

Flow of dynamical causal structures with an application to correlations

Amin Baumeler¹* and Stefan Wolf²¹Facoltà di scienze informatiche, *Università della Svizzera italiana*, 6900 Lugano, Switzerland
²and Facoltà indipendente di Gandria, 6978 Gandria, Switzerland

(Received 23 October 2024; accepted 26 August 2025; published 24 September 2025)

Causal models capture cause-effect relations both qualitatively—via the graphical *causal structure*—and quantitatively—via the *model parameters*. They offer a powerful framework for analyzing and constructing processes. Here, we introduce a tool—the flow of causal structures—to visualize and explore the *dynamical aspect* of classical-deterministic processes, arguably like those present in general relativity. The flow describes all possible ways in which the causal structure of a process can evolve. We also present an algorithm to construct its supergraph—the *superflow*—from the *causal structure only*, without invoking the model parameters. As an application, we show that if all leaves of a flow are trivial, then the corresponding process produces causal correlations only, i.e., correlations where future data cannot influence past events. This strengthens the result that processes, where every directed cycle in their causal structure is chordless, establish causal correlations only. We also discuss the main difficulties for the quantum generalization of the present algorithms.

DOI: [10.1103/jmmm-tq72](https://doi.org/10.1103/jmmm-tq72)

I. INTRODUCTION

A central barrier for the reconciliation of quantum theory with general relativity, as evocatively pointed out by Hardy [1], consists in their diametrically opposed concepts: Quantum theory features *indefinite quantities* that are restricted to evolve along a *static causal structure*, while general relativity contains only *definite* quantities, yet the causal structure is *dynamic*. To overcome this barrier, recently much effort has been put into the study of frameworks [2–4] where the causal connections themselves are governed by quantum theory. The “process-matrix framework” [3], which builds upon the assumption that no deviation from quantum theory is detectable locally, is the most permissive one. It features the “quantum switch” [5], a prime example of a process with *indefinite causal order*. In the “quantum switch,” a qubit controls the order in which two quantum gates are applied to a target system. That causal order may thus be in an indefinite state. This process conceptually exemplifies the general-relativistic trait of dynamical causal order, as in principle attainable via the displacement of matter, and, simultaneously, the quantum-mechanical trait of superposing these possibilities.

Causal connections are conveniently expressed and made explicit using *causal models* [6], a prime tool used in statistics, e.g., for causal reasoning: How may some empirically observed statistics be explained? A causal model consists of two parts: the causal structure, which is a directed graph (*digraph*) $D = (\mathcal{V}(D), \mathcal{E}(D) \subseteq \mathcal{V}(D)^2)$ that encodes cause-effect relations among a set of variables $\mathcal{V}(D)$ via the edges $\mathcal{E}(D)$,

and the model parameters, which functionally express the respective dependencies [see Fig. 1(a)]. Only recently—also for the sake of aligning Reichenbach’s principle [7] with Bell correlations [8]—causal models were generalized to quantum theory [9–12]. The pioneering work by Costa and Shrapnel [10] is a first attempt to relate quantum causal models with processes arising from the “process-matrix framework.” The latest description of quantum causal models by Barrett, Lorenz, and Oreshkov [12], which is sufficiently expressive to represent all unitarily extendible processes [13], proved most useful toward that aim. Quantum causal models [12] are operational (they inherently allow for interventions) by virtue of following the “split-node” approach [14]: Each vertex v of the causal structure represents an agent with two spaces, an input space $\mathcal{I}_v$ and an output space $\mathcal{O}_v$, and the agent may perform an intervention by specifying the state in the output space. In this operational abstraction, dynamical causal order means that the causal relations are affected by the agent’s interventions. The description of Barrett, Lorenz, and Oreshkov [12] led to a series of insights. For instance, with our collaborator Tselentis [15], we show that the admissible quantum causal structures are “siblings-on-cycles” graphs; all other graphs lead to inconsistencies. These “siblings-on-cycles” graphs are such that every directed cycle therein contains no less than one pair of siblings [16].

A. Problem

While causal models were extended to incorporate quantum theory, the *dynamical* aspect of causal relations is only poorly captured. This dynamical aspect manifests itself only in the presence of directed cycles in the causal structure [see Fig. 1(b)]. These cycles, however, solely reflect the *potential* transmission of information. Actually, information traveling along a cycle *must* be interrupted before completing the cycle [15]. In short, the directed cycles are *misleading*; they actually represent only the union of possible causal connections and

*Contact author: amin.baumeler@usi.ch




FIG. 1. (a) Causal models: The model parameters with the directed graph constitute a causal model for the correlations $P_{X,Y,Z}$. (b) Causal structure of the quantum switch: The control qubit specified in the global past (**P**) influences the global future (**F**). Simultaneously, the target system may evolve through **A** and then through **B**, or vice versa. This potentiality of transmitting a signal from **A** to **B** or from **B** to **A** is reflected by the directed cycle in the causal structure. We use bold letters to indicate that the vertices are split nodes.

no “causal loop,” where an event influences itself [17]. This problem is *not* inherent to quantum theory, but is also present in *classical-deterministic* causal models. For example, the “classical” version of the quantum switch (where all qubits are replaced by bits) has the same causal structure [Fig. 1(b)].

In the present treatment, we focus on this classical-deterministic case. This restriction, aside from making the task simpler (the quantum case is at present intractable), is compensated by the following two benefits. First, it allows us to focus on the dynamical aspect without the risk of having quantum effects appear implicitly. Second, by studying this restricted case, we effectively study an “abstraction” of general relativity. Note that general relativity is a classical-deterministic theory. Also, the spacetime structure in general relativity depends on the distribution of matter. An alternation of that distribution thus induces a change in the causal relations. Although it is uncommon to describe such alternations in general relativity—after all, the state of affairs is fully determined by the boundary conditions, there seems to be no conceptual obstacle in doing so. Therefore, the “classical switch” and any classical-deterministic causal model are, in principle, compatible with general relativity. In this context, classical-deterministic causal models could be considered as a toy framework of general relativity.

B. Results

We develop a tool to capture the evolution of the causal structure in classical-deterministic causal models. This tool “unravels” the causal structure via the sequential implementation of interventions at source vertices, i.e., parentless vertices. The resulting *flow graph* [18] (or *flow* for short) encodes the evolution of the causal structure and effectively captures the dynamics induced by the agents’ interventions. This clarifies the nature of the directed cycles in causal models with dynamical causal order. We also show how to generate a *superflow*, an “outer approximation,” from the causal structure only, i.e., without invoking the model parameters. The superflow is thus a purely qualitative object (independent of the dimensions, etc.). It therefore enables statements of qualitative nature, e.g., the characterization of possible correlations (see below). This abstraction, however, comes with a price. For a




FIG. 2. Example of a flow for the model parameters given in Eq. (5) (without dashed part), and the superflow computed by Algorithm 2 (including dashed part).

given causal model, the superflow may describe additional, potentially unrealizable, evolutions of the causal structure: It may be a supergraph of the flow. As an example, Fig. 2 shows the flow and the superflow of some causal model. A specific choice of model parameters, which we present later [Eq. (5)], results in a treelike flow (without dashed part). There, independent of **P**’s intervention, the agents **A** and **B** are causally connected. Logically, however, model parameters for the same initial causal structure may exist, which produce a flow that comprises the dashed part: A specific intervention of **P**, which depends on the model parameters, may causally disconnect **A** and **B**. The cause-effect relations among the agents for any classical-deterministic causal model with that causal structure are restricted to follow one of the paths in the flow. The flow of this example has two leaf nodes. A *leafnode*, or *leaf* for short, is a node without outgoing edges. Here, both leaves hold trivial graphs, i.e., graphs that are composed by a single vertex. Consequently, we call such nodes *trivial leaves*. The above observation has stark consequences. We show that if all leaves of the flow are trivial, then the attainable correlations from the initial causal model decompose causally: Past data may influence future observations only. For more general causal structures, this is not necessarily the case [3, 15, 19–21], and the correlations may be incompatible with any definite causal order. Clearly, that property also holds for superflows. With this, we strengthen a result of Ref. [15]. There, we show that the strictly stronger condition of having chordless cycles only is sufficient. A *chord* in a directed cycle $(v_0 \rightharpoonup \dots \rightharpoonup v_{\ell-1} \rightharpoonup v_0)$ is an edge $v_i \rightharpoonup v_j$ with $j \neq i+1 \pmod{\ell}$. Finally, we discuss the main challenges in “unraveling” quantum causal structures.

II. CAUSAL MODELS

In the scenario with the agents $\mathcal{V}(D)$, the respective input and output spaces $\mathcal{I}_v$ and $\mathcal{O}_v$ of each agent $v \in \mathcal{V}(D)$ are finite sets. An intervention of agent v is a function $\mu_v : \mathcal{X}_v \times \mathcal{I}_v \rightarrow \mathcal{A}_v \times \mathcal{O}_v$, where $\mathcal{X}_v$ is the set of settings and $\mathcal{A}_v$ is the set of results. We use $\mathcal{M}_v$ for the set of all interventions of agent v . The notation we adopt is as follows: By X_S we denote the natural composition of the objects $\{X_s\}_{s \in S}$, and by $\underline{X}$ the special case for $S = \mathcal{V}(D)$. The expression $[\alpha = \beta]$ is 1 whenever the equality holds, and 0 otherwise. For a function $f : \mathcal{Y} \rightarrow \mathcal{Z}$, we write p^f for the induced conditional probability distribution defined through $p^f(z|y) := [f(y) = z]$.

Definition 1 (*Causal model, consistency, and faithfulness* [12,15]). An n -agent classical-deterministic causal model consists of a pair $(D, \mathcal{F})$, where D is a directed graph with $|\mathcal{V}(D)| = n$ (causal structure) and $\mathcal{F}$ is a family of functions $\{\omega_v : \mathcal{O}_{\text{Pa}_D(v)} \rightarrow \mathcal{I}_v\}_{v \in \mathcal{V}(D)}$ (model parameters). The model is called *faithful* whenever, for each $v \in \mathcal{V}(D)$, the function ω_v is signaling from each $\ell \in \text{Pa}_D(v)$ to v :

$$\forall v \in \mathcal{V}(D), \quad \ell \in \text{Pa}_D(v), \quad \exists \delta \in \mathcal{O}_{\text{Pa}_D(v) \setminus \{\ell\}}, \quad (q, r) \in \mathcal{O}_\ell^2, \\ \omega_v(q, \delta) \neq \omega_v(r, \delta). \quad (1)$$

The model is called *consistent* whenever the contraction of $\underline{\omega}$ with any choice of interventions $\underline{\mu}$ is a function from $\underline{\mathcal{X}}$ to $\underline{\mathcal{A}}$:

$$\forall \underline{\mu} \in \underline{\mathcal{M}}, \quad \exists g \in \{\underline{\mathcal{X}} \rightarrow \underline{\mathcal{A}}\}, \quad \forall (\underline{x}, \underline{a}) \in \underline{\mathcal{X}} \times \underline{\mathcal{A}}, \\ \sum_{(i, o) \in \underline{\mathcal{I}} \times \underline{\mathcal{O}}} [\omega(o) = i][\underline{\mu}(\underline{x}, i) = (\underline{a}, o)] = p^g(\underline{a}|\underline{x}). \quad (2)$$

If not otherwise specified, we use the shorthand term “causal model” to refer to a *classical-deterministic causal model*.

A *reduction* over a source agent s is the partial contraction of the causal model with the intervention μ_s , and where every other intervention is left unspecified. In the process of the reduction, the source s also fixes the setting x_s and records the result a_s . Here, we will only consider reductions over source agents.

We will make use of the following two central results. The first states that whenever we reduce over a source, then the reduced causal model remains consistent. This result is essentially an application of Lemma A3 of Ref. [22]. We give a tailored and simplified proof below.

Theorem 1 (*Reduction*). If $(D, \mathcal{F})$ is a consistent causal model and s a source, i.e., $\text{Pa}_D(s) = \emptyset$, then for any $o_s \in \mathcal{O}_s$, the reduced causal model $(D', \mathcal{F}')$ with

$$D' := (\mathcal{V}(D) \setminus \{s\}, \quad \mathcal{E}(D) \setminus \{v \rightarrow s, s \rightarrow v \mid v \in \mathcal{V}(D)\}), \quad (3)$$

$$\mathcal{F}' := \{\omega_v\}_{v \in \mathcal{V}(D) \setminus \text{Ch}_D(s)} \cup \{\omega_v(o_s, \dots)\}_{v \in \text{Ch}_D(s)}, \quad (4)$$

where $\text{Ch}_D(s) := \{v \mid s \rightarrow v \in \mathcal{E}(D)\}$ are the children of s in D and $\omega_v(o_s, \dots)$ is the partial application of the function ω_v , is consistent yet possibly unfaithful.

Proof. The graph D' is a causal structure of the reduced causal model because the reduction over s cannot increase the functional dependencies of the model parameters. The reduced causal model remains consistent for the following reason. Since s is parentless, the model parameter ω_s is a constant c . Moreover, by the specification of $\mu_s \in \mathcal{M}_s$ and $x_s \in \mathcal{X}_s$, the output in the space $\mathcal{O}_s$ is a constant d . The proof is concluded by plugging in these constants in the consistency condition [Eq. (2)]: For any choice of interventions for the remaining parties, the contraction with the reduced model parameters coincides with the function f' defined as the partially evaluated f where the result $\mathcal{A}_s$ is discarded. ■

The second result, as highlighted in the introduction, shows that the causal structure of faithful and consistent causal models is severely restrained.

Theorem 2 (*Admissibility* [15]). If $(D, \mathcal{F})$ is a faithful and consistent causal model, then D is a siblings-on-cycles graph.

ALGORITHM 1. Flow (model-parameter aware).

Algorithm 1 Flow (model-parameter aware)

Input: causal model $(D, \mathcal{F})$

Output: flow $\mathcal{F}$

```

1: Initiate  $\mathcal{G}$  with root  $(D, \mathcal{F})$ 
2:  $\mathcal{L} \leftarrow \text{NONTRIVIALLEAVESWITHSOURCE}(\mathcal{G})$ 
3: while  $\mathcal{L} \neq \emptyset$  do
4:   for  $(L, \mathcal{F}_L) \in \mathcal{L}$  do
5:     for  $s \in \text{SOURCEVERTICES}(L)$  do
6:       for  $o_s \in \mathcal{O}_s$  do
7:          $\mathcal{F}'_L \leftarrow \text{REDUCEMODPARAM}(\mathcal{F}_L, L, s, o_s)$ 
8:          $L' \leftarrow \text{DERIVECAUSALSTRUCTURE}(\mathcal{F}'_L)$ 
9:          $\mathcal{G} \leftarrow \mathcal{G} \cup \{(L, \mathcal{F}_L) \rightarrow (L', \mathcal{F}'_L)\}$ 
10:     $\mathcal{L} \leftarrow \text{NONTRIVIALLEAVESWITHSOURCE}(\mathcal{G})$ 
11:  $\mathcal{F} \leftarrow \text{REMOVEMODELPARAMETERS}(\mathcal{G})$ 
12: return  $\mathcal{F}$ 

```

III. FLOW GRAPH

For the causal model $(D, \mathcal{F})$, the *flow graph* $\mathcal{F}$ contains D as its root (see Fig. 2 and the Appendix for more complex examples) and shows all possible ways in which the causal structure may evolve. Suppose D contains a source vertex s . This means that the observations made by agent s cannot be influenced by any other agent; they are independent of the other agents' choice of intervention. Thus, in order to unravel the causal structure, we may safely implement the intervention of agent s . This intervention may have an effect on the causal connections among the remaining agents: Different interventions may result in different causal connections. The flow $\mathcal{F}$ we construct captures all possibilities that may arise. Thus, $\mathcal{F}$ branches to all distinct causal structures that may follow after an intervention at s . Now, the resulting causal structures may again contain source vertices. The flow $\mathcal{F}$ would then iteratively branch out to causal structures connecting fewer agents. If there are multiple *source* vertices, then the flow also branches to the causal structures after intervening at any other source. Note that branches within the flow might merge (the flow is a rooted directed graph and not necessarily a tree): The evolution of the causal structure along different choices of interventions may result in the same causal structure.

We give a construction of the flow by utilizing the model parameters of the causal model and a construction of the superflow by ignoring the model parameters. Note that a superflow may contain paths unrealizable with actual interventions of the agents: There might exist causal models with the same initial causal structure, but where different parts of the flow are unreachable.

A. Construction: flow of causal structure

The algorithm (see Algorithm 1) consists of two parts: We generate a rooted directed graph $\mathcal{G}$ with causal models as its vertices (lines 1–10), and then we remove all model parameters to arrive at the flow graph $\mathcal{F}$ (line 11). After initiating $\mathcal{G}$, we invoke the function `NONTRIVIALLEAVESWITHSOURCE`, which returns all leaves of $\mathcal{G}$ whose causal structure is non-trivial and contains at least one source. Then, we iterate over

all leaves, sources vertices therein, and interventions for that source (lines 3–6). Since s is a source, the intervention is but the specification of the value on its output space. We then compute the model parameters of the reduced process [see Eq. (4)] and derive the causal structure by determining the functional dependencies. On line 9, we update $\mathcal{G}$ with an edge from that leaf to the reduced one. We repeat the above as long as $\mathcal{G}$ contains nontrivial leaves with source vertices. After this process terminates, we generate the flow $\mathcal{F}$ by discarding the model parameters in all causal models within $\mathcal{G}$. The function `REMOVEDODELPARAMETERS` maps the node set $\{(D_i, \mathcal{F}_i)\}_i$ of $\mathcal{G}$ to the node set $\{D_i\}_i$, and each edge $(D_i, \mathcal{F}_i) \mapsto (D_j, \mathcal{F}_j)$ to $D_i \mapsto D_j$.

1. Correctness

The algorithm terminates after finitely many steps because the initial causal model has a finite number of agents, and the reduction (lines 7 and 8) removes one agent at a time. *Consistency* follows from Theorem 1, and *faithfulness* by construction (line 8). Finally, the iteration on line 6 covers all interventions of agent s . With this, the flow $\mathcal{F}$ describes all possible paths in which the causal structure may evolve.

2. Example

Take the causal structure shown in the root of Fig. 2 and amend it with the model parameters

$$\omega_A(o_P, o_B) := o_P o_B, \quad \omega_B(o_P, o_A) := (1 - o_P) o_A, \quad (5)$$

where all input and all output spaces, but $\mathcal{I}_P$, are $\{0, 1\}$, and where the input space to P is trivial. The resulting flow $\mathcal{F}$ is shown in Fig. 2.

B. Construction: superflow of causal structure

For this construction (see Algorithm 2), we make use of Theorem 1 in conjunction with Theorem 2: Whenever we remove a source from a consistent causal model, then the resulting causal structure must be a siblings-on-cycles graph. The algorithm is as follows: After initiating the graph $\mathcal{S}$ with the causal structure D , we iterate over all leaves that contain

ALGORITHM 2. Superflow (model-parameter agnostic).

Algorithm 2 Superflow (model-parameter agnostic)

Input: causal structure D

Output: superflow $\mathcal{S}$

```

1: Initiate  $\mathcal{S}$  with root  $D$ 
2:  $\mathcal{L} \leftarrow \text{NONTRIVIALLEAVESWITHSOURCE}(\mathcal{S})$ 
3: while  $\mathcal{L} \neq \emptyset$  do
4:   for  $L \in \mathcal{L}$  do
5:     for  $s \in \text{SOURCEVERTICES}(L)$  do
6:        $L' \leftarrow L \setminus \{s\}$ 
7:        $\mathcal{E} \leftarrow \{v \mapsto k \mid k \in \text{Ch}_L(s), v \in \text{Pa}_{L'}(k)\}$ 
8:       for  $\mathcal{R} \subseteq \mathcal{E}$  do
9:          $L'' \leftarrow L' \setminus \mathcal{R}$ 
10:        if  $\text{ISOC}(L'')$  then
11:           $\mathcal{S} \leftarrow \mathcal{S} \cup \{L \mapsto L''\}$ 
12:       $\mathcal{L} \leftarrow \text{NONTRIVIALLEAVESWITHSOURCE}(\mathcal{S})$ 
13: return  $\mathcal{S}$ 

```

source vertices, and over each source (lines 4 and 5). Then, we define L' by removing the source s from the graph L ; i.e., we remove s from the set of vertices and all edges starting from or ending in s [see also Eq. (3)]. After that, we collect all edges of L' that point to any child of s . Then, we iterate over all possible subsets and remove these edges from L' . If the resulting graph is a siblings-on-cycles graph, as decided by the function `ISOC`, we add it to the superflow $\mathcal{S}$.

1. Correctness

For any L on line 4, we will reach line 11 at least once. This together with the finiteness of D implies that Algorithm 2 terminates after finitely many steps. For a causal model $(D, \mathcal{F})$, let $\mathcal{F}$ be the flow computed using Algorithm 1, and $\mathcal{S}$ the superflow using Algorithm 2. Now, any path $\pi = D \mapsto D_2 \mapsto \dots \mapsto D_\ell$ in flow $\mathcal{F}$ is also present in $\mathcal{S}$. The reason for this is that any reduction over a source s may alter the model parameters of the agents $\text{Ch}_L(s)$ only [see line 7 of Algorithm 1 and Eq. (4)], and Algorithm 2 exhausts all possibilities.

C. Run-time

For simplicity, suppose that the maximum cardinality of the output spaces, $\max |O_v|$, is a constant (this is only relevant for the run-time analysis of Algorithm 1). Both algorithms presented have an exponential runtime in the number of vertices n of the causal structure in the root. In particular, the run-time is

$$O(2^{\text{poly}(n)}), \quad (6)$$

where $\text{poly}(n)$ is some polynomial in n . To see this, first note that a flow could coincide with the superflow. The number of potential nodes N in the (super)flow is exponential in n . So, in Algorithm 1 we repeat the lines 5–9 at most N times, and in Algorithm 2 we repeat the lines 8–11 at most N times. In both algorithms, each of these repetitions can be done in at most an exponential number of steps.

We give more details on showing that N is exponential in n . Toward that, let L be a node in the (super)flow, and assume that the digraph L has a single source vertex s . The number of outgoing edges from s is $O(2^{n^2})$. The reason for this is that for each parent vertex of s , we consider all combinations of possibly ingoing edges (lines 7 and 8 in Algorithm 2). For each parent of s , there are $O(2^n)$ such possibilities, and there are $O(n)$ parents of s . In general, L has $O(n)$ source vertices, so we get another exponent of n , and we remain in the form of Eq. (6).

IV. CORRELATIONS

The correlations among the agents $\mathcal{S}$ are expressed by the conditional probabilities $p(a_{\mathcal{S}}|x_{\mathcal{S}})$ of observing $a_{\mathcal{S}}$ under the settings $x_{\mathcal{S}}$. Such correlations are termed *causal* [3,23,24] if and only if they decompose as

$$p(a_{\mathcal{S}}|x_{\mathcal{S}}) = \sum_{v \in \mathcal{S}} p_{\mathcal{S}}(v) p_v(a_v|x_v) p_{a_v}^{x_v}(a_{\mathcal{S} \setminus \{v\}}|x_{\mathcal{S} \setminus \{v\}}), \quad (7)$$

where $p_{\mathcal{S}}$ is a probability distribution over $\mathcal{S}$, p_v a conditional probability distribution of agent v observing a_v under the setting x_v , and $p_{a_v}^{x_v}(a_{\mathcal{S} \setminus \{v\}}|x_{\mathcal{S} \setminus \{v\}})$ are causal correlations




FIG. 3. These causal structures give causal correlations only.

again. Causal correlations are compatible with a dynamical specification of the causal order, where each agent may influence the order of the agents within its future. Single-agent correlations are trivially causal. For a causal model $(D, \mathcal{F})$ and fixed interventions $\underline{\mu} \in \underline{\mathcal{M}}$, the obtained correlations are $p^g(a|x)$, where g is the respective function in Eq. (2). As an application of the developed tool, we present the following result.

Theorem 3 (Causal correlations). If all leaves of a flow $\mathcal{F}$ are trivial, then the correlations attainable from a causal model $(D, \mathcal{F})$ with flow $\mathcal{F}$ are causal.

Proof. For any choice of interventions of the agents, there exists an ordering in which to perform the contractions, such that in every iteration a *source* is contracted. The correlations $p^g(a|x)$ decompose in that order. ■

This theorem trivially also holds for *superflows*. Using this theorem, we can identify causal structures (digraphs) that, when extended to a causal model, always produce causal correlations. So, this theorem strengthens the result [15] that the correlations from a causal model with *chordless* cycles only in its causal structure must be causal. Examples of causal structures with chordal cycles, but that—according to Theorem 3—still produce causal correlations only, are given in Fig. 3. In Fig. 5, we give the superflow of the causal structure shown in Fig. 3(a). In the Appendix, we also list all connected four-vertex digraphs in the “gap” between the result of Ref. [15] and Theorem 3.

V. QUANTUM GENERALIZATION

In the quantum case [12], the input and output spaces of the agents are Hilbert spaces. Also, the model parameters are a family of completely positive trace-preserving maps as opposed to functions that, in the Choi representation [25], pairwise commute. This ensures that the quantum process $W := \prod_{v \in \mathcal{V}(D)} \omega_v$ is well defined. Finally, faithfulness [see Eq. (1)] is defined analogously, and consistency [see Eq. (2)] is

$$\forall \{\mu_v \in \text{CPTP}_v\}_{v \in \mathcal{V}(D)} : \text{Tr} \left[W \cdot \bigotimes_{v \in \mathcal{V}(D)} \mu_v \right] = 1, \quad (8)$$

where CPTP_v is the set of completely positive trace-preserving maps from $\mathcal{L}(\mathcal{I}_v)$ (the set of linear operators on $\mathcal{I}_v$) to $\mathcal{L}(\mathcal{O}_v)$ in the Choi representation.




FIG. 4. All connected four-node cyclic digraphs with at least one source and at least one cycle with a chord.

Take a quantum causal model $(D, \mathcal{F})$, where agent s is a source, and suppose the input Hilbert space $\mathcal{I}_s$ is trivial. If agent s performs as intervention the state preparation μ_s , then it is unclear how to define the *reduced* model parameters, and whether they may be defined in a meaningful way. The reduced process W' is given by $\text{Tr}_{\mathcal{O}_s}[W\mu_s]$. Generally, this partial trace does not extend to $\prod_{v \in \mathcal{V}(D') \setminus \text{Ch}_D(s)} \omega_v \cdot \prod_{v \in \text{Ch}_D(s)} \text{Tr}_{\mathcal{O}_s}[\omega_v \mu_s]$, for D' defined as in Eq. (3).

VI. CONCLUSION

We present a tool—the flow of causal structures—for studying the dynamical aspect of causal models. This tool “linearizes” cycles in causal structures, and hence allows for conventional step-by-step analysis of communication schemes, computation, and general information processing under *dynamical causal order*. In fact, the flow of the directed path graph (a directed line) is just the directed path graph again. We demonstrate the usability of this tool by invoking it in the context of correlations: We show that if all leaves of the flow are trivial, then the possible correlations from the causal model must decompose causally. Possible further applications of this tool are as follows. Flows might be beneficial in bounding the computational power of classical and quantum computation without causal order [26,27]. Also, they might be helpful in resolving the present conjecture [15] that each siblings-on-cycles graph is part of a consistent and faithful causal model. By design, flows allow us to monitor the stream of information in dynamical contexts. Arguably, classical-deterministic processes with dynamical causal order are in principle implementable within general relativity. Thus, on a foundations-of-physics front, flows might facilitate an information-based treatment of general relativity and the design of relativistic protocols for information processing. In a related work, Vilasini and Renner [28] show the limitations of embedding quantum cyclic causal models within a *fixed* relativistic spacetime. Here, we take a different approach and study relativistic causal structures that may be *dynamic*, as guided by the principles of general relativity. We leave open the embedding of quantum cyclic causal models in such *dynamic* spacetimes.

ACKNOWLEDGMENTS

We thank Lefteris Tselentis for helpful discussions and two anonymous referees for their valuable input. This work is supported by the Swiss National Science Foundation (SNSF) through Project No. 214808 and by the Hasler Foundation through Project No. 24010. We have implemented Algorithm 2 as a C program [29]. For the Appendix, we also made use of the “SOC Observation Code” [30].




FIG. 5. Superflow of the example given in Fig. 3(a) [Fig. 4(b)].

DATA AVAILABILITY

The data that support the findings of this article are openly available [29].

APPENDIX: SUPERFLOW EXAMPLES

Figure 4 gives a comprehensive list of all connected four-vertex digraphs that are of relevance for Theorem 3. These digraphs satisfy some properties. First, they are admissible




FIG. 6. Superflow of the example given in Fig. 4(c).




FIG. 7. Superflow of the example given in Fig. 4(g).

as causal structures, i.e., they are siblings-on-cycles graphs. Second, each of them has at least one source vertex. Finally, each of them contains at least one directed cycle with a chord. As by our algorithm, it turns out that all digraphs but (a), (d), and (g) produce causal correlations only.

In Figs. 5 and 6, we present the superflows of the causal structures given in Figs. 4(b) and 4(c). In both cases, the

superflow has trivial leaves. By Theorem 3, the correlations produced by any causal model with such a causal structure are restricted to decompose causally. In Fig. 7, we present the superflow of the causal structure of Fig. 4(g). Here, we have a nontrivial leaf: There might exist model parameters such that the correlations produced by such a causal structure do not decompose causally.

-
- [1] L. Hardy, Probability theories with dynamic causal structure: A new framework for quantum gravity, [arXiv:gr-qc/0509120](#).
 - [2] T. Colnaghi, G. M. D'Ariano, S. Facchini, and P. Perinotti, Quantum computation with programmable connections between gates, *Phys. Lett. A* **376**, 2940 (2012).
 - [3] O. Oreshkov, F. Costa, and Č. Brukner, Quantum correlations with no causal order, *Nat. Commun.* **3**, 1092 (2012).
 - [4] C. Portmann, C. Matt, U. Maurer, R. Renner, and B. Tackmann, Causal boxes: Quantum information-processing systems closed under composition, *IEEE Trans. Inf. Theory* **63**, 3277 (2017).
 - [5] G. Chiribella, G. M. D'Ariano, P. Perinotti, and B. Valiron, Quantum computations without definite causal structure, *Phys. Rev. A* **88**, 022318 (2013).
 - [6] J. Pearl, *Causality* (Cambridge University Press, Cambridge, 2009).
 - [7] H. Reichenbach, The principle of the common cause, in *The Direction of Time*, edited by M. Reichenbach (University of California Press, Berkeley, CA, 1971), Chap. 19, pp. 157–167.
 - [8] J. S. Bell, On the Einstein Podolsky Rosen paradox, *Phys. Phys. Fiz.* **1**, 195 (1964).
 - [9] T. Fritz, Beyond Bell's theorem II: Scenarios with arbitrary causal structure, *Commun. Math. Phys.* **341**, 391 (2016).
 - [10] F. Costa and S. Shrapnel, Quantum causal modelling, *New J. Phys.* **18**, 063032 (2016).
 - [11] J. M. A. Allen, J. Barrett, D. C. Horsman, C. M. Lee, and R. W. Spekkens, Quantum common causes and quantum causal models, *Phys. Rev. X* **7**, 031021 (2017).
 - [12] J. Barrett, R. Lorenz, and O. Oreshkov, Quantum causal models, [arXiv:1906.10726](#).
 - [13] M. Araújo, A. Feix, M. Navascués, and Č. Brukner, A purification postulate for quantum mechanics with indefinite causal order, *Quantum* **1**, 10 (2017).
 - [14] T. Richardson and J. M. Robins, Single world intervention graphs (SWIGs): Unifying the counterfactual and graphical approaches to causality, Working Paper 128, Center for the Statistics and the Social Sciences, University of Washington, Seattle, WA, 2013.
 - [15] E.-E. Tselentis and Ä. Baumeler, Admissible causal structures and correlations, *PRX Quantum* **4**, 040307 (2023).
 - [16] Two vertices u, v in a directed graph D are *siblings*, whenever they share at least one parent, i.e., whenever $\text{Pa}_D(u) \cap \text{Pa}_D(v) \neq \emptyset$, for $\text{Pa}_D(\ell) := \{k \mid k \rightarrow \ell \in \mathcal{E}(D)\}$.
 - [17] Ä. Baumeler and E.-E. Tselentis, Equivalence of grandfather and information antinomy under intervention, in *Proceedings of the 17th International Conference on Quantum Physics and Logic, Electronic Proceedings in Theoretical Computer Science*, edited by B. Valiron, S. Mansfield, P. Arrighi, and P. Panangaden (Open Publishing Association, 2021), Vol. 340, pp. 1–12.
 - [18] A flow graph is a *rooted directed graph*: For every vertex v , there exists a directed path from the root r to v .
 - [19] Ä. Baumeler and S. Wolf, Perfect signaling among three parties violating predefined causal order, in *2014 IEEE International Symposium on Information Theory (ISIT)* (IEEE, Piscataway, NJ, 2014), pp. 526–530.

- [20] C. Branciard, M. Araújo, A. Feix, F. Costa, and Č. Brukner, The simplest causal inequalities and their violation, *New J. Phys.* **18**, 013008 (2015).
- [21] Ä. Baumeler, A. S. Gilani, and J. Rashid, Unlimited non-causal correlations and their relation to non-locality, *Quantum* **6**, 673 (2022).
- [22] Ä. Baumeler, F. Costa, T. C. Ralph, S. Wolf, and M. Zych, Reversible time travel with freedom of choice, *Class. Quantum Grav.* **36**, 224002 (2019).
- [23] O. Oreshkov and C. Giarmatzi, Causal and causally separable processes, *New J. Phys.* **18**, 093020 (2016).
- [24] A. A. Abbott, C. Giarmatzi, F. Costa, and C. Branciard, Multipartite causal correlations: Polytopes and inequalities, *Phys. Rev. A* **94**, 032131 (2016).
- [25] M.-D. Choi, Completely positive linear maps on complex matrices, *Linear Algebra Appl.* **10**, 285 (1975).
- [26] Ä. Baumeler and S. Wolf, Computational tameness of classical non-causal models, *Proc. R. Soc. A* **474**, 20170698 (2018).
- [27] M. Araújo, P. Allard Guérin, and Ä. Baumeler, Quantum computation with indefinite causal structures, *Phys. Rev. A* **96**, 052315 (2017).
- [28] V. Vilasini and R. Renner, Fundamental limits for realizing quantum processes in spacetime, *Phys. Rev. Lett.* **133**, 080201 (2024).
- [29] Ä. Baumeler, Superflow (2025), <https://git.indyfac.ch/amin/superflow>.
- [30] Ä. Baumeler and E.-E. Tselentis, SOC Observation Code (2023), <https://git.indyfac.ch/amin/SOC-Observation-Code>.