

This is the accepted manuscript made available via CHORUS. The article has been published as:

Explaining quantum correlations through evolution of causal models

Robin Harper, Robert J. Chapman, Christopher Ferrie, Christopher Granade, Richard Kueng, Daniel Naoumenko, Steven T. Flammia, and Alberto Peruzzo

Phys. Rev. A **95**, 042120 — Published 17 April 2017

DOI: [10.1103/PhysRevA.95.042120](https://doi.org/10.1103/PhysRevA.95.042120)

Explaining quantum correlations through evolution of causal models

Robin Harper,^{1,2,*} Robert J. Chapman,^{2,3,*} Christopher Ferrie,^{1,2} Christopher Granade,^{1,2}
Richard Kueng,⁴ Daniel Naoumenko,² Steven T. Flammia,^{1,2} and Alberto Peruzzo^{2,3}

¹Centre for Engineered Quantum Systems, University of Sydney, Sydney, NSW, Australia

²School of Physics, University of Sydney, Sydney, NSW, Australia

³Quantum Photonics Laboratory, School of Engineering, RMIT University, Melbourne, Victoria 3000, Australia

⁴Institute for Theoretical Physics, University of Cologne, Germany

(Dated: March 16, 2017)

We propose a framework for the systematic and quantitative generalization of Bell's theorem using causal networks. We first consider the multi-objective optimization problem of matching observed data while minimizing the causal effect of nonlocal variables and prove an inequality for the optimal region that both strengthens and generalizes Bell's theorem. To solve the optimization problem (rather than simply bound it), we develop a novel genetic algorithm treating as individuals causal networks. By applying our algorithm to a photonic Bell experiment, we demonstrate the trade-off between the quantitative relaxation of one or more local causality assumptions and the ability of data to match quantum correlations.

While it seems conceptually obvious that causality lies at the heart of physics, its exact nature has been the subject of constant debate. The fundamental implications of quantum theory shed new light on this debate. It is thought these implications may lead to new insights into the foundations of quantum theory, and possibly even quantum theories of gravity [1–10].

These realizations have their roots in the Einstein-Podolski-Rosen thought experiment [11] and the fundamental theorems of Bell [12] and of Kochen and Specker [13]. A cornerstone of modern physics, Bell's theorem, rigorously excludes classical concepts of causality. Roughly speaking Bell's theorem states that the following concepts are mutually inconsistent: (1) reality; (2) locality; (3) measurement independence; and (4) quantum mechanics.

In philosophical discussions, typically one rejects (1) or (2), which together are often referred to as *local causality*, though the other options have been considered as well. In studies with an operational bent, however, one often considers relaxations of (2) or (3) which is what we concern ourselves with here. These relaxations have been addressed from different perspectives, but only regarding specific causal influences in isolation [14–23], whereas here we wish to study all possible relaxations of the causal assumptions implied by (2) and (3) simultaneously.

The framework of *causal networks* [24, 25] is wildly successful within the field of machine learning and has led some physicists to utilize them to elucidate the tension between causality and Bell's theorem. Recently, Wood and Spekkens have shown that existing principles behind causal discovery algorithms (namely, the absence of fine tuning) still cannot be reconciled with entanglement induced quantum correlations *even if one admits nonlocal models* [9]. However, such results only hold for the exact distributions, and would not necessarily apply to exper-

imental data due to measurement noise, or a relaxation of the demand of reproducing exactly the quantum correlations. Clearly, the further away from the quantum correlations one is allowed to stray, the more likely a locally causal model can be found.

Here we propose a framework for systematic and quantitative generalizations of Bell's theorem by using causal networks. The idea, depicted in Figure 1, is to consider the multi-objective optimization problem of matching the observed data from an experiment while minimizing the causal effect of nonlocal variables. It is in this sense of matching experimental data that we are explaining the quantum correlations. Our first contribution is a rigorous lower bound for this optimization problem, demonstrating a generalization of Bell's theorem. Theorem 1 below establishes that there must exist a tradeoff between the




Figure 1. A sketch of the concept of Pareto optimality for demarcating the boundary between local causality and quantum correlations. In this picture, Bell's theorem rules out the origin only. Our results rule out an entire region of possible models in the presence of relaxations of Bell's assumptions. We rule out this region both rigorously with Theorem 1 and numerically with the evolutionary algorithm that we developed specifically for this task.

* These authors contributed equally to this work.

goodness of fit to experimental data and the quantitative amount of causal influence for *any* model.

This theorem rules out a portion of the space allowed by this new framework, but the bounds are not tight. To solve the optimization problem, and hence numerically find the *optimal* bounds, we develop a type of genetic algorithm called a multi-objective evolutionary algorithm (MOEA) to quantify the relaxations necessary to reproduce the data generated by experiments on entangled quantum systems [26–28]. Our genetic algorithm treats as individuals causal networks and we develop genetic operators which represent the evolution of these networks. By applying our algorithm to a photonic Bell experiment, we show that the tradeoff between the quantitative relaxation of one or more local causality assumptions and the ability to match quantum correlations appears linear.

The outline of the paper is as follows. In [Section I](#) and [Section II](#) we set out the background of the causal models we use and the mathematics required to convert a probability distribution into a fitness function. In [Section III](#) we provide analytic bounds on causal influence. In [Section IV](#) we describe the experiment that provided the input to the algorithm. [Section V](#) briefly describes the result of applying the genetic algorithm to the experimental data. [Section VI](#) describes the process by which we convert the problem into one that can be explored using evolutionary operators and details the construction of the algorithm. We conclude in [Section VII](#) with a discussion.

I. CAUSAL MODELS FOR BELL EXPERIMENTS

The formalism of causal models allows us to quantify the relaxations necessary to avoid the contradiction in Bell’s theorem and, more importantly, explore the trade-offs necessary in minimizing the amount by which the assumptions are violated. Building off the work of Chaves *et al* [29], we make all this concrete through a quantification of the relaxation of each assumption in the context of causal models. The task of minimizing the amount of the relaxation is a multi-objective optimization problem. Bell’s theorem is recast as the statement that all objectives cannot be simultaneously minimized. We explore the trade-offs through the concept of Pareto optimality.

The prototypical “Bell experiment” has two distant parties, often named Alice and Bob. We suppose that Alice and Bob each have devices with binary measurement settings, respectively labeled x and y . Conditioned on these measurement settings, their devices also record binary events, labeled a (Alice) and b (Bob). Suppose it is empirically observed that a and b are correlated. Bell defined a locally causal model of such correlations as follows: there exists a “hidden variable” λ which is the common cause both of a conditioned on x , and of b conditioned on y . We write these random variables as $a | x$ and $b | y$, respectively. Formally, the general conditional




Figure 2. Causal networks for Bell-type experiments. On the left is the local hidden variable model, which respects the assumptions going into Bell’s famous no-go theorem. Such a model cannot account for certain correlations obtained from measuring entangled particles. The graph in the middle contains a causal link between the measurement settings. Such a model exploits the detection loophole and violates measurement independence. Finally, on the right is a superluminal model which contains a causal link between the outcomes of the experiments.

distribution is assumed to satisfy

$$\Pr(a, b | x, y, \lambda) = \Pr(a | x, \lambda) \Pr(b | y, \lambda). \quad (1)$$

Moreover, it is assumed that the choices of settings can be made such that each of x and y can be set independently of the hidden variable λ ,

$$\Pr(x, y | \lambda) = \Pr(x | \lambda) \Pr(y | \lambda) = \Pr(x) \Pr(y). \quad (2)$$

Such an assumption is often motivated by the injection of randomness into the measurement settings or the free-will of Alice and Bob. Bell’s theorem can be stated succinctly as follows: the conditional distributions describing the outcomes of some experiments on quantum systems cannot be factorized as in Eqs. (1) and (2).

A causal network is a directed acyclic graph with nodes representing random variables and edges denoting causal relationships between variables. The defining feature of such networks is the factorization of joint probabilities. Generally, suppose we have nodes $\{x_0, x_1, \dots, x_K\}$, each of which represents a random variable in our model. We will assume that each such random variable is discrete, and without loss of generality, will assume integer labels $x_i \in \{0, \dots, \dim x_i - 1\}$ for its possible values. The edges in the causal network of these variables are defined such that

$$\Pr(x_0, x_1, \dots, x_K) = \prod_{i=0}^K \Pr(x_i | \text{pa}_i), \quad (3)$$

where pa_i denotes the parents of node i .

Take, for example, the causal network in [Figure 2a](#). In general, we can decompose the joint distribution $\Pr(a, b, x, y, \lambda)$ in terms of conditional distributions as

$$\Pr(a, b, x, y, \lambda) = \Pr(a | b, x, y, \lambda) \Pr(b | x, y, \lambda) \Pr(x | y, \lambda) \Pr(y | \lambda) \Pr(\lambda). \quad (4)$$

Using the causal network to eliminate conditionals, (3) implies

$$\Pr(a, b, x, y, \lambda) = \Pr(a|x, \lambda) \Pr(b|y, \lambda) \Pr(x) \Pr(y) \Pr(\lambda), \quad (5)$$

which are identical to Bell’s assumptions on local hidden variable models. Thus, Bell’s theorem is equivalent to the statement that certain quantum correlations cannot be realized by the causal network in Figure 2a.

II. RELAXING BELL’S ASSUMPTIONS

It is known that quantum mechanical correlation arising in a Bell-type experiment *can* however be explained by adding a new causal link to the local hidden variable network [15, 18]. Two examples are shown in Figure 2. In many practical cases, these causal links are not entirely unphysical from the standpoint of respecting relativity and free-will, for example. The reason being that experiments do not actually conform to the exact assumptions Bell made—there are noisy detectors, non-random number generation, losses, inability to space-like separate “Alice” and “Bob,” and so on. When this is the case, such causal models are said to be exploiting *loopholes*.

In Figure 2b, a causal model that allows correlations between the measurement settings is shown. In the same spirit, we could have had either x or y be causally dependent on λ or another hidden variable. Such models are often called *superdeterministic* and are ruled out by the assumption that Alice and Bob are not colluding and have free-will or access to independent randomness. If the experiment only approximately satisfies these assumptions—perhaps due to low detection efficiency—one can still model the data with a local hidden variable said to be exploiting the *detection loophole* [30]. The question of quantifying the amount of independence of the measurement settings necessary has been addressed

from multiple perspectives and has practical quantum cryptographic consequences [15–22].

In Figure 2c, a causal model which allows correlations between the measurement outcomes is shown. This is, and similar models are, called *nonlocal* and could potentially even allow for superluminal signaling. A quintessential example of a nonlocal model which reproduces the predictions of quantum theory is Bohmian mechanics. Toner and Bacon studied the amount of nonlocality necessary to simulate quantum correlation in the context of classical communication costs [14, 23], while Wolf has expressed nonlocality in terms of the compressibility of experimental observations [31].

The current studies, mentioned above, quantifying the relaxations of the causal assumption necessary to replicate quantum correlations are rather disjoint. Recently, Chaves *et al.* placed the question in context of causal networks and found that some measures of these relaxations can be cast as efficiently solvable linear programs [29]. We build on this idea and consider a completely abstract framework amenable to any set of random variables using a single measure of the causal influence of one variable on another. This allows us to consider *all possible relaxations simultaneously* and thus explore the trade-offs necessary to simulate quantum correlations with hidden variable models.

We will now state our model more technically. For consistency we formulate the problem in the context of the two-party Bell experiment, but we emphasize that this approach generalizes in an obvious way to any set of random variables. A model, M , is specified by a joint distribution

$$\Pr(a, b, x, y, \lambda|M). \quad (6)$$

We label the empirical frequencies $F(a, b, x, y)$ and denote the total variational distance (TVD) of a model to these frequencies by

$$\text{TVD}(M) = \|\Pr(a, b, x, y|M) - F(a, b, x, y)\|_1, \quad (7)$$

where the vector being normed is labeled by (a, b, x, y) . Here the 1-norm of a vector x is simply $\|x\|_1 = \sum_i |x_i|$.

The causal influence is defined for a *general* graph as follows:

$$C_{x_i \rightarrow x_j}(M) := \max_{x_i, x'_i, \text{pa}_j}^* \|\Pr(x_j|x'_i, \text{pa}_j \setminus \text{pa}_j^2, M) - \Pr(x_j|x_i, \text{pa}_j \setminus \text{pa}_j^2, M)\|_1, \quad (8)$$

where $\text{pa}_j \setminus \text{pa}_j^2$ is the set of parents of x_j that are not also grandparents of x_j , and where \max^* indicates that the maximization over x_i, x'_i and pa_j is restricted to feasible assignments. That is, the maximization does not consider assignments outside the support of M . In words, the causal influence is non-zero when changing x_i leads to a change in x_j . It is quantified by maximizing over latent variables of the target that are not also latent variables of the control.

For example, if we want to minimize the causal influence between two variables a to b in Figure 2c we consider

$$C_{a \rightarrow b}(M) := \max_{a, a', y}^* \|\Pr(b|a, y, M) - \Pr(b|a', y, M)\|_1. \quad (9)$$

We include the conditions $\Pr(a), \Pr(a') \neq 0$ to prevent the causal influence being maximized by an assignment outside the support of the random variable A ; the maxi-

mization should be taken over all feasible assignments.

Intuitively, this definition represents how distinguishable the different settings of a are when viewed through measurements of b . That is, if a does not causally affect b , then it is not possible for a change in a to be detectable through b alone. We adopt this definition in lieu of the traditional approach of using *interventions*, wherein an external agent imposes a particular value of a while holding all else fixed, effectively cutting out any causal links incident on a other than one originating from the experimentalist themselves. Though some novel experiments have been performed using intervention to reason about quantum mechanics [32], we cannot intervene on quantum mechanical models in general, such that we must instead maximize over conditions for the experiment, here represented by the maximization over a and y .

The task then is to find a model M which minimizes TVD and $C_{\alpha \rightarrow \beta}$ for each $\alpha \rightarrow \beta$ ruled out by local causality and measurement independence. If the empirical frequencies contain some causal dependence between two variables, then either the model must also contain such causal dependence or the observed frequencies from the model must be different from the empirical frequencies. Perhaps interestingly, one might be able to “trade” unwanted causal influence between one pair of variables for another, while maintaining the same TVD. Thus, the problem of determining “how much” relaxation of Bell’s causal assumptions is necessary to match an empirical observed frequency becomes much more interesting and nuanced.

Suppose two models M_1 and M_2 both match the data equally well—i.e. $\text{TVD}(M_1) = \text{TVD}(M_2)$ —but M_2 has some unwanted causal influence $a \rightarrow b$, say, and M_1 does not—that is, $0 = C_{a \rightarrow b}(M_1) < C_{a \rightarrow b}(M_2)$. Clearly, M_1 is preferred and we say M_2 is *dominated* by M_1 . For many objectives, the situation is more complex but can be handled by the concept of Pareto optimality.

Let \mathcal{M} be the set of all models. Let each model’s fitness be represented by the function $f : \mathcal{M} \rightarrow \mathbb{R}^n$, where n is number of objectives. Define the partial order \prec as follows:

$$M \succcurlyeq M' \Leftrightarrow f(M)_k \geq f(M')_k, \quad (10)$$

for all $k \in \{0, n-1\}$. If $M \succcurlyeq M'$, we say M' dominates M (or is equivalent to M , if $M' \succcurlyeq M$ holds as well). The set $\mathcal{P} \subset \mathcal{M}$ of *Pareto optimal* models is now defined as follows:

$$\mathcal{P} = \{M \in \mathcal{M} : \{M' \in \mathcal{M} : M \succcurlyeq M', M' \neq M\} = \emptyset\}. \quad (11)$$

This says that a model is Pareto optimal if the set of other models which dominate it is empty. In other words, the Pareto optimal is the set of non-dominated models.

III. ANALYTICAL BOUNDS

In this section, we provide analytical bounds which relate the amount of causal influences exhibited by any




Figure 3. Photon pairs at 808 nm are emitted via spontaneous parametric down-conversion from a 404 nm pumped bismuth triborate (BiBO) crystal and we prepare polarization enabled photons, labelled “Alice’s Qubit” and “Bob’s Qubit”. The components used are: half wave plates (HWP), quarter wave plates (QWP), polarizing beam splitters (PBS), polarization maintaining fibers (PMF) and an optical delay line (ODL). Once the state is prepared a polarization tomography setup enables projection of each qubit onto any pure polarization state, which is sufficient to perform two-qubit tomography.

model M to its agreement with the empirical frequencies $F(a, b, x, y)$. For the sake of simplicity, we restrict ourselves to analyzing the causal influence between the variables a and b —see Figure 2c. However, we emphasize that analogous statements are valid for causal influences between any two variables. For the variables a and b , the empirical frequencies themselves admit a causal influence

$$C_{a \rightarrow b}(F) = \max_{a, a', y}^* \|F(b|a, y) - F(b|a', y)\|_1 \quad (12)$$

which is defined in complete analogy to (9). To state our theorem, we must define two more quantities. Let $\mathcal{M}_\tau = \mathcal{M}_\tau(F)$ be the set of models having $\text{TVD}(M) \leq \tau$ with respect to the empirical frequencies F , and denote by $f^* = \min_a F(a)$ the minimum empirical marginal frequency.

Theorem 1. For all models $M \in \mathcal{M}_\tau$ and $\tau < 2f^*$,

$$|C_{a \rightarrow b}(F) - C_{a \rightarrow b}(M)| \leq \frac{2\tau(4f^* - \tau)}{f^*(2f^* - \tau)}. \quad (13)$$

We point out that the bound (13) becomes loose and eventually diverges if the minimum empirical marginal frequency f^* approaches zero or if the TVD of the class of models becomes too large relative to f^* .

The proof of Theorem 1 can be found in Appendix A.

IV. BELL EXPERIMENT AND DATA

As input for the MOEA we use data from a polarization photonic Bell experiment, shown in Figure 3 [33]. Indistinguishable horizontally polarized ($|H\rangle$) photon pairs

are generated via type-1 spontaneous parametric down-conversion. Both polarization qubits are rotated into a diagonal state $\frac{1}{\sqrt{2}}(|H\rangle + |V\rangle)$ by a half wave plate (HWP) with fast axis at $\frac{\pi}{8}$ from vertical, where $|V\rangle$ denotes vertical polarization. A polarization phase rotation is applied to photon 1 by two quarter wave plates (QWPs) and a HWP, while photon 2 has its state optimized by a polarizing beamsplitter (PBS). Both photons are collected in polarization maintaining optical fiber (PMF) and are incident on the two input faces of a fibre-coupled PBS, which transmits $|H\rangle$ and reflects $|V\rangle$, preparing Alice's and Bob's qubits. The configuration of the optical fibres results in a σ_x operation applied to Alice's qubit. By measuring in the coincidence basis, we post-select the state

$$\rho = \frac{1+\gamma}{2} |\Phi^+\rangle \langle \Phi^+| + \frac{1-\gamma}{2} |\Phi^-\rangle \langle \Phi^-|, \quad (14)$$

where $|\Phi^\pm\rangle$ are the Bell states $\frac{1}{\sqrt{2}}(|H_A V_B\rangle \pm |V_A H_B\rangle)$ with subscript A (B) corresponding to Alice's (Bob's) qubit. The parameter γ defines the coherence of the state which depends on the overlap of the two photons after the fibre-coupled PBS and is controlled by the optical delay line (ODL). The state prepared when $\gamma = 1$ is a maximally entangled Bell state $|\Phi^+\rangle$ and when $\gamma = 0$ is an incoherent mixture $\frac{1}{2}(|H_A V_B\rangle \langle H_A V_B| + |V_A H_B\rangle \langle V_A H_B|)$. The polarization tomography setup in Figure 3 enables projection onto any pure state and can be used for two-qubit state tomography [34]. The photons are detected with silicon avalanche photo-diodes and coincidence counts recorded by a timing card.

Our input for the MOEA is a normalized frequency distribution $F(a, b, x, y)$ across binary measurement settings for Alice ($x = \{x_1, x_2\}$) and Bob ($y = \{y_1, y_2\}$), and binary measurement outcomes $a = \{|H_A\rangle, |V_A\rangle\}$ and $b = \{|H_B\rangle, |V_B\rangle\}$ respectively. The measurement settings are controlled by wave plate angles in the tomography and the measurement outcome is the collapse of the state onto one of the four basis state $|H_A H_B\rangle$, $|H_A V_B\rangle$, $|V_A H_B\rangle$ or $|V_A V_B\rangle$. A single measurement is the number of photon pairs recorded for a fixed integration time and can be written as $N_{ab}^{xy} = \mathcal{N}\tau \langle ab| U^{xy} \rho U^{xy\dagger} |ab\rangle$ for measurement settings x, y and measurement outcomes a, b . \mathcal{N} is the total photon flux, τ is the integration time and U^{xy} is the operation of the wave plates. We calculate $F(a, b, x, y)$ by measuring all combinations of x, y, a and b , and normalizing by the total number of photon pairs recorded. We note that this experiment is not performed in a loophole-free way, but nonetheless provides us with the quantum correlations we wish to analyse.

Typically, Bell experiments aim to violate the CHSH inequality [35], confirming that quantum mechanical systems cannot be described with local hidden variable models. The CHSH inequality is calculated as

$$|S| \leq 2, \text{ where} \quad (15a)$$

$$S = \mathbb{E}[x_1 y_1] - \mathbb{E}[x_2 y_1] + \mathbb{E}[x_1 y_2] + \mathbb{E}[x_2 y_2], \quad (15b)$$

where $\mathbb{E}[xy]$ defines the correlation between Alice's ($x = \{x_1, x_2\}$) and Bob's ($y = \{y_1, y_2\}$) measurements, given as

$$\mathbb{E}[xy] = \frac{N_{H_A H_B}^{xy} - N_{H_A V_B}^{xy} - N_{V_A H_B}^{xy} + N_{V_A V_B}^{xy}}{\mathcal{N}\tau}. \quad (16)$$

While the CHSH inequality holds for systems which respect local causality, a pair of quantum entangled particles can achieve a maximum value of $|S| = 2\sqrt{2}$. By tuning the γ parameter in (14), then for measurement settings fixed to be optimal for the case $\gamma = 1$, we can prepare states that obey the CHSH inequality when $\gamma \leq \sqrt{2} - 1$ and states that violate it. In order to achieve the maximum violation of the CHSH inequality, it is necessary to choose specific wave plate angles for x and y . Here, we are not interested in violating the CHSH inequality; however, we can use it to benchmark our results from the MOEA.

V. EDGE OF REALITY

Using the experimental data (where $\gamma = 0.984$), we searched for the Pareto optimal models by developing a multi-objective evolutionary algorithm to find the best underlying probability distributions for a causal network. Since this represents a trade-off between a local realistic model and real-world correlations, we call the Pareto optimal surface the ‘‘edge’’ of reality.

An *individual* of the population is a probability distribution over the nodes of a given causal network (each such individual is a causal model, M) and its *multi-objective fitness* depends on how close the model can reproduce the experimental data and the amount of causal influences between nonlocally separated variables.

As an initial step, we examined relaxing one causal edge at a time, beginning with a causal influence from a to b —that is, Alice's outcome is allowed to influence Bob's. The Pareto front (the numerical approximation to the Pareto optimal) is shown below in Figure 5. Like the theoretical bounds (which while not linear, are nearly so in the considered domain), the front appears to be linear (Pearson's ρ^2 value of 0.997, with bisquare robust fitting). That is, increasing locality violations allows observed (quantum) correlations to be more exactly matched, the trade-off being approximately *linear* in nature. Next, we relax the causal edges $a \rightarrow b$ and $x \rightarrow \lambda$ simultaneously. The found Pareto front is shown in Figure 6. Again, we see that the front appears linear ($\rho^2 = 0.9902$). We have also used our algorithm to test other causal networks and found the results to be quantitatively identical to these two cases.

VI. THE EVOLUTIONARY ALGORITHM

In order to find the Pareto front of solutions, it is necessary to find feasible probability distributions that

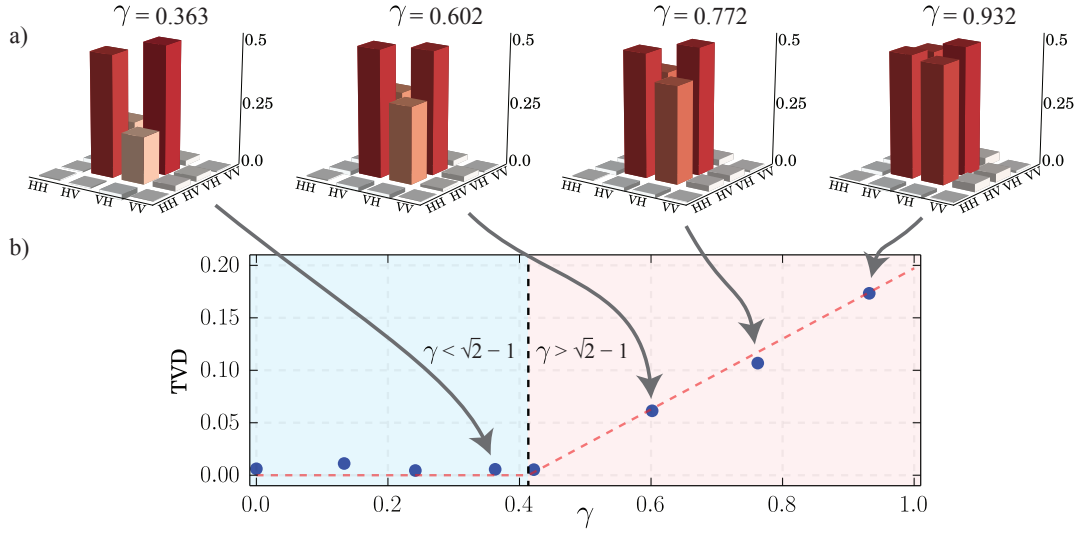


Figure 4. (a) The real component of experimentally measured density matrices for a range γ values measured via two-qubit state tomography. (b) the results of the MOEA (run as an EA) on a local model (see Figure 2a) to generate the best achievable TVD for various values of γ from a Bell-like experiment described in Section IV. As theoretically predicted above a certain threshold ($\gamma = \sqrt{2} - 1$) the local model can no longer explain the measurement results with zero TVD. This threshold corresponds to violating the CHSH inequality in (15). The linear increase in TVD corresponds to the linear increase in S as discussed in Section IV.




Figure 5. The Pareto front for the causal network in Figure 2c using the data from a photonic Bell experiment. The vertical axis labels the causal influence (9) while the horizontal axis labels the closeness to experimental data (7). The blue circles are the values for the non-dominated models found by the evolutionary algorithm. For comparison purposes, the straight line is a linear fit to these data.

give rise to the required TVD with the required causal-violations(s). There is no known way of doing this analytically. Even in simple single edge causal models the search space is prohibitively large and objective non-convex. This search space grows rapidly with additional causal edges. Evolutionary Algorithms are known methods for finding such Pareto fronts where there is only limited knowledge of the underlying search landscape. We wish to numerically find the Pareto optimal set of models representing Bell experiment data. To do so, we use evolutionary computation [36].

Such algorithms are generally well studied for functions of the form $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$. However, here the domain of our objective function f is \mathcal{M} , i.e. the probability distributions on the causal network. Consequently, there are implicit constraints on the relative values these distribution can take (for instance, in each node they need to sum to 1) and so we have devised a set of evolutionary operators that allow the probability distribution of an arbitrary causal network to be evolved. With this we combine several evolutionary computation strategies to evolve and explore the Pareto front of a given arbitrary




Figure 6. The Pareto front for a local causal network with added $a \rightarrow b$ and $x \rightarrow \lambda$ edges using the data from a photonic Bell experiment. The vertical axis labels the closeness to experimental data (7). The two horizontal axes label the causal influences (9) for the added edges. The blue circles are the values for the non-dominated models found by the evolutionary algorithm. The flat surface is a linear fit to these data.

network.

A. Evolutionary Algorithm Overview

As the cornerstone of our multiobjective evolutionary algorithm (MOEA) we utilize the well-known and well-understood NSGA-II algorithm [39]. Although the NSGA-II algorithm specifies both the generation and selection procedures, we utilize the DEAP software library [37] which provides the NSGA-II algorithm only for the “select” stage. The method by which we proceed is to use the $(\mu + \lambda)$ algorithm (detailed in Algorithm 1) where we set $\lambda = \mu$ to be the population size. For the purposes of avoiding confusion we note that the $(\mu + \lambda)$ is more properly an algorithm used with a subset of evolutionary algorithms known as evolutionary strategies, and thus is not part of the toolkit of the separate branch known as genetic algorithms. Consequently, our algorithm is not strictly a genetic algorithm but is an evolutionary algorithm. Although we use an implementation of $(\mu + \lambda)$, by setting $\lambda = \mu$ the algorithm is functionally equivalent to the generation algorithms used in genetic algorithms. In this paper we make no distinction between genetic algorithms and the more general term evolutionary algorithm in the classification of the algorithms used. The overall implementation of the algorithm is thus functionally

Algorithm 1 Evolutionary Algorithm

Input: Population sizes $\mu, \lambda \in \mathbb{N}$, crossover and mutation probabilities p_\times, p_μ
Input: an initial population P_0 , a number of generations N_{gen} .
Input: A genetic operator $\text{EVALUATE}(I)$ that annotates individuals with their fitness $\tilde{f}(I)$.
Input: A genetic operator $\text{MUTATE}(I)$ that mutates an individual in-place.
Input: A genetic operator $\text{CROSSOVER}(I_1, I_2)$ that crosses over two individuals in-place.
Input: A genetic operator $\text{SELECT}(P, \mu)$ that selects μ individuals from the population P .
Output: A Pareto front P_* of individuals with respect to the fitness functions implemented by EVALUATE .
 \triangleright In this Algorithm, we follow the DEAP [37] convention of storing an individuals’ fitness as *metadata*. This prevents having to re-evaluate fitnesses for every comparison.
 $P \leftarrow P_0$
 $P_* \leftarrow \text{KDTree}(\{\})$ \triangleright Initialize the Pareto front to an empty k-d tree [38].
 $\text{EVALUATE}(P)$ \triangleright Evaluate each individual in the initial population.
for $i_{\text{gen}} \leftarrow 1, \dots, N_{\text{gen}}$ **do**
 $P_{\text{offspring}} \leftarrow \{\}$
while $|P_{\text{offspring}}| < \lambda$ **do**
Draw two individuals uniformly at random from P and copy them as I_1 and I_2 .
switch $u \sim \text{Uni}(0, 1)$
case $u \in [0, p_\times)$
 $\text{CROSSOVER}(I_1, I_2)$
case $u \in [p_\times, p_\times + p_\mu)$
 $\text{MUTATE}(I_1)$
case $u \in [p_\times + p_\mu, 1]$
 \triangleright Leave I_1 and I_2 unmodified.
end switch
 $P_{\text{offspring}} \leftarrow P_{\text{offspring}} \cup \{I_1\}$
end while
 $\text{EVALUATE}(P_{\text{offspring}})$
 $P \leftarrow \text{SELECT}(P \cup P_{\text{offspring}}, \mu)$ \triangleright Using the NSGA-II crowding operator, order the individuals and select the next generation from this one and the new offspring.
for $I \in P$ **do**
if there does not exist $I' \in P_*$ such that $I' \succeq I$ **then**
 \triangleright Average time complexity $O(\log |P_*|)$ for k-d trees.
 $P_* \leftarrow P_* \cup \{I\}$
end if
end for
if any individuals were added to P_* this generation **then**
 $P_* \leftarrow \{I | I \in P_* \text{ such that } \forall I' \in P_*, I' \not\succeq I\}$ \triangleright
Remove dominated individuals from the Pareto front.
Rebalance P_* .
end if
end for
return P_*

identical to the original NSGA-II algorithm, save that the selection of parents is random rather than by binary tournament selection.

Consequently this evolutionary algorithm proceeds in *generations*, each of which consists of producing λ offspring from the previous generation’s population, then selecting μ individuals from the combination of the previous population and the new offspring to form the new population. As detailed in Algorithm 1 the $(\mu + \lambda)$ algorithm is expressed abstractly in terms of *genetic operators* that create, crossover, evaluate and select individuals within each population. Thus, we form our algorithm by specifying what an individual is, the fitness functions that we use in evaluating individuals, and by providing suitable genetic operators to create “children” causal networks.

B. Representation of Individuals

Effectively, our genetic algorithm searches for Pareto optimal models $M \in \mathcal{M}$ by representing M as an assignment of conditional distributions to each node in a causal network with a fixed structure. Since the random variables at each node are constrained to be discrete, we represent the conditional distributions by *tensors*, such that finding arbitrary joint, marginal and conditional distributions over subsets of the nodes is then an exercise in standard tensor contractions.

In particular, consider a node x_i with n causal parents $\text{pa}_i = \{x_{i_1}, x_{i_2}, \dots, x_{i_n}\}$. Then, the distribution $\Pr(x_i | \text{pa}_i) = \Pr(x_i | x_{i_1}, \dots, x_{i_n})$ is given by the tensor

$$X[j_0, j_1, \dots, j_n] := \Pr(x_i = j_0 | x_{i_1} = j_1, \dots, x_{i_n} = j_n), \quad (17)$$

where we have used square brackets to indicate indices (similar to C- or Python-style notation).

We can contract repeated indices of two such tensors with the tensor at a corresponding third node to perform expectation values. For example, let A be the tensor for $\Pr(a|x, \lambda)$, B be the tensor for $\Pr(b|y, \lambda)$ and Λ be the tensor for $\Pr(\lambda)$ in the model of Figure 2c. Then, to find $\Pr(a, b|x, y)$, we compute

$$\Pr(a, b|x, y) = \sum_{\lambda} A[a, x, \lambda] B[b, y, \lambda] \Lambda[\lambda]. \quad (18)$$

The general case, allowing for arbitrary numbers of random variables and conditions, is given as Algorithm 2.

C. Fitness Functions

Our algorithm uses two different kinds of fitness functions:

1. The total variational distance (TVD) between the joint distribution computed from an individual and the observed frequencies.

Algorithm 2 Joint and Conditional Distribution Tensors from Individuals

Input: Individual I , random variables x_1, \dots, x_n , random variables y_1, \dots, y_m .
Output: Tensor $J[i_1, \dots, i_n, j_1, \dots, j_m] = \Pr(x_1 = i_1, \dots, x_n = i_n | y_1 = j_1, \dots, y_m = j_m)$ for the distribution represented by I .
 $X' \leftarrow \{x_1, \dots, x_n\}$ \triangleright X' holds those rvs we must still include.
 $F \leftarrow \{\}$ \triangleright F holds those tensor factors we include in the final contraction.
while X' is not empty **do**
 $F \leftarrow F \cup X'$
 $X' \leftarrow \bigcup_{x \in X'} \text{pa}_x \setminus F$ \triangleright Add in any parents that we have not already added.
end while
 $J \leftarrow$ Einstein sum over of F , holding indices $\{x_1, \dots, x_n, y_1, \dots, y_m\}$. \triangleright Marginalize over parents not appearing as x or y .
return J

2. Causal influences along penalized edges, as generalized from the definition given by (9) in Section II.

Dealing with each in turn, the TVD is calculated by taking the vector 1-norm between the flattened joint distribution tensor the observable variables calculated as in Section VI B and the flattened observed frequencies,

$$f_{\text{TVD}}(I) = \|J_{\text{obs}}^b(I) - F^b\|_1, \quad (19)$$

where I is an individual with joint distribution tensor $J_{\text{obs}}(I)$ over all observables, F is the tensor of observed frequencies, and where b indicates flattening—that is, reduction of an arbitrary-rank tensor to a rank-1 tensor.

As discussed in Section II, we adopt a definition of causal influence that allows us to reason even in lieu of interventions. Our definition of the causal influence $C_{a \rightarrow b}(I)$ for an individual I proceeds in three steps. First, we maximize over pairs of settings of a to find which are most distinguishable through observations of b alone. We then maximize over the conditions under which these observations are made, represented by maximizing over feasible assignments to the parents of b . Finally, we *marginalize* over those nodes which are also parents of a to prevent “hiding” causal influence; this is illustrated in Figure 7.

D. Genetic Operators

Having defined the mapping from models to individuals, we complete the specification of our algorithm by detailing the various genetic operators which act on these individuals.

a. Creation In order to create a new individual I , we must specify a new conditional distribution at each




Figure 7. An example of the causal influence measure $C_{a \rightarrow b}(I)$ given by (9) applied to a more complicated graph. The random variable y is conditioned on and maximized over, as it is a parent of b but not of a . By contrast, the variable x is also a parent of a and so it is marginalized over, resulting in the causal influence definition at right.

node of the causal graph. We do so randomly by assigning a tensor $X_i(I)$ with entries drawn uniformly from $[0, 1]$ to each node, then renormalizing to ensure $\sum_{x_i} \Pr(x_i | \text{pa}_i) = 1$. Using the tensor notation defined above,

$$X_i(I)[j_0, j_1, \dots, j_{|\text{pa}_i|}] = \frac{\tilde{X}_i(I)[j_0, j_1, \dots, j_{|\text{pa}_i|}]}{\sum_{j'_0} \tilde{X}_i(I)[j'_0, j_1, \dots, j_{|\text{pa}_i|}]}, \quad (20)$$

where $\tilde{X}_i(I)[j_0, \dots, j_{|\text{pa}_i|}] \sim \text{Uni}(0, 1)$ is the unnormalized tensor of I at X_i .

b. Crossover/Mating Given two individuals I_1 and I_2 , we mate them to produce two new individuals I'_1 and I'_2 by swapping the tensors at each node with probability $p_\chi = 0.5$. That is, for each node x_i , the corresponding tensor $X_i(I'_1)$ of I'_1 is given by

$$X_i(I'_1) = \begin{cases} X_i(I_2) & \text{with probability } p_\chi \\ X_i(I_1) & \text{with probability } 1 - p_\chi \end{cases}. \quad (21)$$

c. Mutation Given a single individual that has been selected for mutation, we proceed by first picking a node x_μ on the causal graph uniformly at random, with corresponding tensor (assuming n parents) $X_\mu[i_0, \dots, i_n]$.

One of the conditional events represented by X_μ (that is, a single element of the tensor) is selected at random and the value (and hence the probability assigned to the selected outcome) is randomly increased or decreased by a sample from a zero-mean Gaussian distribution, where the variance is a user supplied parameter. The mutated element of X_μ is then clipped to the interval $[0, 1]$, and the relevant tensor index renormalized such that (20) holds.

By way of example, if we had a node a with binary values, which in turn had one parent x also with binary values, then the information pertaining to that node would be stored in a 2×2 tensor $A[a, x]$ corresponding to the probability distribution

$$\Pr(A = a | X = x) = A[a, x] = \begin{matrix} & \begin{matrix} x=0 & x=1 \end{matrix} \\ \begin{matrix} a=0 \\ a=1 \end{matrix} & \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} \end{matrix}, \quad (22)$$

where α represents the probability of a being 0 given x is 0 and so on. From this it can be seen that it is necessary that $\alpha + \gamma = \beta + \delta = 1$. One of α , β , γ or δ would be modified as discussed above, and the remaining values renormalized to ensure that the relevant probabilities continue to sum to 1.

d. Selection For selection we used an unmodified version of the NSGA-II algorithm [39]. NSGA-II uses a fast sort algorithm to locate the non-dominated individuals and then applies a crowding distance sorting algorithm to prefer those individuals that explore different parts of the pareto front. The “best” μ individuals are retained for the next generation.

E. Decomposition of the multi-objective optimization (the Island Model)

Here we present an enhancement to the basic genetic algorithm discussed above that aids the discovery of the global Pareto front in multi-dimensional scenarios, where—as is the case here—it is possible to evolve populations to occupy the extremes of any particular front.

As discussed in Section VI G it is well known that the NSGA-II crowding becomes less effective with the exponential increase in the size of the front with the number of dimensions. However, in our case are we able to force the population to start at extreme points of the Pareto front by pre-evolving the population on structurally reduced graphs or with reduced fitness criteria. These populations are able to seed the graph we wish to explore and spread over the front, fleshing it out over multiple runs. This can then be repeated as we increase the dimensions of the fronts. This is not dissimilar to the mechanism used in NeuroEvolution of Augmenting Topologies (NEAT) [40] where populations are evolved on small neural networks prior to allowing additional links to be added. A similar idea of decomposing the objectives is explored in [41]. Effectively, where a multiple dimension Pareto front needs to be explored, different populations are evolved on all permutations of the simpler graphs (on separate “islands”) before being brought together for evolution over the full graph. This is illustrated in Figure 8.

This technique allows us to find a three-dimensional Pareto front based on a graph with two causal edges, from $x \rightarrow \lambda$ as well as $a \rightarrow b$. This was evolved using five runs of the “island model” detailed above. For each run the initial islands had a population of 300 and comprised 4 runs of 400 generations. The initial runs found populations clustered around the three extremes: (1) $\min(\text{TVD})$, hold $C_{x \rightarrow \lambda} = 0$, $C_{a \rightarrow b} = 0$; (2) $\min(C_{a \rightarrow b})$, hold $\text{TVD} = 0$, $C_{x \rightarrow \lambda} = 0$; and (3) $\min(C_{x \rightarrow \lambda})$, hold $\text{TVD} = 0$, $C_{a \rightarrow b} = 0$.

The second set of islands take the relevant individuals generated above, reduce them to the best 400 individuals representing the extreme of the Pareto fronts for that island and transplant them to expanded causal net-




Figure 8. *Island model* diagram showing the steps in evolving a three-dimensional Pareto front. This allows a the edges of the Pareto front to be found by exploring lower dimensional graphs with lower populations. Multiple runs in each of step 1 and step 2 can be done concurrently.

work graphs. In this case there are two second generation islands: one generating the two-dimensional Pareto front for $\{TVD \text{ and } C_{a \rightarrow b}\}$, with $C_{x \rightarrow \lambda}$ held to be 0 (i.e. no causal $x \rightarrow \lambda$ link); and the second the two-dimensional Pareto front $\{TVD \text{ and } C_{x \rightarrow \lambda}\}$, with $C_{a \rightarrow b}$ held to be 0. These populations are then evolved on the respective causal networks generating two dimensional Pareto Fronts similar to Figure 10. These populations are placed in an ϵ -Dominance archive. (In other-words they are only kept if they dominate all previous individuals by at least ϵ , where in this implementation ϵ was $10^{-8} + 10^{-5} * |value|$). The entire process so far is repeated several times (in this experiment 5 times) to ensure we have 2,000 suitable individuals in the archive. These individuals are, effectively, clustered on the two dimensional fronts specified in the second set of islands. This final population is used to generate the 3-dimensional Pareto front shown in Figure 9. The final island had a population of 2,000 individuals (extracted initially from the ϵ -Dominance archive), evolved for 800 generations. This constituted one run of the island model. The Model was run 5 times, with every individual generated by the model being submitted to (but not necessarily accepted by) the global ϵ -Dominance archive.

To illustrate the advantage of using this model, we have also plotted (in red) the best Pareto front found using just the basic algorithm (i.e. evolving only over the full graph). These additional points were collected over 8 runs, using a high population (6,000) and repre-

sented five times the computing power required for the Island Model. As can be seen the global ϵ -Dominance archive for the basic algorithm contained few individuals on the best global Pareto front found by the Island Model. The Pareto front for the Island Model (plotted blue) does appear to be a viable candidate for the actual global front, indicating that - for this model - the trade-offs in the different causality violations considered is linear. The front fits a linear plane with a Pearson's ρ^2 value of 0.9902, the non-fitting points being those with extremely low TVD values (the points which appear on the horizontal part of the mesh). While this still needs to be investigated further we believe it is related to experimental noise which might require increased causal violation to match the noisy data exactly. So far as we are aware this is the first evidence that tradeoffs in multiple causal violations are also linear for such graphs.

F. Previous Work and Design Decisions

Although there has been previous work in using genetic algorithms to explore Bayesian causal networks (e.g. [42], [43] and [44]), the focus of such works has been to create the network and the links therein. For instance, [45] uses a multi-objective genetic algorithm (MOGA) to evolve dynamic Bayesian networks. There the multi-objectives explored were the ability of the evolved networks to explain the data, compared with the complexity of the net-




Figure 9. The 3D Pareto front for $a \rightarrow b$ and $x \rightarrow \lambda$ violations found using the “Island Model” is shown in blue. Beneath it is a linear mesh which fits the Pareto front with a Pearson’s ρ^2 of 0.9902. By Comparison the non-optimal Pareto front found by combining the best individuals from 8 runs using the basic algorithm (i.e. no pre-evolution) on the full graph is shown in red. The front shown in blue took a fraction of the computing time to find compared to the non-optimal red runs.

work in question. The MOGA was used instead of, for instance, a minimum description length (MDL) constraint. In all of these cases the network is being used to model something of interest and then, given some observed values, infer the likely causes. The genetic algorithms are used to construct different models which are then trained, typically the success (or otherwise) of a particular model being its performance on withheld data.

Our work differs because of the way we wish to utilize the Bayesian causal networks, specifically we specify the networks we are interested in, namely those which model a physical view of “reality” with specified local causality violation. Training such networks to replicate observed correlations is of limited interest because successful training results in one specific probability distribution that explains the data. What we are interested in finding are all the relevant probability distributions where the ability to match the observed correlations is contrasted with the strength of the local causality violations. The evolutionary algorithm is used, not to evolve networks, but rather to find these probability distributions given the network. The MOEA is used to guide evolution along these Pareto fronts.

In order to explore the Pareto front some type of MOEA algorithm is required. MOEA on two or three dimensions are relatively well understood. Algorithms to explore large dimensions are still an active area of research (see for example [46]). Since our initial experiments (reported here) would only require causal networks with no penalized edges (a single-value optimization), one penalized edge (a MOEA with a two-dimensional Pareto front) or two penalized edges (a MOEA with a three-dimensional Pareto front), we decided to use the well

understood NSGA-II [39]. Although NSGA-II attempts to return the whole of the Pareto front in a single run it was quite clear that the search space (being the required probability assignments for all the nodes in the network) was not smooth, even though the Pareto fronts may be (and, in fact, turned out to be) smooth. Given this an ϵ -Dominance Archive [47] was maintained and updated through multiple runs. In order to maintain diversity between runs the archive was not used to guide the evolution, but rather served as an updated archive of the best Pareto front found so far. After completing multiple runs, the individuals in the archive thus represent the Pareto front for the entire procedure, rather than for each run taken in isolation.

In Section VIE we describe how the ϵ -Dominance Archive generated from a lower dimension front can be used to seed evolution when a higher dimensional Pareto front is explored, in a manner not dissimilar to the algorithm presented by Liu *et al.* [41] or utilized by Stanley and Miikkulainen [40].

G. Implementation methodology and details

Our initial runs with the EA (i.e. the MOEA with a single objective, being to minimize the TVD) were used to verify that the EA could match known results. In this case we start with a causal network that reflects Bell’s non-locality assumptions (as shown in Figure 2a) and for various values of γ in (14) use the EA to try and match the experimentally observed joint probability distribution. This is a single-objective EA, with fitness being governed solely by the TVD, i.e. by how closely the observed probability distribution of the Model (being the observed joint probability distribution for an individual over the local causal graphs) matches the experimental data. It is known that when γ is less than $\sqrt{2} - 1$ the observed data can be modeled with a local causal network. The TVD values should increase as γ increases to 1 since the empirical distribution no longer factorizes into a locally causal distribution.

An initial population of 300 ($\mu = 300, \lambda = 300$) was chosen, with the probability of crossover, being 0.1 (and mutation 0.9). The mutation operator used a standard deviation of 0.1 (see Section VID). As is typical for experiments using genetic algorithms no systematic attempt was made to find the “best” parameters for the algorithm. Rather during the course of some initial testing runs, runs with variations of parameters were tried and the parameters of the ones that seemed to find solutions quickest were used. Population sizes reflected those minimum populations required to avoid runs being trapped early on in local minima. The parameters reported are not reported in a claim of optimality, but rather are reported for the purposes of reproducibility. In any case, once a solution is obtained, its validity does not depend on the means by which it was found.

Figure 4a shows experimentally measured density ma-

trices for a range of state γ values. The reduction in coherence is observed as decreasing off-diagonal terms. Figure 4b shows the minimum TVD values emerging from 20 runs of the graph for various γ values together with a linear line fitting the data, running from the known y-intercept of 0 TVD for $\gamma = (\sqrt{2} - 1)$. As can be seen the EA fits the expected linear results (Pearson's $\rho^2 = 0.9952$).

Having ensured that the algorithm could correctly match the known results on a causal network consistent with Bell's non-locality assumptions, the next stage is to require a relaxation of local causality to allow the EA (now operating as a MOEA) to match the correlations present in entangled states. As an initial step, we examined relaxing one casual edge at a time, beginning with a causal influence from a to b —that is, Alice's outcome is allowed to influence Bob's. This is illustrated in Figure 2c. This becomes a multi-objective problem with a two-dimensional Pareto front, ostensibly well within the capabilities of NSGA-II. The tensor contractions required are not overly complex but with increasing numbers of hidden variables (as a result of additional causal links) each run takes a non-trivial amount of time. As is typical where the search landscape (being the underlying conditional probability distributions) is not smooth (even though the fitnesses such distributions reduce to are smooth) a number of runs failed to converge to any part of the Pareto front, with most runs finding part, but not all of the Pareto front. In Figure 10a we show the individual results of 40 such runs. As can be seen from the figure just under half of the 40 runs had a large percentage of their front non-optimal, with approximately half the runs being plotted on top of each other on the Pareto Front. To generate the final Pareto front each of the individuals in the ϵ -Dominance archive from each of the 40 runs are submitted to the *global* ϵ -Dominance archive, so that the best estimate of the true Pareto front can emerge, as shown in Figure 10b.

Two points arise from these results. The first is that the front appears to be linear, that is increasing locality violations allows observed (quantum) correlations to be more exactly matched, the tradeoff being *linear* in nature. As far as we are aware this was not previously known. The second arises from the number of failed or only partially successful runs. In particular we note that while in the majority of runs the MOEA was able to find many points on (or close to) the Pareto front, other runs could be trapped and all runs had difficulty at either extreme of the front. It is clear that the search landscape in general is not smooth - the interplay between the conditioning on the hidden local variable and the other probability distributions allow the MOEA to become trapped in some local minima. The larger front (in this case a two-dimensional line) allowed the population to “slide” away from the edge cases. In addition as observed in [9] it is likely the edge cases represent very specific distributions. Whilst some of these observed difficulties could, in part, be ameliorated by using a larger population and

relying on the NSGA-II crowding mechanism to prevent such slippage, as is known this will not be feasible if the front consists of three (or more) dimensions. The front grows exponentially with the number of dimensions, requiring an exponential increase in population size. An alternative MOEA such as NSGA-III may help but each alternative comes with their own difficulties and assumptions. It is, however, possible to use the specifics of the problem space to address these concerns. We know we can evolve the population on a more limited graph (such as the purely local graph) and force the population to find the lowest TVD with a causal violation of zero (i.e. in the local graph $C_{a \rightarrow b} = 0$, since there is no link $a \rightarrow b$). This evolved population can then be “transplanted” on to a graph that does have an $a \rightarrow b$ link (e.g. Figure 2c). The other extreme (i.e. lowest $C_{a \rightarrow b}$ violation for TVD= 0) can be found with a small alteration to the fitness function. To find this point we evolve the population on the graph representing Figure 2c but with a single-objective fitness function, implemented as minimizing the TVD, but where two individuals have the same TVD, the one with the lowest causality violation is preferred. This drives the population towards zero TVD and then minimizes the causality. Even with this the observed correlations were unable to achieve an exact TVD = 0, it is speculated this is a result of experimental noise. The ability to generate populations (distributions) that sat at the extreme points of the Pareto front allowed the entire two-dimensional front to be revealed and, as discussed below, can similarly be utilized to reveal the three-dimensional front created by two simultaneously relaxed local causality constraints.

VII. DISCUSSION

In this work, we have developed a method to allow the study of several relaxations of local hidden variables models simultaneously in a single framework using the tools of casual networks and genetic algorithms.

With further refinement, we hope that our approach can shed light on other scenarios where quantum correlations display richer structure than classical systems would allow. For example, generalizations of the standard Bell scenario to more stations [48, 49] and more outcomes [50, 51], as well as multiple hidden variables [3, 52, 53]. In the latter scenario, very little is known since classical correlations are no longer given by linear constraints. Very recently, Chaves has used the framework of causal networks to systematically study such higher-order constraints [54]. Such measured quantities will be particularly useful to our approach as they can be seen as highly relevant coarse grainings of the exponentially growing data space. Such dimension reduction techniques will be crucial for scaling up our numerical algorithm to the analysis of multi-party quantum correlations.

In addition, there is nothing specifically “quantum”




Figure 10. The results of 40 typical runs of the EA with a causal graph allowing “Alice to Bob” local causality violations (see Figure 2c). (a) For the purpose of producing this graph, each run had its own ϵ -Dominance archive. As can be seen several runs failed to find the correct front at all, indicating that the interplay between hidden nodes and conditioned variables results in a non-trivial search when attempting to match the observed distributions of observable variables. (b) At the conclusion of these 40 runs the ϵ -Dominance archives are combined to form the *global* ϵ -Dominance archive generating the best estimate of the actual Pareto. The front is well fit by a straight line (Pearson’s ρ^2 value of 0.997, with bisquare robust fitting). Exact fitting of the distribution (very low TVD) requires additional causality violation. Experimental noise might be reason for this. Although around 100 runs were conducted to produce the reported results (Figure 5) very few additional points were found on the Pareto Front.

about our core numerical methods. Thus, our approach should find application outside of the problem of understanding quantum correlations. Recently Lee and Spekkens have also used inspiration from the causal analysis of quantum correlations to develop new causal discovery protocols [55]. Like Lee and Spekkens, we depart from the usual considerations of observed correlation to considering the entire joint probability distribution. Our goals differ, however; whereas the aim of Lee and Spekkens is to find all casual models consistent with data, our goal is to find non-dominated models of the plausible correlations. These two approaches are likely to find a harmonious union in the future.

ACKNOWLEDGMENTS

We thank Martin Ringbauer and Andrew White for discussions. This work was supported by the US Army

Research Office grant numbers W911NF-14-1-0098 and W911NF-14-1-0103, and by the Australian Research Council Centre of Excellence for Engineered Quantum Systems CE110001013. STF acknowledges support from an Australian Research Council Future Fellowship FT130101744. AP acknowledges an Australian Research Council Discovery Early Career Researcher Award, Project No. DE140101700 and an RMIT University Vice-Chancellors Senior Research Fellowship. RK acknowledges additional support from the Excellence Initiative of the German Federal and State Governments (Grants ZUK 43 & 81) and the DFG.

-
- [1] M. S. Leifer, “Conditional density operators and the subjectivity of quantum operations,” in *AIP Conference Proceedings*, Vol. 889 (AIP Publishing, 2007) pp. 172–186.
 - [2] M. S. Leifer and D. Poulin, “Quantum graphical models and belief propagation,” *Annals of Physics* **323**, 1899 (2008).
 - [3] T. Fritz, “Beyond Bell’s theorem: correlation scenarios,” *New Journal of Physics* **14**, 103001 (2012).
 - [4] J. Fitzsimons, J. Jones, and V. Vedral, “Quantum correlations which imply causation,” [arXiv:1302.2731](#) (2013).
 - [5] M. S. Leifer and R. W. Spekkens, “Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference,” *Physical Review A* **88**, 052130 (2013).
 - [6] M. S. Leifer and R. W. Spekkens, “A Bayesian approach to compatibility, improvement, and pooling of quantum states,” *Journal of Physics A: Mathematical and Theoretical* **47**, 275301 (2014).
 - [7] Č. Brukner, “Quantum causality,” *Nature Physics* **10**, 259 (2014).
 - [8] E. G. Cavalcanti and R. Lal, “On modifications of Reichenbach’s principle of common cause in light of Bell’s theorem,” *Journal of Physics A: Mathematical and Theoretical* **47**, 424018 (2014).
 - [9] 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 Journal of Physics* **17**, 033002 (2015).
- [10] K. Ried, M. Agnew, L. Vermeyden, D. Janzing, R. W. Spekkens, and K. J. Resch, “A quantum advantage for inferring causal structure,” *Nature Physics* **11**, 414 (2015).
 - [11] A. Einstein, B. Podolsky, and N. Rosen, “Can quantum-mechanical description of physical reality be considered complete?” *Physical Review* **47**, 777 (1935).
 - [12] J. S. Bell, “On the Einstein Podolsky Rosen paradox,” *Physics* **1**, 195 (1964).
 - [13] S. Kochen and E. Specker, “The problem of hidden variables in quantum mechanics,” *Indiana Univ. Math. J.* **17**, 59 (1968).
 - [14] B. F. Toner and D. Bacon, “Communication cost of simulating Bell correlations,” *Physical Review Letters* **91**, 187904 (2003).
 - [15] J. Barrett and N. Gisin, “How much measurement independence is needed to demonstrate nonlocality?” *Physical Review Letters* **106**, 100406 (2011).
 - [16] M. J. W. Hall, “Complementary contributions of indeterminism and signaling to quantum correlations,” *Physical Review A* **82**, 062117 (2010).
 - [17] M. J. W. Hall, “Local deterministic model of singlet state correlations based on relaxing measurement independence,” *Physical Review Letters* **105**, 250404 (2010).
 - [18] M. J. W. Hall, “Relaxed Bell inequalities and Kochen-Specker theorems,” *Physical Review A* **84**, 022102 (2011).
 - [19] D. E. Koh, M. J. W. Hall, Setiawan, J. E. Pope, C. Marletto, A. Kay, V. Scarani, and A. Ekert, “Effects of reduced measurement independence on Bell-based randomness expansion,” *Physical Review Letters* **109**, 160404 (2012).
 - [20] M. Banik, “Lack of measurement independence can simulate quantum correlations even when signaling can not,” *Physical Review A* **88**, 032118 (2013).
 - [21] L. P. Thinh, L. Sheridan, and V. Scarani, “Bell tests with min-entropy sources,” *Physical Review A* **87**, 062121 (2013).
 - [22] G. Pütz, D. Rosset, T. J. Barnea, Y.-C. Liang, and N. Gisin, “Arbitrarily small amount of measurement independence is sufficient to manifest quantum nonlocality,” *Physical Review Letters* **113**, 190402 (2014).
 - [23] K. Maxwell and E. Chitambar, “Bell inequalities with communication assistance,” *Physical Review A* **89**, 042108 (2014).
 - [24] J. Pearl, *Causality: Models, Reasoning, and Inference*, 2nd ed. (Cambridge University Press, 2009).
 - [25] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction and Search* (MIT Press, 2001).
 - [26] B. Hensen, H. Bernien, A. E. Dreau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenbergh, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, “Loophole-free bell inequality violation using electron spins separated by 1.3 kilometres,” *Nature* **526**, 682 (2015), letter.
 - [27] M. Giustina, M. A. M. Versteegh, S. Wengerowsky, J. Handsteiner, A. Hochrainer, K. Phelan, F. Steinlechner, J. Kofler, J.-A. Larsson, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, J. Beyer, T. Gerrits, A. E. Lita, L. K. Shalm, S. W. Nam, T. Scheidl, R. Ursin, B. Wittmann, and A. Zeilinger, “Significant-loophole-free test of bell’s theorem with entangled photons,” *Phys. Rev. Lett.* **115**, 250401 (2015).
 - [28] L. K. Shalm, E. Meyer-Scott, B. G. Christensen, P. Bierhorst, M. A. Wayne, M. J. Stevens, T. Gerrits, S. Glancy, D. R. Hamel, M. S. Allman, K. J. Coakley, S. D. Dyer, C. Hodge, A. E. Lita, V. B. Verma, C. Lambrocco, E. Tortorici, A. L. Migdall, Y. Zhang, D. R. Kumor, W. H. Farr, F. Marsili, M. D. Shaw, J. A. Stern, C. Abellán, W. Amaya, V. Pruneri, T. Jennewein, M. W. Mitchell, P. G. Kwiat, J. C. Bienfang, R. P. Mirin, E. Knill, and S. W. Nam, “Strong loophole-free test of local realism*,” *Phys. Rev. Lett.* **115**, 250402 (2015).
 - [29] R. Chaves, R. Kueng, J. Brask, and D. Gross, “Unifying framework for relaxations of the causal assumptions in Bell’s theorem,” *Physical Review Letters* **114**, 140403 (2015).
 - [30] P. M. Pearle, “Hidden-variable example based upon data rejection,” *Physical Review D* **2**, 1418 (1970).
 - [31] S. Wolf, “Nonlocality without counterfactual reasoning,” *Physical Review A* **92**, 052102 (2015).
 - [32] M. Ringbauer, C. Giarmatzi, R. Chaves, F. Costa, A. G. White, and A. Fedrizzi, “Experimental test of nonlocal causality,” *Science Advances* **2**, e1600162 (2016).
 - [33] R. J. Chapman, M. Santandrea, Z. Huang, G. Corrielli, A. Crespi, M.-H. Yung, R. Osellame, and A. Peruzzo, “Experimental perfect state transfer of an entangled photonic qubit,” *Nature Communications* **7**, 11339 (2016).
 - [34] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White, “Measurement of qubits,” *Phys. Rev. A* **64**, 052312 (2001).
 - [35] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, “Proposed experiment to test local hidden-variable theories,” *Phys. Rev. Lett.* **23**, 880 (1969).
 - [36] J. H. Holland, *Adaptation in natural and artificial systems* (University of Michigan Press, Ann Arbor, 1975).
 - [37] F.-A. Fortin, F.-M. De Rainville, M.-A. Gardner, M. Parizeau, and C. Gagné, “DEAP: Evolutionary algorithms made easy,” *Journal of Machine Learning Research* **13**, 2171 (2012).
 - [38] J. L. Bentley, “Multidimensional binary search trees used for associative searching,” *Commun. ACM* **18**, 509 (1975).
 - [39] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, “A fast and elitist multiobjective genetic algorithm: NSGA-II,” *Evolutionary Computation, IEEE Transactions on* **6**, 182 (2002).
 - [40] K. O. Stanley and R. Miikkulainen, “Evolving neural networks through augmenting topologies,” *Evol. Comput.* **10**, 99 (2002).
 - [41] H. L. Liu, F. Gu, and Q. Zhang, “Decomposition of a multiobjective optimization problem into a number of simple multiobjective subproblems,” *IEEE Transactions on Evolutionary Computation* **18**, 450 (2014).
 - [42] A. G. a. C. de Sá and G. L. Pappa, “Towards a method for automatically evolving Bayesian network classifiers,” in *Proceedings of the 15th Annual Conference Companion on Genetic and Evolutionary Computation, GECCO ’13 Companion* (ACM, New York, NY, USA, 2013) pp. 1505–1512.
 - [43] J. Muruzábal and C. Cotta, “Advances in probabilistic graphical models,” (Springer Berlin Heidelberg, Berlin, Heidelberg, 2007) Chap. A Study on the Evolution of Bayesian Network Graph Structures, pp. 193–213.
 - [44] P. Larraaga, H. Karshenas, C. Bielza, and R. Santana,

- “A review on evolutionary algorithms in bayesian network learning and inference tasks,” *Information Sciences* **233**, 109 (2013).
- [45] B. J. Ross and E. Zuviria, “Evolving dynamic bayesian networks with multi-objective genetic algorithms,” *Applied Intelligence* **26**, 13 (2006).
- [46] Y. Yuan, H. Xu, and B. Wang, “An improved nsga-iii procedure for evolutionary many-objective optimization,” in *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*, GECCO '14 (ACM, New York, NY, USA, 2014) pp. 661–668.
- [47] M. Laumanns, L. Thiele, K. Deb, and E. Zitzler, “Combining convergence and diversity in evolutionary multi-objective optimization,” *Evolutionary computation* **10**, 263 (2002).
- [48] R. F. Werner and M. M. Wolf, “All-multipartite Bell-correlation inequalities for two dichotomic observables per site,” *Physical Review A* **64**, 032112 (2001).
- [49] M. Żukowski, Č. Brukner, W. Laskowski, and M. Wieśniak, “Do all pure entangled states violate Bell’s inequalities for correlation functions?” *Physical Review Letters* **88**, 210402 (2002).
- [50] D. Collins, N. Gisin, N. Linden, S. Massar, and S. Popescu, “Bell inequalities for arbitrarily high-dimensional systems,” *Physical Review Letters* **88**, 040404 (2002).
- [51] D. Collins and N. Gisin, “A relevant two qubit Bell inequality inequivalent to the CHSH inequality,” *Journal of Physics A: Mathematical and General* **37**, 1775 (2004).
- [52] C. Branciard, N. Gisin, and S. Pironio, “Characterizing the nonlocal correlations created via entanglement swapping,” *Physical Review Letters* **104**, 170401 (2010).
- [53] J. Henson, R. Lal, and M. F. Pusey, “Theory-independent limits on correlations from generalized Bayesian networks,” *New Journal of Physics* **16**, 113043 (2014).
- [54] R. Chaves, “Polynomial Bell inequalities,” *arXiv:1506.04325 [quant-ph]* (2015).
- [55] C. M. Lee and R. W. Spekkens, “Causal inference via algebraic geometry: necessary and sufficient conditions for the feasibility of discrete causal models,” *arXiv:1506.03880 [quant-ph, stat]* (2015).

Appendix A: Derivation of the inequality presented in Section III

For the sake of being self-contained, let us start this section with reviewing some basic facts about discrete probability distributions and introduce some notation. Throughout this section, we focus on the empirical frequencies $F(a, b, x, y)$ and the probability distribution $\Pr(a, b, x, y|M)$ associated to a fixed model M . Here a, b are the special nodes whose causal relationship is of interest, y will denote the parents that are not grandparents of b , and x is any set of additional random variables which might include hidden variables λ as well as additional measurement outcomes. Therefore, the discussion is completely general and not specific to the models considered, e.g., in the experiment.

If we marginalize these distributions over any variable, say y , we produce new distributions

$$F(a, b, x) = \sum_y F(a, b, x, y) \quad \text{and} \quad \Pr(a, b, x|M) = \sum_y \Pr(a, b, x, y|M), \quad (\text{A1})$$

respectively. As outlined in (A1), we indicate marginalization over any variable, by simply omitting the corresponding variable in the description. Having such a notation at hand, the *product rule* (for discrete probability distribution) assures that as an immediate consequence of the definition of conditional distributions,

$$F(a, b, x, y) = F(a, b, x|y)F(y) \quad \text{and} \quad \Pr(a, b, x, y|M) = \Pr(a, b, x|y, M) \Pr(y|M) \quad (\text{A2})$$

for the variable y . Analogous formulas are true for any combination of the variables present in the distributions (i.e. $\{a, b, x, y\}$ for $F(\cdot)$ and $\{a, b, x, y\}$ for $\Pr(\cdot|M)$).

With these rules and notational concepts at hand, the following statement is an immediate consequence of the triangle inequality.

Lemma 1. *Let $F(a, b, x, y)$ and $\Pr(a, b, x, y|M)$ be as above. Then*

$$\|\Pr(b|M) - F(b)\|_1 \leq \|\Pr(a, b|M) - F(a, b)\|_1 \leq \|\Pr(a, b, x, y|M) - F(a, b, x, y)\|_1 = \text{TVD}(M). \quad (\text{A3})$$

This Lemma encapsulates two particular instance of the well-known fact that marginalization contracts the total variational distance. Since the latter is a measure of how well two probability distributions can be distinguished and marginalization corresponds to ignoring certain variables, Lemma 1 can be intuitively paraphrased as: “knowing more doesn’t hurt”.

Proof of Lemma 1. Inserting the definitions of marginalization and total variational distance yields

$$\|\Pr(b|M) - F(b)\|_1 = \sum_b \left| \sum_a (\Pr(a, b|M) - F(a, b)) \right| \leq \sum_{a,b} |\Pr(a, b|M) - F(a, b)| = \|\Pr(a, b|M) - F(a, b)\|_1 \quad (\text{A4})$$

upon employing the triangle inequality. The second inequality can be established in complete analogy. \square

We are now ready to establish the main auxiliary result necessary to establish [Theorem 1](#). It requires the concept of the *harmonic mean* for two variables. For $x_1, x_2 > 0$ the harmonic mean is defined as $H(x_1, x_2) = \frac{2x_1x_2}{x_1+x_2}$.

Lemma 2. *Consider two bivariate probability distributions $p(u, v)$ and $q(u, v)$ over finitely many elements labeled by u and v , respectively. Then, the following inequality is valid for any fixed variable v :*

$$\|p(u|v) - q(u|v)\|_1 \leq \frac{\sum_u |p(u, v) - q(u, v)| + |p(v) - q(v)|}{H(p(v), q(v))} \quad (\text{A5})$$

We point out that this estimate is responsible for introducing the on first sight unfavorable scaling of the bounds [\(13\)](#). However, inequality [\(A5\)](#) is actually tight, making the aforementioned behavior essentially unavoidable. To see this, let $u, v \in \{0, 1\}$ be binary variables and let p be the uniform probability distribution over the four possible joint instances. If one chooses q to be a perfectly correlated bivariate distribution—i.e. $q(0, 0) = q(1, 1) = 1/2$ —it is easy to see that equality is attained in the assertion of [Lemma 2](#).

Proof of Lemma 2. Fix an arbitrary label v . Inverting the product rule allows us to rewrite the left hand side of [\(A5\)](#) as

$$\|p(u|v) - q(u|v)\|_1 = \left\| \frac{p(u, v)}{p(v)} - \frac{q(u, v)}{q(v)} \right\|_1 = \frac{1}{p(v)q(v)} \sum_u |q(v)p(u, v) - p(v)q(u, v)|. \quad (\text{A6})$$

For $p(v)$ and $q(v)$ we now define

$$\mu := \frac{1}{2} (p(v) + q(v)) \quad \text{and} \quad \delta := \frac{1}{2} (p(v) - q(v))$$

which obey $p(v) = \mu + \delta$ as well as $q(v) = \mu - \delta$ by construction. Inserting these decompositions into [\(A6\)](#) reveals

$$\begin{aligned} p(v)q(v)\|p(u|v) - q(u|v)\|_1 &= \sum_u |(\mu - \delta)p(u, v) - (\mu + \delta)q(u, v)| \\ &= \sum_u |\mu(p(u, v) - q(u, v)) - \delta(p(u, v) + q(u, v))| \\ &\leq \mu \sum_u |p(u, v) - q(u, v)| + |\delta| \sum_u (p(u, v) + q(u, v)) \\ &= \mu \sum_u |p(u, v) - q(u, v)| + |\delta|(p(v) + q(v)) \\ &= \mu \left(\sum_u |p(u, v) - q(u, v)| + 2|\delta| \right), \end{aligned}$$

where we have employed the triangle inequality and the definition of marginalization. Replacing μ and δ with the original expressions then yields

$$\|p(u|v) - q(u|v)\|_1 \leq \frac{p(v) + q(v)}{2p(v)q(v)} \left(\sum_u |p(u, v) - q(u, v)| + |p(v) - q(v)| \right).$$

The desired statement then follows from this estimate by identifying the pre-factor as $1/H(p(v), q(v))$. \square

We can now show that a bound holds that relates the maximum deviation between the causal influence of any fixed model M and the frequencies F .

Lemma 3. *For any fixed model $M \in \mathcal{M}$ and fixed set of empirical frequencies F , let y denote the parents that are not grandparents of the random variable b . Then the following inequality holds*

$$|C_{a \rightarrow b}(F) - C_{a \rightarrow b}(M)| \leq \frac{4 \text{TVD}(M)}{\min_{a, y}^* H(\Pr(a, y|M), F(a, y))}, \quad (\text{A7})$$

where \min^* denotes the minimization over feasible assignments to the variables a, y .

Proof. Choose an arbitrary model $M \in \mathcal{M}$. To ease notation, denote by y the parents that are not grandparents of the variable b . In order to derive the upper bound presented in (13), we start with inserting the definition (9) of $C_{a \rightarrow b}(M)$ and observe

$$\begin{aligned}
C_{a \rightarrow b}(M) &= \max_{a, a', y}^* \|\Pr(b|a, y, M) - \Pr(b|a', y, M)\|_1 \\
&= \max_{a, a', y}^* \|\Pr(b|a, y, M) - F(b|a, y) - \Pr(b|a', y, M) + F(b|a', y) + F(b|a, y) - F(b|a', y)\|_1 \\
&\leq \max_{a, a', y}^* \|\Pr(b|a, y, M) - F(b|a, y)\|_1 + \max_{a, a', y}^* \|\Pr(b|a', y, M) - F(b|a', y)\|_1 + \max_{a, a', y}^* \|F(b|a, y) - F(b|a', y)\|_1 \\
&= 2 \max_{a, y}^* \|\Pr(b|a, y, M) - F(b|a, y)\|_1 + C_{a \rightarrow b}(F),
\end{aligned} \tag{A8}$$

where we have identified the last term as the empirical average causal effect defined in (12). As a simple bookkeeping device, let us define $v = (a, y)$ to be the cartesian product of the random variables a and y . The first term in (A8) can be bounded by invoking Lemma 2 and Lemma 1. Doing so results in

$$C_{a \rightarrow b}(M) - C_{a \rightarrow b}(F) \leq 2 \max_v^* \|\Pr(b|v, M) - F(b|v)\|_1 \tag{A9}$$

$$\leq 2 \max_v^* \frac{\sum_b |\Pr(v, b|M) - F(v, b)| + |\Pr(v|M) - F(v)|}{H(\Pr(v|M), F(v))} \tag{A10}$$

$$\leq 4 \max_v^* \frac{\|\Pr(v, b|M) - F(v, b)\|_1}{H(\Pr(v|M), F(v))} \tag{A11}$$

$$\leq 4 \max_v^* \frac{\text{TVD}(M)}{H(\Pr(v|M), F(v))} \tag{A12}$$

which is equivalent to the upper bound presented in (A7). The corresponding lower bound can be derived in a completely analogous fashion by starting off with $C_{a \rightarrow b}(F)$ instead of $C_{a \rightarrow b}(M)$. \square

This bound is not yet useful because the right hand side still depends on the unknown model. We seek an inequality that is independent of the model as long as the model has a fixed and sufficiently small value of $\text{TVD}(M)$ with respect to the empirical frequencies. The bound in Section III is a way to avoid this difficulty, and we have now assembled all prerequisites necessary to prove it. We restate the main theorem for completeness.

Theorem 1. *Let M denote any model and let $\text{TVD}(M)$, $C_{a \rightarrow b}(M)$ and $C_{a \rightarrow b}(F)$ be as in (7), (9) and (12), respectively. Denote by \mathcal{M}_τ the set of models having $\text{TVD}(M) \leq \tau$ with respect to the empirical frequencies F , and let $f^* = \min_{a, y} F(a, y)$, where y denotes all parents of the variable b that are not grandparents of b . Then for all $M \in \mathcal{M}_\tau$ and $\tau < 2f^*$ we have*

$$|C_{a \rightarrow b}(F) - C_{a \rightarrow b}(M)| \leq \frac{2\tau(4f^* - \tau)}{f^*(2f^* - \tau)}. \tag{A13}$$

Proof of Theorem 1. Again for the sake of bookkeeping we introduce a variable $v = (a, y)$. We begin with the inequality from Lemma 3 and note that we can simply maximize the righthand side over all $M \in \mathcal{M}_\tau$ to get a universal bound. We have

$$\max_{M \in \mathcal{M}_\tau} \frac{4 \text{TVD}(M)}{\min_v^* H(\Pr(v|M), F(v))} \leq \frac{4 \tau}{\min_{M \in \mathcal{M}_\tau} \min_v^* H(\Pr(v|M), F(v))}. \tag{A14}$$

Therefore we must establish a lower bound on the denominator. Plugging in the definition of the harmonic mean, a simple calculation confirms that $\partial_x H(x, y) = \frac{2y^2}{(x+y)^2} \geq 0$, so the denominator is bounded from below as

$$\min_{M \in \mathcal{M}_\tau} \min_v^* \frac{2 \Pr(v|M) F(v)}{\Pr(v|M) + F(v)} \geq \min_{M \in \mathcal{M}_\tau} \min_v^* \frac{2 \Pr(v|M) f^*}{\Pr(v|M) + f^*}. \tag{A15}$$

Now we relax slightly to allow *all* possible probability distributions (not necessarily ones coming from a causal model M , and denote \mathcal{P}_τ to be the set of all probability distributions p with $\|p - F\|_1 \leq \tau$. Minimizing over a potentially larger set \mathcal{P}_τ may only decrease the function (or keep its minimum unchanged). We find the denominator is now bounded by

$$\min_{M \in \mathcal{M}_\tau} \min_v^* \frac{2 \Pr(v|M) f^*}{\Pr(v|M) + f^*} \geq \min_{p \in \mathcal{P}_\tau} \min_v^* \frac{2p(v) f^*}{p(v) + f^*} \geq \min_{p \in \mathcal{P}_\tau} \frac{2p^* f^*}{p^* + f^*}. \tag{A16}$$

Here in the second inequality we have used the same monotonicity argument for the harmonic mean above (since it is a symmetric function) and replaced the minimum over v with $p^* = \min_v^* p(v)$.

Now we appeal to the monotonicity result of [Lemma 1](#), so that $p \in \mathcal{P}_\tau$ implies that $\|p(v) - F(v)\|_1 \leq \tau$. The claim then follows if we can establish the following result,

$$\min_{p \in \mathcal{P}_\tau} p^* \geq f^* - \frac{\tau}{2}. \quad (\text{A17})$$

A weaker result, that $\min_{p \in \mathcal{P}_\tau} p^* \geq f^* - \tau$, is easy to see if we relax the requirement that p is normalized and add the more stringent requirement that $\tau < f^*$. Begin with the choice $p(a) = F(a)$, and then subtract τ from the smallest component, keeping all other components fixed. This achieves the least value of this relaxed problem. This is a valid solution since the resulting vector is still nonnegative, owing to the constraint $\tau < f^*$. The slightly tighter result follows from reasserting the constraint that the entries of p must sum to 1, and allows us to weaken the constraint on τ to $\tau < 2f^*$. With the normalization condition in place, subtracting any deviation of size δ from a component of p must be compensated by adding δ elsewhere in the vector, and this contributes a total of 2δ to the TVD between these differing vectors. The largest such a deviation can be is half of τ , and to minimize our objective function we put this deviation on the smallest component. This component remains positive because of the condition $\tau < 2f^*$, so this remains a valid probability distribution.

Again by the monotonicity of the harmonic mean, this minimal value can be used to lower bound the denominator. The final inequality is obtained by plugging in the value of [\(A17\)](#) into the denominator expression and simplifying. \square

We remark that the maximum possible value for f^* in a Bell experiment where a takes d possible outcomes is $1/d$. Because this inequality is monotonically decreasing with f^* , the bound becomes weaker as the number of outcomes increases, and the requirement that $\tau < 2f^*$ becomes more demanding.