In each case the required conditional distributions are fully determined by the distributions given in the original RBN.

As long as certain consistency requirements are satisfied (Williamson and Gabbay 2005, §4), the flattenings suffice to determine a joint probability 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 $v_1 \cdots v_m$.

Prediction, explanation and control. With a joint distribution, the RBN determines the probabilities of all combinations of assignments of values to variables, so the model can be used for prediction. For example, the probability that R has value 1 and that the patient will survive 5 years is

$$P(s_1r_1) = P(m_0s_1r_1) + P(m_1s_1r_1) = P(s_1|m_0)P(m_0)P_{m_0}(r_1) + P(s_1|m_1)P(m_1) (P_{m_1}(r_1|d_1)P_{m_1}(d_1) + P_{m_1}(r_1|d_0)P_{m_1}(d_0))$$


and these latter probabilities are all given in the RBN.

More than that, since at each level the arrows in the RBN are interpreted causally, the model can be used for backwards aetiological explanation: one might cite damage response mechanism type 1 as the reason a patient survived 5 years. If the inter-level relations match that of mechanistic composition then the model can be used for mechanistic explanation. Thus the values of the damage and response variables and the link between the two might explain survival.

Finally one can use an RBN to reason about control across levels: by intervening on response R one might change the probability of survival. Interventions in RBNs work in just the same way as they do in standard Bayesian nets. When one intervenes to fix the value of a variable, one creates a new net by deleting all the arrows that go into that variable in the RBN. Then one calculates probabilities in the usual way, instantiating the variable in question to the appropriate value, and using flattenings if necessary. Thus in our example, intervening to set R to value r_1 involves no change at the top level, but both lower-level graphs now have no arrow from D to R :



Both flattenings then have the same graphical structure:



and the new probability of s_1 after intervention to set r_1 ,⁵ is $P(s_1||r_1) = P(s_1 r_1)/P(r_1)$, with these last probabilities calculated using the new network structure:

$$\begin{aligned} P(s_1 r_1) &= P(m_0 s_1 r_1) + P(m_1 s_1 r_1) = \\ &= P(s_1|m_0)P(m_0)P_{m_0}(r_1) + P(s_1|m_1)P(m_1)P_{m_1}(r_1), \end{aligned}$$

and

$$P(r_1) = P(m_0)P_{m_0}(r_1) + P(m_1)P_{m_1}(r_1).$$

While the formal apparatus for handling interventions in RBNs is exactly the same as that for standard Bayesian nets, one needs to be careful now that the level at which the intervention takes place has been correctly identified. Thus while it may be possible to intervene at the lower level to set R to value r_1 , presumably what one wants to do in this case is not to initiate damage response (e.g., cell suicide) in all cells, healthy included. Rather, one wants to intervene to ensure that the response is triggered by damage, i.e., one wants to ensure that D does cause R . This is an intervention to fix the value of M to m_1 , rather than to fix the value of R to r_1 —it is a higher-level intervention.

Note that it is quite common in the mechanisms literature to distinguish interventions at different levels, even when those levels are related by constitutive relations—e.g., in discussions of mutual manipulability of different levels of a mechanism (see §4). So RBNs are not adding anything conceptually here. Also, the formal treatment of interventions that RBNs appeal to is just the formal treatment in ordinary Bayesian nets, so RBNs do not add anything in terms of formal explication. What RBNs do, is allow one to apply the latter formal account of what interventions entail (in terms of the probabilistic inferences one can draw on the basis of interventions) to the former case of multi-level interventions. This is something new and not handled by standard Bayesian nets.

3. Cancer application

In this section the applicability of RBNs to cancer science is illustrated in more detail. We will present a more realistic two-level RBN, comprising a higher, clinical level and a lower, molecular level. At the clinical level (see Figure 2), familial (i.e., hereditary) factors of a certain type of cancer (F) and age (A) cause DNA condition (C), which in turn causes survival in months (S), with A having also a residual influence on S that doesn't go through C . F , A and S are simple variables. C , instead, is a network

⁵ A conditional probability where the condition is an intervention is often represented using a double bar, $P(\cdot||\cdot)$.




Figure 2: F : Familial Factors; A : Age; C : DNA Condition; S : Survival.




Figure 3: Functioning apoptosis mechanism. D : Damaging agents; G : $p53$ gene status; P : $p53$ protein level; O : Apaf1 level; E : caspase 3 level; T : TNF level; I : caspase 8 level.

variable whose two values, c_1 (good) and c_0 (bad), correspond to lower-level networks (Figure 3 and Figure 4) representing the mechanism for *apoptosis* functioning, respectively, correctly or incorrectly.⁶ Let us introduce these lower-level networks whilst describing the mechanism.⁷

DNA damage is responsible (i) for disrupting the cell’s regulatory activities, that is, the cell’s ability to transcribe genes that the affected DNA encodes; and (ii) for modifying the survival ability of the daughter cells, due to the harmful genomic mutations which obtain as the damage is passed from mother cell to daughter cell when the mother cell undergoes mitosis, i.e., cell division. When the DNA is damaged, a well-functioning cell reacts via defence mechanisms known as “DNA repair” mechanisms that heal the cell after damage has arrested its cycle. Depending on the kind of damage (e.g., single-strand, double-strand, mismatch), different enzymes are recruited to fix the damage. As a last resort, if the damage is serious and cannot be effectively repaired, the cell either (i) enters an irreversible state of dormancy (“senescence”) or (ii) commits suicide (“apoptosis”).

⁶ At the higher level, the amount of DNA carried by the cell constitutes evidence for *apoptosis* functioning correctly or incorrectly—in case of malfunctioning of *apoptosis* the cell’s DNA can grow up to 4-5 times larger.

⁷ If one were to draw the flattenings c_1^\downarrow and c_0^\downarrow that contain the lower-level networks represented in figure 3 and figure 4 respectively, one would also need to include nodes for the assignments c_1 and c_0 of the network variable C and arrows from them to their inferiors D, G, P, O, E, T and I in each flattening.




Figure 4: Malfunctioning apoptosis mechanism. D : Damaging agents; G : $p53$ gene status; P : $p53$ protein level; O : Apaf1 level; E : caspase 3 level; T : TNF level; I : caspase 8 level.

However, when none of the above strategies is effective, damaged cells keep growing and dividing and, in so doing, produce mutations, which are the first step toward cancer development.

This is, in short, how errors accumulate irreversibly, from one mitosis to another. During DNA replication, misreading of damaged bases can occur, leading to incorporation of wrong bases opposite damaged ones (e.g., an A:T nucleotide base pair is in the daughter cell's DNA where the mother cell's DNA had a G:C nucleotide base pair, due to a $G \rightarrow A$ mutation). When the cell undergoes division, DNA changes in the mother cell result in mutations, that is, *irreversible* changes in DNA sequence, in the daughter cell. Inherited changes, in fact, cannot be repaired, as they are on both strands of DNA, so template information for correction is lost. As a consequence, they are replicated and inherited through further divisions.

Notice that, although DNA replication mechanisms are very precise, DNA damage due to both internal and external factors can produce a daily number of lesions high enough to be dangerous (King 2000, 125). This is why the mechanisms that allow the cell to correct errors before they are replicated (repair) or prevent mutations (senescence and apoptosis) are so important. In fact, a mutation can start a cascade of mutations, because of its capacity to impair the cell's activities (among which there is the production of enzymes needed in DNA repair itself), so that further mutations occur more easily.

What follows is a summary of the mechanism upon which the cell relies in order to oppose cancer development.⁸ Possible indicators of good or bad functioning of these mechanisms will also be introduced in order to devise the lower-level decomposition of the network variable C .

Although there are ways of measuring DNA damage *in vitro*, there currently is no way to determine this directly *in vivo*. One needs to resort to the expression levels of certain DNA damage response genes as a surrogate for DNA damage; among these genes, $p53$ is considered the “master guardian” of genomic integrity of the cell. After noticing presence of metabolic disorder or genetic damage, protein $p53$ can induce

⁸ This summary develops discussion in Weinberg (2007, ch. 9) and we refer the reader to that text for a more detailed description.

cell-cycle arrest, activate DNA repair proteins, or lead to cell death (apoptosis).⁹ The basic steps leading to cell-cycle arrest are p53's induction of the synthesis of the protein p21^{Cip1}, which prevents cyclin dependent kinases (CDKs) from triggering the process of growth-and-division. Halting the cell cycle, in turn, permits DNA polymerases and other molecules to perform DNA repair. However, if the metabolic derangement or genomic damage is too severe to be cured, p53 may decide—it is still unclear how—that cell death is a better option. For reasons of brevity, only a simplified version of the mechanisms responsible for apoptosis involving p53 will be presented here.

In a well-functioning cell, “wild” (i.e., non-mutant) p53 normally goes through a rapid degradation, due to its being “tagged” by the Mdm2 protein and subsequently “digested” by proteasomes. The amount of p53 increases when, e.g., its phosphorylation due to genotoxic (i.e., DNA damaging) factors (e.g., X-rays), or the phosphorylation of Mdm2 due to ATM kinase, results in Mdm2 being unable to bind to p53. Interestingly enough, p53 promotes synthesis of Mdm2, thereby contributing to its own inhibition in a negative feedback loop. This loop successfully regulates apoptosis unless the gene *p53* mutates. In the latter case, mutation of *p53* prevents Mdm2 from binding to p53 and, as per the wild case, this results in an increase of p53. However, the defective p53 has lost its ability to act as a transcription factor, that is, is unable to bind to the promoters of genes that synthesise pro-apoptotic proteins in the successive stages of the mechanism.

According to available data, gene *p53* is mutated in 30% to 50% of commonly occurring human cancers (Weinberg 2007, 310). The crucial, *causal*, role of the protein p53 is explicitly recognised, as is the possibility of building a mechanistic model around *p53* to explain how alarm signals stop the cell cycle or trigger apoptosis (Weinberg 2007, 316-317). The explanatory value of such a model and the added, predictive, advantage of our RBN will be particularly relevant to cancer types where *p53* is highly mutated.

When modelling, the Mdm2-p53 loop will be simplified so as to represent its overall influence by a single arrow, since RBNs, like standard BNs, do not admit cycles (see also §4). In the RBN, the upstream variables *D* (damaging agents) and *G* (*p53* gene status, which can be wild or mutant) initiate distinct pathways in the case of a healthy cell and a cancerous cell. In both, an edge links *D* and *G* to *P* (p53 protein level). In the network for the healthy cell, *P* is, in turn, directly linked to downstream effects of the regulatory feedback loop p53-Mdm2. This latter link is missing in the network for the cancerous cell. Both networks include, among the values of *D*, both radiotherapy and chemotherapeutic agents (e.g., Cisplatin and the topoisomerase II inhibitors etoposide and mitoxantrone), whose effects on apoptosis are being extensively investigated (see, e.g., <http://www.virtuallumour.co.uk/apoptosis.htm>).

In the well-functioning cell, increased p53 plays an important role in several apoptotic signalling pathways. It is useful to distinguish between an *intrinsic* (internal to the cell) and an *extrinsic* (external) pathway. The apoptotic signal can also be amplified via crosstalk between these pathways.

⁹ Following the lead of the biological literature, we use the same name to refer to a gene and the protein it codes for, and distinguish the former from the latter by italicising it (e.g., *p53* stands for the gene, p53 for the protein).

Internally, p53 acts as transcription factor for the encoding of pro-apoptotic proteins (e.g., Bax) that, by opening the mitochondrial membrane channel, allow release of cytochrome c . Pro-apoptotic proteins belong, together with anti-apoptotic proteins, to a family of proteins named the “Bcl-2 family” (after Bcl-2, the first protein found to contribute to regulation of apoptosis besides p53), due to their sharing a common coding sequence. Their balance determines the opening of the mitochondrial membrane and the release of cytochrome c , that binds a protein called Apaf1 and leads to the formation of a wheel-like heptamer called apoptosome. The apoptosome, then, recruits procaspase 9 and activates it by means of proteolytic cleavage, so that it becomes caspase 9 (“initiator”). Caspase 9, in turn, initiates a cascade of caspases 3, 6 and 7 (“executioners”) that results in the disintegration of the cell. Executioners can be inhibited by IAPs proteins—their action, in turn, being inhibited by another protein, Smac/DIABLO, also released by the mitochondrion together with cytochrome c . Level of apoptosis cannot be easily measured. Expression levels of the caspases 3 and 9 are often used as surrogates.

Also in this case, it will be useful to bypass the intermediary steps of the cascade, as the catalytic action of caspases amplifies the signal via a positive feedback loop. For the sake of simplicity, our RBN will include an edge from P to O , which stands for the effect of p53 on the production of Apaf1 in the healthy cell network, and no such link in the cancerous cell network. Another edge will then depart from O to E to signify the overall activation level of Caspase 3 due to the production and subsequent activation of Apaf1 via the activation of Caspase 9.¹⁰ In both the healthy cell network and the cancerous cell network a direct edge from D to E will stand for all damaging, residual, effects on apoptosis that do not go through P .

Let us turn to the extrinsic pathway. This is due to ligands in the extracellular space (e.g., FasL) belonging to the TNF (tumour necrosis factor) protein family, that bind to death receptors on the surface of the cell (e.g., FasR). The tail of these receptors act in conjunction with the FADD protein to assemble DISC, a complex which, by prior activation of initiator caspases 8 and 10, triggers another cascade of caspases 3, 6 and 7. p53 contributes to this process by promoting the expression of the genes encoding the Fas receptor, thereby increasing the cell’s responsiveness to extracellular

¹⁰ Notice that, although this is a simplified representation of the intrinsic apoptotic pathway, it is sufficiently faithful to reality. As shown by recent studies on the role of Apaf1, XIAP, and Caspases 3 and 9 (Legewie et al. 2006), Caspase 3 activity is bistable and irreversible. Simulations identify hysteretic behaviour, with low active Caspase 3 depending on low active Apaf1 until a threshold point is reached where active Caspase 3 switches irreversibly to a high state. Other simulations explain the role of XIAP in establishing bistability and irreversibility. For low Apaf1, most Apaf1-associated active Caspase 9 is inhibited by XIAP, whereas above the threshold Apaf1 manages to initiate Caspase 3 activation. Active Casp3 then further promotes its own activation by sequestering XIAP away from Apaf1-associated Caspase 9. This results in most of XIAP being bound to Caspase 3, and therefore being unable to inhibit Caspase 9, which is then free to trigger executioner Caspase 3. Furthermore, Caspase 3 activity is maintained even if the stimulus is removed, since Caspase 3, once activated, retains XIAP, thereby preventing full Caspase 9 deactivation. These conclusions are well supported by experimental studies. This makes it plausible to model the influence of p53 on Caspase 3 in a linear fashion, with P contributing to O , and O having a positive net effect on E , and to ignore the inhibiting activity of XIAP.

death ligands, specifically FasL. The extracellular apoptotic signal can be amplified by crosstalk between the two pathways: caspases 8 and 10 cleave Bid, which inhibits the action of Bcl-2 antiapoptotic proteins.

The extrinsic pathway in the RBN will be modelled as follows: an edge will be included to represent the influence of T (TNF proteins) on I (activation level of Caspase 8), which, in turn, has an effect on E . Furthermore, an arrow from P to I will be included to represent the—indirect—effect of p53 on Caspase 8 activation via the expression of genes encoding Fas. Finally, an edge from I to O will be added to model the inhibiting effect of Caspase 8 on Bcl-2 antiapoptotic proteins through the activation of Bid. The latter two edges provide a (simplified) representation of the crosstalk between intrinsic and extrinsic pathways.¹¹

Cancer cells inactivate apoptosis in several ways that enable them to survive and thrive. They can increase the level of anti-apoptotic proteins, change the gene coding for p53 or its upstream regulators, methylate promoters of pro-apoptotic genes, interfere with the release of cytochrome c , inhibit caspases, etc. On the other hand, over-expression of proapoptotic proteins or dysfunction of antiapoptotic proteins due to mutations can result in too much apoptosis and cause other pathological conditions, e.g., neurodegenerative disorders such as Alzheimer's or Parkinson's disease. Thus, an RBN could also prove useful in explaining the relevance of p53 with regard to these latter diseases.

4. RBNs as models of mechanisms

We have described the formalism that allows recursive Bayesian nets to be extended to model mechanistic hierarchy, and illustrated it by applying it to a DNA damage response mechanism. In this section we address philosophical concerns about the interpretation of RBN models. According to the recent mechanisms literature, mechanistic explanation proceeds by identifying mechanisms. Here we argue that RBNs are legitimate descriptions of such mechanisms: BNs, designed to model causal structures, can be extended in the way we suggest to model mechanistic hierarchy. First, we will examine how well the RBN framework can be used to capture the philosophical consensus on mechanisms. In particular, we will illustrate how the network variables in an integrated model—which are the variables that must fit *both* in a causal BN and in the mechanistic hierarchy—can accommodate the relevant constraints. Then we will move on to examine the modelling assumptions that RBNs need to satisfy. We conclude this section by discussing the advantages of representing mechanisms in this way.

¹¹ In this case the simplified representation of the extrinsic pathway in the RBN is also supported by experimental evidence aided by computation. Among other studies, [Mai and Liu \(2009\)](#) have recently produced a 40-node Boolean Network and performed extensive statistical analyses of its state space, which has resulted in the identification of key network components responsible for the stability of the surviving states and the irreversibility of the apoptotic process. Among other things, the study has confirmed the role of TNF as a strong promoter of apoptosis, its effects being only partially offset by GF (growth factor) once apoptosis has been initiated.

Modelling mechanisms. It is natural to have philosophical concerns about integrating BNs and mechanisms. Causal BNs are designed to model causal relations between distinct events, and are explicitly quantitative; the mechanisms literature concerns itself instead with constitutive relations—with how a lower level makes up or composes a higher level that it is not entirely distinct from—and is largely qualitative. We need to explain how a single model can accommodate these differences.

The qualitative-quantitative difference is not a deep one. Attention has been given to qualitative aspects of mechanisms in the philosophical literature on mechanisms because the main worry has been interpretation—the question being what counts as a mechanism and what doesn't—rather than modelling. But there is no reason why relations in mechanisms cannot be modelled quantitatively. The *structure* of this model is another issue: can we represent causal and constitutive structure so readily in the same model?

We will begin to address this question by explaining how RBNs can legitimately be used to describe mechanisms as they are examined in the burgeoning mechanisms literature.¹² The complex-systems approach to mechanisms of Machamer, Darden and Craver (MDC), Glennan, and Bechtel and Abrahamsen (B&A), that we have introduced earlier in §1, is explicitly concerned with explanation. It offers an account of what mechanisms are as they are used in decompositional explanations of higher-level phenomena in terms of the lower-level parts that constitute them. There are disagreements over details within this literature, but a great deal of agreement over core elements vital to a mechanistic explanation. These are:

- Decomposition into entities and activities (always)
- Organisation (always)
- Hierarchy (often)

The first element is that mechanisms have two kinds of parts: the objects in the mechanism, and the things that the objects do. MDC, Glennan, and B&A use different language for the two kinds of parts, but agree on this distinction. We use MDC's language but we use it merely descriptively as has become standard; there is no need to adopt their controversial metaphysical claims of entity-activity dualism in this paper. Entities and activities are not just any old division of a mechanism into pieces. The division must be into functioning components—a mechanistic explanation must identify those entities and activities by which the phenomenon is produced. Put simply, Craver's account of what it is for X to be part of the mechanism for Y is i) X is a part of Y and ii) X and Y are mutually manipulable—wobble X and Y wobbles, and vice versa. (For Craver's more formal statement of this, see Craver (2007, 152-153).)

¹² There is controversy within that literature, of course. We will not address the views of Woodward, and of Glymour and Cheng, on mechanisms (Woodward 2002, S375; Glymour and Cheng 1998). They hold that mechanisms are chains of difference-making relations of the same kind as are already modelled using BNs. Modelling mechanisms of this sort is done merely by adding extra nodes between existing nodes of a simple BN. There is nothing distinctive about *this* kind of mechanism, nothing beyond the BN approach. Modelling this kind of mechanism is no particular challenge because it does not add an extra explanatory dimension to the structure already existing in a standard BN.

The second element is that the entities and activities must be organized, spatially and temporally, in order to produce the phenomenon. This element cannot be ignored. As MDC, for example, write: ‘The organization of these entities and activities determines the ways in which they produce the phenomenon’ (Machamer et al. 2000, 3). See also Bechtel and Abrahamsen (2005, 435) and Darden (2002, S355).

The third element is that mechanistic decomposition is often hierarchical. The decomposition often continues beyond the entities and activities that initially explain a particular phenomenon, to look at the entities and activities that in turn explain one or more of the newly-identified activities, and so on. As Machamer et al. (2000, 13) write: ‘Mechanisms occur in nested hierarchies and the descriptions of mechanisms in neurobiology and molecular biology are frequently multi-level.’ Again, the decomposition is constrained: the entities and activities sought are those responsible for the phenomenon.

Take Glennan’s favourite example of a toilet. We can explain the flushing of the toilet mechanistically. The edge of the handle, and a part of the water are *pieces* of the mechanism of the toilet, but not components. They don’t each have a function that contributes to the functioning of the toilet, so altering them won’t have a recognisable effect on the functioning of the toilet. But the handle itself and the ballcock and valve are components performing specific functions—including triggering, and regulating water flow. Thus, changing one of these will affect the operation of the toilet in a recognisable way. They are component entities and activities. (See Illari and Williamson (2010) for detailed discussion of the importance of function with reference to the mechanisms of protein synthesis and natural selection.) Organization also exists in the toilet. The cistern must be appropriately positioned relative to the bowl, and the pulling of the lever will occur before the flushing of the water. Hierarchy is also present. Presumably, at a lower level the materials that the ballcock and valve are composed of will explain their functioning. For example, the material and construction of the ballcock explains why it floats, so responding to the level of water. The mechanistic explanation can descend one level.

RBNs. Our aim is to use RBNs to describe mechanisms of this and more complex kinds: entities and activities organized to be responsible for a phenomenon, often hierarchically. Broadly, we envisage nodes in the RBN standing for variables describing *either* entities or activities. In our cancer example, variable G describes an entity—whether the $p53$ gene is wild or mutant. Variables P and E measure activities—the rising or falling levels of protein p53 or caspase 3, respectively. Arrows in the RBN, however, will stand only for activities—interactions and influences among the variables. This is because the structure of the arrows in the RBN will represent aspects of the causal organization of the mechanism. Arrows in our cancer example represent, for instance, the influence of familial factors on the patient in question, and the influence of the level of protein p53 in triggering Apaf1 level, and caspase 8 level, which trigger the caspase cascade leading to apoptosis.

RBNs are flexible enough to do this because they use variables, which can represent many things. The interesting question for RBNs concerns the network variables. These are the most constrained since they have to represent the kind of thing that can be

related *causally* to other (same-level) nodes in the model, and *constitutively* to other (higher- or lower-level) nodes in the model. Our network variable is node *C*, standing for DNA condition. Node *C* represents something causally related to age, familial factors, and survival time; and constitutively related to what is represented by the lower-level nodes *D*, *G*, *P*, *O*, *E*, *T* and *I*—the mechanism responsible for DNA condition. This requires some attention, since on the one hand causes conventionally (although not uncontroversially) relate distinct events. On the other hand, mechanisms relate entities and activities and their organization to the explanandum phenomenon they constitute.

Happily, this is no bar to the success of the model. Network variables will have to represent only activities or phenomena that are susceptible to constitutive mechanistic explanations. The activity represented by a particular network variable will be the characteristic activity produced by the lower-level mechanism or mechanisms described in the lower-level nets that are the possible assignments to that variable. In our cancer example, the relevant phenomenon is DNA condition, and the lower-level nets model the mechanisms responsible for correctly functioning apoptosis, yielding good DNA condition, and incorrectly functioning apoptosis, yielding bad DNA condition. In this way, the lower-level BNs describe the mechanisms responsible for the changes in the higher-level variable describing the activity of the mechanisms. Clearly, some nodes will not have a decomposition, since they will not be the kinds of things likely to have a mechanistic explanation. Age, for example, seems unlikely to have a lower-level mechanism. In our RBN, such nodes are simple variables.

The arrows to and from nodes describe relations between the activities produced by the lower-level mechanisms, and variables describing other (same-level) entities and activities in the RBN. In our case, increased age causes worse DNA condition, and worse DNA condition reduces likely survival time. These are modelled by (same-level) arrows in the RBN. Both higher-level and lower-level networks are directed acyclic graphs and so represent causal structures. In so far as they are useful, it will be because they represent something about the causal organization of the mechanisms they describe. The unified model, modelling both causal and constitutive relations, offers valuable insights into the relationships between higher-level clinical variables known to be relevant to survival and lower-level indicators of DNA damage.

Formal modelling assumptions. We move now to consider the formal modelling assumptions that RBNs must satisfy. It is worth clarifying exactly what claim we defend: we claim that the RBN formalism we have presented can usefully represent *some* mechanisms. We are *not* claiming that all RBNs can be interpreted mechanistically, for reasons that will become clear. Nor are we claiming that *all* mechanisms can be represented using an RBN. There are well-known modelling assumptions required for BNs to be applied, and RBNs inherit these. Where the following modelling assumptions are violated, there is no guarantee that RBNs will be useful for modelling mechanisms:¹³

¹³ We do not include stability, faithfulness or minimality assumptions. These assumptions are normally invoked to justify procedures for learning causal networks from data. Since we assume causal relations are given for the purposes of this paper, they are unnecessary.

- Causal Markov Condition: At each level, each variable is independent of its non-effects, conditional on its direct causes.
- Modularity: At each level, an intervention to change the value of a variable does not change the underlying causal structure. It merely breaks the arrow between the altered variable and its parents.
- Acyclicity: The system does not involve any feedback loops.

We have added a new assumption, and there is no guarantee that RBNs will be useful where it is violated:

- Recursive Markov Condition: Each variable is probabilistically independent of those variables that are neither its inferiors nor peers, conditional on its direct superiors.

As explained in §2., the Recursive Markov Condition and Causal Markov Condition are integrated via the Recursive Causal Markov Condition.

The first three assumptions have been thoroughly discussed and are well-understood (see, e.g., Williamson 2005; Cartwright 2007, Part II). There are established strategies for building models that do not violate the Causal Markov Condition or acyclicity (Pearl 1988; Neapolitan 1990; Spirtes et al. 1993; Neapolitan 2004). Problematic cases for modularity are well-known, such as in the social sciences: for example see the discussion of ‘structure altering interventions’ in Steel (2008). We will not discuss these further here. It is sufficient to note that both levels of our cancer RBN separately meet the modelling assumptions of traditional BNs. In particular, the lower-level networks do not include cycles. This is why the Mdm2-p53 negative feedback loop, and the complex intermediate steps of the caspase cascade, which involves positive feedback loops, have each been summarised in a standard way with a single causal arrow. Summarising in this way ensures that the relationships represented in each BN do satisfy these modelling constraints.¹⁴

The *Recursive* Causal Markov Condition, however, is new, and it is worth examining whether every mechanism will satisfy it. Looking again at our cancer example, represented in Figure 2, the network variable is node *C*—DNA condition. The lower-level

¹⁴ In general, handling causal cycles in the right way is of crucial importance for BNs to deliver good predictions under intervention. (Since an RBN is a special kind of BN, the issue of cyclicity is not a distinct, extra problem.) For this purpose, one among the following strategies is commonly employed. One strategy consists in time indexing the variables, which is commonly done when using dynamic BNs (DBNs). For DBNs, it is possible that one variable causes another *and* vice versa—but at different times. DBNs are also useful for inferring the time that elapses between cause and effect. For these reasons, they are often used to model biological mechanisms such as gene expression networks. (For more on DBNs and their biological applications, see, e.g., Friedman et al. (1998, 2000).) However, since the temporal aspect is orthogonal to our primary concern, which is modelling hierarchy, it will not be dealt with here. Another strategy consists in combining the values (e.g., a_1, a_2, b_1, b_2) of two variables (*A* and *B*) that are connected by a causal loop, into a single variable (*AB*, taking the possible values $a_1b_1, a_1b_2, a_2b_1, a_2b_2$). Finally, one can leave out a node if one is not interested in it, provided the node in question is not a common cause of other variables in the net—in the present case, a suitable move is simply to leave out Mdm2. Notice that these strategies do not involve making false assumptions about known cycles, but are instead formal moves to deal with them in a principled way.

nets are represented in Figure 3—correctly-functioning apoptosis—and Figure 4—incorrectly-functioning apoptosis. For this net, the Recursive Causal Markov Condition requires that, conditional on C , all the lower-level variables, D, G, P, O, E, T and I , are probabilistically independent of the other higher-level variables, A, F and S . The condition can be informally linked with Craver’s mutual manipulability criterion for being a component in a mechanism. If D, G, P, O, E, T and I represent components in the mechanism for C , wiggling D, G, P, O, E, T and I should affect C , and vice versa. However, if the value of C is held *fixed*, wiggling any of the higher-level variables A, F and S should have no effect on any of the lower-level variables D, G, P, O, E, T and I , and vice-versa.

There are two possible problems for the Recursive Causal Markov Condition, which are separate. The first problem is the existence of systems such that *no possible* RBN (satisfying the Recursive Causal Markov Condition) could model such a system. An infringement of the condition might be due to a failure of the Causal Markov Condition: for instance, any system where there are two peers that are dependent but not causally connected will create this problem. For example, the flow of the water in the toilet and the sound of the water are dependent, but might not be viewed as causally connected because they do not correspond to spatio-temporally distinct events. (Each seems to be the kind of thing that has a mechanistic explanation at the lower level, and of course the mechanistic explanation of the flow of water in the toilet will at the very least overlap considerably with the mechanistic explanation of the sound of the water.)

The second problem concerns the *usefulness* of RBNs for modelling some systems. The kind of system that might create difficulties (although note that it may not count as a *mechanism*) is a system that is richly integrated so that many of the higher-level variables are related to many of the lower-level variables, and vice-versa. RBNs representing such systems may have to be so complex in order to satisfy the Recursive Causal Markov Condition, as to be not much practical *use*. RBNs are designed to be useful for modelling systems where a subset of the lower-level variables can be selected such that they act on the higher-level net *only* via a subset of the higher-level variables (or a single higher-level variable as in our cancer example). Selecting such a subset may be difficult for some systems.

Consider for example a thermodynamic explanation for the temperature and velocity of a body of air. These two aggregate variables are often related. We have an explanation for the aggregate variable *temperature* in terms of (among other things) the kinetic energy of the individual molecules. We have an explanation for the aggregate *velocity* of the body in terms of (among other things) the velocity of each individual molecule. Perhaps we could produce an RBN that included causal relations between the temperature and aggregate velocity of the body, which would both be network variables, with lower-level nets representing the molecular explanation. However, the kinetic energy of a molecule is *not* independent (causally or probabilistically) of its velocity, so temperature is not independent of the velocity of individual molecules, and aggregate velocity is not independent of the kinetic energy of individual molecules. This is not just because velocity and kinetic energy measure aspects of the same molecule, but because of the meanings of ‘kinetic energy’ and ‘velocity’.

Such a system *can* still be modelled using an RBN that satisfies the Recursive Causal Markov Condition. For the above case, this can be achieved by making temperature *also* a direct superior of the variables describing the average velocity of individual molecules, and making aggregate velocity *also* a direct superior of the variables describing the kinetic energy of individual molecules. Whether this is a model of a mechanism or not depends on whether the superior-inferior relationships have a legitimate interpretation as representing mechanistic hierarchy, as well as the original superior-inferior relationships between temperature and kinetic energy and between aggregate and individual velocity. Either way, this move already makes the RBN messier than we originally wanted. It can be seen that putting further pressure on this example could lead to an RBN where *all* the higher-level variables describing the body of air are direct superiors of *all* the lower-level variables describing aspects of the individual molecules. The Recursive Causal Markov Condition is trivially satisfied in such a net, but this is certainly not the kind of model that would be particularly useful to build. In such a case—and of course if this kind of system counts as a mechanism—then the RBN formalism may be less useful for modelling this mechanism than some of the other formalisms we discuss in the next section.

The advantages of such a representation. By contrast, though, this kind of difficult case highlights situations where the RBN formalism is instead likely to be useful. This happens precisely when lower-level components responsible for particular features of the higher level *can* be identified. Take a machine we have built, such as Glennan's toilet. We have constructed a ballcock and valve to refill the cistern after each flush. It is this system that is responsible for the level of water in the cistern. It is not directly responsible for anything else. We can be confident that the ballcock and valve acts on the flushing of the toilet—say, the flow of water through the toilet bowl—only by being the component that determines the level of water in the cistern. If we hold the level of water in the cistern (the value of the higher-level network variable) fixed, the action of the ballcock and valve is independent of the flow of water through the toilet bowl.

With the idea clear, we can return to our cancer example. This is not as simple as the toilet. Cancer is a complex phenomenon, with many more interacting causes than a toilet has. There are two possible problems. The first is that of many unknown causes. This is clearly a practical concern when modelling cancer systems. However, we are here assuming complete causal knowledge. This simplifies things. If we have complete causal knowledge—if we know all the causal paths from D, G, P, O, E, T and I —we will be in a position to know that they only affect or are affected by the higher-level variables via C . (As we have said, for the purposes of this paper we are setting aside the question of how RBNs might be useful in the absence of complete causal knowledge—in the process of mechanism discovery.) The second problem is whether it is plausible that a lower-level component can be identified that is responsible for the higher-level clinical variable. Since DNA condition can be measured independently of the states of D, G, P, O, E, T and I , it is a legitimate clinical variable. DNA damage and the cell's response is absolutely central to the development, treatment and survival of cancer; it has been extensively studied, and in the construction of this example we draw on a wide consensus that these are some of the lower-level mechanisms responsible for DNA

condition. While of course a great deal is still unknown about cancer, we choose DNA damage response—one of the most intensively studied features of cancer, generating vast amounts of data—because we know enough about these mechanisms to make modelling them using RBNs a plausible strategy, and one that is likely to be useful.

So we believe both that the Recursive Causal Markov Condition will be satisfied in this case, and that our model showing the effect of the lower-level variables D , G , P , O , E , T and I in terms of a single recursive variable describing DNA condition will be fruitful. If the Causal Markov Condition is also satisfied in each net, then the Recursive Causal Markov Condition will be satisfied. The RBN formalism will be applicable, and useful.

There are various general advantages of representing mechanisms in this way. Here, we discuss three.

The first advantage is that this gives a quantitative description of a mechanism. If our intention were to *replace* the more usual qualitative description of a mechanism in words and diagrams with an RBN, there would be both losses and gains. Information from the qualitative description would be lost. Most activities are represented in the same ways using arrows in RBNs, so that richer descriptions of, for example, ‘pushing’ as different from ‘pulling’ would be lost. But our intention is not to replace the descriptions of the biological literature, so the richness of these descriptions is retained. There is only a gain: in the precision and subtleties of quantitative relationships that can now be represented. The two representations put together—diagrams *plus* RBN—are far more powerful than either alone. A flat BN might be useful for these purposes, but a hierarchical RBN will do more. In the next section we discuss other quantitative models which might be used to model mechanisms, and explain the advantages of RBNs.

The second advantage lies in the consistency of the story about explanation, prediction and control that the RBN framework adds to traditional ways of describing mechanisms. The RBN framework enforces consistency. BNs are designed to represent causal relationships and allow prediction and intervention. The RBN formalism extends that to allow intervention and prediction across levels—using precisely those hierarchical relationships that are discovered when mechanisms are investigated. For more complex mechanisms, even the explanatory story given in traditional qualitative descriptions is enhanced by getting these quantitative relationships right. It is increasingly vital in, for example, both proteomics and genomics to represent very precisely not just what proteins are produced, or genes expressed, but exactly what levels and over what time periods.

The third advantage of RBNs is economy. A single RBN encapsulates many BNs. This economy of representation matters, allowing the relationships between the BNs to be represented in a more readily understandable and cognitively manipulable way. Formally equivalent mathematical formalisms are not equal in their representational usefulness, and this economy of representation might lead to fruitful progress.

In sum, RBNs are useful for representing mechanisms as they are discussed in the complex-systems mechanisms literature. There are several benefits to representing mechanisms this way.

5. *Alternative models for prediction, explanation and control*

In this section we further argue for the suitability of the RBN formalism by comparing it with alternative models. Recall, the reason why we opted for Recursive Bayesian Nets was that they satisfy three desiderata:

1. to model causal relations, so that the model can be useful for causal explanation and for control,
2. to model hierarchical mechanisms, so that the model can be useful for mechanistic explanation, and
3. to model probabilistic relations, so that the model can give accurate quantitative predictions.

Alternative models, notably differential equations, agent-based models and multilevel models, certainly perform pretty well in providing quantitative information, as they all go beyond mere qualitative description of phenomena. They look like good potential candidates for our purposes—and consequently like competitors to RBNs—because they are intended to model hierarchies or they can easily be adjusted and interpreted in order to do so. Yet, as we shall show, they all have trouble with at least one of the other two desiderata.

Differential equations. Much can be learnt from a comparison with differential equations. This formalism has been developed in order to model complex systems, and in particular their dynamics, that is how they evolve over time. System analysis (see for instance [Bunge \(2000, 1979\)](#), [Simon \(1969\)](#), [von Bertalanffy \(1968\)](#)) exploits this formalism because the mathematics of differential equations allows the scientists to include in the model a large number of variables at a time, as well as their relations, including back-and-forth interactions between variables.

But system analysis has trouble modelling the kind of mechanisms we are interested in. This is because, in differential equations, variables have no explicit causal role because such equations do not on their own capture any causal asymmetry. Nor is there a natural mechanistic interpretation of differential equations (on this point see, e.g., [Russo \(2010\)](#)). Granted, some differential equations might be interpreted as modelling hierarchies—in principle nothing prevents the inclusion of variables at different levels—but we need *mechanistic* hierarchy. Decomposing and describing the mechanisms in terms of entities and activities makes the model specifically *mechanistic* in that it provides details about the functioning and about the relations between the parts.

Agent-based models. Another type of model that is increasingly receiving attention in the scientific community is the *agent-based model*. Agent-based models are computational models that aim to simulate behaviour of individual agents in a system. The goal is, ultimately, to assess the effects of individual agents on the system as a whole. In particular, agent-based models aim to reproduce emergent behaviour from micro- to macro-level. There has been growing interest in, and more and more applications of, agent-based models since the mid-1990s, especially in business and technology problems. More

recently, there have been attempts to model cellular behaviour in cancer science (see for instance [Chen and Hardoon \(2010\)](#), [Chen \(2009\)](#)). Agent-based models attempt to reproduce the behaviour of individual agents such that the behaviour will give rise to emergent properties/behaviours at a higher level of organisation, thus implying that the system *as a whole* is more than the sum of its individual components. But this methodology has some well-known limitations, restricting its use for prediction, explanation and control.

First, since *simulation* is used to derive the relevant emergent properties, it is typical that so many simplifying assumptions need to be made that it is implausible that the resulting model is an accurate enough model of a complex system for it to make reliable *quantitative* predictions. Second, the direction of explanation is rather limited: agent-based models attempt to simulate behaviour that emerges from the lower level to the higher level, but not the other way round. The lower level is always taken to be at the individual level, and the resulting emergent behaviour to be at a higher, aggregate level. So only two levels are involved and the direction of explanation is always bottom-up. Third, the direction of control is similarly restricted: an agent-based model tells one how to control the emergent phenomenon by changing the parameters or constraints operating on the individuals, but that is the only kind of control question that the model can help answer.

One advantage of agent-based models is that they may be able to capture spatial aspects of phenomena—in our case the spatio-temporal evolution of tumours. RBNs, admittedly, are not able to do this. Yet, RBNs are not meant to be the *only* model for explanation, prediction or control. RBNs are meant to *augment*, rather than replace, pictorial representations of tumour development, which can capture spatial information.

Multilevel models. The most interesting comparison is perhaps with *multilevel* models. The reason is that the primary motivation for developing this formalism was exactly the recognition of a hierarchical structure of systems in virtually all domains.

Multilevel models are an extension of structural equation models (and therefore, in a loose sense, of Bayesian nets too, since Bayesian net models and structural equation models are very similar) that allow one to formalise the following idea: systems, particularly social systems (where this formalism was first developed and applied) have a hierarchical structure. Multilevel analysis posits that society is essentially hierarchically organised (see [Courgeau \(2003, 2007\)](#), [Goldstein \(2003\)](#), [Snijders and Bosker \(2004\)](#)). In the social sciences, various levels of aggregation are possible. For instance, economics is interested in the production, distribution and consumption of wealth; however, there is no a priori restriction on whether analyses concern individuals, markets, firms, or nations. Likewise, demography has no a priori restriction to the level of family, local population, or national population. Multilevel analysis recognises the existence of a multiplicity of levels and tries to specify the relations holding among individuals and/or among different levels of aggregation. The underlying idea is that there is a reciprocal influence of the individual on the group and of the group on the individual. While traditional statistical methods use different models to analyse data at the individual or aggregate level, multilevel models permit the analysis of such hierarchical structures in the framework of a single statistical model. Failing to recognise the hierarchical structure

of social systems leads to incorrect identification of causal relationships and even to the well-known ecological and atomistic fallacies (Robinson 1950).

Multilevel models were first developed in the social sciences. Here is an example of the kind of phenomena social scientists wanted to model. A pioneering case study concerns migration behaviour in Norway (Courgeau (1994, 2003)), where multilevel models were used to explain migration rates of farmers. Multilevel modelling offered a successful explanation of a phenomenon in which an individual-level model and separate aggregate model delivered apparently opposite results. On the one hand, an individual-level model explains individual-level outcomes by individual-level explanatory variables. For instance, in this model we can represent the probability of an individual migrating conditional on the individual characteristic of being/not being a farmer. On the other hand, the aggregate model explains aggregate-level outcomes through aggregate-level variables. For instance, we explain the percentage of migrants in a region through the percentage of people in the population having a certain occupational status (e.g., the percentage of farmers). Multilevel models make claims across the levels—from the aggregate-level to the individual-level and vice versa; that is, the multilevel model of migration in Norway explains the probability of migration of a non-farmer (individual) through the percentage of farmers in the same region (aggregate): in the Norwegian countryside farmers have a low chance of migration, but, as the percentage of farmers increases, they will start to migrate more thus raising the overall probability of migration in the whole population. In other words, the multilevel model allows the proportion of farmers in the population to be a cause of a particular farmer migrating.

Multilevel models go some way toward satisfying the three desiderata we identified: they model probabilistic and causal relations, and they model hierarchical systems. However, multilevel models were not devised to model mechanisms. Now, arguably, causal models in social science, which do include multilevel models, *ought* to be used to model mechanisms insofar as they aim to *explain* a phenomenon, since the explanatory job is carried out by the mechanism being modelled (Russo 2009, ch. 6). Nevertheless, it seems that in cancer science multilevel models are rather used to reconcile results between various levels of aggregation (see e.g., Delsanto et al. (2008)), or to check distribution of disease (see e.g., Short et al. (2002)), but not to model mechanisms that are hierarchical in themselves. In other words, multilevel models are used as means to avoid ecological and atomistic fallacies in drawing causal conclusions from analyses that use data at only one level, and the statistical machinery of multilevel models is not *ipso facto* a legitimate representation of a hierarchical mechanism.

It is also important to bear in mind that multilevel models can model hierarchical structures, but only to a limited extent. In fact, the present state of the art is that multilevel models cannot put *aggregate* variables as *response variables*, while RBNs can. Simply put, this means that, at present, this formalism models how individual-level variables (plus, if needed, aggregate variables) have an influence on aggregate variables.

In sum, RBNs suit our goals because they can account, in a single modelling framework, for the three aspects we took to be essential: modelling causal relations, modelling probabilistic relations, and modelling mechanistic hierarchies. RBNs proved to be more

suitable than BNs for our purposes because they are hierarchical. RBNs are better suited to modelling hierarchical mechanisms than system analysis because they can capture causal connections. RBNs have wider scope than agent-based models which just capture mechanisms involving multi-agent interactions. Finally, RBNs have the advantage of modelling mechanistic constitution relations rather than levels of aggregation, unlike multilevel analysis.

6. Conclusion

The need for quantitative models of hierarchical mechanisms is recognised by biological scientists when they make such claims as:

The best test of our understanding of cells will be to make quantitative predictions about their behaviour and test them. This will require detailed simulations of the biochemical processes taking place within the modules. But making predictions is not synonymous with understanding. We need to develop simplifying, higher-level models and find general principles that will allow us to grasp and manipulate the functions of biological modules. ... Connecting different levels of analysis—from molecules, through modules, to organisms—is essential for an understanding of biology that will satisfy human curiosity. (Hartwell et al. 1999, C52)

This paper has been an attempt to fill this need by introducing the recursive Bayesian network (RBN) formalism and applying it to the modelling of mechanisms. The RBN formalism, we maintain, provides an integrated modelling formalism for explanation, prediction and control. The formalism can be applied to modelling cancer mechanisms, where hierarchy is ubiquitous and vast amounts of data are increasingly available. This kind of model also shows how the current philosophical conception of a mechanism can be further developed, by integrating a quantitative description of the interaction between variables with the philosophically more familiar structural description of hierarchical relations between activities and entities. It is *prima facie* preferable to other kinds of models used in this context, such as agent-based models, differential-equation models, and multilevel models.

But there is much more to do before this kind of model can be routinely applied. First, formal work is needed to develop efficient methods for performing inference in an RBN. While standard Bayesian net methods for inference can be applied to RBNs, there is an added computational consideration for RBNs arising from the multiplicity of flattenings—for a single RBN there are 2^k flattenings, assuming k binary network variables—and work needs to be done to ensure that this extra complexity can be kept under control. Second, it would be helpful to develop the cancer application more fully, in particular on real data, to test the RBN formalism and validate the model. Third, the question of how to build RBNs in the face of incomplete mechanistic, causal and probabilistic knowledge needs thorough investigation. Known causal structure can be imposed on data to help make sense of the remaining relationships, and there are existing methods for doing this using BNs. The RBN formalism offers a parallel way of imposing

known structure on data—by appealing to hierarchical structure, information which is currently normally thrown away. One interesting issue is whether mechanisms currently represented using newly-born representations such as the Systems Biology Graphical Notation (Novere et al. 2009) can be automatically translated into RBN hierarchical structure.

A satisfactory answer to this third question would open up the use of RBNs for mechanism discovery. By treating hierarchical structure in a formally equivalent way to causal structure, this formalism might allow us to extend known methods for extracting unknown causal structure from data to extracting unknown hierarchical structure. This is an exciting possibility for studying cancer, where both causal and hierarchical structure are still to be discovered in the areas opened up by new technology in the last decade. But it is also of relevance to any field in which there is a great deal still to be found about both causal and hierarchical structure, and yet plenty of data available, such as in proteomics and genomics.

Acknowledgements

We are grateful to the Leverhulme Trust and the British Academy for supporting this research. We are also grateful to Amos Folarin, May Yong and Sylvia Nagl of the UCL Cancer Institute for valuable discussions, as well as Jan Lemeire, Dan Mulvihill, Adam White and four anonymous referees.

REFERENCES

- Bechtel, W. and A. Abrahamsen 2005. Explanation: a mechanist alternative. *Studies in the History and Philosophy of the Biological and Biomedical Sciences* 36: 421–441.
- Bunge, M. 1979. *A world of systems. Treatise on basic philosophy, Ontology 2*. Dordrecht: Reidel Publishing Company.
- Bunge, M. 2000. Systemism: the alternative to individualism and holism. *Journal of Socio-Economics* 29: 147–157.
- Cartwright, N. 2007. *Hunting causes and using them*. Cambridge: Cambridge University Press.
- Chen, C. C. 2009. *Complex event types for agent-based simulation*. Ph. D. thesis, University College London.
- Chen, C. C. and D. R. Hardoon 2010. Learning from multi-level behaviours in agent-based simulations: A systems biology application. *Journal of Simulation* 4: 196–203.
- Courgeau, D. 1994. Du groupe à l'individu: l'exemple des comportements migratoires. *Population* 1: 7–26.
- Courgeau, D. (Ed.) 2003. *Methodology and epistemology of multilevel analysis. Approaches from different social sciences*. Dordrecht: Kluwer.
- Courgeau, D. 2007. *Multilevel synthesis: from the group to the individual*. Dordrecht: Springer.
- Craver, C. F. 2007. *Explaining the brain*. Oxford: Oxford University Press.

- Darden, L. 2002, September. Strategies for discovering mechanisms: Schema instantiation, modular subassembly, forward/backward chaining. *Philosophy of Science* 69: S354–S365.
- Darden, L. 2006. *Reasoning in Biological Discoveries*. Cambridge: Cambridge University Press.
- Delsanto, P., C. Condat, N. Pugno, A. Gliozzi, and M. Griffa 2008. A multilevel approach to cancer growth modeling. *Journal of Theoretical Biology* 250: 16–24.
- Friedman, N., M. Linial, I. Nachman, and D. Pe'er 2000. Using Bayesian networks to analyze expression data. *Journal of Computational Biology* 7(3/4): 601–620.
- Friedman, N., K. Murphy, and S. Russell 1998. Learning the structure of dynamic probabilistic networks. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence*, 139–147.
- Glennan, S. 2002. Rethinking mechanistic explanation. *Philosophy of Science. Supplement: Proceedings of the 2000 Biennial Meeting of the Philosophy of Science Association. Part II: Symposia Papers (Sep., 2002)* 69(3): S342–S353.
- Glymour, C. and P. W. Cheng 1998. Causal mechanism and probability: A normative approach. In *Rational Models of Cognition*, eds. M. Oaksford and N. Chater. Oxford: Oxford University Press.
- Goldstein, H. 2003. *Multilevel statistical models*. London: Arnold. Kendall's Library of Statistics (volume 3).
- Gyftodimos, E. and P. Flach 2002. Hierarchical Bayesian networks: a probabilistic reasoning model for structured domains. In *Proceedings of the ICML-2002 Workshop on Development of Representations*, eds. E. de Jong and T. Oates, 23–30. University of New South Wales.
- Hartwell, L. H., J. J. Hopfield, S. Leibler, and A. W. Murray 1999. From molecular to modular cell biology. *Nature* 402: C47–C52.
- Illari, P. M. and J. Williamson 2010. Function and organization: comparing the mechanisms of protein synthesis and natural selection. *Studies in the History and Philosophy of the Biological and Biomedical Sciences* 41: 279–91.
- Jaeger, M. 2001. Complex probabilistic modeling with recursive relational Bayesian networks. *Annals of Mathematics and Artificial Intelligence* 32(1-4): 179–220.
- King, R. J. B. 2000. *Cancer biology* (2nd ed.). Pearson, Singapore.
- Koller, D. and A. Pfeffer 1997. Object-oriented Bayesian networks. In *Proceedings of the 13th Annual Conference on Uncertainty in Artificial Intelligence*, 302–313.
- Lazebnik, Y. 2002. Can a biologist fix a radio?—or, what I learned while studying apoptosis. *Cancer Cell* 2: 179–182.
- Legewie, S., N. Blüthgen, and H. Herzog 2006. Mathematical modeling identifies inhibitors of apoptosis as mediators of positive feedback and bistability. *PLoS Computational Biology* 2(9): 1061–1073.
- Machamer, P., L. Darden, and C. Craver 2000. Thinking about mechanisms. *Philosophy of Science* 67: 1–25.
- Mai, Z. and H. Liu 2009. Boolean network-based analysis of the apoptosis network: Irreversible apoptosis and stable surviving. *Journal of Theoretical Biology* 259: 760–769.

- Neapolitan, R. E. 1990. *Probabilistic reasoning in expert systems: theory and algorithms*. New York: Wiley.
- Neapolitan, R. E. 2004. *Learning Bayesian networks*. Upper Saddle River NJ: Pearson / Prentice Hall.
- Novere, N. L., M. Hucka, H. Mi, S. Moodie, F. Schreiber, A. Sorokin, E. Demir, K. Wegner, M. I. Aladjem, S. M. Wimalaratne, F. T. Bergman, R. Gauges, P. Ghazal, H. Kawaji, L. Li, Y. Matsuoka, A. Villeger, S. E. Boyd, L. Calzone, M. Courtot, U. Dogrusoz, T. C. Freeman, A. Funahashi, S. Ghosh, A. Jouraku, S. Kim, F. Kolpakov, A. Luna, S. Sahle, E. Schmidt, S. Watterson, G. Wu, I. Goryanin, D. B. Kell, C. Sander, H. Sauro, J. L. Snoep, K. Kohn, and H. Kitano 2009. The systems biology graphical notation. *Nature Biotechnology* 27(8): 735–741.
- Pearl, J. 1988. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. San Mateo CA: Morgan Kaufmann.
- Pearl, J. 2000. *Causality: models, reasoning, and inference*. Cambridge: Cambridge University Press.
- Peña, J. M., J. A. Lozano, and P. Larrañaga 2002. Learning recursive Bayesian multinets for clustering by means of constructive induction. *Machine Learning* 47(1): 63–90.
- Robinson, W. S. 1950. Ecological correlations and the behavior of individuals. *American Sociological Review* 15: 351–357.
- Russo, F. 2009. *Causality and causal modelling in the social sciences. Measuring variations*. Methodos Series. New York: Springer.
- Russo, F. 2010. Are causal analysis and system analysis compatible approaches? *International Studies in the Philosophy of Science* 24(1): 67–90.
- Short, M., B. P. Carlin, and S. Bushhouse 2002. Using hierarchical spatial models for cancer control planning in Minnesota (united states). *Cancer Causes and Control* 13: 903–916.
- Simon, H. A. 1982 [1969]. *The sciences of the artificial*. Cambridge, Mass.: MIT Press.
- Snijders, T. A. B. and R. J. Bosker 2004. *Multilevel analysis. An introduction to basic and advanced multilevel modeling* (4th ed.). London: Sage.
- Spirtes, P., C. Glymour, and R. Scheines 1993. *Causation, Prediction, and Search* (Second (2000) ed.). Cambridge MA: MIT Press.
- Steel, D. 2008. *Across the boundaries. Extrapolation in biology and social science*. Oxford University Press.
- von Bertalanffy, L. 1968. *General system theory: foundations, development, applications*. New York: Braziller.
- Weinberg, R. A. 2007. *The biology of cancer*. Garland Science, Taylor & Francis Group, New York.
- Williamson, J. 2005. *Bayesian nets and causality: philosophical and computational foundations*. Oxford: Oxford University Press.
- Williamson, J. and D. Gabbay 2005. Recursive causality in Bayesian networks and self-fibring networks. In *Laws and models in the sciences*, ed. D. Gillies, 173–221. London: King's College Publications. With comments, 223–245.

Woodward, J. 2002. What is a mechanism? A counterfactual account. *Philosophy of Science* 69: S366–S377.

Lorenzo CASINI is a PhD student in Philosophy at the University of Kent, working on theories of causality as applied to complex systems.

ADDRESS: Department of Philosophy, University of Kent, Canterbury, UK. E-mail: l.casini@kent.ac.uk

Phyllis McKay ILLARI is Visiting Research Fellow in Philosophy at the University of Kent, working on mechanisms, causality and the relationship between the two.

ADDRESS: Department of Philosophy, University of Kent, Canterbury, UK. E-mail: p.mckay@kent.ac.uk

Federica RUSSO is Research Associate in Philosophy at the University of Kent, working on causality in the social and health sciences.

ADDRESS: Department of Philosophy, University of Kent, Canterbury, UK. E-mail: f.russo@kent.ac.uk

Jon WILLIAMSON is Professor of Reasoning, Inference and Scientific Method in Philosophy at the University of Kent, working on causality, probability and logic.

ADDRESS: Department of Philosophy, University of Kent, Canterbury, UK. E-mail: j.williamson@kent.ac.uk