

## **Efficient computation of counterfactual bounds**

Zaffalon, M., Antonucci, A., Cabañas, R., Huber, D., & Azzimonti, D.

*Published in:*

International Journal of Approximate Reasoning

*DOI (link to publication from Publisher):*

<https://doi.org/10.1016/j.ijar.2023.109111>

*Publication date:*

2024

*Document Version:*

Accepted author manuscript, peer reviewed version

*Citation for published version (APA):*

Zaffalon, M., Antonucci, A., Cabañas, R., Huber, D., & Azzimonti, D. (2024). Efficient computation of counterfactual bounds. *International Journal of Approximate Reasoning*, 109111.

# Efficient Computation of Counterfactual Bounds

Marco Zaffalon<sup>a</sup>, Alessandro Antonucci<sup>a,\*</sup>, Rafael Cabañas<sup>b</sup>, David Huber<sup>a</sup>,  
Dario Azzimonti<sup>a</sup>

<sup>a</sup>*IDSIA, Lugano (Switzerland)*

<sup>b</sup>*Department of Mathematics, University of Almería, Almería (Spain)*

---

## Abstract

We assume to be given structural equations over discrete variables inducing a directed acyclic graph, namely, a *structural causal model*, together with data about its internal nodes. The question we want to answer is how can we compute bounds for partially identifiable counterfactual queries from such an input. We start by giving a map from structural causal models to *credal networks*. This allows us to compute exact counterfactual bounds via algorithms for credal nets on a subclass of structural causal models. Exact computation is going to be inefficient in general given that, as we show, causal inference is NP-hard even on polytrees. We target then approximate bounds via a causal EM scheme. We evaluate their accuracy by providing credible intervals on the quality of the approximation; we show through a synthetic benchmark that the EM scheme delivers accurate results in a fair number of runs. In the course of the discussion, we also point out what seems to be a neglected limitation to the trending idea that counterfactual bounds can be computed without knowledge of the structural equations. We also present a real case study on palliative care to show how our algorithms can readily be used for practical purposes.

**Keywords:** Causal analysis, structural causal models, partial identifiability, imprecise probability, counterfactuals, credal networks, expectation maximisation.

---

---

\*Corresponding author

*Email addresses:* zaffalon@idsia.ch (Marco Zaffalon),  
alessandro.antonucci@idsia.ch (Alessandro Antonucci), rcabanas@ual.es (Rafael Cabañas), david.huber@idsia.ch (David Huber), dario.azzimonti@idsia.ch (Dario Azzimonti)

## 1. Introduction

Since early times, dealing with causality has been—and under many respects, still is—a true challenge for scientists and philosophers (Hume, 1739). Nowadays, causality represents an important direction for data science, with many applications to machine learning (e.g., Schölkopf 2022), reinforcement learning (e.g., Zhang and Bareinboim 2020) and explainable AI (e.g., Galhotra et al. 2021).

Structural causal models are a natural formalism for causal modelling and inference, in particular for their appealing graphical representation (Pearl, 2009). They are also very general and equivalent to the prominent alternative formalisms proposed to handle causality (see Ibeling and Icard 2023, for instance, for a discussion on their relation with the Neyman-Rubin potential-outcome framework).

We focus in particular on Pearl’s non-parametric structural causal models with discrete variables. We start by showing that they can be represented by *credal networks* (Cozman, 2000), which are a class of imprecise-probabilistic graphical models originally proposed as tools for sensitivity analysis in Bayesian networks. The representation is exact: every query in the causal model can be reformulated as a query in the credal network and solved by standard algorithms for the latter. The output is made of a lower and an upper bound in the case of partially identifiable queries, this typically being the case for counterfactuals, or a sharp value for identifiable ones such as the interventional queries we might compute by Pearl’s do calculus. Yet, credal network inference remains a challenging task: exact inference belongs to a complexity class higher than that of Bayesian networks and the existing approximate schemes can in practice be directly applied only to a specific class of structural models. In fact, we use the relation between structural causal models and credal networks to prove that inference on the former is NP-hard even when the causal graph is a poly-tree.

To bypass such limitations, we derive an *expectation-maximisation* (EM) scheme that reduces the credal network inference to an iterated EM with multiple initialisations. It can be understood as a sampling approach that yields a set of points inside the exact counterfactual bounds, via inferences on Bayesian networks with the same topology of the original model. Lower and upper values of these points define an approximate *range* of values for the counterfactual that is encompassed by the actual bounds (we say it is an *inner approximation*). We derive credibility intervals to evaluate the quality of such a range in the form

of distance from the exact bounds.

As a matter of fact, relatively few works have considered so far the case of partially identifiable problems, for which bounds need to be computed. One of the first attempts in this direction is from Balke and Pearl (1997), where the computation of bounds on interventional queries is reduced to a linear program. This allows to compute tight bounds, but the size of the program grows exponentially large, being in practice feasible only for very small models. Kang and Tian (2006) present a more systematic technique to derive constraints on the query, but still with exponential growth and no explicit methods to compute bounds. Sachs et al. (2023) detect a class of models allowing for an efficient reduction to linear programs, but no guarantees are provided in the general case. More recently, both Zhang et al. (2022) and Duarte et al. (2023) consider an exact reduction to polynomial programming, with the former giving in addition an approximate sampling scheme.

The paper is organised as follows: after providing background material in Section 2, we discuss the notion of identifiability for causal queries, as well as that of M-compatibility, i.e., the logical consistency between a structural causal model and an empirical distribution, in Section 3. An algorithm to convert a causal model into a credal network, whose quantification is defined by the observational data, is provided in Section 4. The approximate EM solution is instead reported in Section 5. Experiments on synthetic and real data are discussed in Section 6 along with a case study. Conclusions are presented in Section 7, proofs are gathered in Appendix A, while Appendix B contains additional technical results. Finally, Appendix C provides a list of questions and answers intended to further clarify various aspects of the paper.<sup>1</sup>

## 2. Background

In this section we review the necessary background on probabilistic graphical models, namely Bayesian and credal networks (Section 2.1) and structural causal models (Section 2.2), while also discussing typical inference tasks on all

---

<sup>1</sup>The results we present here as a journal contribution are a revised and extended version, with new experiments on synthetic and real data and a case study, of material originally presented in three conference papers. The credal network mapping has been introduced by Zaffalon et al. (2020), the EM scheme by Zaffalon et al. (2021), and the general formulation for the credibility intervals by Zaffalon et al. (2022). The main focus of the last paper is the analysis of biased data, which is not included here.

these models (Section 2.3). For more extensive discussions on these topics we point the reader to the books of Koller and Friedman (2009) and Pearl (2009).

### 2.1. Bayesian and Credal Networks

Variable  $X$  is assumed to take values from the set  $\Omega_X$ . The generic element of  $\Omega_X$  is denoted as  $x$ . Here we only consider discrete variables, i.e.,  $|\Omega_X| < +\infty$ . Denote as  $P(X)$  a probability mass function (PMF) over  $X$ , and as  $K(X)$  a *credal* set (CS), which is a set of PMFs over  $X$ . Given variables  $X$  and  $Y$ , a conditional probability table (CPT)  $P(X|Y)$  is a collection of (conditional) PMFs indexed by the values of  $Y$ , i.e.,  $\{P(X|y)\}_{y \in \Omega_Y}$ . If all PMFs in a CPT are *degenerate*, i.e., there is a state receiving probability mass one and hence all the other ones receive zero, we say that also the CPT is degenerate. A credal CPT (CCPT)  $K(X|Y)$  is similarly a collection of CSs over  $X$  indexed by the values of  $Y$ . With a small abuse of terminology, we might call CPT (CCPT) also a single PMF (CS).

Consider a joint variable  $\mathbf{X} := (X_1, \dots, X_n)$  and a directed acyclic graph  $\mathcal{G}$  whose nodes are in a one-to-one correspondence with the variables in  $\mathbf{X}$ . Note that we use a node in  $\mathcal{G}$  and its corresponding variable interchangeably. Given  $\mathcal{G}$ , a Bayesian network (BN) is a collection of CPTs  $\{P(X_i|\text{Pa}_{X_i})\}_{i=1}^n$ , where  $\text{Pa}_{X_i}$  denotes the *parents* of  $X_i$ , i.e., the direct predecessors of  $X_i$  according to  $\mathcal{G}$ . A BN induces a joint PMF  $P(\mathbf{X})$  that factorises as follows:

$$P(\mathbf{x}) = \prod_{i=1}^n P(x_i|\text{pa}_{X_i}), \quad (1)$$

for each  $\mathbf{x} \in \Omega_X$ , where  $(x_i, \text{pa}_{X_i}) \sim \mathbf{x}$ , which we use to denote that  $x_i$  and  $\text{pa}_{X_i}$  are the values of  $X_i$  and  $\text{Pa}_{X_i}$  consistent with  $\mathbf{x}$  for each  $i = 1, \dots, n$ . A credal network (CN) is similarly intended as a collection of CCPTs. A CN defines a joint CS  $K(\mathbf{X})$  whose elements are PMFs factorising as those of a BN with CPT values taken from the corresponding CCPTs. Some authors also require CS convexity, but this is irrelevant for the inferences we consider in this paper.

### 2.2. Structural Causal Models

Let us first define a *structural equation* (SE)  $f_Y$  associated with variable  $Y$  and based on the input variable(s)  $X$  as a surjective function  $f_Y : \Omega_X \rightarrow \Omega_Y$  that determines the value of  $Y$  from that of  $X$ . We shall represent such a SE via a degenerate CPT  $P(Y|X)$  such that  $P(y|x) := \mathbb{I}[f_Y(x) = y]$  for each  $y \in \Omega_Y$  and  $x \in \Omega_X$ , where  $\mathbb{I}[\cdot]$  denote the Iverson brackets that take the value one if the statement inside the brackets is true and zero otherwise.

Consider two sets of variables  $\mathbf{U}$  and  $\mathbf{V}$ , to be called, respectively, *exogenous* and *endogenous*. A collection of SEs  $\{f_V\}_{V \in \mathbf{V}}$  such that the input variables of  $f_V$  are in  $(\mathbf{U}, \mathbf{V})$  for each  $V \in \mathbf{V}$  is called a *partially specified structural causal model* (PSCM) over  $(\mathbf{U}, \mathbf{V})$ . It coincides with the notion of a ‘functional causal model’ in Pearl (2009, Chapter 1.4).

A PSCM  $M$  induces the specification of a directed, so-called *causal*, graph  $\mathcal{G}$  whose nodes are in a one-to-one correspondence with the variables in  $(\mathbf{U}, \mathbf{V})$  and such that there is an arc between two variables if and only if the first variable is an input variable for the SE of the second. The exogenous variables are therefore root nodes of  $\mathcal{G}$ . We focus on *semi-Markovian* PSCMs, i.e., such that their graph is acyclic. Moreover, if there is no exogenous variable with more than one endogenous child, we call the PSCM *Markovian*.

In a (semi-Markovian) PSCM  $M$ , a joint state of  $\mathbf{V}$  is obtained from a (joint) state of  $\mathbf{U}$  by applying the SEs of  $M$  consistently with a topological order for  $\mathcal{G}$ . A *fully specified structural causal model* (FSCM, see Pearl, 2009, top of p. 69) is just a PSCM  $M$  paired with a collection of marginal PMFs, one for each exogenous variable. As SEs induce (degenerate) CPTs, overall, an FSCM provides a BN specification based on  $\mathcal{G}$  whose joint PMF factorises according to Equation (1), i.e.:

$$P(\mathbf{u}, \mathbf{v}) = \prod_{U \in \mathbf{U}} P(u) \cdot \prod_{V \in \mathbf{V}} P(v | \text{pa}_V), \quad (2)$$

where, for each  $\mathbf{u} \in \Omega_{\mathbf{U}}$  and  $\mathbf{v} \in \Omega_{\mathbf{V}}$ ,  $(u, v, \text{pa}_V) \sim (\mathbf{u}, \mathbf{v})$ ,  $\text{Pa}_V$  are the parents of  $V$  according to  $\mathcal{G}$  (i.e., the inputs of SE  $f_V$ ).

Let us clarify the above concepts by means of a small example.

**Example 1.** Given the endogenous variables  $\mathbf{V} := (V_1, V_2, V_3, V_4)$  and the exogenous ones  $\mathbf{U} := (U_1, U_2, U_3)$ , SEs  $f_{V_1}(U_1)$ ,  $f_{V_2}(U_2, V_1)$ ,  $f_{V_3}(U_3, V_2)$ , and  $f_{V_4}(U_2, V_3)$  define a semi-Markovian PSCM based on the causal graph in Figure 1. An FSCM based on the same SEs induces a joint PMF that factorises as in Equation (2). By expressing the (degenerate) probabilities in the endogenous CPTs through the SEs we can therefore write the joint probability as:

$$\begin{aligned} P(u_1, u_2, u_3, v_1, v_2, v_3, v_4) &= P(u_1)P(u_2)P(u_3) \\ &\cdot \mathbb{I}[f_{V_1}(u_1) = v_1] \mathbb{I}[f_{V_2}(u_2, v_1) = v_2] \mathbb{I}[f_{V_3}(u_3, v_2) = v_3] \mathbb{I}[f_{V_4}(u_2, v_3) = v_4], \end{aligned} \quad (3)$$

for each  $v_i \in \Omega_{V_i}$  and  $u_j \in \Omega_{U_j}$  with  $i = 1, 2$  and  $j = 1, 2, 3, 4$ .

Given the causal graph  $\mathcal{G}$  of a PSCM (or FSCM)  $M$ , obtain  $\mathcal{G}'$  by removing from  $\mathcal{G}$  any arc connecting pairs of endogenous variables. Let  $\{\mathcal{G}_c\}_{c \in \mathcal{C}}$  denote




Figure 1: The causal graph of a semi-Markovian PSCM.

the connected components of  $\mathcal{G}'$  with  $\mathcal{C} = \{1, \dots, n_C\}$  and  $n_C$  being the number of connected components of  $\mathcal{G}'$ . The  $c$ -components of  $M$  are the elements of the partition  $\{V^c\}_{c \in \mathcal{C}}$  of  $V$ , where  $V^c$  denotes the endogenous nodes in  $\mathcal{G}_c$ , for each  $c \in \mathcal{C}$  (Tian, 2002). This procedure also induces a partition of  $U$ , similarly denoted as  $\{U^c\}_{c \in \mathcal{C}}$ . Moreover, for each  $c \in \mathcal{C}$ , let  $W^c$  denote the union of the endogenous parents of the nodes in  $V^c$  and  $V^c$  itself. Finally, for each  $V \in V^c$ , obtain  $W_V$  by removing from  $W^c$  the nodes topologically following  $V$  and  $V$  itself (we dropped the superscript  $c$  as this can be implicitly retrieved from  $V$ ).

Tian (2002) shows that the joint PMF  $P(V)$  obtained by marginalising the exogenous variables out of the joint PMF in Equation (2) is a BN that factorises as follows:

$$P(\mathbf{v}) = \prod_{V \in V} P(v | \mathbf{w}_V), \quad (4)$$

for each  $\mathbf{v} \in \Omega_V$  with  $(v, \mathbf{w}_V) \sim \mathbf{v}$  and  $V \in V$  (remember that symbol  $\sim$  is used to denote consistent instances of variables as from Section 2.1). In the following we call such a BN the *endogenous* BN of a FSCM.

For Markovian models, where the partitions induced by the  $c$ -components are made of singletons, we trivially have  $W_V = \text{Pa}_V \cap V$  for each  $V \in V$ , i.e., the parents of an endogenous node in the endogenous BN are its endogenous parents in the original model. For non-Markovian models, we illustrate the procedure by the following example.

**Example 2.** *The FSCM in Example 1 has three  $c$ -components inducing the endogenous partition  $\{V^c\}_{c=1}^3$  with  $V^1 := \{V_1\}$ ,  $V^2 := \{V_2, V_4\}$ , and  $V^3 := \{V_3\}$ , and the exogenous partition  $\{U^c\}_{c=1}^3$  with  $U^1 := \{U_1\}$ ,  $U^2 := \{U_2\}$ , and  $U^3 := \{U_3\}$ . We easily obtain  $W^1 = \{V_1\}$ ,  $W^2 = \{V_1, V_2, V_3, V_4\}$ , and  $W^3 = \{V_2, V_3\}$ . Finally, by considering the unique topological order over the four endogenous variables, we get  $W_{V_1} = \emptyset$ ,  $W_{V_2} = \{V_1\}$ ,  $W_{V_3} = \{V_2\}$ , and  $W_{V_4} = \{V_1, V_2, V_3\}$ . The endogenous BN obeying the factorisation in Equation (4) therefore corresponds to the graph in Figure 2.*




Figure 2: The graph of the endogenous BN for the FSCM in Figure 1.

In practice, the CPTs in the right-hand side of Equation (4) can be computed through standard BN inference algorithms by just regarding the FSCM as a BN. With PSCMs, assuming the availability of a dataset  $\mathcal{D}$  of endogenous observations, we might also quantify the endogenous BN by learning the probabilities in the CPTs from  $\mathcal{D}$ .

Finally, let us note that PSCMs assume SEs to be given. However, in the case where endogenous variables take on finitely many values, recent results allow one to specify a PSCM by only giving a causal graph; SEs are automatically defined, without loss of generality (yet adding a likely excess of caution to the model), via a *canonical specification*. Let us introduce it in the simple case of Markovian models: we say that SE  $f_V$  is canonical if the states of the (single because of Markovianity) exogenous parent  $U$  of  $V$  index all the deterministic relations between the endogenous parents (i.e.,  $\mathbf{W}_V$  because of Markovianity) and  $V$ . This requires:

$$|\Omega_U| = |\Omega_V|^{\prod_{W \in \mathbf{W}_V} |\Omega_W|}. \quad (5)$$

A Markovian PSCM whose SEs are all canonical is also called canonical. We refer the reader to the works of Duarte et al. (2023) and Zhang et al. (2022) for a generalisation of such a concept to non-Markovian models.

### 2.3. Updating, Interventions and Counterfactuals

Computing the posterior probability for the state of a queried variable given an evidence with respect to the joint PMF of a BN, as well as the bounds of this probabilities with respect to the joint CS of a CN are NP-hard tasks (e.g., Mauá et al., 2014). Yet, polynomial algorithms computing approximate inferences for the general case (e.g., Antonucci et al., 2015) or exact ones for classes of sub-models (e.g., Fagioli and Zaffalon, 1998) are available for CNs, not to mention the copious tools for BNs (e.g., Koller and Friedman, 2009).

Observational queries in FSCMs can be addressed in the endogenous BN, and the same can be done for PSCMs by assuming the availability of the dataset  $\mathcal{D}$  of endogenous observations.

To perform causal inference, *interventions* denoted as  $\text{do}(\cdot)$  should be considered instead. In an FSCM or PSCM  $M$ , given  $V \in \mathbf{V}$  and  $v \in \Omega_V$ ,  $\text{do}(V = v)$



simulates a physical action on  $M$  forcing  $V$  to take the value  $v$ . The original SE  $f_V$  should be consequently replaced by a constant map  $V = v$ . Notation  $M_v$  is used for such a modified model, whose graph is obtained by removing from  $\mathcal{G}$  the arcs entering  $V$ , and for which evidence  $V = v$  is considered. In an FSCM  $M$ , given  $V, W \in \mathbf{V}$  and  $v \in \Omega_V$ ,  $P(w|\text{do}(v))$  denotes the conditional probability of  $W = w$  in the post-intervention model, i.e.,  $P'(w|v)$ , where  $P'$  is the joint PMF induced by  $M_v$ . As interventions commute, there are no ordering issues when coping with multiple interventions. If evidence is also available, i.e., some variables have been observed, it is customary to assume that observations take place after the interventions. Note that interventional queries assume the set of observed variables and that of intervened variables disjoint.

A more general setup is provided by *counterfactual* queries, where the same variable may be observed as well as subject to intervention, albeit in distinct ‘worlds’. In mathematical parlance, if  $\mathbf{W}$  are the queried variables,  $\mathbf{V}'$  the observed ones and  $\mathbf{V}''$  the intervened ones, we write the query by  $P(\mathbf{W}_{\mathbf{V}''}|\mathbf{v}')$  with possibly  $\mathbf{V}' \cap \mathbf{V}'' \neq \emptyset$ . A popular counterfactual query involving two endogenous Boolean variables  $X$  and  $Y$  of an FSCM is the *probability of necessity* (PN), i.e., the probability that event  $Y$  would not have occurred by disabling  $X$ , given that  $X$  and  $Y$  did in fact occur. This corresponds to  $P(Y_{X=0} = 0|X = 1, Y = 1)$ . Similarly, the probability of *sufficiency* (PS) is the probability that  $Y$  would have occurred by activating  $X$ , given that  $X$  and  $Y$  did not occur, i.e.,  $P(Y_{X=1} = 1|X = 0, Y = 0)$ . Finally, the probability of *necessity and sufficiency* (PNS), corresponding to  $P(Y_{X=1} = 1, Y_{X=0} = 0)$ , is the probability that  $Y$  would respond to  $X$  both ways, thus measuring both the sufficiency and necessity of  $X$  to produce  $Y$ . A characterisation of these three counterfactual probabilities has been provided by Pearl (1999).

Computing counterfactual queries in an FSCM may be achieved via an auxiliary structure called a *twin network* (Balke and Pearl, 1994). This is simply an FSCM where the original endogenous nodes (and their SEs) have been duplicated, while remaining affected by the same exogenous variables. As an example, Figure 3 depicts the twin network for the model in Figure 1. More general (e.g., involving more than two copies of the same endogenous node) and compact structures can be also considered (Shpitser and Pearl, 2007). Computing a counterfactual in the twin network of an FSCM is analogous to what is done with interventional queries provided that interventions and observations are associated with distinct copies of the same variable. BN inference eventually allows one to compute the counterfactual query in such an augmented model.

In the next section we discuss the computation of causal queries in PSCMs.




Figure 3: The graph of the twin network for the model in Figure 1.

### 3. Identifiability and M-Compatibility

As shown in the previous section, FSCMs allow to compute any causal query by standard BN inference algorithms. Yet, as the exogenous variables are typically latent, the marginal PMFs over them are rarely available and we should instead cope with a PSCM (i.e., the SEs) only. We assume the PSCM  $M$  comes paired up with a dataset  $\mathcal{D}$  of endogenous observations.

In the following we first formalise the concept of identifiability and discuss how to address partially identifiable queries (Section 3.1). This is based on the central notion of M-compatibility and its likelihood-based reformulation that becomes our main operational tool (Section 3.2).

#### 3.1. Partial Identifiability and M-Compatibility

As noticed in Section 2.2, given a PSCM, the endogenous observations in  $\mathcal{D}$  can be used to quantify the CPTs of the related endogenous BN (Tian, 2002). In this paper we use maximum likelihood estimation, whence we set:

$$P(v|\mathbf{w}_V) := \frac{n(v, \mathbf{w}_V)}{n(\mathbf{w}_V)}, \quad (6)$$

for each  $v \in \Omega_V$ ,  $\mathbf{w}_V \in \Omega_{\mathbf{w}_V}$ , and  $V \in \mathbf{V}$ , with  $n(\cdot)$  denoting the frequencies in  $\mathcal{D}$  of its argument.

We take the endogenous BN obtained in this way as our *ground truth*: it represents the joint mass function  $P(\mathbf{V})$  on the observable variables that results from the SCM we are given in input and the available data. Of course, this is not the same BN that we would obtain in the limit of infinite data; let us remark that we are not concerned with limit considerations in this paper, and we

rather work with finite samples and the related estimates. Consequently,  $P(\mathbf{V})$  is taken as our ‘objective’ piece of information on the domain, which is not under question.  $P(\mathbf{V})$  is not sufficient to draw causal inference, though. For this we need to go back to our SCM.

But the question is that there are in general many FSCMs that lead to the very same  $P(\mathbf{V})$ . Using Equation (2), we can characterise them as follows:

$$\sum_{\mathbf{u} \in \Omega_{\mathbf{U}}} \left[ \prod_{U \in \mathbf{U}} P(u) \cdot \prod_{V \in \mathbf{V}} P(v|\text{pa}_V) \right] = P(\mathbf{v}), \quad (7)$$

for each  $\mathbf{v} \in \Omega_{\mathbf{V}}$ , with  $(v, \text{pa}_V) \sim \mathbf{v}$  and  $u \sim \mathbf{u}$ . We call this *M-compatibility* to emphasise that it implicitly defines all and only the models (FSCMs) that are compatible with the PSCM given in input as well as with the data, in the sense that they all yield  $P(\mathbf{V})$ ; we denote the set of all these models by  $\mathcal{M}_{M,P}$ .

Let us illustrate this multiplicity with a simple example.

**Example 3.** Consider a PSCM with a Boolean endogenous variable  $V$ , a ternary exogenous parent  $U$  with  $\Omega_U := \{0, 1, 2\}$ , and a SE such that  $f_V(U = 0) = 0$  and  $f_V(U = 1) = f_V(U = 2) = 1$ . A dataset  $\mathcal{D}$  of observations for  $V$  allows to assess the endogenous PMF  $P(V)$ . Equation (7) rewrites therefore as:

$$\sum_{u=0}^2 P(u) \cdot P(v|u) = P(v), \quad (8)$$

for  $v = 0, 1$ . This is a linear system with solution:

$$P(U) = \begin{bmatrix} P(V = 0) \\ \alpha \\ P(V = 1) - \alpha \end{bmatrix}, \quad (9)$$

for each  $\alpha \in [0, 1 - P(V = 0)]$ .

That  $\mathcal{M}_{M,P}$  is not a singleton in general is the reason why one usually talks of ‘partial identifiability’ when it comes to causal inference: namely, that it is not possible in general to identify a single causal model that is consistent with the data; the best one can do is to rather consider the set  $\mathcal{M}_{M,P}$ . From this it also follows the notion of ‘partially identifiable query’: this is a causal query, such as an intervention or a counterfactual, whose outcome is partially indeterminate due to the multiplicity of FSCMs in  $\mathcal{M}_{M,P}$  we have to query to get a answer. If we ask for an expectation, for instance, each FSCM  $M'$  will yield a number  $q_{M'}$ ,

and all these numbers will eventually be summarised by a lower and an upper value, namely the shortest interval that contains all the delivered expectations:

$$\left[ \min_{M' \in \mathcal{M}_{M,P}} q_{M'}, \max_{M' \in \mathcal{M}_{M,P}} q_{M'} \right]. \quad (10)$$

Sometimes the interval will naturally collapse to a number, thus embodying the case of an identifiable query. But this will just be a special case. It occurs, for example, when the *do calculus* of Pearl (2009) and its extensions (see, e.g., Bareinboim and Pearl, 2012) can reduce interventional queries on PSCMs to observational ones (since the *do calculus* is sound and complete, it can be used to have an equivalent condition to identifiability in the interventional case).

Our focus in the rest of the paper is however on partial identifiability, which in practice means the computation of the bounds in Equation (10).

### 3.2. *M-Compatibility and Likelihood Maximisation*

To address causal queries by Equation (10), we need the set  $\mathcal{M}_{M,P}$  to be non-empty, this meaning that Equation (7) should admit at least one FSCM solution. If this was not the case, the endogenous distribution  $P(\mathbf{V})$  would be incompatible with the PSCM  $M$  under consideration, in the sense that no FSCM based on the same SEs of the PSCM could generate such a distribution when the latent variables get marginalised out.

Remember that  $M$ -compatibility is about the possibility of reconstructing the uncertainty about  $\mathbf{U}$  from the PSCM together with the endogenous PMF  $P(\mathbf{V})$ . If it fails, we know that the task is hopeless, because there is no  $P(\mathbf{U})$  that can eventually lead to  $P(\mathbf{V})$ . This may happen either because the sample is too small, or because the PSCM is a wrong model of the phenomenon under study. In either case, one should refrain from making inferences using jointly the PSCM and  $P(\mathbf{V})$  as they would be logically contradicting each other.

The notion of  $M$ -compatibility is central to this paper. Now we proceed to reformulate it in a way that makes it easier to use it as an operational tool.

By exploiting the factorisation in Equation (2), while also grouping together the terms corresponding to the different  $c$ -components, the log-likelihood of  $\mathcal{D}$  from an FSCM can be written as follows:

$$l(\theta_{\mathbf{U}}) := \sum_{c \in \mathcal{C}} \sum_{\mathbf{w}^c \in \Omega_{W_C}} n(\mathbf{w}^c) \cdot \log \sum_{\mathbf{u}^c \in \Omega_{U_C}} \left[ \prod_{U \in \mathbf{U}^c} \theta_u \cdot \prod_{V \in \mathbf{V}^c} P(v|\text{pa}_V) \right], \quad (11)$$

with  $(u, v, \text{pa}_V) \sim (\mathbf{u}^c, \mathbf{w}^c)$  for each  $c \in \mathcal{C}$  and  $\theta_{\mathbf{U}} := (\theta_u)_{U \in \mathbf{U}, u \in \Omega_U}$ , with  $\theta_u$  being the unknown ‘true’ chance of  $U = u$ . For the endogenous BN of an FSCM, the log-likelihood of  $\mathcal{D}$  is instead:

$$\lambda(\theta_V) := \sum_{V \in \mathbf{V}} \sum_{v, \mathbf{w}_V} n(v, \mathbf{w}_V) \cdot \log \theta_{v|\mathbf{w}_V}, \quad (12)$$

where, for the sake of a light notation, the domains of the sums over  $v$  and  $\mathbf{w}_V$  are left implicit and  $\theta_V := (\theta_{v|\mathbf{w}_V})_{v \in \Omega_V, \mathbf{w}_V \in \Omega_{\mathbf{w}_V}, V \in \mathbf{V}}$ . Note that the conditional chances in  $\theta_V$  can be directly obtained from those in  $\theta_{\mathbf{U}}$ .

Equation (12) exhibits the decomposable structure of a multinomial likelihood. Such a concave function has a unique global maximum achieved where the relative frequencies in Equation (6) are attained by the arguments  $\theta_V$ , and no local maxima (see, e.g., Koller and Friedman, 2009). Let us denote as  $\lambda^*$  the value of global maximum of the function in Equation (12).

**Theorem 1.** *Let  $\mathcal{K}$  denote the set of joint mass functions  $P(\mathbf{U})$  consistent with Equation (7). For  $\mathcal{K} \neq \emptyset$ , the log-likelihood  $l$  in Equation (11) achieves its global maximum, equal to  $\lambda^*$ , if and only if  $\theta_{\mathbf{U}} \in \mathcal{K}$ . For  $\mathcal{K} = \emptyset$ , instead,  $l$  takes only values strictly smaller than  $\lambda^*$ .*

In other words, global optimality of the log-likelihood is tantamount to finding an FSCM based on the given PSCM compatible with the data, this also allowing to decide M-compatibility, as formalised by the following result.

**Corollary 1.** *The function in Equation (11) achieving its global maximum  $\lambda^*$  is an equivalent condition for M-compatibility.*

This corollary is important because it allows us to narrow the attention to the M-compatible models only via likelihoods, thus enabling us to use a simple numerical test of M-compatibility.

Before using such a test, in the next section we show how CNs can be used to address the exact computation of a partially identifiable query as well as the test of M-compatibility in a way alternative to the one provided by Corollary 1.

#### 4. Credal Networks for Partial Identifiability

CNs offer a suitable formalism to address partial identifiability as in Equation (10). Exactly as a single FSCM is equivalent to a BN, the FSCMs in  $\mathcal{M}_{M,P}$  (all with the same SEs) can be regarded as a CN. To achieve that, a trivial but highly

inefficient approach, inspired by Antonucci and Zaffalon (2008), would consist in adding an auxiliary parent variable whose states index all the exogenous PMF specifications in  $\mathcal{K}$ , these clearly being in a one-to-one correspondence with the elements of  $\mathcal{M}_{M,P}$ . We derive instead a compact CN specification for Markovian models (Section 4.1), and then extend it to more general cases (Section 4.2). Such a mapping allows to virtually reduce causal inference to CN inference, while also leading to a new characterisation of the computational complexity of the latter. Those results are demonstrated by a number of examples leading to a deeper discussion on the importance of testing M-compatibility (Section 4.3).

#### 4.1. Credal Network Mapping in the Markovian Case

We already noticed how the process of specifying the endogenous BN for Markovian PSCMs is particularly simple:  $\mathbf{W}_V$ , i.e., the parents of  $V$  in the endogenous BN, are just the endogenous parents of  $V$  in the original model, for each  $V \in \mathbf{V}$ . This allows for a decomposition of the non-linear constraints in Equation (7), which are mixing the marginal PMFs for the different exogenous variables. We consequently obtain a separate set of linear constraints for each exogenous variable. This corresponds to a CS specification for each  $U \in \mathbf{U}$  and, overall, a proper CN with no need of auxiliary variables. The procedure is shown in Algorithm 1, where, just for the sake of readability, we ignore the case of endogenous nodes without exogenous parents and we do not explicitly write the normalisation and non-negativity constraints of CSs.

---

**Algorithm 1:** CN map for the Markovian case.

---

**input :** A Markovian PSCM  $M$  over  $(\mathbf{U}, \mathbf{V})$  and a BN  $P(\mathbf{V})$   
**output:** A credal set  $K(\mathbf{U})$

```

1 for  $V \in \mathbf{V}$  do
2    $U \leftarrow \text{Pa}_V \cap \mathbf{U}$ ;
3    $\mathbf{W}_V \leftarrow \text{Pa}_V \setminus \{U\}$ ;
4    $K(U) \leftarrow \left\{ P(U) \left| \begin{array}{l} \sum_{u \in \Omega_U} P(u) \cdot \mathbb{I}[f_V(u, \mathbf{w}_V) = v] = P(v|\mathbf{w}_V) \\ \forall v \in \Omega_V, \mathbf{w}_V \in \Omega_{\mathbf{W}_V} \end{array} \right. \right\}$ 
5 end
```

---

Note that we try to keep notation simple by denoting the collection of credal sets  $\{K(U)\}$ , with  $U \in \mathbf{U}$ ,  $K(\mathbf{U})$ —while also abusing terminology by calling the latter a credal set.

The algorithm correctness is guaranteed by the following result.

**Theorem 2.** *For a Markovian PSCM  $M$ , the FSCMs of  $\mathcal{M}_{M,P}$  can be represented as the BNs of a CN whose CCPTs are those induced by the SEs of  $M$  and the CS  $K(\mathbf{U})$  returned by Algorithm 1.*

The result implies that in the Markovian case we can regard the computation of the bounds of a PSCM query in Equation (10) as an inference task on the CN based on the output of Algorithm 1. In spite of the hardness of CN inference, the mapping is not increasing the complexity of PSCM queries as shown by the following result.

**Theorem 3.** *The computation of post-interventional queries for single variables in polytree-shaped PSCMs is NP-hard.*

The proof of the above result, reported in Appendix A, is based on the analogous complexity result for CNs derived by da Rocha and Cozman (2002). Other complexity results for CNs might be similarly applied to PSCMs by exploiting the CN formulation proposed by Cozman and Mauá (2017).

In the next section we discuss how to cope with non-Markovian PSCMs.

#### 4.2. Beyond Markovianity

In order to extend the procedure discussed in the previous section, let us first consider a simple non-Markovian example.

**Example 4.** *Consider a PSCM  $M$  over the graph in Figure 4. The endogenous variables  $V_1$  and  $V_2$  are Boolean, while their common exogenous parent  $U$  is such that  $\Omega_U := \{0, 1, \dots, 4\}$ . For  $V_1$ , we have  $f_{V_1}(U = u)$  equal to zero for  $U = 0, 3, 4$  and one for the other states of  $U$ . For  $V_2$ , we have instead  $f_{V_2}(U = u, V_1 = 0) = 0$  for  $u = 0, 2$  and  $f_{V_2}(U = u, V_1 = 1) = 0$  for  $U = 2, 3$ , while the SEs return one in the other cases. Equation (7) rewrites as:*

$$\sum_{u \in \Omega_U} P(u) \cdot P(v_1|u) \cdot P(v_2|u, v_1) = P(v_1, v_2), \quad (13)$$

to be considered for each  $v_1 \in \Omega_{V_1}$  and  $v_2 \in \Omega_{V_2}$ . In the sum on the left-hand side, the terms corresponding to values of  $u$  that are not simultaneously consistent, through SEs  $f_{V_1}$  and  $f_{V_2}$ , with both  $v_1$  and  $v_2$ , are zero. We thence rewrite Equation (13) as:

$$\sum_{u \in \Omega_U: \substack{f_{V_1}(u)=v_1, \\ f_{V_2}(u, v_1)=v_2}} P(u) = P(v_1, v_2). \quad (14)$$

Equation (14) defines a linear system analogous to the one in Example 3, whose solutions define a CS  $K(U)$ . As an example, if the joint endogenous PMF is such that  $P(V_1 = 0, V_2 = 0) = \frac{1}{5}$ ,  $P(V_1 = 1, V_2 = 0) = \frac{2}{5}$ , and  $P(V_1 = 0, V_2 = 1) = \frac{4}{15}$ , the corresponding linear constraints becomes  $P(U = 0) = \frac{1}{5}$ ,  $P(U = 1) = \frac{2}{15}$  and  $P(U = 2) = \frac{2}{5}$ , and  $P(U = 3) + P(U = 4) = \frac{4}{15}$ . The elements of the corresponding CS  $K(U)$  can be therefore parametrised as  $P(U) = [\frac{1}{5}, \frac{2}{15}, \frac{2}{5}, \alpha, \frac{4}{15} - \alpha]$  with  $\alpha \in [0, \frac{4}{15}]$ .




Figure 4: The causal graph of a semi-Markovian, but not Markovian, PSCM.

In the non-Markovian case, the common exogenous parents of two or more endogenous variables are called *confounders*. Confounders express the lack of Markovianity also at the PSCM level, being input variables common to two or more SEs. In these cases, the surjectivity we assume for single SEs (see Section 2.2) is extended to the joint SE involving all the SEs with the same confounder in input. For instance, the joint SE  $(f_{V_1}, f_{V_2})$  in Example 4 is *jointly surjective* as all the four joint states of  $(V_1, V_2)$  can be obtained by (at least) a value of the input variable  $U$ .

Let us call *quasi-Markovian* a PSCM (or FSCM) such that  $|\mathbf{U}^c| = 1$  for each  $c \in \mathcal{C}$ , this meaning that no endogenous variable has more than one exogenous parent. This is for instance the case of the PSCM in Figure 1, as well as the one in Figure 4. In a quasi-Markovian PSCM, notation  $\text{Pa}'_V$  is used for the endogenous parents of  $V$ , i.e.,  $\text{Pa}'_V := \text{Pa}_V \cap \mathbf{V}$ . It is easy to check that  $\text{Pa}'_V \subseteq \mathbf{W}_V$ .

A procedure analogous to that in Example 4 can be derived for any quasi-Markovian PSCM. As in the previous section, the key point is that the constraints on the marginal PMF of an exogenous variable  $U \in \mathbf{U}$  imposed by the consistency with the endogenous PMF can be specified separately from those of the other exogenous variables. This corresponds to Algorithm 2 that allows to derive the CS  $K(\mathbf{U})$  and hence obtain a CN from a quasi-Markovian model.

It is not difficult to check that Algorithm 2 coincides with Algorithm 1 in the case of Markovian inputs. The relation between the two procedures is even stronger as shown by the following result, which can be regarded as a generalisation of Theorem 2.



---

**Algorithm 2:** CN map for the quasi-Markovian case.

---

**input :** A quasi-Markovian PSCM  $M$  over  $(\mathbf{U}, \mathbf{V})$  and a BN  $P(\mathbf{V})$   
**output:** A credal set  $K(\mathbf{U})$

```

1 for  $U \in \mathbf{U}$  do
2    $c \leftarrow$  index of the c-component of  $U$ ;
3    $\gamma \leftarrow \emptyset$ ;
4   for  $w^c \in \Omega_{W^c}$  do
5      $\Omega_U^{w^c} \leftarrow \left\{ u \in \Omega_U \mid \begin{array}{l} f_V(u, \text{pa}'_V) = v \\ \forall V \in V^c, (v, \text{pa}'_V) \sim w^c \end{array} \right\}$ ;
6      $\gamma \leftarrow \gamma \cup \left\{ \sum_{u \in \Omega_U^{w^c}} P(u) = \prod_{V \in V^c} P(v | w_V), (v, w_V) \sim w^c \right\}$ 
7   end
8    $K(U) \leftarrow \{P(U) \mid \gamma\}$ 
9 end

```

---

**Theorem 4.** *For a quasi-Markovian PSCM  $M$ , the FSCMs of  $\mathcal{M}_{M,P}$  are the BNs associated with a CN whose CCPTs are those induced by the SEs of  $M$  and the CS  $K(\mathbf{U})$  returned by Algorithm 2.*

Regarding complexity, unlike Algorithm 1, where the number of constraints defining the CSs of the CN roughly corresponds to the size of the degenerate CPTs defining the SEs in the input PSCM, the bottleneck of Algorithm 2 is the loop over lines 4–7, which takes  $O(2^{|W^c|})$  time. For quasi-Markovian models,  $W^c$  is to the union of the children of  $U$  (i.e.,  $V^c$ ) and their parents, i.e., the Markov blanket of  $U$ . Setting a bound to this number (e.g., having that confounders only act on pairs of endogenous variables and a bounded indegree) would therefore make the algorithm polynomial.

The above result allows for addressing the computation of the bounds in Equation (10) through CN inference even in the quasi-Markovian case. In principle, any semi-Markovian model can be turned into quasi-Markovian, e.g., by clustering all  $U$  variables of a c-component into a single one; yet this neglects the exponential blowup in the computation that follows as a consequence.<sup>2</sup>

Overall we mapped a hard task (cf. Theorem 3) to another hard task (Mauá et al., 2014). Compared to analogous efforts of Duarte et al. (2023), where a

---

<sup>2</sup>Note also that Theorem 2.4 of Zhang et al. (2022) allows to extend our results to models with continuous exogenous variables, provided the endogenous ones remain discrete.

mapping to general polynomial programming has been derived, CN inference represents a more specific field that has been subject of intense investigation in the last three decades (Mauá and Cozman, 2020), and for which dedicated and stable solvers are freely available (Huber et al., 2020).

Moreover, as discussed in the next section, Algorithm 2 allows to exactly test the M-compatibility (see Section 3.2) with quasi-Markovian PSCMs, this corresponding to the feasibility of the linear constraints of each  $U \in \mathbf{U}$ . In the next section we discuss this procedure by also advocating the importance of checking M-compatibility and the role of SEs to achieve that.

#### 4.3. M-Compatibility and the Limits of Rung 2 of Pearl's Hierarchy

Let us first consider a simple example to be used along the section to clarify our findings.

**Example 5.** *Consider the setup from Mueller et al. (2022, Section 4.1) referring to a study about the recovery  $Y$  of patients of gender  $Z$  possibly subject to treatment  $X$ . A sample of 700 patients is considered (the corresponding frequencies are in Table 1). The authors show that for a PSCM whose endogenous BN has the same graph of that induced from the causal graph in Figure 5, given the sample, the PNS for treatment ( $X$ ) on effect ( $Y$ ) is no greater than 0.01 (this is the value rounded to the second decimal place; a more precise estimate is 0.015). They tell this for any PSCM with that graph, irrespectively of the specific SEs it uses; they only require that the endogenous PMF factorises according to the (endogenous BN induced by that) graph.*

*Our CN approach needs instead the PSCM specification, i.e., the SEs. Let us assume  $M$  canonical. Because of Equation (5), this corresponds to  $|\Omega_{U_Z}| = 2$ ,  $|\Omega_{U_X}| = 4$  and  $|\Omega_{U_Y}| = 16$ . To obtain canonical SEs, we set  $f_Z$  equal to the identity map, while for  $f_X$ :*

$$f_X(Z, U_X = 0) := \neg Z, \quad (15)$$

$$f_X(Z, U_X = 1) := 0, \quad (16)$$

$$f_X(Z, U_X = 2) := Z, \quad (17)$$

$$f_X(Z, U_X = 3) := 1. \quad (18)$$

*The relations between  $Y$  and  $(X, Z)$  according to  $f_Y$  induced by the sixteen states of  $U_Y$  can be similarly listed. M-compatibility with data in Table 1 corresponds to the constraints in line 4 of Algorithm 1 (the model is Markovian). For  $U_Z$ , this*

trivially means  $P(U_Z = 0) = P(Z = 0) = \frac{470}{700}$  and  $P(U_Z = 1) = P(Z = 1) = \frac{230}{700}$ . For  $P(U_X)$  we have instead the two following independent constraints:

$$P(U_X = 1) + P(U_X = 2) = P(X = 0|Z = 0) = \frac{116}{470}, \quad (19)$$

$$P(U_X = 0) + P(U_X = 1) = P(X = 0|Z = 1) = \frac{120}{230}. \quad (20)$$

Four independent linear constraints can be similarly obtained for  $P(U_Y)$ .

By standard linear programming tools, we might trivially check that these constraints are feasible and define the non-empty CSs  $K(U_X)$ ,  $K(U_Y)$ , and  $K(U_Z)$ . This proves the M-compatibility of the endogenous PMF based on the data in Table 1. CN inference eventually provides the PNS bounds. These bounds are computed in the twin network of Figure 6 by first performing the interventions in the two copies of the  $X$  variables and then jointly querying the  $Y$  variables. The resulting bounds coincides with the one presented by Mueller et al. (2022) on the basis of their formulae.

Gender (Z)	Treatment (X)	Recovery (Y)	#
0	0	0	2
0	0	1	114
0	1	0	41
0	1	1	313
1	0	0	107
1	0	1	13
1	1	0	109
1	1	1	1

Table 1: Data from an observational study involving three Boolean variables (Mueller et al., 2022, Section 4.1). The state equal to one means *female* for  $Z$ , *treated* for  $X$  and *recovered* for  $Y$ .

The M-compatibility with the endogenous PMF induced by the data in Table 1 for the Markovian model in Example 5 reflects a general result about M-compatibility and Markovianity.

**Theorem 5.** *In a canonical Markovian PSCM  $M$  over  $(\mathbf{U}, \mathbf{V})$  any  $P(\mathbf{V})$  is M-compatible.*

In practice, by Theorem 5, if the SEs of a Markovian SCM are not available and we adopt a canonical specification, compatibility is guaranteed.




Figure 5: The causal graph of a Markovian PSCM.




Figure 6: The twin network of the of the PSCM in Figure 5.

The bounds in Mueller et al. (2022) correctly predict the right numbers we get in Example 5. But we are not only assuming that the data factorise according to the graph in Figure 5, as the mentioned paper does. We are also assuming that the specification is canonical (i.e., all the mechanisms are possible). What happens if we do not—considered that the latter is not an assumption required to derive the bounds. This is discussed in the next example, still based on the setup of Example 5.

**Example 6.** *Assume that, thanks to some expert knowledge, the state  $U_X = 3$  is deemed impossible, this meaning that a deterministic mechanism forcing the treatment of all the patients is considered unrealistic. Such an assumption preserves the separate surjectivity of  $f_X$  and also the joint surjectivity of the model. In other words there is at least a joint state of  $(U_X, U_Y, U_Z)$  inducing any possible joint state of  $(X, Y, Z)$ . Under this additional information, the negation of Equation (19) rewrites as:*

$$P(U_X = 0) = P(X = 1|Z = 0) = \frac{354}{470}, \quad (21)$$

*and this is clearly inconsistent with Equation (20).*

This means that the model obtained by dropping state  $U = 3$  cannot generate the distribution that we see in the data: namely, the data are incompatible with such a non-canonical model. Note also that this shows that the endogenous distribution of the data can factorise according to the graph, while at the same time, it may be incompatible with the related model. This may happen because the latter form of incompatibility depends on the SEs, not only on the graph.

In fact, as shown by the next example, M-incompatibility is subtle in that it may be based on an endogenous probability that is inconsistent with the very model used to produce the data!

**Example 7.** *In the same setup of Example 5 assume that some expert tells us that, out the sixteen states of  $U_Y$ , only three of them are possible, namely those indexing the following logical relations:*

$$f_Y(X, Z, U_Y = 0) = X \vee \neg Z, \quad (22)$$

$$f_Y(X, Z, U_Y = 1) = \neg X \vee \neg Z, \quad (23)$$

$$f_Y(X, Z, U_Y = 2) = \neg X \wedge Z. \quad (24)$$

*Note that this preserves surjectivity. Call  $M'$  the reduced SCM that we obtain from the canonical specification (in  $M$ ) by keeping only the values 0, 1, 2 both for  $U_X$  and  $U_Y$ . Consider the following distributions for the exogenous variables in  $M'$ :  $P(U_Y) := [0.47, 0.439, 0.091]$ ,  $P(U_X) := [0.677, 0.000, 0.323]$ , and  $P(U_Z)$  taking the same values of  $P(Z)$  as in the case of  $M$ . Because of the previous discussion, the data in Table 1 are incompatible with  $M'$ . And yet they can be produced by  $M'$ , because  $P(X, Y, Z)$  is a positive distribution under  $M'$ . So any data can be generated—with different log-likelihoods. In particular, the ratio between the log-likelihood of  $M'$  and that of  $M$  is 0.71: thus it turns out that it is not at all unlikely to produce the data in Table 1 with model  $M'$ . And since those data factorise according to the graph, as before, we should be allowed to apply the bounds as before, claiming that the upper bound of PNS for  $M'$  is 0.015. But the PNS for  $M'$  is 0.15, 10 times above the bound. And so the bounds of Mueller et al. (2022, Section 4.1) fail here.*

They fail because ‘factorising according to the graph’ is not strong enough an assumption to derive bounds; we should ask for  $M$ -compatibility. Stated differently, whenever we produce formulae, or algorithms, for computing counterfactuals, we should make sure that  $M$ -compatibility holds, otherwise the results will be unwarranted.

At this point, one might be tempted to escape the problem of testing  $M$ -compatibility in absence of SEs (i.e., in presence of data and the causal graph alone) by using a canonical specification of the SEs, given that the resulting model is always compatible with the data (cf. Theorem 5). Even more so that recent efforts (Duarte et al., 2023; Zhang et al., 2022) appear to provide a canonical specification for general PSCMs; this would seem to enable doing general counterfactual inference without SEs. Yet, canonical specifications do not seem to provide us with a safe way out, as shown next.

Given a canonical Markovian PSCM  $M$ , any other Markovian PSCM  $M'$  over the same endogenous variables and graph can be defined by dropping some values of the exogenous variables, this being equivalent to remove from the SEs the corresponding mechanisms. That is,  $M'$  is defined via sets  $\{\Omega'_U\}_{U \in \mathbf{U}}$  such that, for each  $U \in \mathbf{U}$ ,  $\Omega'_U \subseteq \Omega_U$  are the states of  $U$  indexing the deterministic relations of the corresponding SE of  $M'$ . As a consequence we can regard a canonical PSCM as the set of all the possible non-canonical PSCMs. Some of these may be incompatible with the data, though. So how does inference in the canonical PSCM and in the set of non-canonical PSCMs relate to each other?

To answer this question, let us say that  $M$  *embeds*  $M'$  if and only there is at

least a  $P(\mathbf{U}) \in \mathcal{K}$  (see Theorem 1 for the definition of  $\mathcal{K}$ ) assigning  $P(U = u) = 0$  to all values  $u$ , of all exogenous variables  $U$ , corresponding to SEs that should be dropped to obtain  $M'$  from  $M$ .<sup>3</sup> In other words, embedding a sub-model  $M'$  means to ‘soft’-drop the SEs that do not belong to it via zero probabilities. The following holds:

**Theorem 6.** *A Markovian canonical PSCM  $M$  cannot embed incompatible models.*

As a consequence, running an inference algorithm in a canonical model corresponds to *automatically discarding* the incompatible models while using only the compatible ones. The important implication here is that if the true underlying model is not compatible with the available data, the results obtained by using the canonical model will be unwarranted as an approximation to the actual one. This is in fact the reason why in Example 7 the PNS interval obtained by the canonical model does not contain the actual value of PNS.

Overall, the lesson appears to be that we cannot have guaranteed bounds without knowing the SEs of the underlying SCM. The relation of M-compatibility to sample size appears therefore to be a subject that needs to be studied more deeply in order to at least provide probabilistic guarantees on the delivered bounds.

## 5. Causal Expectation Maximisation

In spite of the hardness of causal inference (cf. Theorem 3), the mapping to CN inference derived in the previous section may allow to take advantage of existing CN algorithms. The limitations are those related to non-quasi-Markovian models and to the intrinsic hardness of CN inference. Nevertheless, we show in this section that it is possible to leverage the particular optimisation required by causal queries, which leads to an iterative EM approach that provides an approximate range for partially identifiable queries for general PSCMs (Section 5.1). We also derive credible intervals to characterise the quality of such an approximation (Section 5.2).

---

<sup>3</sup>The analogous definition in Zaffalon et al. (2021) mistakenly reported  $P(U \in \Omega'_U) > 0$  instead of  $P(U \in \Omega'_U) = 1$  as in this paper.

### 5.1. EM for Causal Computations

Consider a partially identifiable query as in Equation (10). We already noticed that the set of FSCMs  $\mathcal{M}_{M,P}$  is in a one-to-one correspondence with the elements of the set  $\mathcal{K}$  considered by Theorem 1. Thus, in principle, if  $\mathcal{K}$  is available, one could compute the bounds of a query by optimising the corresponding function of  $P(\mathbf{U})$  over  $\mathcal{K}$ .

In practice coping with partial identifiability is demanding even with simple topologies and queries (cf. Theorem 3) and we should generally consider methods to compute approximate bounds. As a matter of terminology, here and in the following we shall call *range* the values spanned by the counterfactual query inside the approximate bounds. As mentioned already in the Introduction, the range is an inner approximation of the actual counterfactual interval as identified by the exact bounds.

Corollary 1 provides a good match w.r.t. developing approximations: it tells us that sampling the global optimum points of the log-likelihood corresponds to sampling from, and hence approximating,  $\mathcal{K}$ . And the crucial observation is that we can easily sample the optimum points of the log-likelihood with the *expectation-maximisation* (EM) scheme (Dempster et al., 1977): in fact, exogenous variables are missing at random in  $\mathcal{D}$ , just because they are latent (missing with probability one).

In particular, given an initialisation  $P_0(\mathbf{U})$ , the EM algorithm consists in regarding the posterior probability  $P_0(u|\mathbf{v})$  as a *pseudo-count* for  $(u, \mathbf{v})$ , for each  $\mathbf{v} \in \mathcal{D}$ ,  $u \in \Omega_U$  and  $U \in \mathbf{U}$  (E-step). A new estimate is consequently obtained as  $P_1(u) := |\mathcal{D}|^{-1} \sum_{\mathbf{v} \in \mathcal{D}} P_0(u|\mathbf{v})$  (M-step). This scheme, called EMCC (*EM for Causal Computation*) is iterated until convergence. Algorithm 3 depicts the EMCC pseudo-code.

Subroutine `initialise` (line 1) provides a random initialisation of the exogenous PMFs, while `components` (line 2) finds the c-components of  $M$ . A restriction of the dataset to the variables in the  $c$ -th component is achieved in line 4. In line 8 d-separation properties allow to replace  $P(U|\mathbf{v})$  with  $P(U|\mathbf{w}^c)$  when  $U \in \mathbf{U}^c$ . Since  $P_{t+1}$  gets a higher log-likelihood than  $P_t$  (Koller and Friedman, 2009, Theorem 19.3), we adopt likelihood stationarity as a stopping criterion to decide convergence (line 10). As a side remark, note that each iteration of the loop in line 3 of Algorithm 3 can be executed in parallel to the others because of the d-separation among c-components. Something similar can be done at the dataset level when computing, by standard BN algorithms, the queries in line 8.



---

**Algorithm 3:** EM for Causal Computation (EMCC).

---

**input :** A PSCM  $M$  over  $(\mathbf{U}, \mathbf{V})$  and a dataset  $\mathcal{D}$  of observations of  $\mathbf{V}$   
**output:**  $P(\mathbf{U})$

```
1  $P_0(\mathbf{U}) \leftarrow \text{initialise}(M);$ 
2  $\{\mathbf{U}^c, \mathbf{V}^c\}_{c \in \mathcal{C}} \leftarrow \text{components}(M);$ 
3 for  $c \in \mathcal{C}$  do
4    $\mathcal{D}^c \leftarrow \mathcal{D} \downarrow \mathbf{W}^c;$ 
5    $t \leftarrow 0;$ 
6   repeat
7     for  $U \in \mathbf{U}^c$  do
8        $P_{t+1}(U) \leftarrow |\mathcal{D}|^{-1} \sum_{\mathbf{w}^c \in \mathcal{D}^c} P_t(U | \mathbf{w}^c)$   $t \leftarrow t + 1$ 
9     end
10  until  $l(P_{t+1}(\mathbf{U}^c)) = l(P_t(\mathbf{U}^c));$ 
11 end
```

---

Multiple EMCC runs, started with different seeds, will yield an approximating subset of  $\mathcal{K}$ . It is known (Wu, 1983) that the EM algorithm only converges to stationary points of the likelihood. In practice, after the convergence of EM we evaluate the log-likelihood and, if a value smaller than the global maximum is achieved, we reject the point. Empirically we have never observed a case of convergence to a stationary point that is either a global maximum or a saddle point.

### 5.2. Deriving Credible Intervals for EMCC Inferences

In this section we characterise the accuracy of our EMCC procedure in terms of credible intervals. Let us consider a, possibly partially identifiable, query, whose exact bounds defined as in Equation (10) are  $[a^*, b^*]$ . In this section we assume that  $[a^*, b^*] \subseteq [0, 1]$ , i.e., the query results in a probability value, to simplify the notation. The results listed below, however can be easily extended to the general case of any bounded query  $[a^*, b^*]$ .

Say that  $k$  EMCC runs have been executed to approximate the query and denote the result as  $\rho := \{\pi_i\}_{i=1}^k$ . Let  $a := \min_{i=1}^k \pi_i$  and  $b := \max_{i=1}^k \pi_i$  denote their approximations and let  $L := b - a$  be the size of the range. By construction, we have  $a^* \leq a \leq b \leq b^*$ . To evaluate the quality of such an inner approximation we compute the probability of covering  $[a^*, b^*]$  if we assume that both  $a$  and  $b$  are subject to an error of  $\epsilon$  relative to  $L$ . The probability depends on how the EMCC outputs are distributed:

**Theorem 7.** Assume that the EMCC runs in  $\rho$  are distributed as a four parameter beta distribution, i.e.,  $\pi_i \sim \text{Beta}(\alpha, \beta, a^*, b^*)$ , for each  $i = 1, \dots, k$ . The following equality holds:

$$P\left(a - \varepsilon L \leq a^* \leq b^* \leq b + \varepsilon L \mid \rho\right) = \frac{\int_0^{\delta_b/2} \int_0^{\delta_a/2} P(x, y; L, \alpha, \beta, k) dx dy}{\int_0^{a+(1-b)} \int_0^{a+(1-b)-y} P(x, y; L, \alpha, \beta, k) dx dy}, \quad (25)$$

where  $P(x, y; L, \alpha, \beta, k)$  is equal to

$$\left( \frac{(L+x)^\alpha {}_2F_1(\alpha, 1-\beta, \alpha+1, \frac{L+x}{L+x+y}) - x^\alpha {}_2F_1(\alpha, 1-\beta, \alpha+1, \frac{x}{L+x+y})}{\alpha(L+x+y)^\alpha B(\alpha, \beta)} \right)^k, \quad (26)$$

${}_2F_1$  is the Gaussian (ordinary) hypergeometric function,  $B(\alpha, \beta)$  is the beta function evaluated at  $\alpha, \beta > 0$  and:

$$\delta_a := \begin{cases} 2L\varepsilon & \text{if } \varepsilon \leq \frac{a}{L} \\ 2a & \text{otherwise} \end{cases} \quad (27)$$

$$\delta_b := \begin{cases} 2L\varepsilon & \text{if } \varepsilon \leq \frac{1-b}{L} \\ 2(1-b) & \text{otherwise.} \end{cases} \quad (28)$$

Note that Equation (25) is a valid probability if  $\varepsilon$  is such that:  $\delta_a \leq 2(1-L) - \delta_b$  and  $\delta_b \leq 1-L$ .

There are some notable special cases to the above procedure for which a simplified version of Equation (25) can be derived.

**Corollary 2.** For  $a = 0$ , we have  $a^* = 0$ ,  $L = b$ , and  $\pi_i \sim \text{Beta}(\alpha, \beta, 0, b^*)$ , for each  $i = 1, \dots, k$ , and hence:

$$P\left(b^* \leq b + \varepsilon L \mid \rho\right) = \frac{\int_0^{\delta/2} P(y; L, \alpha, \beta, k) dy}{\int_0^{(1-b)} P(y; L, \alpha, \beta, k) dy}, \quad (29)$$

with  $\delta = 2L\varepsilon$  for  $\varepsilon \leq \frac{1-b}{L}$ , where  $P(y; L, \alpha, \beta, k)$  is equal to

$$\left( \frac{L^\alpha {}_2F_1(\alpha, 1-\beta, \alpha+1, \frac{L}{L+y})}{\alpha(L+y)^\alpha B(\alpha, \beta)} \right)^k. \quad (30)$$

Similarly, for  $b = 1$ , we have  $b^* = 1$ ,  $L = 1 - a$ , and  $\pi_i \sim \text{Beta}(\alpha, \beta, a^*, 1)$ , for each  $i = 1, \dots, k$  and hence:

$$P\left(a^* \geq a - \varepsilon L \mid \rho\right) = \frac{\int_0^{\delta/2} P(x; L, \alpha, \beta, k) dx}{\int_0^a P(x; L, \alpha, \beta, k) dx}, \quad (31)$$

with  $\delta = 2L\varepsilon$  for  $\varepsilon \leq \frac{a}{L}$ , where  $P(x; L, \alpha, \beta, k)$  is equal to

$$\left( \frac{(L+x)^\alpha {}_2F_1(\alpha, 1-\beta, \alpha+1, \frac{L+x}{1+L+x}) - x^\alpha {}_2F_1(\alpha, 1-\beta, \alpha+1, \frac{x}{1+L+x})}{\alpha(1+L+x)^\alpha B(\alpha, \beta)} \right)^k. \quad (32)$$

Finally, for  $a = b$ , i.e., all  $k$  runs in  $\rho$  are equal, then:

$$P(a^* = b^* \mid \rho) = 1 + 9/3^k - 8/2^k. \quad (33)$$

The last case in the corollary above implies that nine equal runs make identifiability probable with 99% confidence.

Theorem 7 is a proper extension of the original EMCC characterisation (Zafalon et al., 2021), as proved by the following result.

**Corollary 3.** *If the EMCC runs in  $\rho$  are uniformly distributed in  $[a^*, b^*]$ , i.e., for each  $i = 1, \dots, k$ ,  $\pi_i \sim \text{Beta}(1, 1, a^*, b^*)$ , then we get:*

$$P\left(a - \varepsilon L \leq a^* \leq b^* \leq b + \varepsilon L \mid \rho\right) = \frac{1 + (1 + 2\varepsilon)^{2-k} - 2(1 + \varepsilon)^{2-k}}{(1 - L^{k-2}) - (k-2)(1-L)L^{k-2}}. \quad (34)$$

In practice we can use the theorems above as stopping criterion for the EMCC procedure. After any  $k$  EMCC runs, we compute  $a$  and  $b$  and, given a relative error  $\varepsilon$  we regard as acceptable, compute the probability in Equation (25). If the corresponding probability is sufficiently high, we stop iterating EMCC, otherwise we keep iterating the procedure to collect new points for  $\rho$  and achieve greater probabilities.

Note that the computation of the probability in Equation (25) requires estimating the parameters  $\alpha$  and  $\beta$  of the distribution  $\text{Beta}(\alpha, \beta, a^*, b^*)$ . We estimate them via a maximum likelihood procedure over the  $k$  values collected in  $\rho$ . Moreover, since  $a^*, b^*$  are unknown, we also need to choose the range of the  $\text{Beta}$  distribution. Here we set  $\text{Beta}(\alpha, \beta, a - \varepsilon L, b + \varepsilon L)$ . The integrals in Equations (25), (29) and (31) cannot be computed analytically and are thus estimated numerically.

## 6. Empirical Validation

We present here a numerical validation of our techniques to estimate the bounds of partially identifiable queries. This is achieved through simulations on an extensive benchmark of synthetic models (Section 6.1), a case study on palliative care (Section 6.2), and the analysis of a classical BN for medical diagnosis (Section 6.3).

### 6.1. Tests on Synthetic Models

Let us test both the CN mapping described in Section 4 and the EMCC algorithm introduced in Section 5. To do that we consider the computation of counterfactual, partially identifiable, queries on a benchmark of synthetic PSCMs.

*Sampling FSCMs.* For each experiment we first sample a ‘ground-truth’ FSCM  $M^*$ . The causal graph of  $M^*$  is sampled with the Erdős-Rényi model specialised to the case of directed acyclic graphs (Ide et al., 2004). We set the maximum in-degree equal to three and the maximum out-degree equal to two, with the total number of nodes specified as an input. We cope with semi-Markovian models, this implying that only the root nodes are associated with the exogenous variables. Thus, if the sampled graph is such that there are non-root nodes without root parents, we add to the graph a new root node as a parent. Let  $\mathcal{G}$  denote the causal graph obtained in this way. The root nodes of  $\mathcal{G}$  correspond to the exogenous variables  $\mathbf{U}$  of  $M^*$ , while the non-root nodes refer to the endogenous variables in  $\mathbf{V}$ . Just for the sake of simplicity, the variables in  $\mathbf{V}$  are assumed Boolean. Regarding the cardinality of the exogenous variables and the SEs, for the c-components including a single exogenous variable we adopt a canonical specification. If two or more exogenous variables are in the same c-component, we set instead  $|\Omega_U| = 16$ , for each  $U$  in the component. The SE of each endogenous variable in the c-component is obtained by sampling a surjective map determining the value of the endogenous from that of its endogenous parents, for each joint state of the exogenous parents. Let  $M$  denote the PSCM obtained in this way. FSCM  $M^*$  is eventually obtained by adding to  $M$  a (uniformly sampled) random PMF  $P^*(U)$  such that  $P(U = u) > 0$  for each  $u \in \Omega_U$ , for each  $U \in \mathbf{U}$ . Overall, we create a benchmark of 101 FSCMs, 39 of them being quasi-Markovian. The number of variables over the different models in the benchmark ranges from 5 to 19 (average 9.9), while the exogenous cardinality ranges from 3 to 256 (average 30.2) and the treewidth of  $\mathcal{G}$  from 2 to 4 (average 2.9).

*M-compatibility.* A dataset  $\mathcal{D}$  of endogenous observations is sampled from  $M^*$ . The PSCM  $M$  underlying FSCM  $M^*$  and the dataset  $\mathcal{D}$  (together with the endogenous BN quantified from  $\mathcal{D}$ ) are the inputs of each experiment. To check the M-compatibility of the (joint PMF associated with the) endogenous BN, we first use the BN to compute  $\lambda^*$  (see Equation (12)). After that, we perform a single EMCC run with  $M$  and  $\mathcal{D}$  as inputs to check whether, after convergence, the maximum value  $\lambda^*$  is achieved by the likelihood (cf. Corollary 1). If this is not case, we sample more observations until M-compatibility is guaranteed.

*Queries and Ground Truth.* As a partially identifiable (counterfactual) query, we consider a PNS having as *cause* and *effect* respectively the first and the last of the variables in  $V$ , assuming them sorted in a topological order. The bounds  $[a^*, b^*]$  can be computed only for quasi-Markovian models by exact CN inference. In practice this is feasible only with small models ( $< 10$  nodes). For the other models in the benchmark we use the EMCC (Algorithm 3) with a high ( $> 300$ ) number  $r$  of runs. Let  $[a_r, b_r]$  denote the range induced by those  $r$  runs. Taking a threshold probability  $P^* := 0.99$ , we check the minimum value  $\epsilon^*$  of  $\epsilon$  such that the probability in Equation (29) is greater than or equal to  $P^*$ . This allows to further approximate the true bounds as  $a^* \simeq a_r - \epsilon^* L$  and  $b^* \simeq b_r + \epsilon^* L$ . We take these as our ‘ground-truth’ bounds to be compared with the inner approximation obtained by the EMCC with a smaller number of runs. For quasi-Markovian models, we also use the CN mapping (Algorithm 2) linked with the approximate *ApproxLP* algorithm for CN inference. Note that also this procedure provides inner approximations (Antonucci et al., 2015). The quality of the (inner) approximation of  $[a^*, b^*]$  provided by the range  $[a, b]$  in output of one of our algorithms is described by the following relative root mean square error:

$$\text{RRMSE} := \sqrt{\frac{(a - a^*)^2 + (b - b^*)^2}{2(b^* - a^*)^2}}, \quad (35)$$

which is computed for each experiment.

*Implementation and Results.* Both the ECMM and the CN mapping have been implemented within the CREDICI library for causal inference (Cabañas et al., 2020).<sup>4</sup> CN inferences are computed instead by the CREMA library (Huber et al., 2020), which can be directly imported in CREDICI. Note that the CREDICI

---

<sup>4</sup>[github.com/idsia/credici](https://github.com/idsia/credici).

library also includes scripts for model and data generation, as well as the computation of the integrals to obtain the credible intervals discussed in Section 5.2. The scripts to reproduce the results discussed here and in Section 6.3 are stored in a separate repository, which also includes the case studies discussed in the next sections.<sup>5</sup> The simulations are sequentially executed on an AMD EPYC-7542 32-Core Processor with 256 GB.

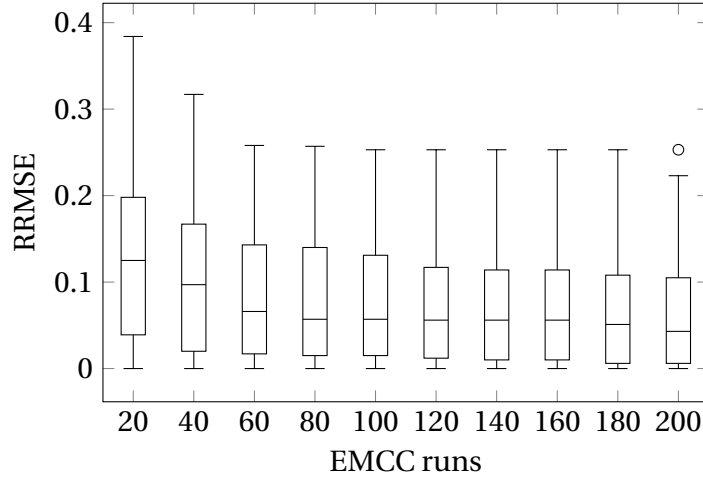


Figure 7: RRMSE vs. EMCC runs.

The boxplots in Figure 7 describe the RRMSE of the ranges obtained by the EMCC for an increasing number of runs. As expected, the accuracy of our ranges increases with the number of EMCC runs. With less than 200 runs, we already obtain an  $RRMSE < 0.1$  in most of the cases. With many runs and small errors, it becomes more difficult for the EMCC to further expand its ranges: this might explain the slower improvements in the right part of Figure 7.

Figure 8 depicts instead a comparison between EMCC and CN approaches. The ApproxLP algorithm invoked by the CN method is slightly more accurate, but, contrary to the EMCC, its direct application is restricted to quasi-Markovian models. Thus, we might reasonably regard the EMCC as the algorithm of choice to compute the ranges of partially identifiable queries.

<sup>5</sup>[github.com/IDSIA-papers/2023-IJAR-efficient-bounding](https://github.com/IDSIA-papers/2023-IJAR-efficient-bounding).




Figure 8: EMCC ( $r = 200$ ) vs. approximate CN approach on quasi-Markovian models.

## 6.2. A Counterfactual Analysis in Palliative Care

Figure 9 represents a causal model for the study of terminally ill cancer patients’ preferences with respect to their place of death: home or hospital. In fact most patients prefer to die at home, but the majority actually die in institutional settings. The study aimed at understanding interventions by health care professionals that can facilitate dying at home. The graph corresponds to the network proposed by Kern et al. (2020) reduced to the subset of variables for which data were available—variables have been binarised too.

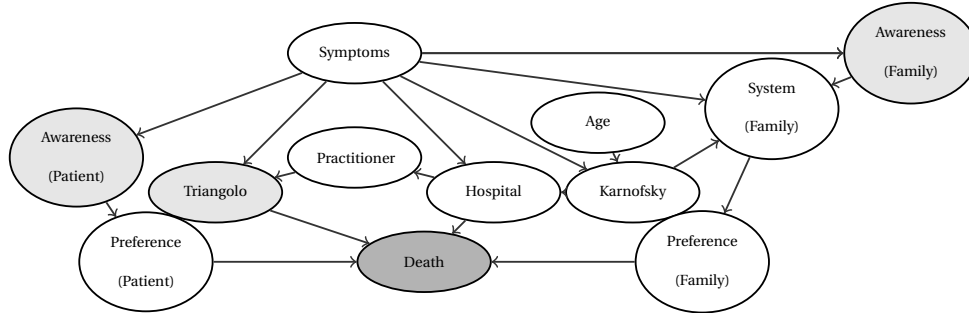


Figure 9: The model to study preferences about the place of death in oncology patients.

One can intervene on three variables in the network (light grey nodes in Figure 9): the patient’s and the family’s *awareness* of death (which involves communication with the doctors); and home assistance (provided by the *Triangolo* association). Our goal is to use the tools of causal analysis to understand what is the most important variable on which to act, and we measure importance by the PNS values having the variable *death* as effect variable (dark grey node

in Figure 9). The idea is that the one with highest PNS will lead to the highest increase of people that can die at home by its very reason.

To this end, we turn the casual graph into a Markovian SCM by adding an exogenous variable to each node. The model is taken Markovian as a consequence of the fact that all the potential confounders have been explicitly represented in the causal graph. As for SEs, we stick to the canonical representation since we want to be least-committal w.r.t. the true underlying mechanisms. Note that this induces a high cardinality for the exogenous variable associated with node *death* as, by Equation (5), we have  $|\Omega_{U_{\text{death}}}| = 2^{16}$ . We generate a sample of endogenous data from the original network model (we have no access to the patients' original data). The high cardinality of  $\Omega_{U_{\text{death}}}$  prevents an application of the CN algorithms, while the EMCC can be executed and we obtain our PNSs based on 160 EMCC runs (EM convergence achieved in around 500 iterations on average) in approximately 75 minutes. More specifically,  $[0.30, 0.31]$  is the resulting range when *Triangolo* is the intervened variables; similarly, *patient's awareness* gives  $[0.03, 0.10]$  and *family's awareness*  $[0.06, 0.10]$ . The conclusion is that one should clearly act on *Triangolo* first: for instance, by making *Triangolo* available to all patients, we should expect a reduction of people at the hospital by 30%. This would save money too, and would allow politicians to do economic considerations as to which amount it is even economically profitable to fund *Triangolo*, and have patients die at home, rather than spending more to have patients die at the hospital.

### 6.3. A Counterfactual Analysis with the Asia Network

In line with the preceding section, we perform a causal analysis based on a classical BN model, which is also related to a medical domain. We consider the *Asia* BN, whose graph is depicted in Figure 10. In this case, the aim is to determine the foremost causes to the occurrence of *Dyspnoea*. For this purpose, we compute the PNS values having the variable *Dyspnoea* as effect variable, and the variables *Bronchitis*, *Lung Cancer*, *Tuberculosis*, *Smoker* and *Asia* as causes.

As in the study on palliative care, we obtain from the causal graph a Markovian SCM by introducing an exogenous variable to each node and considering canonical SEs. A dataset of 5,000 instances is sampled from the initial BN. PNS ranges are obtained by means of the CN mapping (Algorithm 1) linked to exact and approximate (ApproxLP) CN algorithms. EMCC ranges are also computed with  $r = 71$  runs. The PNS ranges rounded to the second decimal place are in Table 2. Both approximate methods are very accurate. *Bronchitis* appears as






Figure 10: The graph of the Asia Bayesian network used for the causal analysis.

the most relevant variable for an intervention aiming to avoid *Dyspnoea*, while *Asia* is the least relevant.

Cause	CN (Exact)	CN (Approx)	EMCC
Bronchitis	[0.65, 0.76]	[0.65, 0.76]	[0.65, 0.76]
Lung Cancer	[0.37, 0.48]	[0.37, 0.48]	[0.37, 0.48]
Tuberculosis	[0.36, 0.46]	[0.36, 0.46]	[0.36, 0.46]
Smoker	[0.23, 0.53]	[0.24, 0.52]	[0.24, 0.51]
Asia	[0.01, 0.02]	[0.01, 0.02]	[0.01, 0.02]

Table 2: PNS ranges about the effect on *Dyspnoea* of different causes for the model in Figure 10.

## 7. Conclusions

We have presented two algorithms to do counterfactual inference in partially specified structural causal models: the first is based on a mapping to credal networks while the second on an EM scheme. The CN one is limited to quasi-Markovian models; it delivers exact inference for relatively small models and approximate inference otherwise. The EMCC algorithm is just approximate but works on any semi-Markovian SCM. Both algorithms appear to work rather efficiently and accurately. We have empirically verified this via a structured experimental validation based also on credible intervals to establish the quality of solutions obtained. In spite of a few alternative approaches in the literature to solve this kind of problems, we seem to be the first to do a systematic evaluation of the proposed algorithms as well as the first to publicly release the related code. We regard this as a relevant contribution of this work in order to start establishing fair comparisons of diverse approaches; in particular

the NP-hardness of the involved problems, proved here, implies that solutions will have to be pursued via heuristics in general, which should be compared on standard benchmarks.

Both algorithms assume that endogenous nodes are categorical, and both assume that structural equations are given—along with a dataset for the endogenous variables. Let us stress that the requirement that SEs are given is not as stringent as it might seem, given that we can produce the needed SEs by a preprocessing step if the actual ones are not available: this is possible thanks to recent work (Duarte et al., 2023; Zhang et al., 2022) that has introduced a ‘canonical’ specification of the SEs. This can be understood as a least-committal specification that, loosely speaking, can be used without loss of generality; the implication is that the output intervals will tend to be weaker compared to the case where the actual SEs are given. In this sense, our work is therefore as general as the works that do not assume the SEs to be given.

Another contribution of the present work is a theoretical analysis about the potential logical incompatibility of an SCM with (the empirical distribution inferred from) a dataset; we have called this M-incompatibility. We have observed that there are cases where M-incompatibility prevents one from obtaining guaranteed counterfactual bounds in absence of knowledge about the SEs, even when one uses the canonical specification. It is a problem that in principle can affect any method that tries to deliver bounds without knowing the actual SEs. It is unclear to us how severe this problem is in practice. For the time being we are pointing to a problem that does not seem to have obtained the due attention in the literature. We leave to future research a dedicated analysis that clarifies the problem in detail.

We regard the EMCC in particular as a very practical algorithm to compute counterfactual inference, which—in this respect—does not seem to have many competitors at the moment. There are however many directions to improve it starting from the work in this paper. Some recent papers have already explored a few of these: Zaffalon et al. (2023) have extended the EMCC to handle multiple data sources, of experimental and observational data, along with possible selection biases; Huber et al. (2023) have instead sped the EMCC up by one order of magnitude compared to the implementation described here, by a circuital compilation and some parallelisation steps. Note that the EMCC is query-agnostic at the moment: it aims at quantifying the uncertainty of the exogenous variables without targeting a specific counterfactual query. This means that once the uncertainty is quantified, any counterfactual query can be approximately solved without having to reconsider the data. However in

some situations it may be more efficient to use query-specific algorithms that evaluate the uncertainty of the exogenous variables only in relation to a given query. This would entail using some local optimisation that would more efficiently find the counterfactual range. Finally, another useful direction would be to extend the EMCC to continuous domains.

## Appendix A. Proofs

In this technical appendix we gather the proofs of the theorems presented in the paper together with some necessary lemma.

**Lemma 1.** *Equation (7) is equivalent to have the following equation:*

$$\sum_{\mathbf{u}^c \in \Omega_{\mathbf{U}}^c} \left[ \prod_{U \in \mathbf{U}^c} P(u) \cdot \prod_{V \in \mathbf{V}^c} P(v|\text{pa}(V)) \right] = \prod_{V \in \mathbf{V}^c} P(v|\mathbf{w}_V), \quad (\text{A.1})$$

satisfied for each  $c \in \mathcal{C}$ .

*Proof.* The result has been proved by Tian (2002). Equation (A.1) corresponds to Equation (4.41) in that work.  $\square$

**Proof of Theorem 1.** Let us first consider the case  $\mathcal{K} \neq \emptyset$ . Take  $\theta_{\mathbf{U}} \in \mathcal{K}$ . We have Equation (A.1) by Lemma 1. Putting this in Equation (11) we get:

$$l(\theta_{\mathbf{U}}) = \sum_{c \in \mathcal{C}} \sum_{\mathbf{w}^c} n(\mathbf{w}^c) \sum_{V \in \mathbf{V}^c} \log P(v|\mathbf{w}_V). \quad (\text{A.2})$$

Equation (A.2) corresponds to Equation (12) with the probabilities in Equation (6) that are giving the global maximum  $\lambda^*$ . This proves the sufficient condition. To prove the necessary condition, consider again Equation (A.1), which is equivalent to Equation (7). Assume, *ad absurdum*, that there is a  $\theta_{\mathbf{U}} \notin \mathcal{K}$  such that the log-likelihood in Equation (11) attains its global maximum. This means that Equation (A.1) should be violated for at least a value of  $\mathbf{W}^c$  in a c-components. In this case, putting Equation (A.1) in Equation (12) produces a value smaller than the global maximum  $\lambda^*$ , as the maximum is achieved if and only if the values in Equation (6) are used. The same deduction can be applied to any  $\theta_{\mathbf{U}}$  when  $\mathcal{K} = \emptyset$ .  $\square$

**Proof of Corollary 1.** The result trivially follows from Theorem 1.  $\square$

**Proof of Theorem 2.** To prove the theorem, let us first take an FSCM  $M' \in \mathcal{M}_{M,\bar{P}}$  and prove that this is also a BN of the CN based on the SEs of M and the CSs returned by Algorithm 1. By the definition of  $\mathcal{M}_{M,\bar{P}}$ ,  $M'$  satisfies Equation (7), which is equivalent to Equation (A.1) thanks to Lemma 1. As in a Markovian model  $U$  is the unique element of  $\mathbf{U}^c$  and  $\text{Pa}(V) = (U, \mathbf{W}_V)$ , this means:

$$\sum_{u \in \Omega_U} P(u) \cdot P(v|u, \mathbf{w}_V) = P(v|\mathbf{w}_V), \quad (\text{A.3})$$

and hence:

$$\sum_{u \in \Omega_U} P(u) \cdot \mathbb{I}[f_V(u, \mathbf{w}_V) = v] = P(v|\mathbf{w}_V). \quad (\text{A.4})$$

Equation (A.4) is the constraint in line 4 of Algorithm 1. This proves  $P(U) \in K(U)$  for each  $U \in \mathbf{U}$ , and this means that  $\mathcal{M}_{M,\bar{P}}$  is included in the (set of BNs associated with the) CN.

To prove the inverse inclusion, let us take a PMF  $P(U) \in K(U)$  for each  $U \in \mathbf{U}$ . This induces a (Markovian) FSCM  $M'$  based on M that should satisfy Equation (A.4) because of the definition of  $K(U)$  as in line 4 of Algorithm 1. But this means that also Equation (A.3) and hence Equation (A.1) is satisfied. Lemma 1 eventually implies  $M' \in \mathcal{M}_{M,\bar{P}}$  and hence the thesis follows.  $\square$

**Proof of Theorem 3.** The proof follows from the analogous result for CNs derived by da Rocha and Cozman (2002). The polytree-shaped CN used in that proof (Figure A.11) has degenerate CCPTs for the non-root nodes. We can intend those CCPTs as the SEs of a PSCM M and regard the variables associated with the non-root nodes of the CN as the endogenous variables of M. These endogenous nodes form a chain, and each node has a single exogenous parent apart from the first one in the chain, which has also a second exogenous parent. Even if such a model is non-Markovian, we can add an auxiliary exogenous parent to one of the two parents of the first node of the chain and obtain an equivalent Markovian model such that Algorithm 1 returns the CSs on the root node of the polytree-shaped CN.

As a query, the authors consider the upper bound of a marginal query in the last node of the chain. The task amounts to the identification of the upper bound of the causal effect on the last node of the chain given an intervention in an additional endogenous parent of the first node of the chain. An algorithm to bound interventional queries in PSCMs would therefore solve inference in polytree-shaped CN. This contradicts the result of the authors.  $\square$




Figure A.11: The polytree used in the complexity proof of da Rocha and Cozman (2002).

**Proof of Theorem 4.** For a quasi-Markovian model, Equation (A.1) rewrites as:

$$\sum_{u \in \Omega_U} \left[ P(u) \cdot \prod_{V \in V^c} P(v|u, \text{pa}'_V) \right] = \prod_{V \in V^c} P(v|\mathbf{w}_V), \quad (\text{A.5})$$

and hence:

$$\sum_{u \in \Omega_U} \left[ P(u) \cdot \prod_{V \in V^c} \mathbb{I}_{f_V(u, \text{pa}'_V) = v} \right] = \prod_{V \in V^c} P(v|\mathbf{w}_V), \quad (\text{A.6})$$

which coincides with the linear constraint for  $P(U)$  in line 6 of Algorithm 2. Such a remark allows to prove the thesis by a scheme perfectly analogous to that considered for the proof of Theorem 2.  $\square$

An alternative proof of Theorem 4 based on an arc reversal strategy can be found in Zaffalon et al. (2020).

**Proof of Theorem 5.** Because of Corollary 1, the thesis is equivalent to the fact that the constraints in Equation (7) can be satisfied for at least a specification of  $P(\mathbf{U})$ . Lemma 1 allows to replace Equation (7) with Equation (A.1), to be considered for each  $c \in \mathcal{C}$ . As  $M$  is Markovian, as already discussed for the proof of Theorem 2, we can rewrite Equation (A.1) as:

$$\sum_{u \in \Omega_U} P(u) \cdot \mathbb{I}_{f_V(u, \mathbf{w}_V) = v} = P(v|\mathbf{w}_V). \quad (\text{A.7})$$

to be satisfied for each  $v \in \Omega_V$  and  $\mathbf{w}_V \in \Omega_{\mathbf{w}_V}$ . The representation result provided by Druzdzel and Simon (1993, Theorem 1) can be used to prove that the linear constraints in Equation (A.7) can be always satisfied if the SEs are canonical. In their proof the authors consider a setup analogous to the current one but, instead of  $U$ , a continuous  $U' \in [0, 1]$  with a uniform density  $P(U')$  is considered. Say that  $\Omega_V := \{v_1, \dots, v_q\}$ . For each  $\mathbf{w}_V \in \Omega_{\mathbf{w}_V}$ , define the vector  $H_{\mathbf{w}_V} := \{h_{\mathbf{w}_V}^{(i)}\}_{i=0}^q$  such that  $h_{\mathbf{w}_V}^{(0)} := 0$  and  $h_{\mathbf{w}_V}^{(i)} := \sum_{j=1}^i P(v_j|\mathbf{w}_V)$  for each

$i = 1, \dots, q$ . Note that  $h_{\mathbf{w}_V}^{(i)} \leq h_{\mathbf{w}_V}^{(i+1)}$  for each  $i = 1, \dots, q-1$  and  $h_{\mathbf{w}_V}^{(q)} = 1$ . The authors show that Equation (A.7) is satisfied if the SE  $V = f'_V(U', \mathbf{W}_V)$  is such that:

$$f'_V(u', \mathbf{w}_V) := \left\{ v_i \in \Omega_V : u' \in \left[ h_{\mathbf{w}_V}^{(i)}, h_{\mathbf{w}_V}^{(i+1)} \right] \right\}, \quad (\text{A.8})$$

for each  $u' \in [0, 1]$  and  $\mathbf{w}_V \in \Omega_{\mathbf{W}_V}$ . A partition of  $[0, 1]$  is obtained by removing the left endpoints from the intervals in Equation (A.8) apart from the first one. Thus, for a given  $\mathbf{w}_V \in \Omega_{\mathbf{W}_V}$ , Equation (A.8) can be regarded as a discrete SE, mapping to  $V$  the values of a discretisation of  $U'$  based on the partition of  $[0, 1]$  induced by  $H_{\mathbf{w}_V}$ .

A *least common partition* is obtained from a set of partitions by putting together and sorting the endpoints of the intervals of all the partitions. Take the discretisation of  $U'$  induced by such a least common partition when considering all the partitions of  $[0, 1]$  induced by each  $\mathbf{w} \in \Omega_{\mathbf{W}_V}$ . In practice, the SE  $V = f'_V(\mathbf{W}_V, U')$  can be equivalently described by  $V = \hat{f}_V(\mathbf{W}_V, \hat{U})$ , where  $\hat{U}$  is a discrete variable whose states are in correspondence with the above considered set of discretisation intervals for  $U$ . The uniform density  $P(U')$  is consequently mapped to a PMF  $P(\hat{U})$  such that  $P(\hat{U} = \hat{u})$  is equal to the integral of  $P(U')$  on the interval associated with  $\hat{u}$  and hence it is equal to its width. Finally, observe that, for each  $\hat{u} \in \Omega_{\hat{U}}$ ,  $\hat{f}_V$  defines a deterministic relation between  $\mathbf{W}_V$  and  $V$  and this should correspond to a state of  $U$  in  $\mathbf{M}$ , as  $U$  is enumerating all these possible relations because of the canonical specification. This defines a map  $\mu : \Omega_{\hat{U}} \rightarrow \Omega_U$  with  $\mu(\hat{u}) := \{u \in \Omega_U : f_V(u, \mathbf{W}_V) = \hat{f}_V(\hat{u}, \mathbf{W}_V)\}$ . For  $P(U)$  we have:

$$P(u) = \sum_{\hat{u} \in \Omega_{\hat{U}} : \mu(\hat{u}) = u} P(\hat{u}), \quad (\text{A.9})$$

and  $P(u) = 0$  for the states of  $u$  that are not in the domain of  $\mu$ . This is the exogenous quantification that proves the thesis.  $\square$

**Proof of Theorem 6.** Note that  $M$  embeds  $M'$  if and only if there is at least a  $P(\mathbf{U}) \in \mathcal{K}$  so that  $P(U \in \Omega'_{U'}) = 1$ , for all exogenous variables  $U$ .

Assume, *ad absurdum*, that  $M$  embeds an incompatible  $M'$ . For each  $U \in \mathbf{U}$ , let  $\mathcal{K}'(U)$ ,  $U \in \mathbf{U}$ , denote the sets of compatible specifications of  $M'$  relative to variable  $U$ . By Corollary 1, the incompatibility of  $M'$  implies that  $\mathcal{K}'(U)$  should be empty for at least a  $U \in \mathbf{U}$ . By the definition of embedding, for each  $U \in \mathbf{U}$ , we should have at least a  $P(U) \in \mathcal{K}(U)$  such that  $P(U \in \Omega'_{U'}) = 1$ .

Obtaining  $M'$  from  $M$  can be regarded as the result of conditioning the exogenous PMFs on the events  $U \in \Omega'_{U'}$ , for each  $U \in \mathbf{U}$ . As  $P(U \in \Omega'_{U'}) = 1$ , such a

conditioning is well defined for  $P$ . Let  $P'(U)$  denote the resulting PMF for  $M'$ , which simply corresponds to the restriction of  $P$  to  $\Omega'_U$ .

Consider the compatibility constraints of  $M$  involving PMF  $P(U)$  as in Equation (7). As  $P(U) \in \mathcal{K}(U)$ , these constraints should be satisfied by  $P(U)$ . With Markovian models, the constraints are linear. Since by  $P(U \in \Omega'_U) = 1$ , the same constraints hold for  $P'(U)$ . But this means  $P'(U) \in \mathcal{K}'(U)$  and hence  $\mathcal{K}'(U) \neq \emptyset$  (for each  $U \in \mathbf{U}$ ), which is a contradiction.  $\square$

**Proof of Theorem 7.** Consider the l.h.s. of Equation (25). The corresponding joint density is:

$$P\left(\Delta_a \leq \frac{\delta_a}{2}, \Delta_b \leq \frac{\delta_b}{2}, \rho\right) = \gamma \int_0^{\frac{\delta_b}{2}} \int_0^{\frac{\delta_a}{2}} P(\rho | \Delta_a = x, \Delta_b = y) dx dy, \quad (\text{A.10})$$

where a uniform prior is considered for  $\Delta_a := (a - a^*)$  and  $\Delta_b := (b^* - b)$ .

As  $\pi_i \sim \text{Beta}(\alpha, \beta, a^*, b^*)$ :

$$P(\pi_i \in [a, b] | \Delta_a = x, \Delta_b = y) = P(\pi_i \in [a, b] | a^* = a - x, b^* = b + y), \quad (\text{A.11})$$

and hence:

$$P(\pi_i \in [a, b] | \Delta_a = x, \Delta_b = y) = \int_a^b \frac{(\rho - a + x)^{\alpha-1} (b + y - \rho)^{\beta-1}}{(b - a + y + x)^{\alpha+\beta-1} B(\alpha, \beta)} d\rho. \quad (\text{A.12})$$

We solve (A.12) analytically and obtain that  $P(\rho | \Delta_a = x, \Delta_b = y)$  equals to:

$$\left( \frac{(b - a + x)^\alpha {}_2F_1(\alpha, 1 - \beta, \alpha + 1, \frac{b-a+x}{b-a+x+y}) - x^\alpha {}_2F_1(\alpha, 1 - \beta, \alpha + 1, \frac{x}{b-a+x+y})}{\alpha(b - a + x + y)^\alpha B(\alpha, \beta)} \right)^k. \quad (\text{A.13})$$

The joint  $P(\rho, \Delta_a \leq \frac{\delta_a}{2}, \Delta_b \leq \frac{\delta_b}{2})$  can thus be obtained by the following integral:

$$\gamma \int_0^{\delta_b/2} \int_0^{\delta_a/2} \left( \frac{(b - a + x)^\alpha {}_2F_1(\alpha, 1 - \beta, \alpha + 1, \frac{b-a+x}{b-a+x+y}) - x^\alpha {}_2F_1(\alpha, 1 - \beta, \alpha + 1, \frac{x}{b-a+x+y})}{\alpha(b - a + x + y)^\alpha B(\alpha, \beta)} \right)^k dx dy. \quad (\text{A.14})$$

Note that the condition on the l.h.s. of Equation (25) is  $\Delta_a = a - a^* \leq \varepsilon L$ , but the condition in Equation (25) is only meaningful if  $\Delta_a \geq 0$ , i.e. if  $a \geq \varepsilon L$ . This means that  $\frac{\delta_a}{2} = \varepsilon L$  if  $a \geq \varepsilon L$  and  $\frac{\delta_a}{2} = a$  otherwise. By the same reasoning on  $\Delta_b$ , if we rearrange the conditions we obtain the values for  $\delta_a, \delta_b$  outlined in Equations (27) and (28).

The marginal distribution for  $\rho$  can also be obtained by solving the following integral:

$$P(\rho) = \gamma \int_0^{a+(1-b)} \int_0^{a+(1-b)-y} P(\rho|\Delta_a = x, \Delta_b = y) dx dy.$$

The l.h.s. of Equation (25) is just the ratio between  $P(\rho, \Delta_a \leq \frac{\delta_a}{2}, \Delta_b \leq \frac{\delta_b}{2})$  and  $P(\rho)$ .  $\square$

**Proof of Corollary 2.** Let us discuss separately the three cases.

*Case  $a = 0$ .* If  $a = 0$  then  $a^* = 0$  and  $\Delta_a = 0$ , then  $P(\pi_i \in [a, b] | a^* = 0, b^* = b + y)$  becomes

$$\int_a^b \frac{(\rho)^{\alpha-1} (b+y-\rho)^{\beta-1}}{(b+y)^{\alpha+\beta-1} B(\alpha, \beta)} d\rho. \quad (\text{A.15})$$

This integral can be solved analytically and, by exploiting the independence of  $\pi_i$ , we obtain the expression in Equation (30). By following the same steps as in the proof of Theorem 7 we obtain the result.

*Case  $b = 1$ .* If  $b = 1$ , then  $b^* = 1$  and  $\Delta_b = 0$ , then  $P(\pi_i \in [a, b] | a^* = a - x, b^* = 1)$  becomes

$$\int_a^b \frac{(\rho - a + x)^{\alpha-1} (1-\rho)^{\beta-1}}{(1-a+x)^{\alpha+\beta-1} B(\alpha, \beta)} d\rho. \quad (\text{A.16})$$

This integral can be solved analytically and, by exploiting the independence of  $\pi_i$ , we obtain the expression in Equation (32). By following the same steps as in the proof of Theorem 7 we obtain the result.

*Case  $a = b$ .* If  $a = b$  then  $a^* = b^*$  and  $L \rightarrow 0$ . Moreover we assume here that  $\delta \rightarrow 0$  and  $\epsilon \rightarrow 1$ . We assume that the distribution of the outputs of the EMCC iterations is uniformly distributed in  $[a^*, b^*]$ .  $P(a = b | \rho)$  is obtained from Equation (34) in the limit  $L \rightarrow 0$  and  $\epsilon \rightarrow 1$  and the thesis follows trivially.  $\square$

**Proof of Corollary 3.** For  $\alpha = \beta = 1$ ,  ${}_2F_1(\alpha, 1 - \beta, \alpha + 1, \mu) = 1$ ,  $B(\alpha, \beta) = 1$  and hence:

$$P(\rho | \Delta_a = x, \Delta_b = y) = \left( \frac{b - a + x - x}{b - a + x + y} \right)^n = \left( \frac{b - a}{b - a + x + y} \right)^k. \quad (\text{A.17})$$

Equation (34) is finally obtained by computing the integrals as in Zaffalon et al. (2021).  $\square$



## Appendix B. Characterising the Likelihood in the Joint Case

Equation (11) contains the product of the chances of the exogenous variables in  $\mathbf{U}$ ; this is consequence of their being mutually independent (that is, being root nodes in an SCM). In this appendix, we study the log-likelihood function in the simpler setting in which we work in the space of the joint chances  $\theta_{\mathbf{U}}$ . This is equivalent to considering an SCM with a single exogenous node; we call the resulting log-likelihood  $l'$ . The latter contains the linear term  $\theta_{\mathbf{U}}$  where  $l$  had instead the product of the chances. This makes  $l'$  easier to characterise than  $l$ , in the way that follows (we assume we are under M-compatibility to make sure that a global maximum exists).

**Lemma 2.** *The function  $l'$  is concave.*

*Proof.* Consider a dataset made of a single observation. In this case,  $l'$  is the logarithm of a linear function of  $\theta_{\mathbf{U}}$ , whence it is concave. For more general datasets,  $l'$  is then a sum of concave functions, which is also concave.  $\square$

**Theorem 8.** *The function  $l'$  cannot have saddle points or local maxima.*

*Proof.* This is a direct consequence of Lemma 2.  $\square$

**Theorem 9.** *Denote by  $l^*$  the global maximum value of  $l'$ . The global maximum points of  $l'$  form a convex region, i.e., for any two points  $\theta'_{\mathbf{U}}, \theta''_{\mathbf{U}}$  such that  $l'(\theta'_{\mathbf{U}}) = l'(\theta''_{\mathbf{U}}) = l^*$ , then  $l'(\theta_{\mathbf{U}}) = l^*$  for each  $\theta_{\mathbf{U}} := \alpha\theta'_{\mathbf{U}} + (1 - \alpha)\theta''_{\mathbf{U}}$  with  $0 \leq \alpha \leq 1$ .*

*Proof.* Lemma 2 implies  $l'(\theta_{\mathbf{U}}) \geq \alpha l'(\theta'_{\mathbf{U}}) + (1 - \alpha)l'(\theta''_{\mathbf{U}})$  and hence  $l'(\theta_{\mathbf{U}}) \geq l^*$ . But, as the maximum of  $l'(\theta_{\mathbf{U}})$  is  $l^*$ , we have  $l'(\theta_{\mathbf{U}}) = l^*$ , i.e., the thesis.  $\square$

The characterisations above are not immediate to extend to the case of  $l$ , due to the non-linearity introduced by the product of the chances. However we conjecture that in such a case there should not be local maxima either, since we have not encountered that case in our experiments. We leave this proof for future work.

## Appendix C. Questions and Answers

In this section, we discuss various aspects of the paper in an attempt to clarify them via a more direct approach. We are taking inspiration from our rebuttal to the reviewers' reports to lay down a section that should hopefully address the main doubts a reader might be concerned with.

**Q1.** Are the bounds you compute credible intervals? That is, are you after the ‘true’ distribution obtained in the limit of infinite data? What do you mean by saying that bounds are ‘exact’?

This is perhaps one of the biggest potential sources of confusion, therefore we will discuss it at some length.

Let us consider for a moment a Bayesian network. When we estimate its parameters from data, however we do that, we are after some kind of, frequentist or Bayesian or some other type of, expectation (in particular, the expectation of indicator functions  $I$ ). We plug these expectations into the network as parameters and then we use it to do inference, like belief updating or *maximum a posteriori* explanation. We typically do not use credible intervals or the like to estimate the net’s parameters, we just use expectations.

Now imagine that the Bayesian net embeds one latent root node  $U$ . We still want to estimate the net’s parameters from data; but we have data only about the manifest variables. In this case we typically use the EM or some other algorithms to estimate the parameters of the unobserved node, that is, its unconditional probability mass function  $P(U)$ . What we get in the end is again an expectation  $E[I_{U=u}] = P(U = u)$ , for all values  $u$  that  $U$  can take on. And in particular, if the problem is identifiable, the estimated  $P(U)$  corresponds to the global maximum  $\theta_{\mathbf{U}}^*$  of the EM. The point is that also in this case, we are after expectations.

Now consider the case of unidentifiable problems. In this case there are multiple unconditional mass functions  $P(U)$  for the latent node that lead to the distribution we estimate from data (the empirical distribution) once we marginalise  $U$  out. Each one of them corresponds to a global maximum of the EM (a result in this sense was already given by Redner 1981). We take the set of all these unconditional mass functions and call it  $\mathcal{K} := \{P(U)\}$ . We have no information as to whether one element of  $\mathcal{K}$  is more probable than another. We just have no second order probability on top of  $\mathcal{K}$ —which is the essence of non-identifiability.

The best we can do, at this point, is to consider the (possibly infinite) set of Bayesian networks that we obtain by considering all the mass functions for  $U$  that are in  $\mathcal{K}$ ; we run our inference (say, updating) on each of them and we summarise the results by the lower and the upper values of the updating over such a set of Bayesian nets. These are unidentifiability bounds; all we know is where the updated probability lies (we have no probability distribution on top of such a probability, even less so a uniform one). Let us stress that those are

indeed bounds on an expectation (of an indicator function). They are not determining a credible interval. More specifically, those bounds have nothing to do with inferring population values from sample outcomes. The procedure is just the same that we follow for Bayesian nets without latent nodes; we estimate expectations and deliver expectations. The difference is ‘only’ that now we have interval-valued expectations, because of unidentifiability (which means that our model is incomplete). What we are saying so far is that non-identifiability has not necessarily to be subject to question of statistical inference; it is rather like identifiable problems: we use our models via expectations.

And here is the crucial point of potential confusion: whenever in the paper we refer to the ‘exact’ bounds, we refer to the interval-valued expectation that we would obtain if we could run all the models in  $\mathcal{K}$  one after the other and then summarise the (observational, interventional, counterfactual) results by the lower and upper bounds obtained across all those models. Usually we rather try to compute directly exact bounds via an optimisation problem—which is our credal-net based approach in this paper. But that is not possible in general and we need approximations. In this case we resort to our EMCC, which will deliver ranges, i.e., inner approximations of the bounds.

Let us stress once more that we never aim in this paper to compute the counterfactual bounds that we would get if we had an infinite sample to estimate parameters. In the same way as in Bayesian nets one does not refer to the outcome of an updating as the actual value of the probability that one would obtain if the parameters of the Bayesian nets were estimated from an infinite sample. We just know this is impossible and we stick to the best we can: using expectations. For us ‘exact’ refers to exactness with respect to optimising over  $\mathcal{K}$ .

An additional note is that also our credible intervals should be understood in the same light. That is, we yield an approximating interval-valued expectation (probability)  $[a, b]$ , which we know that by construction is included in the exact interval-valued expectation  $[a^*, b^*]$ : where the latter is made by the exact yet unknown bounds that result from an optimisation problem over  $\mathcal{K}$ . Our credible interval tells us then, in probability, how much smaller is  $[a, b]$  compared to  $[a^*, b^*]$ . Note that the former gets closer and closer to the latter the more EM runs we do, with the same data fixed (not with increasing data).

**Q2.** M-compatibility seems to be referring to the empirical distribution. Shouldn’t it refer to the ‘true’ one? Moreover, the empirical distribution is estimated via maximum likelihood and obtained via

the factorisation in (4). Is your notion of compatibility restricted only to these specific choices?

Our definition of compatibility indeed applies to the empirical distribution; this paper is not focused on statistical inference. It is true that there are different ways to define the empirical distribution (we use MLE, but one could use Bayes; we use (4)—because that is induced by the very marginal DAG that generates the data—, but others might want to use other factorisations); and yet we are not concerned about the way by which one defines the empirical distribution. What matters to us is that once it is defined, it complies with the logical constraints imposed by the SCM, otherwise we would be using a model that embeds a logical contradiction and that would not allow us to make sensible calculations.

Let us also note that defining compatibility with the ‘true’ distribution (obtained in the limit of infinite data) can be questionable, because the compatibility problem vanishes in such a case: in fact, one could use a canonical specification (that is always compatible with the data—as we show for Markovian SCMs), thus eventually (i.e., in the limit) automatically retaining only the distributions that are compatible with the given causal graph.

### Q3. What is Theorem 1 actually proving?

We start by having assessed in some way the empirical mass function  $P(\mathbf{V})$ . If our problem were identifiable, the global maximum of the EM would correspond to the  $P(\mathbf{U})$  that leads to  $P(\mathbf{V})$  when the  $U$  variables are marginalised out of the given SCM. In reality the problem is not identifiable and hence there are (possibly infinitely) many  $P(\mathbf{U})$  that lead to  $P(\mathbf{V})$  when we marginalise the  $U$  variables out. The theorem’s claim is just that those many  $P(\mathbf{U})$ , which we gather in set  $\mathcal{K}$  (representing them as collections of marginals over the exogenous nodes), are in a one-to-one correspondence with the global maxima of the likelihood.

Stated differently, the problem of the logical compatibility of the structural equations with a given (empirical) probability distribution, can be regarded as one of feasibility: that is, we could write it down as the problem determining whether the feasible region of an optimisation problem is empty or not. The theorem allows us to sample such a region by multiple runs of the EM (instead of solving the NP-hard optimisation problem).

**Q4.** The paper advocates testing for  $M$ -compatibility, which is a valid point; Examples 5–7 are interesting and useful. But do they justify the claims the paper makes, such as:

- if the true underlying model is not compatible with the available data, the results obtained by using the canonical model will be unwarranted as an approximation to the actual one.
- We cannot have guaranteed bounds without knowing the structural equations of the underlying SCM.

Does estimated  $P$  being compatible with  $M$  or not really make a fundamental difference (like a phase transition) in the validity of the inferred bounds? Or the quality of the inferred bounds simply depends on the distance between  $P$  and the ‘true’ distribution of the endogenous variables? Is it just a problem of statistical inference then?

Let us stay for a moment in the case of finite sample sizes. Our point is that many results published for that case assume that a causal graph (that is, one without structural equation) together with data about the exogenous nodes can yield us valid bounds. Our examples show that this is not the case; it is particularly disturbing that the empirical distribution can be incompatible with the very SCM that has generated it.

Part of the questions asked seems to imply the following: can this be avoided with your notion of compatibility? In a sense, yes: let us imagine that we can test all the PSCMs with the given causal graph, and that each of them is compatible with the empirical distribution. This test makes sure that our expectation bounds (which are of the same type of those by Mueller et al. 2022, for instance) make sense: because they have passed a basic test of rationality. In fact, even if only one of those PSCMs were incompatible, we would be leaving open the possibility that such a PSCM is the actual one generating the data; and our bounds would be at risk of being wrong, because they would not cover such a case (as in our examples).

One point that we make is that the canonical specification is not a way out to such a problem. In principle, it seems to allow us to use all the possible SEs at once (which is also the reason why it is always compatible with the empirical distribution), thus bypassing the problem. But in reality the incompatible SEs/PSCMs (implicitly considered in the canonical specification) are just put

under the carpet by internally never assigning probability zero to the structural equations that do not belong to them, which is equivalent to actually neglecting those sub-models; and the problem would still be there.

Yet a sufficient condition to address the problem exists: we should use the canonical specification and check whether all sub-models are embedded. In the Markovian case this is readily possible to do since we can exactly represent the set  $\mathcal{K}$  via linear constraints as Theorem 2 proves (with an increase in complexity, the same holds for quasi-Markovian models). In case they were, our expectation bounds would not be at risk. In the opposite case, we should give the disclaimer that the bounds are not guaranteed to be correct; or we should just refrain from yielding them.

Now let us consider inferential arguments, that is, how this relates to the similarity of the ‘true’ distribution to the empirical one. Let us consider again the canonical specification. Since this implicitly considers all the possible SEs, in the limit of infinite data, all and only the compatible models will be embedded. And since the true SCM that has generated the data must be compatible with the ‘true’ distribution, by definition, it will be included among them, and our inferences will be correct.

The problem here is that all this happens in the limit, and we do not know yet how this should be related to questions of finite samples: how large will a sample have to be to declare that the true SCM is for sure embedded?

We seem to have to connect the distance between the distributions to the transition into compatibility of the true SCM. But we do not know whether this is a continuous map, and there could indeed be a kind of phase-transition (in particular because SEs are deterministic, sharp, equations). As long as this map is not studied in detail, we do not seem to be in the conditions to give probabilities of compatibility along these lines.

**Q5.** One stated advantage of the proposed EMCC approach over the existing work is that it provides ranges that approximate the exact ones from inside (i.e., a so-called inner approximation). Where does this claim comes from?

It looks as if  $[a^*, b^*]$  here are not the ‘true’ exact bounds due to partial identifiability but those one obtains by assuming  $P$  in (6) is the ‘true’ distribution. Thus  $[a^*, b^*]$  approach the ‘true’ bounds only given infinite samples; otherwise, their relation with the ‘true’ bounds is unclear as well as it is unclear how EMCC can escape the finite sample issue.

The question in the first paragraph relates to some of our previous replies: ‘exact’ bounds does not mean in this paper the bounds one gets with infinite data; rather it is the exact solution of an optimisation problem over  $\mathcal{K}$ . In the second paragraph the question is similar: we do not aim at escaping the finite sample issue by the EMCC.

**Q6.** It looks as if EMCC vs. Zhang et al. (2022) is basically frequentist vs. Bayesian.

Actually we would not say that this is the main issue at stake. Let us try to clarify the main differences between our EMCC and the alternative sampling method proposed by Zhang et al. (2022) to approximate partially identifiable queries.

The latter is indeed based on Bayesian ideas, in the sense that a prior is put over the chances of the exogenous variables that is later updated to a posterior via the dataset  $\mathcal{D}$  of endogenous observations. With the posterior they can get an expected value over the chances that in turn leads to an expected value over the query of interest. That is to say that the partially identifiable counterfactual query is reduced to a point estimate when it comes to expectation. This should be contrasted with the EMCC, which delivers approximate lower and upper expectations (bounds) for the query of interest.

However Zhang et al. (2022) are after the expectation bounds as we do. To this end, they use a different strategy than ours and that is based on credible intervals. Their Theorem 3.2 shows that the 100% credible interval contains the expectation bounds (i.e., it is an *outer* approximation) and coincides with it almost surely in the limit. When dealing with finite samples, they provide further results giving probabilistic guarantees that the computed interval is in fact a 100% credible interval. This result is based on a tolerance rate  $\varepsilon$  that is somewhat similar in its aim to the  $\varepsilon$  we use in Section 5.2 for our own credible intervals. In fact the two, quite different, types of credible intervals, theirs and ours, can be regarded to be after the same goal: that is, to state with which probability the delivered interval contains the exact expectation interval.

It would be interesting in the future to compare the relative power of the two types of credible intervals. For the time being, let us just note that the EMCC provides an inner approximation to the actual expectation bounds, besides a probabilistic guarantee on the fact that the credible interval is an outer approximation. Furthermore, Zhang et al. (2022) appears to be oblivious to the question of M-compatibility, which, as we saw, has the power to invalidate what might look like safe conclusions.

## Acknowledgements

The reviewers of this paper have been very careful in checking it, thus helping us to identify some of its unclear and imprecise parts; we gratefully acknowledge their dedication and stubbornness, which the paper has definitely benefited from. We are also grateful to Heidi Kern from the Triangolo association for her support with the palliative care problem discussed in Section 6.2. This research was partially funded by MCIN/AEI/10.13039/501100011033 with FEDER funds for the projects PID2019-106758GB-C32 and PID2022-139293NB-C31. Finally, we would like to thank the “María Zambrano” grant (RR\_C\_2021\_01) from the Spanish Ministry of Universities and funded with NextGenerationEU funds.

## References

- Antonucci, A., de Campos, C.P., Huber, D., Zaffalon, M., 2015. Approximate credal network updating by linear programming with applications to decision making. *International Journal of Approximate Reasoning* 58, 25–38.
- Antonucci, A., Zaffalon, M., 2008. Decision-theoretic specification of credal networks: A unified language for uncertain modeling with sets of Bayesian networks. *International Journal of Approximate Reasoning* 49, 345–361.
- Balke, A., Pearl, J., 1994. Counterfactual probabilities: Computational methods, bounds and applications, in: de Mantaras, R.L., Poole, D. (Eds.), *Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann. p. 4654.
- Balke, A., Pearl, J., 1997. Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association* 92, 1171–1176.
- Bareinboim, E., Pearl, J., 2012. Causal inference by surrogate experiments: Z-identifiability, in: de Freitas, N., Murphy, K.P. (Eds.), *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, AUAI Press.
- Cabañas, R., Antonucci, A., Huber, D., Zaffalon, M., 2020. CREDICI: a Java library for causal inference by credal networks, in: Yeager, M., Nielsen, T.D. (Eds.), *Proceedings of the Tenth International Conference on Probabilistic Graphical Models*. JMLR.org. volume 138 of *PMLR*, pp. 597–600.



- Cozman, F.G., 2000. Credal networks. *Artificial intelligence* 120, 199–233.
- Cozman, F.G., Mauá, D.D., 2017. On the complexity of propositional and relational credal networks. *International Journal of Approximate Reasoning* 83, 298–319.
- Dempster, A.P., Laird, N.M., Rubin, D.B., 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society* 39, 1–22.
- Druzdzel, M.J., Simon, H.A., 1993. Causality in Bayesian belief networks, in: Heckerman, D., Mamdani, A. (Eds.), *Proceedings of the Ninth International Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann. p. 311.
- Duarte, G., Finkelstein, N., Knox, D., Mummolo, J., Shpitser, I., 2023. An automated approach to causal inference in discrete settings. *Journal of the American Statistical Association* 0, 1–16. doi:10.1080/01621459.2023.2216909.
- Fagioli, E., Zaffalon, M., 1998. 2U: An exact interval propagation algorithm for polytrees with binary variables. *Artificial Intelligence* 106, 77–107.
- Galhotra, S., Pradhan, R., Salimi, B., 2021. Explaining black-box algorithms using probabilistic contrastive counterfactuals, in: Li, G., Li, Z., Idreos, S., Srivastava, D. (Eds.), *Proceedings of the 2021 International Conference on Management of Data*, Association for Computing Machinery. p. 577590.
- Huber, D., Cabañas, R., Antonucci, A., Zaffalon, M., 2020. Crema: A Java library for credal network inference, in: Yeager, M., Nielsen, T.D. (Eds.), *Proceedings of the Tenth International Conference on Probabilistic Graphical Models*. JMLR.org. volume 138 of *PMLR*, pp. 613–616.
- Huber, D., Chen, Y., Antonucci, A., Darwiche, A., Zaffalon, M., 2023. Tractable bounding of counterfactual queries by knowledge compilation. URL: <https://tractable-probabilistic-modeling.github.io/tpm2023/papers>. Sixth Workshop on Tractable Probabilistic Modeling @ UAI 2023.
- Hume, D., 1739. *A Treatise of Human Nature*. Oxford University Press.

- Ibeling, D., Icard, T., 2023. Comparing causal frameworks: Potential outcomes, structural models, graphs, and abstractions. arXiv:2306.14351 .
- Ide, J.S., Cozman, F.G., Ramos, F.T., 2004. Generating random Bayesian networks with constraints on induced width, in: de Mantaras, R.L., Saitta, L. (Eds.), Proceedings of the Sixteenth European Conference on Artificial Intelligence, IOS Press. p. 353357.
- Kang, C., Tian, J., 2006. Inequality constraints in causal models with hidden variables, in: Dechter, R., Richardson, T. (Eds.), Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence, AUAI Press. p. 233240.
- Kern, H., Corani, G., Huber, D., Vermes, N., Zaffalon, M., Varini, M., Wenzel, C., Fringer, A., 2020. Impact on place of death in cancer patients: a causal exploration in southern Switzerland. BMC Palliative Care 19, 1–10.
- Koller, D., Friedman, N., 2009. Probabilistic Graphical Models: Principles and Techniques. MIT.
- Mauá, D.D., De Campos, C.P., Benavoli, A., Antonucci, A., 2014. Probabilistic inference in credal networks: New complexity results. Journal of Artificial Intelligence Research 50, 603–637.
- Mauá, D.D., Cozman, F.G., 2020. Thirty years of credal networks: Specification, algorithms and complexity. International Journal of Approximate Reasoning 126, 133–157.
- Mueller, S., Li, A., Pearl, J., 2022. Causes of effects: learning individual responses from population data, in: De Raedt, L. (Ed.), Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, International Joint Conferences on Artificial Intelligence Organization. pp. 2712–2718.
- Pearl, J., 1999. Probabilities of causation: Three counterfactual interpretations and their identification. Synthese 121, 93–149.
- Pearl, J., 2009. Causality. Cambridge University Press.
- Redner, R., 1981. Note on the consistency of the maximum likelihood estimate for nonidentifiable distributions. The Annals of Statistics 9, 225–228.

- da Rocha, J.C.F., Cozman, F.G., 2002. Inference with separately specified sets of probabilities in credal networks, in: Darwiche, A., Friedman, N. (Eds.), *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann. p. 430437.
- Sachs, M.C., Gabriel, E.E., Sölander, A., Gabriel, E.E., 2023. Symbolic computation of tight causal bounds. *Journal of Computational and Graphical Statistics* 32, 567–576.
- Schölkopf, B., 2022. Causality for machine learning, in: *Probabilistic and Causal Inference: The Works of Judea Pearl*. Association for Computing Machinery, pp. 765–804.
- Shpitser, I., Pearl, J., 2007. What counterfactuals can be tested, in: Parr, R., van der Gaag, L. (Eds.), *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*, AUAI Press. p. 352359.
- Tian, J., 2002. *Studies in Causal Reasoning and Learning*. Ph.D. thesis. UCLA.
- Wu, C.F.J., 1983. On the convergence properties of the EM algorithm. *The Annals of Statistics* 11, 95–103.
- Zaffalon, M., Antonucci, A., Cabañas, R., 2020. Structural causal models are (solvable by) credal networks, in: Jaeger, M., Nielsen, T.D. (Eds.), *Proceedings of the Tenth International Conference on Probabilistic Graphical Models*. JMLR.org. volume 138 of *PMLR*, pp. 581–592.
- Zaffalon, M., Antonucci, A., Cabañas, R., 2021. Causal expectation-maximisation. *Why-21@NeurIPS* URL: [https://why21.causalai.net/papers/WHY21\\_52.pdf](https://why21.causalai.net/papers/WHY21_52.pdf).
- Zaffalon, M., Antonucci, A., Cabañas, R., Huber, D., 2023. Approximating counterfactual bounds while fusing observational, biased and randomised data sources. *International Journal of Approximate Reasoning* 162, 109023.
- Zaffalon, M., Antonucci, A., Cabañas, R., Huber, D., Azzimonti, D., 2022. Bounding counterfactuals under selection bias, in: Salmerón, A., Rumí, R. (Eds.), *Proceedings of the Eleventh International Conference on Probabilistic Graphical Models*. JMLR.org. volume 186 of *PMLR*, pp. 289–300.

- Zhang, J., Bareinboim, E., 2020. Designing optimal dynamic treatment regimes: a causal reinforcement learning approach, in: Daumé, H., Singh, A. (Eds.), Proceedings of the Thirty-Seventh International Conference on Machine Learning, JMLR.org. pp. 11012–11022.
- Zhang, J., Tian, J., Bareinboim, E., 2022. Partial counterfactual identification from observational and experimental data, in: Chaudhuri, K., Jegelka, S., Song, L., Szepesvári, C., Niu, G., Sabato, S. (Eds.), Proceedings of the Thirty-Ninth International Conference on Machine Learning, JMLR.org. pp. 26548–26558.