Credal networks under epistemic irrelevance using sets of desirable gambles

Jasper De Bock
Ghent University, Belgium
jasper.debock@ugent.be

Gert de Cooman
Ghent University, Belgium
gert.decooman@ugent.be

Abstract

We present a new approach to credal networks, which are graphical models that generalise Bayesian nets to deal with imprecise probabilities. Instead of applying the commonly used notion of strong independence, we replace it by the weaker notion of epistemic irrelevance. We show how assessments of epistemic irrelevance allow us to construct a global model out of given local uncertainty models, leading to an intuitive expression for the so-called irrelevant natural extension of a network. In contrast with Cozman [2], who introduced this notion in terms of credal sets, our main results are presented using the language of sets of desirable gambles. This has allowed us to derive a number of useful properties of the irrelevant natural extension. It has powerful marginalisation properties and satisfies all graphoid properties but symmetry, both in their direct and reverse forms.

Keywords. Credal networks, epistemic irrelevance, sets of desirable gambles, graphoid properties, irrelevant natural extension, lower previsions, coherence.

1 Introduction

In his overview paper [2], Cozman discussed and compared a number of different extensions for so-called credal networks, which generalise standard Bayesian networks to allow for imprecise probability assessments.

One of these extensions is the so-called irrelevant natural extension, which captures that the non-parent non-descendants of any variable in the network are epistemically irrelevant to that variable given the value of its parents. Cozman argues that of all the possible extensions, this irrelevant natural extension is perhaps the most appealing one. Nevertheless, it has thus far received little attention.

The present paper tries to remedy this situation by providing a firm theoretical foundation for the irrelevant natural extension of a network, leading to, amongst other things, a powerful marginalisation property and a proof that it satisfies all graphoid properties but symmetry.

The main results are stated using the theory of sets of desirable gambles, which we introduce in Section 2. We go on to introduce and discuss important concepts such as directed acyclic graphs and epistemic irrelevance in Section 3, and use these in Section 4 to show how assessments of epistemic irrelevance can be combined with given local sets of desirable gambles to construct a joint model. We call this the *irrelevant natural extension* of the credal network and prove that it is the most conservative coherent model that extends the local models and expresses all conditional irrelevancies encoded in the network. In Section 5 we present a powerful marginalisation property, and in Section 6, we use an asymmetric version of D-separation to show that the irrelevant natural extension satisfies all graphoid properties except symmetry, both in their direct and reverse forms. Finally, Section 7 establishes a connection between the sets of desirable gambles approach to credal networks under epistemic irrelevance that we presented in this paper, and a similar approach using coherent lower previsions.

2 Sets of desirable gambles

Consider a variable X taking values in some non-empty and finite set $\mathscr{X}$. Beliefs about the possible values this variable may assume can be modelled in various ways: probability mass functions, credal sets and coherent lower previsions are only a few of the many options. We choose to adopt a different approach, using sets of desirable gambles. We will model a subject's beliefs regarding the value of a variable X by means of his behaviour: which gambles (or bets) on the unknown value of X would our subject strictly prefer to the status quo (the zero gamble).

Although they are not as well known as other (imprecise) probability models, sets of desirable gambles have definite advantages. To begin with, they are more expressive than both credal sets and lower previsions. For example, they are easily able to deal with such things as conditioning on events with probability zero, which tends to be much more involved when using other imprecise probability models. Secondly, they have the advantage of being operational, meaning that there is a practical way of constructing a model that represents the subject's beliefs. For sets of desirable gambles this can be done by offering the subject certain gambles and asking him whether or not he strictly prefers them to the status quo. And finally, our experience tells us that it is usually easier to construct proofs in the language of coherent sets of desirable gambles than in other, perhaps more familiar languages. We give a brief survey of the basics of sets of desirable gambles and refer to Refs. [7, 1, 12] for more details and further discussion.

2.1 Desirable gambles

A gamble f is a real-valued map on $\mathscr{X}$ that is interpreted as an uncertain reward. If the value of the variable X turns out to be x, the (possibly negative) reward is $f(x)$. A non-zero gamble is called *desirable* to a subject if he strictly prefers to zero the transaction in which (i) the actual value x of the variable is determined, and (ii) he receives the reward $f(x)$. The zero gamble is therefore not considered to be desirable.

We model a subject's beliefs regarding the possible values $\mathscr{X}$ that a variable X can assume by means of a set $\mathscr{D}$ of desirable gambles—some subset of the set $\mathscr{G}(\mathscr{X})$ of all gambles on $\mathscr{X}$. For any two gambles f and g in $\mathscr{G}(\mathscr{X})$, we say that $f \geq g$ if $f(x) \geq g(x)$ for all x in $\mathscr{X}$ and $f > g$ if both $f \geq g$ and $f \neq g$. We use $\mathscr{G}(\mathscr{X})_{>0}$ to denote the set of all gambles $f \in \mathscr{G}(\mathscr{X})$ for which $f > 0$ and $\mathscr{G}(\mathscr{X})_{\leq 0}$ to denote the set of all gambles $f \in \mathscr{G}(\mathscr{X})$ for which $f \leq 0$. As a special kind of gambles we consider *indicators* $\mathbb{I}_A$ of events $A \subseteq \mathscr{X}$. $\mathbb{I}_A$ is equal to 1 if the event A occurs—the variable X assumes a value in A—and zero otherwise.

2.2 Coherence

In order to represent a rational subject's beliefs about the values a variable can assume, a set $\mathscr{D} \subseteq \mathscr{G}(\mathscr{X})$ of desirable gambles should satisfy some rationality requirements. If these requirements are met, we call the set $\mathscr{D}$ *coherent*. We require that for all f, f_1, $f_2 \in \mathscr{G}(\mathscr{X})$ and all real $\lambda > 0$:

D1. if $f \leq 0$ then $f \notin \mathscr{D}$;

D2. if $f > 0$ then $f \in \mathscr{D}$;

D3. if $f \in \mathscr{D}$ then $\lambda f \in \mathscr{D}$; [scaling]

D4. if $f_1, f_2 \in \mathscr{D}$ then $f_1 + f_2 \in \mathscr{D}$. [combination]

Requirements D3 and D4 turn $\mathscr{D}$ into a convex cone: $\operatorname{posi}(\mathscr{D}) = \mathscr{D}$, where we use the positive hull operator 'posi' that generates the set of finite strictly positive linear combinations of elements of its argument set:

$$\operatorname{posi}(\mathscr{D}) := \left\{ \sum_{k=1}^{n} \lambda_k f_k \colon f_k \in \mathscr{D}, \lambda_k \in \mathbb{R}_0^+, n \in \mathbb{N}_0 \right\}.$$

Here $\mathbb{R}_0^+$ is the set of all (strictly) positive real numbers, and $\mathbb{N}_0$ the set of all natural numbers (zero not included).

3 Credal networks

3.1 Directed acyclic graphs

A directed acyclic graph (DAG) is a graphical model that is well known for its use in Bayesian networks. It consists of a finite set of nodes (vertices), joined into a network by a set of directed edges, each edge connecting one node with another. Since this directed graph is assumed to be acyclic, it is not possible to follow a sequence of edges from node to node and end up at the same node one started out from.

We will call G the set of nodes s associated with a given DAG. For two nodes s and t, if there is a directed edge from s to t, we denote this as $s \to t$ and say that s is a *parent* of t and t is a *child* of s. A single node can have multiple parents and multiple children. For any node s, its set of parents is denoted by $P(s)$ and its set of children by $C(s)$. If a node s has no parents, $P(s) = \emptyset$, and we call s a *root node*. If $C(s) = \emptyset$, then we call s a *leaf*, or *terminal node*.

Two nodes s and t are said to have a *path* between them if one can start from s, follow the edges of the DAG regardless of their direction and end up in t. In other words: one can find a sequence of nodes $s = s_1, \dots, s_n = t$, $n \geq 1$, such that for all $i \in \{1, \dots, n-1\}$ either $s_i \to s_{i+1}$ or $s_i \leftarrow s_{i+1}$. If this sequence is such that $s_i \to s_{i+1}$ for all $i \in \{1, \dots, n-1\}$ (all edges in the path point away from s), we say that there is a *directed path* from s to t and write $s \sqsubseteq t$. In that case we also say that s *precedes* t. If $s \sqsubseteq t$ and $s \neq t$, we say that s *strictly precedes* t and write $s \sqsubset t$. For any node s, we denote its set of *descendants* by $D(s) := \{t \in G \colon s \sqsubset t\}$ and its set of *non-parent non-descendants* by $N(s) := G \setminus (P(s) \cup \{s\} \cup D(s))$. We also use the shorthand notation $PN(s) := P(s) \cup N(s) = G \setminus (\{s\} \cup D(s))$ to refer to the so-called *non-descendants* of s.

We extend these notions to subsets of G in the following way. For any $K \subseteq G$, $P(K) := (\bigcup_{s \in K} P(s)) \setminus K$ is its set of parents and $D(K) := (\bigcup_{s \in K} D(s)) \setminus K$ is its set of descendants. The non-parent non-descendants of K are given by $N(K) := G \setminus (P(K) \cup K \cup D(K)) = \bigcap_{s \in K} N(s)$, and we also define $PN(K) := P(K) \cup N(K)$. This last set cannot be referred to as the non-descendants of K since $P(K)$ and $D(K)$ are not necessarily disjoint.

Special subsets of G that we will consider, are the closed ones: we call a set $K \subseteq G$ *closed* if for all $s, t \in K$ and any $k \in G$ such that $s \sqsubseteq k \sqsubseteq t$, it holds that $k \in K$. For closed $K \subseteq G$, $P(K) \cap D(K) = \emptyset$ and therefore $PN(K) = G \setminus (K \cup D(K))$, which means that for closed K, $PN(K)$ can rightfully be referred to as the non-descendants of K.

With any subset K of G, we can associate a so-called *sub-DAG* of the DAG that is associated with G. The nodes of this sub-DAG are the elements of K and the directed edges of this sub-DAG are those edges in the original DAG that

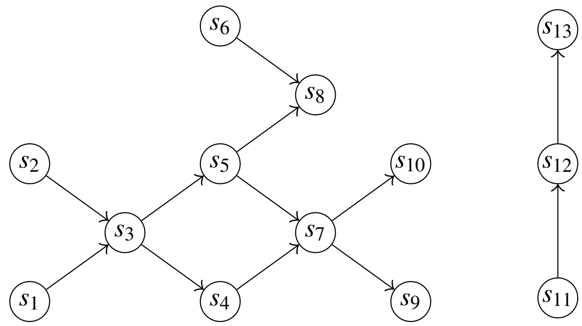

$$G = \{s_1, s_2, s_3, s_4, s_5, s_6, s_7, s_8, s_9, s_{10}, s_{11}, s_{12}, s_{13}\}$$

Figure 1: Example of a directed acyclic graph (DAG)

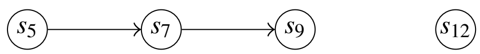

Figure 2: Example of a sub-DAG

connect elements in K. For a sub-DAG that is associated with some subset K of G, we will use similar definitions as those for the original DAG, adding the subset K as an index. As an example: for all $k \in K$, we denote by $P_K(k)$ the parents of k in the sub-DAG that is associated with the nodes in K. For all $K \subseteq G$ and $k \in K$, we have $P_K(k) = P(k) \cap K$ and $P(k) \setminus P_K(k) = P(k) \cap P(K)$.

Example 1. *Consider the DAG in Figure 1. For the node $s_7 \in G$, we find that $P(s_7) = \{s_4, s_5\}$, $D(s_7) = \{s_9, s_{10}\}$ and $N(s_7) = \{s_1, s_2, s_3, s_6, s_8, s_{11}, s_{12}, s_{13}\}$. For the closed subset $K = \{s_5, s_7, s_9, s_{12}\} \subset G$, we have $P(K) = \{s_3, s_4, s_{11}\}$, $D(K) = \{s_8, s_{10}, s_{13}\}$ and $N(K) = \{s_1, s_2, s_6\}$. The sub-DAG that corresponds to K is drawn in Figure 2. We find that $P_K(s_7) = \{s_5\}$, $D_K(s_7) = \{s_9\}$ and $N_K(s_7) = \{s_{12}\}$.* $\diamondsuit$

3.2 Variables and gambles on them

With each node s of the network, we associate a variable X_s assuming values in some non-empty finite set $\mathscr{X}_s$. We denote by $\mathscr{G}(\mathscr{X}_s)$ the set of all gambles on $\mathscr{X}_s$. We extend this notation to more complicated situations as follows. If S is any subset of G, then we denote by X_S the tuple of variables whose components are the X_s for all $s \in S$. This new joint variable assumes values in the finite set $\mathscr{X}_S := \times_{s \in S} \mathscr{X}_s$ and the corresponding set of gambles is denoted by $\mathscr{G}(\mathscr{X}_S)$. When $S = \emptyset$, we let $\mathscr{X}_\emptyset$ be a singleton. The corresponding variable $X_\emptyset$ can then only assume this single value, so there is no uncertainty about it. $\mathscr{G}(\mathscr{X}_\emptyset)$ can then be identified with the set $\mathbb{R}$ of real numbers. Generic elements of $\mathscr{X}_s$ are denoted by x_s or z_s and similarly for x_S and z_S in $\mathscr{X}_S$. Also, if we mention a tuple z_S, then for any $t \in S$, the corresponding element in the tuple will be denoted by z_t. We assume all variables in the network to be logically independent, meaning that the variable X_S may assume *all* values in $\mathscr{X}_S$, for all $\emptyset \subseteq S \subseteq G$.

We will use the simplifying device of identifying a gamble f_S on $\mathscr{X}_S$ with its *cylindrical extension* to $\mathscr{X}_U$, where $S \subseteq U \subseteq G$: the gamble f_U on $\mathscr{X}_U$ defined by $f_U(x_U) := f_S(x_S)$ for all $x_U \in \mathscr{X}_U$. For instance, if $\mathscr{K} \subseteq \mathscr{G}(\mathscr{X}_G)$, this allows us to consider $\mathscr{K} \cap \mathscr{G}(\mathscr{X}_S)$ as the set of those gambles in $\mathscr{K}$ that depend only on the variable X_S.

3.3 Modelling our beliefs about the network

Throughout, we consider sets of desirable gambles as models for a subject's beliefs about the values that certain variables in the network may assume. One of the main contributions of this paper, further on in Section 4, will be to show how to construct a joint model for our network, being a coherent set $\mathscr{D}_G$ of desirable gambles on $\mathscr{X}_G$.

From such a joint model, one can derive both conditional and marginal models [7, 6]. Let us start by explaining how to condition the global model $\mathscr{D}_G$. Consider an event $A_I \subseteq \mathscr{X}_I$, with $I \subseteq G$, and assume that we want to update the model $\mathscr{D}_G$ with the information that $X_I \in A_I$. This leads to the following updated set of desirable gambles:

$$\mathscr{D}_G \rfloor A_I := \left\{ f \in \mathscr{G}(\mathscr{X}_{G \setminus I}) \colon \mathbb{I}_{A_I} f \in \mathscr{D}_G \right\},$$

which represents our subject's beliefs about the value of the variable $X_{G \setminus I}$, conditional on the observation that X_I assumes a value in A_I. This definition is very intuitive, since $\mathbb{I}_{A_I} f$ is the unique gamble that is called off (is equal to zero) if $X_I \notin A_I$ and equal to f if $X_I \in A_I$. Since $\mathbb{I}_{\{x_\emptyset\}} = 1$, the special case of conditioning on the certain variable $X_\emptyset$ yields no problems: it amounts to not conditioning at all.

Marginalisation too is very intuitive in the language of sets of desirable gambles. Suppose we want to derive a marginal model for our subject's beliefs about the variable X_O, where O is some subset of G. This can be done by using the set of desirable gambles that belong to $\mathscr{D}_G$ but only depend on the variable X_O:

$$\mathrm{marg}_O(\mathscr{D}_G) := \left\{ f \in \mathscr{G}(\mathscr{X}_O) \colon f \in \mathscr{D}_G \right\} = \mathscr{D}_G \cap \mathscr{G}(\mathscr{X}_O).$$

Now let I and O be *disjoint* subsets of G and let A_I be any subset of $\mathscr{X}_I$. By sequentially applying the process of conditioning and marginalisation we can obtain conditional marginal models for our subject's beliefs about the value of the variable X_O, conditional on the observation that X_I assumes a value in A_I:

$$\mathrm{marg}_O(\mathscr{D}_G \rfloor A_I) = \left\{ f \in \mathscr{G}(\mathscr{X}_O) \colon \mathbb{I}_{A_I} f \in \mathscr{D}_G \right\}. \qquad (1)$$

Conditioning and marginalisation are special cases of Eq. (1); they can be obtained by letting $O = G \setminus I$ or $I = \emptyset$. If A_I is a singleton $\{x_I\}$, with $x_I \in \mathscr{X}_I$, we will use the shorthand notation $\mathrm{marg}_O(\mathscr{D}_G \rfloor x_I) := \mathrm{marg}_O(\mathscr{D}_G \rfloor \{x_I\})$.

Since coherence is trivially preserved under both conditioning and marginalisation, we find that if the joint model $\mathscr{D}_G$ is coherent, all the derived models will also be coherent.

3.4 Epistemic irrelevance

We now have the necessary tools to introduce one of the most important concepts for this paper, that of epistemic irrelevance. We describe the case of conditional irrelevance, as the unconditional version of epistemic irrelevance can easily be recovered as a special case.

Consider three disjoint subsets C, I, and O of G. When a subject judges X_I to be *epistemically irrelevant to X_O conditional on X_C*, denoted as $\mathrm{IR}(I,O|C)$, he assumes that if he knew the value of X_C, then learning in addition which value X_I assumes in $\mathscr{X}_I$ would not affect his beliefs about X_O. More formally put, he assumes for all $x_C \in \mathscr{X}_C$ and $x_I \in \mathscr{X}_I$ that:

$$\mathrm{marg}_O(\mathscr{D}_G \rfloor x_{C\cup I}) = \mathrm{marg}_O(\mathscr{D}_G \rfloor x_C).$$

Alternatively, a subject can make the even stronger statement that he judges X_I to be epistemically *subset-irrelevant* to X_O conditional on X_C, denoted as $\mathrm{SIR}(I,O|C)$. In that case, he assumes that if he knew the value of X_C, then receiving the additional information that X_I is an element of any non-empty subset A_I of $\mathscr{X}_I$ would not affect his beliefs about X_O. In other words, he assumes for all $x_C \in \mathscr{X}_C$ and all non-empty $A_I \subseteq \mathscr{X}_I$ that:

$$\mathrm{marg}_O(\mathscr{D}_G \rfloor \{x_C\} \times A_I) = \mathrm{marg}_O(\mathscr{D}_G \rfloor x_C).$$

Making a subset-irrelevance statement $\mathrm{SIR}(I,O|C)$ implies the corresponding irrelevance statement $\mathrm{IR}(I,O|C)$. Even stronger, it implies for all $I' \subseteq I$ that $\mathrm{IR}(I',O|C)$. The converse does not hold in general. However, as we will show further on, credal networks under epistemic irrelevance are a useful exception: although we define the joint model by imposing irrelevance, it will also satisfy subset-irrelevance. For the unconditional irrelevance case it suffices, in the discussion above, to let $C = \emptyset$. This makes sure the variable X_C has only one possible value, so conditioning on that variable amounts to not conditioning at all.

Irrelevance and subset-irrelevance can also be extended to cases where I, O and C are not disjoint, but $I \setminus C$ and $O \setminus C$ are. We then call X_I epistemically (subset-)irrelevant to X_O conditional on X_C provided that $X_{I\setminus C}$ is epistemically (subset-)irrelevant to $X_{O\setminus C}$ conditional on X_C. Although these cases are admittedly artificial, they will help us state and prove some of the graphoid properties further on.

3.5 Local uncertainty models

We now add *local uncertainty models* to each of the nodes s in our network. These local models are assumed to be given beforehand and will be used further on in Section 4 as basic building blocks for constructing a joint model for a given network.

If s is not a root node of the network, i.e. has a non-empty set of parents $P(s)$, then we have a conditional local model for every instantiation of its parents: for each $x_{P(s)} \in \mathscr{X}_{P(s)}$, we have a coherent set $\mathscr{D}_{s\rfloor x_{P(s)}}$ of desirable gambles on $\mathscr{X}_s$. It represents our subject's beliefs about the variable X_s conditional on its parents $X_{P(s)}$ assuming the value $x_{P(s)}$.

If s is a root node, i.e. has no parents, then our subject's local beliefs about the variable X_s are represented by an unconditional local model. It should be a coherent set of desirable gambles and will be denoted by $\mathscr{D}_s$. As was explained in Section 3.3, we can also use the common generic notation $\mathscr{D}_{s\rfloor x_{P(s)}}$ in this unconditional case, since for a root node s, its set of parents $P(s)$ is equal to the empty set $\emptyset$.

3.6 The interpretation of the graphical model

In classical Bayesian nets, the graphical structure is taken to represent the following assessments: for any node s, conditional on its parent variables, the associated variable is independent of its non-parent non-descendant variables

When generalising this interpretation to credal networks, the classical notion of independence gets replaced by a more general, imprecise-probabilistic notion of independence, which in the existing literature is usually chosen to be strong independence; see Ref. [3] for an overview of different approaches, including relevant references. Here, we will not do so: we choose to use the weaker, asymmetric notion of epistemic irrelevance, introduced in Section 3.4. In the special case of precise uncertainty models, both epistemic irrelevance and strong independence reduce to the classical notion of independence and the corresponding interpretations of the graphical network are equivalent to the one used in classical Bayesian networks.

In the present context, we therefore assume that the graphical structure of the network embodies the following conditional irrelevance assessments, turning the network into a *credal network under epistemic irrelevance*. Consider any node s in the network, its set of parents $P(s)$ and its set of non-parent non-descendants $N(s)$. Then *conditional on $X_{P(s)}$, $X_{N(s)}$ is assumed to be epistemically irrelevant to X_s*:

$$\mathrm{IR}(N(s),\{s\}|P(s)).$$

For a coherent set of desirable gambles $\mathscr{D}_G$ that describes our subject's global beliefs about all the variables in the network, this has the following consequences. For every $s \in G$ and all $x_{PN(s)} \in \mathscr{X}_{PN(s)}$, $\mathscr{D}_G$ must satisfy:

$$\mathrm{marg}_s(\mathscr{D}_G \rfloor x_{PN(s)}) = \mathrm{marg}_s(\mathscr{D}_G \rfloor x_{P(s)}). \tag{2}$$

4 Constructing a joint model

We now show how to construct a joint model for the variables in the network, and argue that it is the most conservative coherent model that extends the local models and expresses all conditional irrelevancies encoded in the network. But before we do so, let us provide some motivation. Suppose we have a global set of desirable gambles $\mathscr{D}_G$, how do we express that such a model is compatible with the assessments encoded in the network?

4.1 Defining properties of the joint model

We will require our joint model to satisfy the following four properties. First of all, we require that our global model should extend the local ones. This means that the local models derived from the global one by marginalisation should be equal to the given local models:

G1. The joint model $\mathscr{D}_G$ marginalises to the given local uncertainty models: $\operatorname{marg}_s(\mathscr{D}_G \rfloor x_{P(s)}) = \mathscr{D}_{s\rfloor x_{P(s)}}$ for all $s \in G$ and $x_{P(s)} \in \mathscr{X}_{P(s)}$.

The second requirement is that our model should reflect all epistemic irrelevancies encoded in the graphical structure of the network:

G2. $\mathscr{D}_G$ satisfies all equalities that are imposed by Eq. (2). In these equalities, the right hand side can be replaced by $\mathscr{D}_{s\rfloor x_{P(s)}}$ due to requirement G1.

The third requirement is that our model should be coherent:

G3. $\mathscr{D}_G$ is coherent (satisfies requirements D1–D4).

Since requirements G1–G3 do not determine a unique global model, we impose a final requirement to ensure that all inferences we make on the basis of our global models are as conservative as possible, and are therefore based on no other considerations than what is encoded in the network:

G4. $\mathscr{D}_G$ is the smallest set of desirable gambles on $\mathscr{X}_G$ satisfying requirements G1–G3: it is a subset of any other set that satisfies them.

We will now show how to construct the unique global model $\mathscr{D}_G$ that satisfies all of the four requirements G1–G4.

4.2 An intuitive expression for the joint model

Let us start by looking at a single given marginal model $\mathscr{D}_{s\rfloor z_{P(s)}}$ and investigate some of its implications for the joint model $\mathscr{D}_G$. Consider any node s and fix values $z_{P(s)}$ and $z_{N(s)}$ for its parents and non-parent non-descendants. Due to requirements G1 and G2, any gamble $f \in \mathscr{D}_{s\rfloor z_{P(s)}}$ should also be an element of $\operatorname{marg}_s(\mathscr{D}_G \rfloor z_{PN(s)})$, which by definition means that $\mathbb{I}_{\{z_{PN(s)}\}} f \in \mathscr{D}_G$. Inspired by this observation, we introduce the following set of gambles on $\mathscr{X}_G$:

$$\mathscr{A}_G^{\mathrm{irr}} := \left\{ \mathbb{I}_{\{z_{PN(s)}\}} f \colon s \in G,\ z_{PN(s)} \in \mathscr{X}_{PN(s)},\ f \in \mathscr{D}_{s\rfloor z_{P(s)}} \right\}.$$

It should now be clear that $\mathscr{A}_G^{\mathrm{irr}}$ must be a subset of our joint model $\mathscr{D}_G$.

Proposition 1. *$\mathscr{A}_G^{\mathrm{irr}}$ is a subset of any joint model $\mathscr{D}_G$ that satisfies requirements* G1 *and* G2*.*

Since our eventual joint model should also be coherent (satisfy requirement G3), and thus in particular should be a convex cone, we can derive the following corollary.

Corollary 2. $\operatorname{posi}(\mathscr{A}_G^{\mathrm{irr}})$ *is a subset of any joint model $\mathscr{D}_G$ that satisfies requirements* G1–G3*.*

We now suggest the following expression for the joint model describing our subject's beliefs about the variables in the network:

$$\mathscr{D}_G^{\mathrm{irr}} := \operatorname{posi}(\mathscr{A}_G^{\mathrm{irr}}). \tag{3}$$

We will refer to $\mathscr{D}_G^{\mathrm{irr}}$ as the *irrelevant natural extension* of the local models $\mathscr{D}_{s\rfloor x_{P(s)}}$. Since we know from Corollary 2 that it is guaranteed to be a subset of the joint model we are looking for, we propose it as a candidate for the joint model itself. In the next section, we set out to prove that $\mathscr{D}_G^{\mathrm{irr}}$ is indeed the unique joint model satisfying all four requirements G1–G4.

We would like to point out that $\mathscr{D}_G^{\mathrm{irr}}$ is a generalisation of the so-called *independent natural extension* of a number of unconditional marginal models [6, Section 7]. This special case corresponds to a DAG that has no edges, consisting of a finite amount of disconnected nodes [6, Section 10]. Quite a few of the results obtained further on can therefore be regarded as generalisations of those in Ref. [6].

4.3 Justifying our expression for the joint model

We start by proving a number of useful properties of $\mathscr{D}_G^{\mathrm{irr}}$.

Proposition 3. *A gamble $f \in \mathscr{G}(\mathscr{X}_G)$ is an element of $\mathscr{D}_G^{\mathrm{irr}}$ if and only if it can be written as:*

$$f = \sum_{s \in G} \sum_{z_{PN(s)} \in \mathscr{X}_{PN(s)}} \mathbb{I}_{\{z_{PN(s)}\}} f_{s,z_{PN(s)}},$$

where $f_{s,z_{PN(s)}} \in \mathscr{D}_{s\rfloor z_{P(s)}} \cup \{0\}$ for every $s \in G$ and all $z_{PN(s)} \in \mathscr{X}_{PN(s)}$, and at least one of them is non-zero.

Proposition 4. *$\mathscr{G}(\mathscr{X}_G)_{>0}$ is a subset of $\mathscr{D}_G^{\mathrm{irr}}$.*

These two propositions serve as a first step towards the following coherence result, which states that our joint model $\mathscr{D}_G^{\mathrm{irr}}$ satisfies requirement G3.

Proposition 5. *$\mathscr{D}_G^{\mathrm{irr}}$ satisfies requirement* G3*: it is a coherent set of desirable gambles.*

Our proof for this result has an interesting feature that deserves to be borne out. The crucial step hinges on the assumption that if the local models of our network were precise probability mass functions, we would be able to construct a joint probability mass function that satisfies all irrelevancies (in that case independencies) that are encoded in our network. Since the precise version of a credal net under epistemic irrelevance is a classical Bayesian network, this assumption is indeed true. What we believe is useful about this approach, is that it can be extended to credal networks with irrelevance assumptions that differ from the ones we impose in the present article, as long as the assumption above is satisfied. In this way, it enables us to use existing coherence results for precise networks to prove their counterparts for credal networks.

Next, we turn to an important factorisation result that is essential in order to prove that our joint model extends the local models and expresses all conditional irrelevancies encoded in the network, and therefore satisfies G1 and G2.

Proposition 6. *Fix arbitrary $s \in G$, $x_{P(s)} \in \mathscr{X}_{P(s)}$ and $g \in \mathscr{G}(\mathscr{X}_{N(s)})_{>0}$. For every $f \in \mathscr{G}(\mathscr{X}_s)$:*

$$g\mathbb{I}_{\{x_{P(s)}\}}f \in \mathscr{D}_G^{\mathrm{irr}} \Leftrightarrow f \in \mathscr{D}_{s\rfloor x_{P(s)}}.$$

Corollary 7. *$\mathscr{D}_G^{\mathrm{irr}}$ satisfies requirements* G1 *and* G2*: it holds for every $s \in G$ and all $x_{PN(s)} \in \mathscr{X}_{PN(s)}$ that*

$$\mathrm{marg}_s(\mathscr{D}_G^{\mathrm{irr}}\rfloor x_{PN(s)}) = \mathrm{marg}_s(\mathscr{D}_G^{\mathrm{irr}}\rfloor x_{P(s)}) = \mathscr{D}_{s\rfloor x_{P(s)}}.$$

We now have all tools necessary to formulate our first important result. It is one of the main contributions of this paper and provides a justification for the joint model $\mathscr{D}_G^{\mathrm{irr}}$ that was proposed in Eq. (3).

Theorem 8. *The irrelevant natural extension $\mathscr{D}_G^{\mathrm{irr}}$ is the unique set of desirable gambles on $\mathscr{X}_G$ that satisfies all four requirements* G1–G4.

It is already apparent from Proposition 6 that the properties of the irrelevant natural extension $\mathscr{D}_G^{\mathrm{irr}}$ are not limited to G1–G4. As a first example, Proposition 6 implies that for any node s, conditional on its parent variables $X_{P(s)}$, the non-parent non-descendant variables $X_{N(s)}$ are not only epistemically irrelevant, but also subset-irrelevant to X_s.

Corollary 9. *All nodes $s \in G$ satisfy the subset-irrelevance statement $\mathrm{SIR}(N(s),\{s\}|P(s))$: for any $x_{P(s)} \in \mathscr{X}_{P(s)}$ and non-empty $A_{N(s)} \subseteq \mathscr{X}_{N(s)}$, it holds that*

$$\mathrm{marg}_s(\mathscr{D}_G^{\mathrm{irr}}\rfloor\{x_{P(s)}\}\times A_{N(s)}) = \mathrm{marg}_s(\mathscr{D}_G^{\mathrm{irr}}\rfloor x_{P(s)}).$$

In the next two sections, we establish a number of even stronger properties of $\mathscr{D}_G^{\mathrm{irr}}$.

5 Additional marginalisation properties

As explained in Section 3.1, a subset K of G can be associated with a so-called sub-DAG of the original DAG. Similarly to what we have done for the original DAG, we can use Eq. (3) to construct a joint model for this sub-DAG. All we need to do is provide, for every $s \in K$ and $z_{P_K(s)} \in \mathscr{X}_{P_K(s)}$, a local model $\mathscr{D}_{s\rfloor z_{P_K(s)}}$.

One particular way of providing these local models is to derive them from the ones of the original DAG. The starting point to do so is fixing a value $x_{P(K)} \in \mathscr{X}_{P(K)}$ for the parent variables of K. This provides us, for every $s \in K$, with a value $x_{P(s)\setminus P_K(s)} \in \mathscr{X}_{P(s)\setminus P_K(s)}$ because $P(s)\setminus P_K(s) \subseteq P(K)$. For every $s \in K$ and $z_{P_K(s)} \in \mathscr{X}_{P_K(s)}$, we can then identify the local model $\mathscr{D}_{s\rfloor z_{P_K(s)}}$ of the sub-DAG with the local model $\mathscr{D}_{s\rfloor z_{P(s)}}$ of the original DAG, where $z_{P(s)\setminus P_K(s)} = x_{P(s)\setminus P_K(s)}$. In other words, for every $s \in K$ and $z_{P_K(s)} \in \mathscr{X}_{P_K(s)}$

$$\mathscr{D}_{s\rfloor z_{P_K(s)}} = \mathscr{D}_{s\rfloor(z_{P_K(s)},x_{P(s)\setminus P_K(s)})}.$$

Example 2. *Consider again the DAG in Figure 1 and the sub-DAG in Figure 2 that corresponds to the closed subset $K = \{s_5, s_7, s_9, s_{12}\} \subset G$. In order to provide this sub-DAG with local models, we fix a value $x_{P(K)} \in \mathscr{X}_{P(K)}$. Using Eq.* (5)*, this provides us with unconditional local models $\mathscr{D}_{s_5} = \mathscr{D}_{s_5\rfloor x_{s_3}}$ and $\mathscr{D}_{s_{12}} = \mathscr{D}_{s_{12}\rfloor x_{s_{11}}}$, for all $z_{s_5} \in \mathscr{X}_{s_5}$, a conditional local model $\mathscr{D}_{s_7\rfloor z_{s_5}} = \mathscr{D}_{s_7\rfloor(z_{s_5},x_{s_4})}$ and, for all $z_{s_7} \in \mathscr{X}_{s_7}$, a conditional local model $\mathscr{D}_{s_9\rfloor z_{s_7}}$.* ◊

For every $K \subseteq G$ and all $x_{P(K)} \in \mathscr{X}_{P(K)}$, the resulting joint model for the sub-DAG that is associated with K is given by

$$\mathscr{D}^{\mathrm{irr}}_{K\rfloor x_{P(K)}} := \mathrm{posi}(\mathscr{A}^{\mathrm{irr}}_{K\rfloor x_{P(K)}}),$$

where

$$\mathscr{A}^{\mathrm{irr}}_{K\rfloor x_{P(K)}} := \Big\{\mathbb{I}_{\{z_{PN_K(s)}\}}f\colon s \in K, z_{PN_K(s)} \in \mathscr{X}_{PN_K(s)}, \\ f \in \mathscr{D}_{s\rfloor(z_{P_K(s)},x_{P(s)\setminus P_K(s)})}\Big\}.$$

A question that now naturally arises is whether these joint models for sub-DAGs can be related to the original joint model $\mathscr{D}_G^{\mathrm{irr}}$. It turns out that, for subsets K of G that are closed, this is indeed the case.

Theorem 10. *If K is a closed subset of G, then for any $x_{P(K)} \in \mathscr{X}_{P(K)}$, $g \in \mathscr{G}(\mathscr{X}_{N(K)})_{>0}$ and $f \in \mathscr{G}(\mathscr{X}_K)$:*

$$g\mathbb{I}_{\{x_{P(K)}\}}f \in \mathscr{D}_G^{\mathrm{irr}} \Leftrightarrow f \in \mathscr{D}^{\mathrm{irr}}_{K\rfloor x_{P(K)}}.$$

The proof, although complex and elaborate, is essentially a simple separating hyperplane argument. We consider this result to be the main technical achievement of this paper. It is a significant generalisation of Proposition 6 [with $K = \{s\}$] and has a number of interesting consequences. As a first example, it implies the following generalisations of Corollaries 7 and 9.

Corollary 11. *For all closed $K \subseteq G$, $x_{P(K)} \in \mathscr{X}_{P(K)}$ and non-empty $A_{N(K)} \subseteq \mathscr{X}_{N(K)}$, we have that*

$$\mathrm{marg}_K(\mathscr{D}_G^{\mathrm{irr}} \rfloor \{x_{P(K)}\} \times A_{N(K)}) = \mathscr{D}^{\mathrm{irr}}_{K\rfloor x_{P(K)}}.$$

Corollary 12. *All closed sets $K \subseteq G$ satisfy the subset-irrelevance statement* $\mathrm{SIR}(N(K), K|P(K))$*: for any $x_{P(K)} \in \mathscr{X}_{P(K)}$ and non-empty $A_{N(K)} \subseteq \mathscr{X}_{N(K)}$, it holds that*

$$\mathrm{marg}_K(\mathscr{D}_G^{\mathrm{irr}} \rfloor \{x_{P(K)}\} \times A_{N(K)}) = \mathrm{marg}_K(\mathscr{D}_G^{\mathrm{irr}} \rfloor x_{P(K)}).$$

In the next section, we will extend this subset-irrelevance result to even more general cases.

6 AD-Separation and graphoid properties

In credal networks that are defined by means of a symmetrical independence concept, the notion of D-separation is a very powerful tool [9]. For asymmetrical independence concepts such as epistemic (subset-)irrelevance, D-separation has been modified to take this asymmetry into account. Moral [8] speaks of *asymmetrical D-separation* (AD-separation) and Vantaggi [10] has introduced the very similar *L-separation* criterion. Here, we choose not to use one of these existing concepts, but to introduce a slightly modified version of AD-separation. We do so because our definition is weaker (more general) than both Moral's AD-separation and L-separation and yet has stronger properties.

Consider any path $s_1, \ldots, s_n$ in G, with $n \geq 1$. We say that this path is *blocked* by a set of nodes $C \subseteq G$ whenever at least one of the following four conditions holds:

B1. $s_1 \in C$;

B2. there is some $1 < i < n$ such that $s_i \to s_{i+1}$ and $s_i \in C$;

B3. there is some $1 < i < n$ such that $s_{i-1} \to s_i \leftarrow s_{i+1}$, $s_i \notin C$ and $D(s_i) \cap C = \emptyset$;

B4. $s_n \in C$.

Now consider (not necessarily disjoint) subsets I, O and C of G. We say that O is *AD-separated* from I by C, denoted as $\mathrm{AD}(I, O|C)$, if every path $i = s_1, \ldots, s_n = o$, $n \geq 1$, from a node $i \in I$ to a node $o \in O$, is blocked by C. Our version of AD-separation satisfies a number of useful properties.

Theorem 13. *For any subsets I, O, S and C of G, the following properties hold:*

Direct redundancy: $\mathrm{AD}(I, O|I)$

Reverse redundancy: $\mathrm{AD}(I, O|O)$

Direct decomposition: $\mathrm{AD}(I, O \cup S|C) \Rightarrow \mathrm{AD}(I, O|C)$

Reverse decomposition: $\mathrm{AD}(I \cup S, O|C) \Rightarrow \mathrm{AD}(I, O|C)$

Direct weak union: $\mathrm{AD}(I, O \cup S|C) \Rightarrow \mathrm{AD}(I, O|C \cup S)$

Reverse weak union: $\mathrm{AD}(I \cup S, O|C) \Rightarrow \mathrm{AD}(I, O|C \cup S)$

Direct contraction:

$$\mathrm{AD}(I, O|C) \;\&\; \mathrm{AD}(I, S|C \cup O) \Rightarrow \mathrm{AD}(I, O \cup S|C)$$

Reverse contraction:

$$\mathrm{AD}(I, O|C) \;\&\; \mathrm{AD}(S, O|C \cup I) \Rightarrow \mathrm{AD}(I \cup S, O|C)$$

Direct intersection: *if $O \cap S = \emptyset$, then*

$$\mathrm{AD}(I, O|C \cup S) \;\&\; \mathrm{AD}(I, S|C \cup O) \Rightarrow \mathrm{AD}(I, O \cup S|C)$$

Reverse intersection: *if $I \cap S = \emptyset$, then*

$$\mathrm{AD}(I, O|C \cup S) \;\&\; \mathrm{AD}(S, O|C \cup I) \Rightarrow \mathrm{AD}(I \cup S, O|C)$$

This result (and our proof for it) is very similar to, and heavily inspired by, the work of Vantaggi [10, Theorem 7.1]. The main difference is that Vantaggi does not include the two redundancy properties, since L-separation is defined only for *disjoint* subsets I, O and C of G. Moral's version of AD-separation [8] does not require I, O and C to be disjoint, but it does not satisfy direct redundancy, and proofs for a number of other properties are not given [8, Theorem 4]. We therefore prefer our version of AD-separation.

Example 3. *Consider the sets of nodes $I = \{s_2, s_3, s_4, s_{11}\}$, $O = \{s_5, s_6, s_9, s_{13}\}$, $C = \{s_4, s_6, s_{12}\}$, $S_{\mathrm{d}} = \{s_8, s_{10}\}$ and $S_{\mathrm{r}} = \{s_1\}$ in the DAG that is depicted in Figure 1. The direct properties in Theorem 13 are illustrated by I, O, C and S_{d} and the reverse ones by I, O, C and S_{r}.* ◊

Theorem 10 implies a very general factorisation result.

Theorem 14. *If $I, O, C \subseteq G$ are such that* $\mathrm{AD}(I, O|C)$ *then for all $x_C \in \mathscr{X}_C$, $g \in \mathscr{G}(\mathscr{X}_{I\setminus C})_{>0}$ and $f \in \mathscr{G}(\mathscr{X}_{O\setminus C})$:*

$$g\mathbb{I}_{\{x_C\}} f \in \mathscr{D}_G^{\mathrm{irr}} \Leftrightarrow \mathbb{I}_{\{x_C\}} f \in \mathscr{D}_G^{\mathrm{irr}}.$$

This result can be combined with Theorem 13 to derive a collection of (subset-)irrelevance statements that are fulfilled by the irrelevant natural extension $\mathscr{D}_G^{\mathrm{irr}}$.

Corollary 15. *For any $I, O, C \subseteq G$ such that* $\mathrm{AD}(I, O|C)$ *we have that* $\mathrm{SIR}(I, O|C)$ *(and thus also* $\mathrm{IR}(I, O|C)$*): for all $x_C \in \mathscr{X}_C$ and non-empty $A_{I\setminus C} \subseteq \mathscr{X}_{I\setminus C}$ it holds that*

$$\mathrm{marg}_{O\setminus C}(\mathscr{D}_G^{\mathrm{irr}} \rfloor \{x_C\} \times A_{I\setminus C}) = \mathrm{marg}_{O\setminus C}(\mathscr{D}_G^{\mathrm{irr}} \rfloor x_C).$$

This family of subset-irrelevance statements satisfies all graphoid properties except symmetry: it satisfies redundancy, decomposition, weak union, contraction and intersection, both in their direct and reverse form.

We leave it to the reader to show that Theorem 14 is a generalisation of Theorem 10 and that Corollary 15 generalises the first part of Corollary 12. In other words: for any closed subset K of G, it holds that $\mathrm{AD}(N(K), K|P(K))$.

Readers who are familiar with the work in Ref. [8] might have noticed the similarity between Ref. [8, Theorem 5] and the first part of Corollary 15. The main difference between our approach and Moral's approach [8], besides the fact that we use a slightly different separation criterion, is that he enforces a more stringent version of epistemic irrelevance than we do. He calls X_I epistemically irrelevant to X_O if and only if the joint model $\mathscr{D}_{I\cup O}$ is the so-called irrelevant natural extension of $\mathscr{D}_I$ and $\mathscr{D}_O$ and refers to our concept of irrelevance as weak epistemic irrelevance. Consequently, if we understand his work correctly, his results are not applicable to all directed acyclic networks. As a simple example: his concept of irrelevance does not seem to allow for two variables to be mutually irrelevant, except in some degenerated uninformative cases. Therefore, it appears to us his results cannot be applied to a network consisting of two unconnected nodes.

As far as the second part of Corollary 15 is concerned, some clarification is perhaps in order. We do not claim that epistemic irrelevance satisfies the graphoid axioms that are stated in Theorem 13. As was proven in Ref. [4], epistemic irrelevance can violate direct contraction and both direct and reverse intersection. In fact, we believe that this negative result might even be one of the main reasons why a result such as Corollary 15 has thus far not appeared in any literature.

Indeed, in Bayesian networks, proving the counterpart to Corollary 15—with AD-separation replaced by D-separation and epistemic irrelevance replaced by stochastic independence—is usually done by using the fact that stochastic independence satisfies the graphoid axioms [9]. By applying these axioms to the independence assessements that are used to define a Bayesian network, one can infer new independencies, namely those that correspond to D-separations in the DAG of that network.

If one tries to mimic this approach in our context, then since epistemic irrelevance can fail some of the graphoid axioms, one might suspect that Corollary 15 cannot be proven. However, it is not necessary to use the axioms: our proof for Theorem 14—of which the the first part of Corollary 15 is a straightforward consequence—uses only Theorem 10 and a number of properties of AD-separation. At no point does it invoke graphoid properties of epistemic irrelevance. The second part of Corollary 15 is then but a mere consequence of the first part and Theorem 13. It states that the family of irrelevance statements that are proven to hold in the first part, are closed under the graphoid properties in Theorem 13.

So in order to conclude this section: epistemic irrelevance can fail a number of graphoid axioms, which implies that the irrelevance statements that are proven in Corollary 15 do not necessarily hold for every joint model $\mathscr{D}_G$ that satisfies requirements G1–G3. However for the unique one that also satisfies G4, being the irrelevant natural extension $\mathscr{D}_G^{\mathrm{irr}}$ of the network, this family of irrelevance statements does hold, the reason being that for this specific model, one can provide a direct proof that does not invoke any graphoid axioms of epistemic irrelevance.

7 Credal nets under epistemic irrelevance using coherent lower previsions

Credal networks under epistemic irrelevance can also be defined using imprecise probability concepts other than coherent sets of desirable gambles. In this section, we describe an approach that uses coherent lower previsions, and we show how it is related to the desirable gambles approach of the previous sections.

7.1 Coherent lower previsions

For any subset O of G, we define a *coherent lower prevision* $\underline{P}_O$ as a real-valued functional on $\mathscr{G}(\mathscr{X}_O)$ that satisfies the following three conditions. For all $f,g \in \mathscr{G}(\mathscr{X}_O)$ and all real $\lambda \geq 0$:

C1. $\underline{P}_O(f) \geq \min f$;

C2. $\underline{P}_O(\lambda f) = \lambda \underline{P}_O(f)$; [non-negative homogeneity]

C3. $\underline{P}_O(f+g) \geq \underline{P}_O(f) + \underline{P}_O(g)$. [super-additivity]

Now consider two disjoint subsets O and I of G and suppose that we have, for all $x_I \in \mathscr{X}_I$, a coherent lower prevision $\underline{P}_O(\cdot|x_I)$ on $\mathscr{G}(\mathscr{X}_O)$. The corresponding *coherent conditional lower prevision* $\underline{P}_{O\cup I}(\cdot|X_I)$ is then a special two-place function that is defined, for all $f \in \mathscr{G}(\mathscr{X}_{O\cup I})$ and $x_I \in \mathscr{X}_I$, by $\underline{P}_{O\cup I}(f|x_I) := \underline{P}_O(f(\cdot,x_I)|x_I)$.

7.2 Defining a credal network

Suppose now that the local models of our credal network under epistemic irrelevance are coherent lower previsions: for all $s \in G$ and $x_{P(s)} \in \mathscr{X}_{P(s)}$, we have a coherent lower prevision $\underline{P}_{s\rfloor x_{P(s)}}$ on $\mathscr{G}(\mathscr{X}_s)$.

The irrelevance assessments that are encoded in the network can then be expressed as follows. For all $s \in G$, $I \subseteq N(s)$, $x_{P(s)\cup I} \in \mathscr{X}_{P(s)\cup I}$ and $f \in \mathscr{G}(\mathscr{X}_s)$, we require that:

$$\underline{P}_{\{s\}}(f|x_{P(s)\cup I}) := \underline{P}_{s\rfloor x_{P(s)}}(f).$$

For all $s \in G$ and $I \subseteq N(s)$, the corresponding conditional lower prevision $\underline{P}_{\{s\}\cup P(s)\cup I}(\cdot|X_{P(s)\cup I})$ is then given, for all $f \in \mathscr{G}(\mathscr{X}_{\{s\}\cup P(s)\cup I})$ and $x_{P(s)\cup I} \in \mathscr{X}_{P(s)\cup I}$, by

$$\underline{P}_{\{s\}\cup P(s)\cup I}(f|x_{P(s)\cup I}) := \underline{P}_{s\rfloor x_{P(s)}}(f(\cdot,x_{P(s)\cup I})).$$

We will denote the set consisting of all these conditional lower previsions as $\mathscr{I}(\underline{P}_{s\rfloor x_{P(s)}}, s \in G, x_{P(s)} \in \mathscr{X}_{P(s)})$.

The global model $\underline{E}_G^{\mathrm{irr}}$ is now defined as the smallest coherent lower prevision on $\mathscr{G}(\mathscr{X}_G)$ that is (strongly) coherent with this set of conditional lower previsions. We will refer to it as the *irrelevant natural extension* of the local models $\underline{P}_{s\rfloor x_{P(s)}}$. We will not get into the details of what strong coherence means, but one can very roughly think of it as requiring that the conditional lower previsions in the set $\mathscr{I}(\underline{P}_{s\rfloor x_{P(s)}}, s \in G, x_{P(s)} \in \mathscr{X}_{P(s)})$ (i) are compatible with one another and (ii) can be obtained by conditioning the global model $\underline{E}_G^{\mathrm{irr}}$; see Ref. [5, Section 2.4] for more details on strong coherence.

We know from Walley's Finite Extension Theorem [11, Theorem 8.1.9] that if $\underline{E}_G^{\mathrm{irr}}$ exists, then it is equal to the *natural extension* of the collection $\mathscr{I}(\underline{P}_{s\rfloor x_{P(s)}}, s \in G, x_{P(s)} \in \mathscr{X}_{P(s)})$ to an unconditional lower prevision on $\mathscr{G}(\mathscr{X}_G)$. In that case, by applying a derivation that is similar to the one for [5, Eq.(10), Section 5.2], we find for all $f \in \mathscr{G}(\mathscr{X}_G)$ that

$$\begin{aligned}\underline{E}_G^{\mathrm{irr}}(f) = \sup_{\substack{g_{\{s\}\cup P(s)\cup I}\\ \in\mathscr{G}(\mathscr{X}_{\{s\}\cup P(s)\cup I})}} \Bigg\{ \min_{z_G \in \mathscr{X}_G} \Big[f(z_G) \\ - \sum_{s\in G, I\subseteq N(s)} [g_{\{s\}\cup P(s)\cup I}(z_s, z_{P(s)\cup I}) \\ - \underline{P}_{s\rfloor z_{P(s)}}(g_{\{s\}\cup P(s)\cup I}(\cdot, z_{P(s)\cup I}))]\Big]\Bigg\}. \quad (4)\end{aligned}$$

7.3 Connections with our approach

For every $s \in G$ and $x_{P(s)} \in \mathscr{X}_{P(s)}$, the local coherent set of desirable gambles $\mathscr{D}_{s\rfloor x_{P(s)}}$ uniquely defines a corresponding coherent lower prevision $\underline{P}_{s\rfloor x_{P(s)}}$. For all $f \in \mathscr{G}(\mathscr{X}_s)$

$$\underline{P}_{s\rfloor x_{P(s)}}(f) := \sup\{\mu \in \mathbb{R} \colon f - \mu \in \mathscr{D}_{s\rfloor x_{P(s)}}\}. \quad (5)$$

Conversely, every local coherent lower prevision $\underline{P}_{s\rfloor x_{P(s)}}$ has at least one coherent set of desirable gambles $\mathscr{D}_{s\rfloor x_{P(s)}}$ from which it can be derived by Eq. (5). These sets are however not unique since coherent sets of desirable gambles are generally more expressive than coherent lower previsions. Using any such family of corresponding local sets of desirable gambles, we can then apply Eq. (3) to obtain their irrelevant natural extension $\mathscr{D}_G^{\mathrm{irr}}$. This joint set also has a corresponding coherent lower prevision. It is denoted as $\underline{P}_G^{\mathrm{irr}}$ and given for all $f \in \mathscr{G}(\mathscr{X}_G)$ by

$$\underline{P}_G^{\mathrm{irr}}(f) := \sup\{\mu \in \mathbb{R} \colon f - \mu \in \mathscr{D}_G^{\mathrm{irr}}\}. \quad (6)$$

The coherent lower prevision $\underline{P}_G^{\mathrm{irr}}$ that is constructed in this way from given local models $\underline{P}_{s\rfloor x_{P(s)}}$ might depend on the particular choice for the sets $\mathscr{D}_{s\rfloor x_{P(s)}}$ in its construction. We will show in Theorem 17 that such is not the case, however.

Proposition 16. *Choose, for all $s \in G$ and $x_{P(s)} \in \mathscr{X}_{P(s)}$, any coherent local set of desirable gambles $\mathscr{D}_{s\rfloor x_{P(s)}}$ on $\mathscr{X}_s$ such that the given local coherent lower prevision $\underline{P}_{s\rfloor x_{P(s)}}$ satisfies Eq.* (5)*. Construct the irrelevant natural extension $\mathscr{D}_G^{\mathrm{irr}}$ by applying Eq.* (3) *and let $\underline{P}_G^{\mathrm{irr}}$ be the coherent lower prevision on $\mathscr{G}(\mathscr{X}_G)$ as given by Eq.* (6)*. Then $\underline{P}_G^{\mathrm{irr}}$ is strongly coherent with $\mathscr{I}(\underline{P}_{s\rfloor x_{P(s)}}, s \in G, x_{P(s)} \in \mathscr{X}_{P(s)})$.*

Proposition 16 shows that it is possible to construct at least one coherent lower prevision $\underline{P}_G^{\mathrm{irr}}$ on $\mathscr{G}(\mathscr{X}_G)$ that is strongly coherent with $\mathscr{I}(\underline{P}_{s\rfloor x_{P(s)}}, s \in G, x_{P(s)} \in \mathscr{X}_{P(s)})$, implying that the irrelevant natural extension $\underline{E}_G^{\mathrm{irr}}$ is always well defined and given by Eq. (4).

The following result now establishes the final connection between the irrelevant natural extensions $\mathscr{D}_G^{\mathrm{irr}}$ and $\underline{E}_G^{\mathrm{irr}}$ that were outlined in this paper. We show that $\underline{P}_G^{\mathrm{irr}}$ is always equal to the irrelevant natural extension $\underline{E}_G^{\mathrm{irr}}$, regardless of the local sets $\mathscr{D}_{s\rfloor x_{P(s)}}$ that are chosen to construct it.

Theorem 17. *Let $\mathscr{D}_G^{\mathrm{irr}}$ be the irrelevant natural extension of local coherent sets of desirable gambles $\mathscr{D}_{s\rfloor x_{P(s)}}$, $s \in G$ and $x_{P(s)} \in \mathscr{X}_{P(s)}$, as given by Eq.* (3)*. Construct local coherent lower previsions $\underline{P}_{s\rfloor x_{P(s)}}$ by applying Eq.* (5) *and let $\underline{E}_G^{\mathrm{irr}}$ be their irrelevant natural extension, as given by Eq.* (4)*. It then holds for all $f \in \mathscr{G}(\mathscr{X}_G)$ that*

$$\underline{E}_G^{\mathrm{irr}}(f) = \sup\{\mu \in \mathbb{R} \colon f - \mu \in \mathscr{D}_G^{\mathrm{irr}}\} = \underline{P}_G^{\mathrm{irr}}(f).$$

We believe that this connection between the two approaches can be used to translate at least some of our results for sets of desirable gambles into the language of coherent lower previsions. We intend to explore this further in future work.

8 Summary and conclusions

This paper has developed the notion of a credal network under epistemic irrelevance using sets of desirable gambles. We have proven that the resulting irrelevant natural extension of a network has a number of interesting properties. It marginalises in an intuitive way and satisfies all graphoid properties except symmetry. Finally, we have established a connection with an approach to credal networks under epistemic irrelevance that uses coherent lower previsions.

Future goals that we intend to pursue are to derive counterparts to the marginalisation and graphoid properties in this paper, expressed in terms of coherent lower previsions rather than sets of desirable gambles. By exploiting these properties, we would like to develop algorithms for credal networks under epistemic irrelevance that are able to perform inferences in an efficient manner.

Acknowledgements

Jasper De Bock is a Ph.D. Fellow of the Fund for Scientific Research – Flanders (FWO) and wishes to acknowledge the financial support of the FWO. The authors also wish to thank three anonymous referees for their helpful comments.

References

[1] I. Couso, S. Moral: Sets of desirable gambles: conditioning, representation, and precise probabilities. *International Journal of Approximate Reasoning*, 52(7):1034–1055, 2011.

[2] F. G. Cozman. Credal networks. *Artificial Intelligence*, 120:199–233, 2000.

[3] F. G. Cozman. Graphical Models for Imprecise Probabilities. *International Journal of Approximate Reasoning*, 39(2–3):167–184, 2005.

[4] F. G. Cozman and P. Walley. Graphoid properties of epistemic irrelevance and independence. *Annals of Mathematics and Artificial Intelligence*, 45(1–2):173–195, 2005.

[5] G. de Cooman, E. Miranda, M. Zaffalon. Independent natural extension. *Artificial Intelligence*, 175(12–13):1911-1950, 2011.

[6] G. de Cooman, E. Miranda: Irrelevant and Independent Natural Extension for Sets of Desirable Gambles. *Journal of Artificial Intelligence Research*, 45:601–640, 2012.

[7] G. de Cooman, E. Quaeghebeur: Exchangeability and sets of desirable gambles. *International Journal of Approximate Reasoning*, 53(3):363–395, 2012. Special issue in honour of Henry E. Kyburg, Jr.

[8] S. Moral: Epistemic irrelevance on sets of desirable gambles. *Annals of Mathematics and Artificial Intelligence* 45:197–214, 2005.

[9] J. Pearl: *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA, 1988.

[10] B. Vantaggi: The L-separation criterion for description of cs-independence models. *International Journal of Approximate Reasoning*, 29:291–316, 2002.

[11] P. Walley: *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London, 1991.

[12] P. Walley: Towards a unified theory of imprecise probability. *International Journal of Approximate Reasoning*, 24:125–148, 2000.