Research Article

Miguel Navascués* and Elie Wolfe

The Inflation Technique Completely Solves the Causal Compatibility Problem

<https://doi.org/10.1515/jci-2018-0008>

Received Mar 01, 2018; accepted Jun 13, 2020

Abstract: The causal compatibility question asks whether a given causal structure graph — possibly involving latent variables — constitutes a genuinely plausible causal explanation for a given probability distribution over the graph's observed categorical variables. Algorithms predicated on merely *necessary* constraints for causal compatibility typically suffer from false negatives, i.e. they admit incompatible distributions as *apparently* compatible with the given graph. In DOI: 10.1515/jci-2017-0020, one of us introduced the *inflation technique* for formulating useful relaxations of the causal compatibility problem in terms of linear programming. In this work, we develop a formal hierarchy of such causal compatibility relaxations. We prove that inflation is asymptotically tight, i.e., that the hierarchy converges to a zero-error test for causal compatibility. In this sense, the inflation technique fulfills a longstanding desideratum in the field of causal inference. We quantify the rate of convergence by showing that any distribution which passes the n^{th} -order inflation test must be $O(n^{-1/2})$ -close in Euclidean norm to some distribution genuinely compatible with the given causal structure. Furthermore, we show that for many causal structures, the (unrelaxed) causal compatibility problem is faithfully formulated already by either the first or second order inflation test.

Keywords: causal compatibility problem, inflation technique

2010 Mathematics Subject Classification: 62D20, 81P13, 81P42, 90C90

1 Introduction

A Bayesian network or **causal structure** is a directed acyclic graph (DAG) where vertices represent random variables, each of which is generated by a non-deterministic function depending on the values of its parents. Nowadays, causal structures are commonly used in bioinformatics, medicine, image processing, sports betting, risk analysis, and experiments of quantum nonlocality. In this work we consider causal structures with two distinct types of vertices: categorical variables which may be directly observed, and variables which cannot be observed, referred to as latent variables.¹ We make no assumption whatsoever on the state spaces of the latent variables; they can be discrete or continuous. Nevertheless, every causal structure encodes a possible *hypothesis* of causal explanation for statistics over its observed variables.

Naturally, understanding how different causal structures give rise to different sets of compatible distributions is a fundamental goal within the field of causal inference. Many prior works are ultimately concerned with the **causal discovery problem**, which asks to enumerate (or graphically characterize) all legitimate hypotheses of causal structure which are capable of explaining some observed probability distribution [2–9]. For computational tractability, practical causal discovery algorithms typically exclude causal explanations which are unfaithful (fine-tuned). Fundamentally, however, the faithfulness assumption is not an essential

*Corresponding Author: Miguel Navascués: Institute for Quantum Optics and Quantum Information (IQOQI) Vienna, Austrian Academy of Sciences, Boltzmannngasse 3, 1090 Vienna, Austria

Elie Wolfe: Perimeter Institute for Theoretical Physics, 31 Caroline St. N, Waterloo, Ontario, Canada, N2L 2Y5

¹ Pearl [1] refers to such graphical models as “latent structures”.

criterion for causal discovery. Demanding faithfulness can be thought of as a *second* filtering step, where the *fundamental* filtering of causal discovery is the exclusion of any causal structure which cannot explain the observed probability distribution, even granting fine-tuning. In this manuscript, therefore, causal discovery refers to the foundational problem of returning *all* causal structures compatible with the given distribution. Selecting a single “best” causal model — or even *scoring* the quality of the different causal explanations [8, 9] — constitute refinements to the causal discovery problem which we do not address here.

The **causal characterization problem** is the focus of a distinct line of research. It concerns characterizing the set of statistics compatible with a single given causal structure, that is, the derivation of causal compatibility constraints [9–18]. Quantum information theorists have recently joined this research effort [19–26]. Causal characterization is useful for proving the impossibility of simulating certain quantum optics experiments with classical devices [27–29], or for confirming the nonclassicality of quantumly realizable statistics in novel hypothetical scenarios [30–34].

We must note that the causal characterization problem has also been tackled in scenarios where the state-spaces of the latent variables are prescribed [35–38]. Critically, Ref. [39] provides upper bounds on the cardinalities of a causal structure’s latent variables without any loss of generality (whenever the observed variables have discrete state spaces). Consequently, the set of multi-variate categorical distributions compatible with any given causal structure is always a semi-algebraic set, admitting characterization in terms of a finite number of polynomial equality and inequality constraints. Nevertheless, *identifying* the full set of causal compatibility constraints via exploiting the constrained state-spaces of the latent variables is often intractable [35–38]. The inflation technique considered herein, by contrast, has no dependence on the latent variables’ state-spaces. Hereafter, therefore, we consider only causal structures with *unconstrained* latent variables.

Causal discovery relates a single distribution to many structures; causal characterization relates many distributions to a single structure. Both such efforts, therefore, are oracle-wise equivalent, and hinge fundamentally on the **causal compatibility problem (CCP)**, which simply asks a yes-or-no question: Is the given distribution compatible with the given causal structure? The **inflation technique** [40] is a way of relating approximations of the causal compatibility problem to linear programming (LP) problems. Every LP *satisfaction* problem can be dualized and recast as an equivalent *optimization* problem. Inspired by LP duality, we will formulate a dual notion of causal compatibility, through which we will be able to rigorously upper-bound the error introduced by approximating the CCP as an LP problem via inflation. Our main result here is that this error asymptotically tends to zero when inflation is expressed as a hierarchy of ever-higher-order tests of causal compatibility. This implies that the inflation criterion — far from being a relaxation — is meta-equivalent to the causal compatibility problem, and hence constitutes an alternative way of understanding general causal structures.

In contrast with Ref. [40], in this paper we define the inflation technique as a hierarchy of causal compatibility tests applicable *exclusively* to the special class of causal structures (introduced by Fritz [19]) called **correlation scenarios**. At the same time, however, we also introduce a *graphical preprocessing* which precisely recasts the general causal compatibility problem in terms of causal compatibility with correlation scenarios, such that there is no loss of generality in our approach.

This paper is organized as follows: In Section 2 we introduce the concept of a correlation scenario, we define primal and dual notions of the causal compatibility problem and their approximations. In Section 3 we review the inflation technique as a means for approximately solving either form of the causal compatibility problem. In Section 4 we state our main theorems concerning the convergence of inflation for correlation scenarios, though the formal proofs are deferred to the appendices. In Section 5 we build upon existing causal inference techniques to describe a natural graphical preprocessing which maps general causal structures into correlation scenarios. This preprocessing — fairly useful in its own right — implies the universal applicability of the inflation technique as defined here. Finally, Section 6 presents our conclusions.

2 Preliminary Definitions

The graphical models we study are fully general causal structures. A causal structure is represented by a directed acyclic graph imbued with some distinction among the vertices to clarify if a node in the graph represents either an observable or a latent variable. In this work, we use a pink color and subscripts of the letter “ U ” (“ U ” from “Unobserved”) to indicate the latent variables in a graph. We follow the convention of Refs. [14, 18, 41] and depict exogenous (i.e., non-root) observable variables as square-shaped nodes in their graphs.

Correlation scenarios are a special type of causal structures. The graph of a correlation scenario has just two layers: a bottom layer of independently distributed latent random variables $\{U_1, U_2, \dots, U_L\}$ and a top layer of observable random variables $\{A_1, A_2, \dots, A_m\}$, see Figure 1. The observable distribution $P(A_1, A_2, \dots, A_m)$ is generated via non-deterministic functions $A_x = A_x(\bar{U}^{L_x})$, with $\bar{U} = (U_1, \dots, U_L)$ and $L_x \subset \{1, \dots, L\}$. Here (and in the following) the notation $\bar{v}^S$, where $\bar{v}$ is a vector with N entries and $S \subset \{1, \dots, N\}$, will represent the vector with entries v^s , $s \in S$. Readers familiar with **d-separation** may appreciate that although the implication $(A_i \perp_d A_j | A_k) \Rightarrow (A_i \perp_d A_j | A_k)$ does *not hold* for general causal structures, it is *true* for correlation scenarios. Later on in Section 5, we will relate distributions over general causal structures to distributions over some correlation scenarios associated with them. As we will see, correlation scenarios are the atomic constituents upon which the inflation hierarchy acts.

Figure 1: A generic correlation scenario. The independent latent variables $U_1, U_2, \dots$ influence the observed variables $A_1, A_2, \dots$

2.1 The Causal Compatibility Problem, its dual and their approximate versions

A distribution P over the observable variables of a causal structure $\mathcal{G}$ is said to be **compatible** with $\mathcal{G}$ if P is the observable *marginal* of some distribution P' over *all* the variables of $\mathcal{G}$, and where P' can be factored into a product of singleton-variable conditional probability distributions associated with every individual vertex in $\mathcal{G}$ (conditioned on all of the vertex’s parents, if any). A diverse vocabulary of phrases synonymous with “ P is compatible with $\mathcal{G}$ ” can be found in conventional literature, such as “ P can be *realized* in $\mathcal{G}$ ”, “ P can *arise* from $\mathcal{G}$ ”, “ $\mathcal{G}$ *gives rise to* P ”, “ $\mathcal{G}$ *explains* P ”, “ $\mathcal{G}$ can *simulate* P ”, and “ $\mathcal{G}$ is a *model* for P ”.

Figure 2: The three-on-line correlation scenario.

Figure 3: The triangle scenario.

Consider, for example, the correlation scenario dubbed **the triangle scenario**, with $m = L = 3$, see Figure 3. Denoting A_1, A_2, A_3 respectively by A, B, C , we have that a probability distribution $P(A, B, C)$ is realizable in the triangle scenario if A, B, C are generated via the non-deterministic functions $A(U_1, U_2), B(U_2, U_3), C(U_3, U_1)$. Alternatively, $P(A, B, C)$ is realizable in the triangle scenario if it admits a decomposition of the form:

$$P(A, B, C) = \sum_{U_1, U_2, U_3} P(A|U_1, U_2)P(B|U_2, U_3)P(C|U_3, U_1)P(U_1)P(U_2)P(U_3). \quad (1)$$

For causal structures which are *not* correlation scenarios, however, the non-deterministic functions giving rise to the observed variables will also depend on other observed variables. An example is given by the **instrumental scenario** [33, 42] (Figure 4, up), where X and U are, respectively, a free observable and a latent variable, and the observed variables A and B are generated via the non-deterministic functions $A = A(X, U), B = B(A, U)$. Alternatively, $P(A, B|X)$ is realizable in the instrumental scenario if it admits a decomposition of the form:

$$P(A, B|X) = \sum_U P(A|X, U)P(B|A, U)P(U). \quad (2)$$

Per Ref. [39], the set of distributions P compatible with a given causal structure $\mathcal{G}$ is a semi-algebraic set whenever the observable random variables are categorical, i.e., when they have finite cardinality. This implies that it can be characterized in terms of a finite number of polynomial inequalities. Unfortunately, the computational complexity of deriving such a characterizing set of inequalities makes the problem intractable already for networks of a very small size [43]. Furthermore, within the context of quantum foundations, there exist fairly natural causal structures for which the total number of such inequalities grows exponentially with the dimensionality of P [44]. We must resort thus to partial characterizations of the original set of distributions. This notion is better formalized by the following problem.

Figure 4: The instrumental scenario. (Not a correlation scenario.)

Problem 1. Approximate Causal Compatibility

INPUT: $\epsilon > 0$, a causal structure $\mathcal{G}$ and a particular probability distribution P over the observed variables.

OUTPUT: If there does not exist a probability distribution $\tilde{P}$ over the observed variables, such that $\|P - \tilde{P}\|^2 \leq \epsilon$ and $\tilde{P}$ is compatible with $\mathcal{G}$, then return a function F such that $F(P) < 0$ and $F(\hat{P}) \geq 0$ for all distributions $\hat{P}$ compatible with $\mathcal{G}$.

OBJECTIVE: Determine if P is “approximately compatible” with $\mathcal{G}$; if not, provide a witness F to prove incompatibility of P .

Note that, if P is not approximately compatible with $\mathcal{G}$, the function F witnessing its incompatibility is not required to be *universal*. Namely, there could exist other distributions P' incompatible with $\mathcal{G}$ such that $F(P') \geq 0$.

The goal of this paper is to provide a solution to this problem. Note that, since for any $\mathcal{G}$ the set of compatible distributions is closed [39], it follows that any distribution P that is ϵ -compatible for all $\epsilon > 0$ must be compatible with $\mathcal{G}$. The analog problem for $\epsilon = 0$ will be simply referred to as *Causal Compatibility*.

A related problem that we will also solve is the following:

Problem 2. Approximate Causal Optimization

INPUT: $\epsilon > 0$, a causal structure $\mathcal{G}$ and a real function F of the probabilities of the observed events.

OUTPUT: A real value f such that $f \leq f_* = \min_{\hat{P}} F(\hat{P}) \leq f + \epsilon$, where the minimum is over all distributions $\hat{P}$ compatible with $\mathcal{G}$.

OBJECTIVE: Given a function F , find a good lower bound on its minimum value over all distributions compatible with $\mathcal{G}$.

This problem is dual to Approximate Causal Compatibility, and it is interesting in its own right. In quantum optics experiments, we test and quantify non-classicality via the violation of inequalities of the form $F(P) \geq K$. Identifying values of K for which the above holds for all distributions P compatible with the considered causal structure $\mathcal{G}$ is a must before any experiment is actually carried out. Similarly as before, for $\epsilon = 0$, we name the analog problem *Causal Optimization*.

Coming back to the triangle scenario, an instance of Approximate Causal Optimization would be minimizing

$$(P(0, 0, 0) - 1/2)^2 + (P(1, 1, 1) - 1/2)^2 \quad (3)$$

over all distributions $P(A, B, C)$ with $A, B, C \in \{0, 1\}$ compatible with the triangle scenario. In any experimental setup where bipartite optical sources play the role of U_1, U_2, U_3 in the triangle scenario, any observed distribution $P(A, B, C)$ for which the value of (3) is smaller than the lower bound f provided by Approximate Causal Optimization evidences the presence of quantum effects.

There exist a number of algorithms which provide outer approximations for the set of distributions compatible with a given causal structure $\mathcal{G}$ [3, 14, 25, 26, 41, 45, 46]. By minimizing functions over such outer approximations, existing algorithms can provide lower bounds on the true minimum and thus solve Approximate Causal Optimization, as long as ϵ exceeds some threshold (determined by the mismatch between the aforementioned relaxations and the original set of compatible distributions). As we will see, the Inflation Technique can be used to solve both Approximate causal compatibility and Approximate Causal Optimization for arbitrarily small values of ϵ .

3 The Inflation Hierarchy for Correlation Scenarios

3.1 Some examples

Let $P(A, B, C)$ be a distribution realizable in the triangle scenario, and suppose that we generate n independently distributed copies of U_1, U_2, U_3 , that is, the variables $\{U_1^i, U_2^i, U_3^i : i = 1, \dots, n\}$. Then we could define

the random variables

$$A^{ij} \equiv A(U_1^i, U_2^j), B^{ij} \equiv B(U_2^i, U_3^j), C^{ij} \equiv C(U_3^i, U_1^j). \quad (4)$$

The causal structure associated with the independently distributed copies of U_1, U_2, U_3 and their observable children $\{\{A^{ij}\}, \{B^{kl}\}, \{C^{pq}\}\}$ is termed an **inflation graph**. The inflation graph of a correlation scenario is also a correlation scenario; as an example, Figure 5 depicts the $n = 2$ inflation graph for the triangle scenario. These observable variables follow a probability distribution $Q_n(\{A^{ij}\}, \{B^{kl}\}, \{C^{pq}\})$ with the property

Figure 5: The second-order inflation graph of the triangle scenario.

$$Q_n(\{A^{ij}=a^{ij}, B^{kl}=b^{kl}, C^{pq}=c^{pq}\}) = Q_n(\{A^{ij}=a^{\pi(i)\pi'(j)}, B^{kl}=b^{\pi'(k)\pi''(l)}, C^{pq}=c^{\pi''(p)\pi(q)}\}), \quad (5)$$

for all permutations of n elements π, π', π'' . Expanded out for $n = 2$, Eq. (5) becomes

$$\begin{aligned} & Q_2(a^{11}, a^{12}, a^{21}, a^{22}, b^{11}, b^{12}, b^{21}, b^{22}, c^{11}, c^{12}, c^{21}, c^{22}) \\ &= Q_2(a^{21}, a^{22}, a^{11}, a^{12}, b^{11}, b^{12}, b^{21}, b^{22}, c^{12}, c^{11}, c^{22}, c^{21}) \\ &= Q_2(a^{12}, a^{11}, a^{22}, a^{21}, b^{21}, b^{22}, b^{11}, b^{12}, c^{11}, c^{12}, c^{21}, c^{22}) \\ &= Q_2(a^{11}, a^{12}, a^{21}, a^{22}, b^{12}, b^{11}, b^{22}, b^{21}, c^{21}, c^{22}, c^{11}, c^{12}) \end{aligned}$$

We treat with special distinction the **diagonal variables** $\{A^{ii}, B^{ii}, C^{ii}\}_i$. Given the global distribution Q_n , we denote by Q_n^g the marginal distribution of the diagonal variables with indices up to $g \leq n$, i.e.

$$Q_n^g := Q_n\left(\bigwedge_{i=1}^g \{A^{ii}=a^i, B^{ii}=b^i, C^{ii}=c^i\}\right) \quad (6)$$

In the following, we call Q_n^g the **diagonal marginal of degree- g** .

A related concept is the **degree- g lifting** of a distribution P , consisting of the statistics of g independent and identically distributed copies of P , that is

$$P^{\otimes g}\left(\bigwedge_{i=1}^g \{A^{ii}=a^i, B^{ii}=b^i, C^{ii}=c^i\}\right) := \prod_{i=1}^g P(a^i, b^i, c^i).$$

Taking the random variables in the inflation graph to arise per Eq. (4) implies that the diagonal marginals associated with the inflation graph must be related to the lifted distributions of the original distribution over observed variables per

$$Q_n^g = P^{\otimes g}, \text{ for } g = 1, \dots, n. \quad (7)$$

Expanded out for $g = n = 2$, Eq. (7) identifies the diagonal marginal in this scenario

$$Q_{n=2}^{g=2} = \sum_{A^{12}, A^{21}, B^{12}, B^{21}, C^{12}, C^{21}} Q_{n=2} \text{ as} \\ Q_{n=2}^{g=2}(a^{11}, a^{22}, b^{11}, b^{22}, c^{11}, c^{22}) = P(a^{11}, b^{11}, c^{11})P(a^{22}, b^{22}, c^{22}) \quad (8)$$

Note that there exist additional relations between Q_n and the original distribution $P(A, B, C)$, some of which involve polynomials of the probabilities $P(A, B, C)$ with degree greater than n . For instance,

$$Q_2(\{A^{12}=a, B^{12}=b, C^{12}=c\}) = P_A(a)P_B(b)P_C(c). \quad (9)$$

In this paper we will not exploit such higher degree relations, though they are quite useful in practical implementations.²

Given an arbitrary distribution $P(A, B, C)$, the **inflation technique** consists in demanding the degree- n lifting of $P(A, B, C)$ be the degree- n diagonal marginal of a distribution Q_n over the inflated variables satisfying (5). When condition (7) is met, we call the associated distribution Q_n an **n^{th} -order inflation** of P . Clearly, if $P(A, B, C)$ does not admit an n^{th} -order inflation for some n , then it cannot be realized in the triangle scenario. Deciding if the degree- n lifting of $P(A, B, C)$ is a member of the set of degree- n diagonal marginals can be cast as a linear program [47].

If the linear program is infeasible, i.e., if no n^{th} -order inflation exists for $P(A, B, C)$, then the program will find a witness to detect its incompatibility. Such a witness will be of the form

$$\bar{F} \cdot P^{\otimes n} < \min_{Q_n} \bar{F} \cdot Q_n^n, \quad (10)$$

where $\bar{F}$ is a real vector and the minimum on the right hand side is taken over all distributions Q_n satisfying Eq. (5). Call F the n^{th} -degree polynomial such that $F(Q) = \bar{F} \cdot Q^{\otimes n}$ for all Q . For some distributions P , the inflation technique will thus output a polynomial witness of incompatibility F , hence solving the corresponding (Approximate) Causal Compatibility problem.

Note that, for $n \geq n'$, any distribution P admitting an n^{th} -order inflation also admits an $(n')^{\text{th}}$ -order inflation. This suggests that we might be able to detect the incompatibility of a distribution P via the inflation technique just by taking the order n high enough.

Since any polynomial of a probability distribution can be lifted to a linear function acting on g -liftings, we can also use the inflation technique to attack Approximate Causal Optimization, as long as the function F to minimize happens to be a polynomial. Suppose that this is the case and that F has degree g . We wish to minimize $F(P)$ over all distributions compatible with the triangle scenario. Our first step would be to express F as a vector $\bar{F}$, such that $F(P) = \bar{F} \cdot P^{\otimes g}$, for all distributions P . Our second step consists in solving the linear program

$$f_n \equiv \min_{Q_n} \bar{F} \cdot Q_n^g, \\ \text{where } Q_n^g \text{ is defined by Eq. (6),} \\ \text{and such that } Q_n \text{ is a distribution} \\ \text{satisfying condition (5).} \quad (11)$$

² For instance, the result in Example 2 of Ref. [40] excluding the so-called “W-distribution” can be recovered using second order inflation if the higher degree relation of Eq. (9) is imposed. Without that relation, however, the W-distribution is excluded only at *third* order inflation.

Since the g -lifting of any distribution P compatible with the triangle scenario can be viewed as the diagonal marginal of degree g of a distribution Q_n satisfying (5), we thus have that $f_n \leq f_* = \min_P F(P)$, for all n , just as in the definition of Approximate Causal Optimization. Moreover, $f_n \geq f_{n'}$, for $n \geq n'$, i.e., as we increase the order n of the inflation, we should expect to obtain increasingly tighter lower bounds on f_* . If, by whatever means, we were to obtain an upper bound f_+ , then we would have solved Approximate Causal Optimization for all $\epsilon \geq f_+ - f_n$.

In the triangle scenario, the inflation technique can therefore be used to tackle both Approximate Causal Compatibility and Approximate Causal Optimization.

For further elucidation, consider another correlation scenario. In the **three-on-line scenario**, Figure 2, we again have three random variables A, B, C which are defined, respectively, via the non-deterministic functions $A(U_1), B(U_1, U_2), C(U_2)$. As always, the exogenous latent variables $\{U_1, U_2\}$ are independently distributed. The $n = 2$ inflation graph for the three-on-line scenario is depicted in Figure 6.

In this scenario, an n^{th} -order inflation corresponds to a distribution Q_n over the variables $\{A^i\}, \{B^{jk}\}, \{C^l\}$, where i, j, k, l range from 1 to n . Q_n must satisfy the linear constraints:

$$Q_n(\{A^i = a^i, B^{jk} = b^{jk}, C^l = c^l\}) = Q_n(\{A^i = a^{\pi(i)}, B^{jk} = b^{\pi(j)\pi'(k)}, C^l = c^{\pi'(l)}\}), \quad (12)$$

for all permutations of n elements π, π' . Expanded out for $n = 2$, Eq. (12) becomes

$$\begin{aligned} Q_2(a^1, a^2, b^{11}, b^{12}, b^{21}, b^{22}, c^1, c^2) \\ = Q_2(a^2, a^1, b^{21}, b^{22}, b^{11}, b^{12}, c^1, c^2) \end{aligned} \quad (13a)$$

$$= Q_2(a^1, a^2, b^{12}, b^{11}, b^{22}, b^{21}, c^2, c^1) \quad (13b)$$

Additionally, relating the degree- g liftings of P to the diagonal marginal in this scenario requires

$$Q_n^g(\bigwedge_{i=1}^g \{A^i = a^i, B^{ii} = b^i, C^i = c^i\}) = \prod_{i=1}^g P(a^i, b^i, c^i), \quad (14)$$

for any choice of integer g such that $g \leq n$. Expanded out for $g=n=2$, Eq. (14) identifies the diagonal marginal for this scenario $Q_{n=2}^{g=2} = \sum_{B^{12}, B^{21}} Q_{n=2}$ as

$$Q_{n=2}^{g=2}(a^1, a^2, b^{11}, b^{22}, c^1, c^2) = P(a^1, b^{11}, c^1)P(a^2, b^{22}, c^2) \quad (15)$$

The above ideas are easy to generalize to arbitrary correlation scenarios (remember, though, that correlation scenarios are just a special class of causal structures).

3.2 Inflation of an Arbitrary Correlation Scenario

To set up the n^{th} -order inflation of an arbitrary correlation scenario, first imagine n independent copies of all the latent variables, and then consider all the observable variables which are children of these, following the prescription of the original correlation scenario. Each observable variable in the inflation graph has as many superindices as latent variables it depends on. Then, one must impose symmetry restrictions on the total probability distribution Q_n , demanding that it be invariant under any relabeling-permutations applied to the index of any one latent variable, i.e.,

$$Q_n(\{A_1^{\bar{i}_1} = a_1^{\bar{i}_1}, \dots, A_m^{\bar{i}_m} = a_m^{\bar{i}_m} : \bar{i}_1, \dots, \bar{i}_m\}) = Q_n(\{A_1^{\bar{i}_1} = a_1^{\bar{\pi}^{L_1}(\bar{i}_1)}, \dots, A_m^{\bar{i}_m} = a_m^{\bar{\pi}^{L_m}(\bar{i}_m)} : \bar{i}_1, \dots, \bar{i}_m\}), \quad (16)$$

for all vectors $\bar{\pi} = (\pi^1, \dots, \pi^L)$ of L independent permutations (one for each latent variable or index type). Here $\bar{i}_x$ denotes the tuple of superindices on which variable A_x depends. It should go almost without saying that we demand non-negativity and normalization of the inflation probabilities

$$Q_n(\{A_1^{\bar{i}_1} = a_1^{\bar{i}_1}, \dots, A_m^{\bar{i}_m} = a_m^{\bar{i}_m} : \bar{i}_1, \dots, \bar{i}_m\}) \geq 0, \quad \sum_{\vec{a}} Q_n(\vec{a}) = 1. \quad (17)$$

The central object we consider, then, is the set of all diagonal marginals consistent with such an n^{th} -order inflation. We denote such a generic diagonal marginal by

$$Q_n^g := Q_n \left(\bigwedge_{i=1}^g \{A_1^{i\dots i} = a_1^i, \dots, A_m^{i\dots i} = a_m^i\} \right). \quad (18)$$

The compatibility conditions

$$Q_n^g \left(\bigwedge_{i=1}^g \{A_1^{i\dots i} = a_1^i, \dots, A_m^{i\dots i} = a_m^i\} \right) = \prod_{i=1}^g P(a_1^i, \dots, a_m^i) \quad (19)$$

require the degree- g lifting of P to be consistent with such a degree- g diagonal marginal.

Notice that any distribution Q_n subject to the constraints (16-19) must be such that the marginals associated with relabellings of the indices of the diagonal variables obey the same compatibility conditions as the canonical diagonal marginals do, i.e.,

$$Q_n \left(\bigwedge_{i=1}^g \{A_1^{\bar{\pi}^{l_1}(i\dots i)} = a_1^i, \dots, A_m^{\bar{\pi}^{l_m}(i\dots i)} = a_m^i\} \right) = \prod_{i=1}^g P(a_1^i, \dots, a_m^i), \quad (20)$$

for all $\bar{\pi}$.

Actually, the original description of the inflation technique in Ref. [40] imposes the constraints (20) rather than (16-19) over the distribution Q_n , as demanding the existence of a distribution Q_n satisfying condition (20) can be shown to enforce over $P(a_1, \dots, a_m)$ exactly the same constraints as demanding the existence of a distribution satisfying (16-19). Indeed, as noted in Ref. [40, App. C], any distribution Q_n satisfying (20) can be twirled or symmetrized (see Appendix) to a distribution $\tilde{Q}_n$ satisfying Eqs. (16-19). For convenience, from now on we will just refer to the formulation of the inflation technique involving the symmetries (16). This formulation has the added advantage that the symmetry constraints can be exploited to reduce the time and memory complexity of the corresponding linear program, see for instance Ref. [48].

It isn't hard to see how this general notion of inflation can also be used to tackle Approximate Causal Compatibility and Approximate Causal Optimization in general correlation scenarios $\mathcal{G}$:

Problem 3. Inflation for Causal Compatibility

INPUT: A positive integer n , a causal structure $\mathcal{G}$ and a particular probability distribution over the observed variables P .

PRIMAL LINEAR PROGRAM:

$$\begin{aligned} & \min_{Q_n} 0, \\ & \text{where } P \text{ relates to } Q_n \text{ by Eqs. (18,19),} \\ & \text{and such that } Q_n \text{ satisfies conditions (16,17).} \end{aligned} \quad (21)$$

DUAL LINEAR PROGRAM:

$$\begin{aligned} & \min_{\bar{F}} \bar{F} \cdot P^{\otimes n}, \\ & \text{such that } 0 \leq \bar{F} \cdot Q_n^{g=n} \leq 1, \\ & \text{where } Q_n^g \text{ is defined by Eq. (18),} \\ & \text{and such that } Q_n \text{ satisfies conditions (16,17).} \end{aligned} \quad (22)$$

SUMMARY: If the degree- n lifting of P is not in the set of degree- n diagonal marginals consistent with an n^{th} -order inflation of $\mathcal{G}$, then the returned dual variable $\bar{F}$ will witness the incompatibility of P per $\bar{F} \cdot P^{\otimes n} < 0$ while $\bar{F} \cdot P'^{\otimes n} \geq 0$ for all distributions P' compatible with $\mathcal{G}$.

Similarly,

Problem 4. Inflation for Causal Optimization

INPUT: A positive integer n , a causal structure $\mathcal{G}$ and a degree- g polynomial function F of the probabilities of the observed events.

LINEAR PROGRAM:

$$\begin{aligned} f_n &\equiv \min_{Q_n} \bar{F} \cdot Q_n^g, \\ \text{where } Q_n^g &\text{ is defined by Eq. (18),} \\ \text{and such that } Q_n &\text{ satisfies conditions (16,17).} \end{aligned} \quad (23)$$

SUMMARY: The program returns a degree- g diagonal marginals consistent with an n^{th} -order inflation of $\mathcal{G}$ which minimizes the input function F . Since such diagonal marginals contain all degree- g liftings of distributions compatible with $\mathcal{G}$, it follows that f_n is a lower bound on the minimum value of F over all distributions compatible with $\mathcal{G}$.

In certain practical cases, we may not know the full probability distribution of the observable variables, but only the probabilities of a restricted set E of *observable events*. As we will see in Section 5, this often happens when we map the causal compatibility problem from a general causal structure to a correlation scenario. To apply the inflation technique to those cases, rather than fixing the value of all probability products, like in Eq. (19), we will impose the constraints

$$\sum_{\vec{a}^1 \in e^1, \dots, \vec{a}^n \in e^n} Q_n(\{A_1^{i \dots i} = a_1^i, \dots, A_m^{i \dots i} = a_m^i\}_i) = \prod_{i=1}^n P(e^i), \quad (24)$$

for all $e^1, \dots, e^n \in E$. Any distribution Q_n satisfying both (16) and (24) will be dubbed an n^{th} order inflation of the distribution of observable events.

For example, consider again the three-on-line scenario (Fig. 2), and assume that our experimental setup just allows us to detect events of the form $e(a) \equiv \{(A, B, C) : A = B = C = a\}$. Then our set of observable events is $E = \cup_a \{e(a)\}$ and the input of the causal inference problem is the distribution $\{P(e), e \in E\}$. An n^{th} order inflation Q_n of $P(e)$ would satisfy Eq. (12) and the linear conditions

$$Q_n(\{A^i = a^i, B^{ii} = a^i, C^i = a^i\}_i) = \prod_{i=1}^n P(e(a^i)) = \prod_{i=1}^n P(a^i, a^i, a^i). \quad (25)$$

4 Convergence of Inflation

The main result of this article is that the inflation technique can be used to solve Approximate Causal Compatibility and Approximate Causal Optimization for arbitrarily small values of ϵ , just by taking the order n of the inflation high enough. Depending on which of the two problems we wish to solve and which causal structures are involved, we will have either finite-order convergence or asymptotic convergence.

4.1 On finite-order convergence

Even at low orders, the Inflation Technique has been shown to provide very good outer approximations to the set of distributions compatible with the triangle scenario [40]. Furthermore, for certain correlation scenarios, a second-order inflation can be shown to fully characterize the set of compatible distributions.

Figure 6: The second-order inflation graph of the three-on-line scenario.

Consider, for instance, the three-on-line scenario (Figure 2), whose second-order inflation was depicted in Figure 6. Note that condition (15) implies that

$$Q_{n=2}^{g=2}(A^1=a^1, C^2=c^2) = P(A=a^1)P(C=c^2) \quad (26)$$

and condition (13b) implies that

$$Q_{n=2}(A^1=a^1, C^1=c^1, C^2=c^2) = Q_{n=2}(A^1=a^1, C^1=c^2, C^2=c^1). \quad (27)$$

From the last condition, it follows that $Q_{n=2}(A^1=a, C^1=c) = Q_{n=2}(A^1=a, C^2=c)$. Invoking condition (26), we thus have that

$$Q_{n=2}(A^1=a^1, C^1=c^1) = P(A=a^1)P(C=c^1). \quad (28)$$

This is sufficient to ensure that $Q_{n=2}^{g=1}$ is realizable in the three-on-line scenario, since then $P(A, B, C) = P(A, C)P(B|A, C) = P(A)P(C)P(B|A, C)$. This last expression represents a realization of $P(A, B, C)$ in the three-on-line scenario, where the hidden variables U_1, U_2 are, respectively, A and C .

This example can be generalized to prove convergence at order $n=2$ of any star-shaped correlation scenario. Star-shaped scenarios with N observable variables have the defining property that, in some subset of $N-1$ observable variables, every pair of variables share no latent parents [28, 29], see Figs. 2 and 7. (This definition assumes that every set of variables in a correlation scenario all of which have the *same* set of latent parents are implicitly merged into a single vector-value variable.) Given an arbitrary star-shaped correlation scenario with N random variables, call $B_1, \dots, B_{N-1}$ any set of $N-1$ random variables without a common ancestor; and A , the remaining variable. Using the same trick as before, one can prove that, for any $i \neq j$, $P(B_i, B_j) = P(B_i)P(B_j)$. Similarly, one can group variables B_i, B_j and argue that, for any $l \neq i, j$, $P(B_i, B_j, B_l) = P(B_i, B_j)P(B_l) = P(B_i)P(B_j)P(B_l)$. Iterating this argument, we show that $P(B_1, \dots, B_{N-1})$ factors into $N-1$ products. Analogously, it is proven that $P(A, B_{i_1}, \dots, B_{i_m}) = P(A)P(B_{i_1}, \dots, B_{i_m})$, for any set of indices $i_1, \dots, i_m$ such that $B_{i_1}, \dots, B_{i_m}$ do not share parents with A . This is enough to prove compatibility.

In these examples, using the inflation technique is an overkill, as compatibility can be determined solely by checking the satisfaction of all independence relations. There are many correlation scenarios, however, where distribution compatibility is also determined by *inequality* constraints. Examples of such “interesting”³ correlations scenarios include the triangle scenario, as well as the four-on-line scenario depicted in Figure 8. And actually, in the former scenario, the problem of compatibility of distributions is *not* completely solved by second-order inflation.

³ Here we use “interesting” in precisely the meaning of Refs. [31, 32].

Figure 7: A star-shaped correlation scenario.

Figure 8: The four-on-line correlation scenario.

Indeed, all binary variable distributions of the form

$$P_v(A=a, B=b, C=c) = \begin{cases} v & \text{if } abc=111 \\ \frac{1-v}{7} & \text{otherwise.} \end{cases} \quad (29)$$

pass the second-order inflation test for triangle scenario causal compatibility. On the other hand, the Finner inequality applied to the triangle scenario in Ref. [49, Thm. 1] certifies the incompatibility of all P_v for which $v > 57/64$.

To be clear, however, there exists situations where, conversely, a second order inflation outperforms the Finner inequality. For probability distributions of the form

$$P_{q,r}(A=a, B=b, C=c) = \begin{cases} q & \text{if } abc=000 \\ r & \text{if } abc=111 \\ \frac{1-q-r}{6} & \text{otherwise.} \end{cases} \quad (30)$$

compatibility with second order inflation for the triangle scenario requires

$$q \leq \frac{13r + 8 - 3\sqrt{16r(r+2) + 1}}{5} \quad (31)$$

This bound is strictly tighter than bounds implied by the Finner inequality [49] or the semidefinite causal compatibility constraints involving covariances of Refs. [25, 50] throughout the parameter region $0.0283 \lesssim r \leq q$.

For arbitrary correlation scenarios $\mathcal{G}$ with observed variables of specified cardinality, we inquire whether some finite-order inflation is always sufficient to characterize the set of compatible distributions. Are there causal structures for which inflation converges only asymptotically? Could the triangle scenario be such an example?

Open Question. For any correlation scenario $\mathcal{G}$, does there exist n such that n^{th} -order inflation solves exact Causal Compatibility?

Interestingly, we can prove that, when used to solve Approximate Causal Optimization, the inflation technique does not converge, in general, in a finite number of steps. Indeed, consider the trivial correlation scenario consisting of a single observable variable A and its single latent parent U . We wish to use the inflation technique to minimize the function $-P(A=0)P(A=1)$. Clearly the solution of this problem is $-\frac{1}{4}$. An n^{th} -order function inflation assessment (starting at $n \geq 2$), however, would effectively reduce this problem to the LP

$$\begin{aligned} \min_{Q_n} \quad & -Q_n^{g=2}(0, 1) \equiv -\sum_{a^3, \dots, a^n} Q_n(0, 1, a^3, \dots, a^n) \quad \text{s.t.} \\ & Q_n(a^1, \dots, a^n) \geq 0, \quad \sum_{a^1, \dots, a^n} Q_n(a^1, \dots, a^n) = 1, \\ & Q_n(a^1, \dots, a^n) = Q_n(a^{\pi(1)}, \dots, a^{\pi(n)}), \quad \forall \pi \in S^n. \end{aligned}$$

For $n = 2n'$, consider the symmetric probability distribution Q_n given by randomly choosing without replacement n bits $a^1, \dots, a^n$ from a pool of n' 0's and n' 1's. Then it can be verified that

$$-\sum_{a^3, \dots, a^n} Q_n(0, 1, a^3, \dots, a^n) = -\frac{1}{2} \frac{n'}{2n' - 1} < -\frac{1}{4}, \quad (32)$$

overshooting the magnitude of the true minimum for all n' . Nonetheless, note that the above quantity converges to the correct result of $-\frac{1}{4}$ asymptotically as $O(1/n)$.

4.2 Asymptotic convergence

In this section we will prove that, for any correlation scenario $\mathcal{G}$, the inflation technique characterizes the set of compatible correlations asymptotically. More precisely, we will show that any distribution P admitting an n^{th} order inflation is $O(1/\sqrt{n})$ -close in Euclidean norm to a compatible distribution $\tilde{P}$. Similarly, we will show that f_n , as defined in Eq. (11), satisfies $f_\star - f_n \leq O(1/n)$. In order to solve Approximate Causal Compatibility and Approximate Causal Optimization for a given value of ϵ , we just have to use the Inflation Technique up to orders $O(1/\epsilon^2)$, $O(1/\epsilon)$, respectively. Since the set of compatible distributions is closed [39], this implies that, for any incompatible distribution P , there exists n such that P does not admit an n^{th} order inflation.

Before we proceed with the proof, a note on the scope of our results is in order. The inflation technique is fairly expensive in terms of time and memory resources. At order n , it involves optimizing over probability distributions of $\sum_{i=1}^m n^{|L_i|}$ variables (we remind the reader that $L_i \subset \{1, \dots, L\}$ denotes the set of indices j such that the hidden variable U_j influences A_i). If each of these variables can take d possible values, then the number of free variables in the corresponding linear program is $N \equiv d^{\sum_{i=1}^m n^{|L_i|}}$. That is, the memory resources required by the inflation technique are superexponential on n . Add to this the fact that the best LP solvers in the market have a running time of $O(N^3)$ [51], and you will come to the conclusion that a brute-force implementation of the inflation technique in the triangle scenario is already unrealistic for $n = 4$, even in the simplest case of $d = 2$. What is the relevance, then, of proving asymptotic convergence?

For us, it is a matter of principle. Even at low orders, the inflation technique has proven itself very useful at identifying non-trivial constraints on observable probability distributions. It is therefore natural to ask whether the inflation technique just provides a partial characterization of compatibility, or, on the contrary, it reflects an alternative way of comprehending the latter. Our work settles this question completely: by proving that any unfeasible distribution must violate one of the inflation conditions, we refute the first hypothesis and validate the second one.

The key to deriving the asymptotic convergence of the Inflation Technique is the following theorem, proven in the Appendix.

Theorem 1. Let $\mathcal{G}$ be a correlation scenario with L latent variables, and let Q_n^g be the degree- g diagonal marginal of a distribution Q_n satisfying the symmetry conditions (16). Then, there exist normalized probability distributions P_μ compatible with $\mathcal{G}$ and probabilities $p_\mu \geq 0$, $\sum_\mu p_\mu = 1$ such that

$$D\left(Q_n^g, \sum_\mu p_\mu P_\mu^{\otimes g}\right) \leq O\left(\frac{Lg^2}{n}\right), \quad (33)$$

where $D(q, r) = \sum_x |q(x) - r(x)|$ denotes the total variation distance between the probability distributions $q(X), r(X)$.

This theorem can be regarded as an extension of the finite de Finetti theorem [52], that states that the marginal $P(a^1, \dots, a^g)$ of a symmetric distribution $P(a^1, \dots, a^n)$ is $O(g^2/n)$ -close in total variation distance to a convex combination of degree- g liftings.

The solvability of Approximate Causal Optimization through the Inflation Technique follows straightforwardly from Theorem 1. Let F be a polynomial of degree g , with $f_\star = \max_P F(P)$, and let Q_n be the symmetric distribution achieving the value f_n in Eq. (11). Then, by the previous theorem, we have that

$$f_n = \bar{F} \cdot Q_n^g = \sum_\mu p_\mu \bar{F} \cdot P_\mu^{\otimes g} - O\left(\frac{Lg^2}{n}\right) = \sum_\mu p_\mu F(P_\mu) - O\left(\frac{Lg^2}{n}\right) \geq f_\star - O\left(\frac{Lg^2}{n}\right). \quad (34)$$

It follows that

$$f_n \leq f_\star \leq f_n + O\left(\frac{Lg^2}{n}\right).$$

Proving the analog result for Approximate Causal Compatibility is only slightly more complicated. Let P be a probability distribution over the observed variables, and suppose that P admits an n^{th} -order inflation Q_n . Define the second-degree polynomial $N(R) = \sum_{\bar{a}} (R(\bar{a}) - P(\bar{a}))^2$, and let $\bar{N}$ be a linear functional such that $\bar{N} \cdot q^{\otimes 2} = N(q)$ for all distributions q . Note that, due to conditions (19), $\bar{N} \cdot Q_n^2 = N(P) = 0$. Thus the minimum value f_n of $\bar{N} \cdot Q_n^2$ over all diagonal marginals of degree 2 of a distribution Q_n satisfying the symmetry conditions (16) is such that $f_n \leq 0$. Invoking Eq. (34) for $g = 2$, we have that

$$f_n \leq f_\star \leq f_n + O\left(\frac{L}{n}\right), \quad (35)$$

where $f_\star$ is the minimum value of $N(Q)$ over all compatible distributions Q . This implies that there exists a compatible distribution $\tilde{P}$ such that

$$\sqrt{N(\tilde{P})} = \|P - \tilde{P}\|_2 \leq O\left(\sqrt{\frac{L}{n}}\right). \quad (36)$$

This proof of convergence easily extends to scenarios where we only know the probabilities of set E of observable events. Indeed, choosing the polynomial N such that $N(R) = \sum_{e \in E} (P(e) - R(e))^2$, and following the same derivation as in Eq. (35), we conclude that a distribution of observable events admitting an n^{th} order inflation is $O\left(\sqrt{\frac{L}{n}}\right)$ -close in Euclidean norm to a compatible distribution.

5 Unpacking Causal Structures

So far we have just been referring to correlation scenarios, i.e., those causal structures where all observed variables only depend on a number of independent latent variables. However, in a general causal model, the value of a given variable can depend, not only on latent variables, but also on the values of other observed

variables. In the following, we define procedures call **exogenization** and **unpacking** which cumulatively map the problem of causal compatibility with an arbitrary causal structure to problems of causal compatibility with the structure's implicit constituent correlation scenarios. Consequently, these procedures enable application of the inflation technique to general causal structure via preprocessing into correlation scenarios.

Suppose $\mathcal{G}$ is a DAG with latent variables. If U is an endogenous (non-root) latent variable in $\mathcal{G}$, one can **exogenize** U by first adding all possible directed edges originating from a parent of U and terminating at a child of U , and then deleting from $\mathcal{G}$ all directed edges which terminated at U . The resulting graph admits precisely the same set of feasible observed distributions as $\mathcal{G}$, per Ref. [53, Sec. 3.2]. Hereafter, therefore, we restrict our attention to causal compatibility problems involving causal structures where all latent variables are exogenous.

In addition, we will always consider distributions as implicitly conditional on the values of any exogenous observable variables. Of course, this mapping from raw probability distributions to conditional probability distributions only makes sense if the distribution of exogenous observable variables factorizes, i.e., if all exogenous observable variables are independent from each other. As an example, the sorts of distributions we consider for the Bell scenario depicted in Figure 9 are of the form $P(A, B|X, Y)$, as opposed to $P(A, B, X, Y)$.

To go from general causal structures to correlation scenarios, we introduce counterfactual variable sets, in which we consider all the different ways a variable can respond to its observable parents as distinct variables. We call the procedure for eliminating all dependencies between observed variables **unpacking**. Unpacking is related to — but distinct from — the *single world intervention graphs* introduced in Ref. [17] and the *e-separation* method developed in Ref. [16]. As quantum physicists, we understand unpacking as a manifestation of *counterfactual definiteness*, which is a natural assumption mysteriously inconsistent with quantum theory [54–56]. Since counterfactual definiteness *does* hold in the “classical” causal models considered in this paper, we promote maximally exploiting this assumption as a first step towards resolving any causal compatibility problem.

By way of example, consider the structure $\mathcal{G}^1$ depicted in Figure 10. The correlation scenario which results from unpacking $\mathcal{G}^1$ — assuming that all observable variables are discretely valued in the range $[0, 1]$ — is shown in Figure 11. The unpacked scenario can be thought of as having either seven binary-valued vari-

Figure 9: The Bell scenario. (Not a correlation scenario.)

Figure 10: The example structure $\mathcal{G}^1$. (Not a correlation scenario.)

Figure 11: The unpacking of $\mathcal{G}^1$ in Figure 10, which we denote by $\mathcal{G}^2$.

ables $\{A^{X=0}, A^{X=1}, B, C^{A=0,B=0}, C^{A=0,B=1}, C^{A=1,B=0}, C^{A=1,B=1}\}$ or simply three variables, two of which are vector valued. We use the latter interpretation for the visualization of the unpacked scenario, but the former interpretation is convenient to explicitly relate the packed distributions to the unpacked distributions. A distribution over the observable variables in $\mathcal{G}^1$ (conditioned on the exogenous observable variable X) is compatible with $\mathcal{G}^1$ iff there exists another distribution compatible with $\mathcal{G}^2$ (over $\mathcal{G}^2$'s observable variables) such that the first distribution is recovered via suitable *varying* marginals of the latter. Explicitly,

$$P_{\text{original}}(A=a, B=b, C=c|X=x) = P_{\text{unpacked}}(A^{X=x}=a, B=b, C^{A=a,B=b}=c). \quad (37)$$

Note that for each of the eight distinct choices of $\{a, b, x\} \in \{0, 1\}^3$, the marginal of P_{unpacked} referenced in Eq. (37) is distinct. The *subset of variables* specifying the relevant marginal of P_{unpacked} does *not* vary depending on the value of c , however.

We now describe how to unpack an arbitrary causal structure. Let A be an observed variable. We denote the observable parents of A as $\text{paOBS}[A]$. Suppose the (set) $\text{paOBS}[A]$ can take d different values, e.g.: $\text{paOBS}[A] \in \{1, \dots, d\}$. The different values of $\text{paOBS}[A]$ are generally vector-valued; we may nevertheless denote such value-tuples by a single scalar index, for compactness of notation. To unpack the vertex A , we break all edges between $\text{paOBS}[A]$ and A , unpacking A into the counterfactual variable set $\{A^{\text{paOBS}[A]=1}, \dots, A^{\text{paOBS}[A]=d}\}$. Unpacking all the endogenous observed variables, and regarding the resulting counterfactuals as observed variables themselves, we arrive at the associated correlation scenario.

The probabilities of the observed variables in $\mathcal{G}$ can be obtained from the probabilities of a set of measurable events in the associated correlation scenario $\mathcal{G}'$. To be clear, let $\bar{A}(\bar{X})$ denote all the observable endogenous (exogenous) variables in $\mathcal{G}$. Then,

$$P_{\text{original}}(\bar{A}=\bar{a}|\bar{X}=\bar{x}) = P_{\text{unpacked}}\left(\bigwedge_i A_i^{\text{paOBS}[A_i]=\{\bar{a}, \bar{x}\}_{A_i}}=a_i\right) \quad (38)$$

where $\{\bar{a}, \bar{x}\}_{A_i}$ denotes selecting those elements out of the set $\bar{a} \cup \bar{x}$ which corresponding to the values of $\text{paOBS}[A_i]$.

The original Approximate Causal Compatibility (Approximate Causal Optimization) problem in $\mathcal{G}$ is thus mapped to an Approximate Causal Compatibility (Approximate Causal Optimization) problem in the correlation scenario $\mathcal{G}'$, with a non-trivial set of observable events. The inflation technique can then be applied on $\mathcal{G}'$ to solve either problem on the original structure $\mathcal{G}$ up to arbitrary precision ϵ .

Note that unpacking can be valuable even without further inflation. For instance, unpacking the instrumental scenario of Figure 4 leads to an associated correlation scenario which is trivial, such as depicted in Figure 12. Any distribution over the four variables $\{A^{X=0}, A^{X=1}, B^{A=0}, B^{A=1}\}$ is compatible with Figure 12. Nevertheless, demanding that the distributions $P(A, B|X)$ admit such an unpacking leads to nontrivial constraints. We can formulate the admission of an unpacked distribution as a linear program, via Eq. (38). Ex-

Figure 12: The unpacking of the instrumental scenario.

explicitly formulating this linear program for the instrumental scenario looks like

$$\begin{aligned} P_{\text{original}}(A=a, B=b|X=x) &= P_{\text{unpacked}}(A^x=a, B^a=b), \\ P_{\text{unpacked}}(A^0=a^0, A^1=a^1, B^0=b^0, B^1=b^1) &\geq 0, \\ \sum_{a^0 a^1 b^0 b^1} P_{\text{unpacked}}(A^0=a^0, A^1=a^1, B^0=b^0, B^1=b^1) &= 1, \end{aligned}$$

and leads to the famous instrumental inequalities [42], such as

$$P(A=0, B=0|X=0) + P(A=0, B=1|X=1) \leq 1. \quad (39)$$

This example substantially generalizes. Unpacking alone also completely solves the causal compatibility problem for any single-district causal structure (see [18] for a definition) containing one (or fewer) latent variables. This includes all Bell scenarios and the entire hierarchy of their relaxations as described in Ref. [57].

Furthermore, one can take advantage of known results concerning **observationally equivalent** causal structures. We say that $\mathcal{G}$ and $\mathcal{G}'$ are observationally equivalent whenever both structures admit precisely the same set of compatible distributions over their observed variables. Prop. 5 in Ref. [53] is a prescription for replacing latent variables with sets of directed edges while preserving observational equivalence. We encourage aggressive application of that prescription in order to convert (unpacked) correlation scenarios into observationally equivalent structures which can be unpacked further. For instance, it can be invoked to convert the four-on-line correlation scenario of Figure 8 into the observationally equivalent Bell scenario of Figure 9, to convert the three-on-line correlation scenario of Figure 2 into the observationally equivalent graph $A \rightarrow B \leftarrow C$ with no latent variables, or to replace all latent variables in star scenarios such as Figure 7 with inwards-pointing directed edges. Interestingly, all the challenging causal structures collected in Figure 14 of Ref. [53] unpack to the four-on-line correlation scenario. One can readily demonstrate the non-saturation of those structures by converting their unpacked forms to the Bell scenario à la Prop. 5 of Ref. [53], and then unpacking a second time.

Of course, unpacking supplemented with inflation is far more powerful than unpacking alone. Unpacking and inflation are both naturally formulated as linear programs, and hence can be easily combined into a single composite linear program to solve Causal Compatibility or Causal Optimization (over polynomials of conditional distributions).

6 Conclusions

The inflation technique was first proposed by Wolfe *et al.* [40] as a means to obtain strong causal compatibility constraints for arbitrary causal structures. Here, we have formulated inflation as a formal hierarchy of problems for assessing causal compatibility relative to correlation scenarios. We have proven the inflation hierarchy to be complete, in the sense that any distribution incompatible with a given correlation scenario will be detected as incompatible by inflation. More quantitatively, we showed that any distribution P passing the n^{th} -order inflation test is $O\left(\frac{1}{\sqrt{n}}\right)$ -close in Euclidean norm to some other distribution which *can* be realized within the considered scenario.

The inflation technique is fully applicable to any causal structure, since unpacking allows one to map any causal assessment problem (for either distributions or functions) to an equivalent assessment problem relative to a correlation scenario. The observed distribution in the original structure is mapped to probabilities pertaining to restricted sets of measurable events in the unpacked correlation scenario. Since, however, our proof of the convergence of inflation allowed for restricted sets of measurable events, the convergence theorems are still applicable when using inflation to assess compatibility relative to general causal structures.

We have therefore shown that the inflation technique is much more than a useful machinery to derive statistical limits; it is an alternative way to define causal compatibility!

For the purpose of practical causal discovery, we envision the inflation technique being used as *final refinement*. That is, inflation (and unpacking) should be employed as a postprocessing, after first filtering the set of candidate causal explanations by means of computationally-cheaper but less-sensitive algorithms. Indeed, our attitude concerning the primacy of single-district graphs reflects our implicit assumption that all the kernels of a multi-district graph will have been identified. In other words, whatever distribution is being assessed for causal compatibility via inflation, we are presuming that it has already been verified to satisfy the nested Markov property (NMP) relative to the considered graph [14, 41]. Thus, we envision testing for compatibility via inflation only *after* first testing for compatibility via NMP algorithms. This is not strictly necessary, as our results here imply that the inflation technique alone can recover all the constraints implied by NMP, though we imagine it is relatively inefficient to impose NMP only indirectly through inflation.

Alternatively, inflation could be used to estimate the distances of a distribution P from the sets of distributions compatible with various causal structures. We speculate that such distances could prove valuable in helping compute scores for the *ranked* causal discovery problem [8, 9], though we defer further analysis to future research.

Acknowledgement: We thank T. Fritz, T.C. Fraser, A. Acín, and A. Pozas-Kerstjens for interesting discussions. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Colleges and Universities. This work was not supported by the European Research Council.

References

- [1] J. Pearl, *Causality* (Cambridge University Press, 2009).
- [2] D. M. Chickering, “Optimal structure identification with greedy search,” *J. Mach. Learn. Res.* **3**, 507 (2002).
- [3] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning* (The MIT Press, 2009).
- [4] P. Spirtes and K. Zhang, “Causal discovery and inference: concepts and recent methodological advances,” *Applied Informatics* **3**, 3 (2016).
- [5] J. D. Ramsey and D. Malinsky, “Comparing the Performance of Graphical Structure Learning Algorithms with TETRAD,” *arXiv:1607.08110* (2016).
- [6] N. Wermuth, “Graphical Markov Models, Unifying Results and Their Interpretation,” in *Wiley StatsRef* (American Cancer Society, 2015) pp. 1–29.
- [7] S. Lauritzen and K. Sadeghi, “Unifying Markov properties for graphical models,” *Ann. Statist.* **46**, 2251 (2018).
- [8] S. Magliacane, T. Claassen, and J. M. Mooij, “Ancestral Causal Inference,” (2016).
- [9] R. J. Evans, “Model selection and local geometry,” *arXiv:1801.08364* (2018).
- [10] T. Verma and J. Pearl, “An Algorithm for Deciding if a Set of Observed Independencies Has a Causal Explanation,” in *Proc. 8th Conf. Uncertainty in Artificial Intelligence* (1992) pp. 323–330.
- [11] D. Geiger and C. Meek, “Graphical Models and Exponential Families,” in *Proc. 14th Conf. Uncertainty in Artificial Intelligence* (1998) pp. 156–165.
- [12] J. Tian and J. Pearl, “On the Testable Implications of Causal Models with Hidden Variables,” in *Proc. 18th Conf. Uncertainty in Artificial Intelligence* (2002) pp. 519–527.
- [13] C. Kang and J. Tian, “Inequality Constraints in Causal Models with Hidden Variables,” in *Proc. 22nd Conf. Uncertainty in Artificial Intelligence* (2006) pp. 233–240.

- [14] T. S. Richardson, J. M. Robins, and I. Shpitser, “Nested Markov Properties for Acyclic Directed Mixed Graphs,” in *Proc. 28th Conf. Uncertainty in Artificial Intelligence* (2012) p. 13.
- [15] B. Steudel and N. Ay, “Information-Theoretic Inference of Common Ancestors,” *Entropy* **17**, 2304 (2015).
- [16] R. J. Evans, “Graphical methods for inequality constraints in marginalized DAGs,” in *IEEE International Workshop on Machine Learning for Signal Processing* (2012).
- [17] T. S. Richardson and J. M. Robins, *Single World Intervention Graphs (SWIGs) : A Unification of the Counterfactual and Graphical Approaches to Causality* (Now Publishers Inc, 2013).
- [18] R. J. Evans, “Margins of discrete Bayesian networks,” *Ann. Stat.* **46**, 2623 (2018).
- [19] T. Fritz, “Beyond Bell’s theorem: correlation scenarios,” *New J. Phys.* **14**, 103001 (2012).
- [20] R. Chaves, L. Luft, T. O. Maciel, D. Gross, D. Janzing, and B. Schölkopf, “Inferring latent structures via information inequalities,” in *Proc. 30th Conf. Uncertainty in Artificial Intelligence* (2014) pp. 112–121.
- [21] D. Rosset, C. Branciard, T. J. Barnea, G. Pütz, N. Brunner, and N. Gisin, “Nonlinear Bell Inequalities Tailored for Quantum Networks,” *Phys. Rev. Lett.* **116**, 010403 (2016).
- [22] A. Tavakoli, “Bell-type inequalities for arbitrary noncyclic networks,” *Phys. Rev. A* **93**, 030101 (2016).
- [23] R. Chaves, “Polynomial Bell inequalities,” *Phys. Rev. Lett.* **116**, 010402 (2016).
- [24] N. Miklin, A. A. Abbott, C. Branciard, R. Chaves, and C. Budroni, “The entropic approach to causal correlations,” *New J. Phys.* **19**, 113041 (2017).
- [25] A. Kela, K. V. Prillwitz, J. Aberg, R. Chaves, and D. Gross, “Semidefinite Tests for Latent Causal Structures,” *IEEE Trans. Info. Theo.* **66**, 339 (2020).
- [26] M. Weilenmann and R. Colbeck, “Analysing causal structures with entropy,” *Proc. Roy. Soc. A* **473**, 20170483 (2017).
- [27] C. Branciard, D. Rosset, N. Gisin, and S. Pironio, “Bilocal versus nonbilocal correlations in entanglement-swapping experiments,” *Phys. Rev. A* **85**, 032119 (2012).
- [28] A. Tavakoli, M. O. Renou, N. Gisin, and N. Brunner, “Correlations in star networks: from Bell inequalities to network inequalities,” *New J. Phys.* **19**, 073003 (2017).
- [29] F. Andreoli, G. Carvacho, L. Santodonato, R. Chaves, and F. Sciarrino, “Maximal violation of n -locality inequalities in a star-shaped quantum network,” *New J. Phys.* **19**, 113020 (2017).
- [30] C. J. Wood and R. W. Spekkens, “The lesson of causal discovery algorithms for quantum correlations: causal explanations of Bell-inequality violations require fine-tuning,” *New J. Phys.* **17**, 033002 (2015).
- [31] J. Henson, R. Lal, and M. F. Pusey, “Theory-independent limits on correlations from generalized Bayesian networks,” *New J. Phys.* **16**, 113043 (2014).
- [32] J. Pienaar, “Which causal structures might support a quantum–classical gap?” *New J. Phys.* **19**, 043021 (2017).
- [33] Van Himbeec *et al.*, “Quantum violations in the Instrumental scenario and their relations to the Bell scenario,” *Quantum* **3**, 186 (2019).
- [34] T. C. Fraser and E. Wolfe, “Causal compatibility inequalities admitting quantum violations in the triangle structure,” *Phys. Rev. A* **98**, 022113 (2018).
- [35] D. Mond, J. Smith, and D. van Straten, “Stochastic factorizations, sandwiched simplices and the topology of the space of explanations,” *Proc. Roy. Soc. A* **459**, 2821 (2003).
- [36] T. Kocka and N. Zhang, “Dimension Correction for Hierarchical Latent Class Models,” in *Proc. 18th Conf. Uncertainty in Artificial Intelligence* (2002) pp. 267–274.
- [37] E. Allman, J. Rhodes, and A. Taylor, “A Semialgebraic Description of the General Markov Model on Phylogenetic Trees,” *SIAM J. Disc. Math.* **28** (2012).
- [38] D. Geiger and C. Meek, “Quantifier elimination for statistical problems,” *CoRR* **abs/1301.6698** (2013), arXiv:1301.6698 .
- [39] D. Rosset, N. Gisin, and E. Wolfe, “Universal bound on the cardinality of local hidden variables in networks,” *Quant. Info. & Comp.* **18**, 910 (2018).
- [40] E. Wolfe, R. W. Spekkens, and T. Fritz, “The Inflation Technique for Causal Inference with Latent Variables,” *J. Caus. Inf.* **7** (2019).
- [41] I. Shpitser, R. J. Evans, T. S. Richardson, and J. M. Robins, “Introduction to nested Markov models,” *Behaviormetrika* **41**, 3 (2014).
- [42] J. Pearl, “On the Testability of Causal Models with Latent and Instrumental Variables,” in *Proc. 11th Conf. Uncertainty in Artificial Intelligence* (1995) pp. 435–443.
- [43] S. Basu, R. Pollack, and M. Roy, *Algorithms in Real Algebraic Geometry*, Algorithms and Computation in Mathematics (Springer Berlin Heidelberg, 2006).
- [44] R. F. Werner and M. M. Wolf, “All-multipartite Bell-correlation inequalities for two dichotomic observables per site,” *Phys. Rev. A* **64**, 032112 (2001).
- [45] R. J. Evans and T. S. Richardson, “Marginal log-linear parameters for graphical markov models,” *J. Roy. Stat. Soc. B* **75**, 743 (2013).
- [46] R. J. Evans and T. S. Richardson, “Smooth, identifiable supermodels of discrete dag models with latent variables,” *Bernoulli* **25**, 848 (2019).
- [47] D. Alevras and M. W. Padberg, *Linear Optimization and Extensions* (Springer Berlin Heidelberg, 2001).

- [48] I. P. Gent, K. E. Petrie, and J.-F. Puget, “Chapter 10 - Symmetry in Constraint Programming,” in *Handbook of Constraint Programming*, Vol. 2 (Elsevier, 2006) pp. 329 – 376.
- [49] M.-O. Renou, Y. Wang, S. Boreiri, S. Beigi, N. Gisin, and N. Brunner, “Limits on Correlations in Networks for Quantum and No-Signaling Resources,” *Phys. Rev. Lett.* **123**, 070403 (2019).
- [50] J. Aberg, R. Nery, C. Duarte, and R. Chaves, “Semidefinite tests for quantum network topologies,” (2020).
- [51] K. Anstreicher, “Linear Programming in $O([n^3/\ln n]L)$ Operations,” *SIAM J. Optimization* **9**, 803 (1999).
- [52] P. Diaconis and D. Freedman, “Finite exchangeable sequences,” *Ann. Prob.* **8**, 745–764 (1980).
- [53] R. J. Evans, “Graphs for margins of Bayesian networks,” *Scandinavian J. Stat.* **43**, 625 (2016), Note: the arXiv version numbers the propositions differently than the published version. In particular, Prop. 5 in the published version corresponds to Prop. 6.1 in the arXiv version.
- [54] R. D. Gill, “Statistics, Causality and Bell’s Theorem,” *Statist. Sci.* **29**, 512 (2014).
- [55] M. S. Leifer and R. W. Spekkens, “Pre- and Post-Selection Paradoxes and Contextuality in Quantum Mechanics,” *Phys. Rev. Lett.* **95**, 200405 (2005).
- [56] Y.-C. Liang, R. W. Spekkens, and H. M. Wiseman, “Specker’s parable of the overprotective seer: A road to contextuality, nonlocality and complementarity,” *Phys. Rep.* **506**, 1 (2011).
- [57] R. Chaves, D. Cavalcanti, and L. Aolita, “Causal hierarchy of multipartite Bell nonlocality,” *Quantum* **1**, 23 (2017).

Appendix: Generalizing de Finetti's Theorem

The purpose of this Appendix is to prove Theorem 1.

We will first prove the result for the triangle scenario; the generalization to arbitrary correlation scenarios will be obvious. Given a distribution of $3n^2$ variables $Q(\{A^{i,j}, B^{k,l}, C^{p,q}\})$, consider its symmetrization $\tilde{Q}$, defined by

$$\tilde{Q}(\{A^{ij}=a^{ij}, B^{kl}=b^{kl}, C^{pq}=c^{pq}\}) = \frac{1}{n!^3} \left(\sum_{\pi, \pi', \pi'' \in S_n} Q(\{A^{ij}=a^{\pi(i)\pi'(j)}, B^{kl}=b^{\pi'(k)\pi''(l)}, C^{pq}=c^{\pi''(p)\pi(q)}\}) \right). \quad (A1)$$

Note that $\tilde{\tilde{Q}} = \tilde{Q}$. In addition, any distribution Q satisfying the symmetry condition (5) fulfills $\tilde{Q} = Q$ and any symmetrized distribution satisfies (5).

Let $\mathbb{I}_{\{\hat{a}(i,j), \hat{b}(k,l), \hat{c}(p,q)\}}$ be the deterministic distribution assigning the values $\hat{a}(i,j)$, $\hat{b}(k,l)$, $\hat{c}(p,q)$ to the random variables $A^{i,j}$, $B^{k,l}$, $C^{p,q}$, for $i, j, k, l, p, q \in \{1, \dots, n\}$. Since any distribution is a convex combination of deterministic points, it follows that any distribution satisfying Eq. (5) can be expressed as a convex combination of symmetrized distributions of the form $\tilde{\mathbb{I}}_{\{\hat{a}(i,j), \hat{b}(k,l), \hat{c}(p,q)\}}$.

For clarity of notation, let us assume that the values $\{\hat{a}(i,j), \hat{b}(k,l), \hat{c}(p,q)\}$ are fixed and denote the symmetrization of $\mathbb{I}_{\{\hat{a}(i,j), \hat{b}(k,l), \hat{c}(p,q)\}}$ by $\tilde{P}$. Call $\tilde{P}^1$ its diagonal marginal of degree 1, i.e., $\tilde{P}(A^{1,1}, B^{1,1}, C^{1,1})$. It can be verified, by symmetry, that $\tilde{P}^1$ is given by the formula:

$$\tilde{P}^1(a, b, c) = \frac{1}{n^3} \sum_{i,j,k=1}^n \delta(\hat{a}(i,j), a) \delta(\hat{b}(j,k), b) \delta(\hat{c}(k,i), c), \quad (A2)$$

where $\delta(i,j)$ denotes the Kronecker delta function, i.e., $\delta(i,j) = 1$ if $i = j$ or zero otherwise. Notice that $\tilde{P}^1(a, b, c)$ can be reproduced in the triangle scenario. Indeed, the latent variables are i, j, k , they can take values in $\{1, \dots, n\}$ and are uniformly distributed. The observed variables a, b, c are deterministic functions of (i,j) , (j,k) and (k,i) , respectively.

Consider now the diagonal marginal of degree g , $\tilde{P}^g \equiv \tilde{P}(A^{1,1}, B^{1,1}, C^{1,1}, \dots, A^{g,g}, B^{g,g}, C^{g,g})$. By symmetry, it is expressed as:

$$\begin{aligned} \tilde{P}^g(a^1, b^1, c^1, \dots, a^g, b^g, c^g) = \\ \frac{1}{n^3(n-1)^3 \dots (n-g+1)^3} \sum_{\bar{i}, \bar{j}, \bar{k}} \prod_{x=1}^g \delta(\hat{a}(i^x, j^x), a^x) \delta(\hat{b}(j^x, k^x), b^x) \delta(\hat{c}(k^x, i^x), c^x), \end{aligned} \quad (A3)$$

where the sum is taken over all tuples $\bar{i}, \bar{j}, \bar{k} \in \{1, \dots, n\}^g$ with no repeated indices, i.e., such that $i^x \neq i^y$, $j^x \neq j^y$, $k^x \neq k^y$ for $x \neq y$.

Now, compare $\tilde{P}^g$ with the degree- g lifting $(\tilde{P}^1)^{\otimes g}$. It is straightforward that

$$\begin{aligned} (\tilde{P}^1)^{\otimes g}(a^1, b^1, c^1, \dots, a^g, b^g, c^g) &= \prod_{x=1}^g \tilde{P}^1(a^x, b^x, c^x) \\ &= \frac{1}{n^{3g}} \sum_{\bar{i}, \bar{j}, \bar{k}} \prod_{x=1}^g \delta(\hat{a}(i^x, j^x), a^x) \delta(\hat{b}(j^x, k^x), b^x) \delta(\hat{c}(k^x, i^x), c^x), \end{aligned} \quad (A4)$$

where, this time, the sum contains all possible tuples $\bar{i}, \bar{j}, \bar{k} \in \{1, \dots, n\}^g$. The total variation distance between the two distributions is bounded by $\frac{1}{n^3(n-1)^3 \dots (n-g+1)^3} - \frac{1}{n^{3g}}$ times the number of tuples with non-repeated indices (namely, $n^3(n-1)^3 \dots (n-g+1)^3$), plus $1/n^{3g}$ times the number of tuples with repeated indices (namely, $n^{3g} - n^3(n-1)^3 \dots (n-g+1)^3$). The result is

$$D(\tilde{P}^g, (\tilde{P}^1)^{\otimes g}) \leq 2 \left(1 - \frac{n^3(n-1)^3 \dots (n-g+1)^3}{n^{3g}} \right). \quad (A5)$$

Finally, let Q_n be any distribution satisfying Eq. (5). Then, $Q_n = \sum_{\mu} p_{\mu} \tilde{P}_{\mu}$, where $p_{\mu} \geq 0$, $\sum_{\mu} p_{\mu} = 1$, and $\tilde{P}_{\mu}$ is the result of symmetrizing $P_{\mu} = \mathbb{I}_{\{\hat{a}_{\mu}(i,j), \hat{b}_{\mu}(k,l), \hat{c}_{\mu}(p,q)\}}$, for some values $\{\hat{a}_{\mu}(i,j), \hat{b}_{\mu}(k,l), \hat{c}_{\mu}(p,q)\}$. By convexity of the total variation distance, we have that

$$\begin{aligned} D\left(Q_n^g, \sum_{\mu} p_{\mu} (\tilde{P}_{\mu}^1)^{\otimes g}\right) &\leq \sum_{\mu} p_{\mu} D\left(\tilde{P}_{\mu}^g, (\tilde{P}_{\mu}^1)^{\otimes g}\right) \leq 2 \left(1 - \frac{n^3(n-1)^3 \dots (n-g+1)^3}{n^{3g}}\right) \\ &= 2 \left(1 - \left[\frac{n^g - n^{g-1}(1+2+\dots+g-1) + O(n^{g-2})}{n^g}\right]^3\right) = 2 \left(1 - \left[1 - \frac{O(g^2)}{n}\right]^3\right) = O\left(\frac{3g^2}{n}\right). \end{aligned} \quad (\text{A6})$$

Extending this result to general correlation scenarios is straightforward, so we will just sketch the proof. First, the action of the corresponding symmetrization over a deterministic distribution equals a distribution $\tilde{P}$ whose 1-marginal $\tilde{P}^1(a_1, \dots, a_m)$ is a uniform mixture over the tuple of indices $\vec{i}$ of deterministic distributions of the form $\prod_{x=1}^m \delta(a_x, a_x(\vec{i}^{L_x}))$. Again, we remind the reader that $L_x \subset \{1, \dots, L\}$ denotes the indices of the hidden variables on which A_x depends. It thus follows that $\tilde{P}^1$ is realizable within the correlation scenario. The diagonal marginal $\tilde{P}_g$ is also a uniform mixture of deterministic distributions of a similar type, but where no repeated indices are allowed between the different blocks of variables. The statistical difference between $\tilde{P}^g$ and $(\tilde{P}^1)^{\otimes g}$ is thus bounded by

$$2 \left(1 - \frac{n^L(n-1)^L \dots (n-g+1)^L}{n^{gL}}\right) = O\left(\frac{Lg^2}{n}\right). \quad (\text{A7})$$

As before, the general result follows from the convexity of the total variation distance.