



HAL
open science

Embedding cyclic information-theoretic structures in acyclic space-times: No-go results for indefinite causality

Vilasini Venkatesh, Renato Renner

► To cite this version:

Vilasini Venkatesh, Renato Renner. Embedding cyclic information-theoretic structures in acyclic space-times: No-go results for indefinite causality. *Physical Review A*, 2024, 110 (2), pp.022227. 10.1103/PhysRevA.110.022227 . hal-04720652

HAL Id: hal-04720652

<https://inria.hal.science/hal-04720652v1>

Submitted on 9 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Embedding cyclic information-theoretic structures in acyclic spacetimes: no-go results for indefinite causality

V. Vilasini^{1,2,*} and Renato Renner^{2,†}

¹*Université Grenoble Alpes, Inria, 38000 Grenoble, France*

²*Institute for Theoretical Physics, ETH Zurich, 8093 Zürich, Switzerland*

(Dated: 8th October, 2024)

The notions of causality adopted within the quantum information and spacetime physics communities are distinct. Although experience tells us that these notions play together in a compatible manner in physical experiments, their general interplay is little understood in theory. Therefore, we develop a theoretical framework that connects the two causality notions, while also clearly distinguishing them. The framework describes a composition of quantum operations through feedback loops, and the embedding of the resulting, possibly cyclic information-theoretic structure in an acyclic spacetime structure. Relativistic causality (which forbids superluminal communication) then follows as a graph-theoretic compatibility condition between the two structures. Demonstrating that indefinite causal order (ICO) processes, widely studied in the quantum information community, can be formulated within our framework, we shed light on the links between indefinite and cyclic causality, and on questions regarding their physicality. In particular, there are several experiments that claim to implement ICO processes in Minkowski spacetime, presenting an apparent theoretical paradox: how can an *indefinite* information-theoretic causal structure be consistent with a *definite* spacetime structure? We address this through no-go theorems, showing that as a consequence of relativistic causality, (a) realisations of ICO processes necessarily involve the non-localisation of systems in spacetime and (b) will nevertheless admit an explanation in terms of a definite and acyclic causal order process, at a more fine-grained level. These results are made possible by introducing the concept of fine-graining that allows causal structures to be analysed at different levels of detail. This fully resolves the apparent paradox and bears implications for the physical interpretation of ICO experiments. Our work also sheds light on the limits of quantum information processing in spacetime and offers concrete insights on the operational meaning of indefinite causality, both within and beyond the context of a fixed spacetime.

Contents

I. Introduction and summary of contributions	2
II. Abstract causal structures	6
III. Information-theoretic causal structures	7
A. Cyclic quantum networks	7
B. Signalling and causal structures of a network	10
C. Compatibility with an abstract causal structure	12
D. Fine-graining of quantum networks	13
IV. Spatio-temporal causal structures	15
V. Linking information-theoretic and spatio-temporal causal structures	17
A. Realisations of cyclic quantum networks in fixed acyclic spacetimes	17
B. Relativistic causality	18
VI. Indefinite causal structures in the process matrix framework	18
VII. Indefinite causal structures as instances of cyclic quantum networks	20

*Electronic address: vilasini@inria.fr

†Electronic address: renner@ethz.ch

VIII. No-go results for spacetime realisations of quantum processes	23
A. Localisation of information in spacetime	23
B. Existence of a fine-grained explanation in terms of a fixed order process	25
IX. Causality in the quantum switch	26
A. The QS process	26
B. No go result for QS	27
C. Implications for experimental realisations	29
D. A realisation of QS in Minkowski spacetime with systems localised in time	31
X. Towards demystifying indefinite causation	32
A. Different notions of events	33
B. Relating the process matrix assumptions to those of our no-go theorems	33
C. Operational meaning of causal inequalities	34
D. Cyclic causal structures and insights from causal inference	34
E. Beyond classical spacetimes and classical reference frames	35
XI. Conclusions and outlook	36
A. Overview of symbols and notation	38
B. Further details of loop composition	38
C. Illustrative examples of fine-graining and link to causal models	40
D. Physical spacetime realisations are causal boxes	43
E. Further details of the process matrix framework	45
F. Equivalence of notions of signalling in process networks	46
G. Fine-grained acyclic network describing a spacetime realisation of QS	48
1. Causal box description of the fine-grained network	48
2. Encoders and decoders relating the fine and coarse grained networks	50
H. Comparing fixed spacetime and quantum gravitational realisations	53
I. Relation to causal loops without superluminal signalling	54
J. Proofs of all results from the main text	56
References	64

I. INTRODUCTION AND SUMMARY OF CONTRIBUTIONS

The notion of *causality*, although fundamental to science, appears in several different forms across scientific disciplines. A defining feature of a *causal structure* is that it specifies an *order relation* between certain *events*, the meaning of the events and the relation varies depending on the notion of causality considered. In relativity theory, causality is a property of the geometry of spacetime, the events are points (or regions) in spacetime, and the causal relation is specified by the metric tensor together with a time direction.¹ We refer to this as *spacetime causality*. In the context of quantum theory, causality is more commonly used to describe how interventions, such as the choices of parameters by an agent (such as an experimenter), influence observations. The events are associated to random variables (or, more generally, quantum Hilbert spaces) and the order relation is specified by information-theoretic

¹ Considering the fact that the metric tensor can be recovered from the causal structure up to a scaling factor, one can basically identify the causal relation with the spacetime geometry.

channels connecting them, which can generate correlations between them. This enables analyses of the flow of information in quantum communication protocols between multiple agents. We refer to this as *information-theoretic causality*. The information-theoretic approach also aligns with how causality is defined within causal modelling and causal inference approaches originally developed in classical statistics [1], which have found diverse applications across data driven fields such as machine learning, economics and medicine.

In physical experiments, information-theoretic structures are *embedded* in spacetime and these two notions of causality play together in a compatible manner. In fact, *relativistic causality principles* such as the impossibility of superluminal signalling impose such compatibility conditions by requiring the direction of information flow to align with the light-cone structure and time direction of the spacetime (see also Fig. 1a), and this has been formalised in general physical theories, in recent works involving one of the authors [2–4]. As a consequence of Bell’s theorem [5], the interplay between spacetime and information-theoretic causality becomes particularly subtle in quantum experiments where freely chosen measurements are performed on entangled quantum systems. The theorem indicates a fundamental incompatibility between classical information-theoretic notions of causality and the spatio-temporal notion of causality [6, 7] in explaining such experiments. In light of this, there has been significant progress in the development of quantum causal modelling frameworks [7–20], which reveal that free choice along with relativistic principles of causality in the background spacetime can be simultaneously preserved if we replace our classical intuitions with a quantum information-theoretic notion of causation. These formalisms provide a faithful and compatible explanation of both notions of causation in quantum experiments where information-theoretic events are localised in an acyclic spacetime.

However, we can have physical scenarios, even in a classical background spacetime, where quantum systems (associated with the in and outputs of quantum channels) are not localised in space and time (see e.g., [21–23]). More generally, in quantum gravitational settings, a fixed background spacetime structure may no longer be available [24–26]. This has motivated information-theoretic formalisms for analysing multi-agent information-processing protocols where a definite acyclic ordering of the agents’ operations is not assumed [24, 27, 28]. A prominent example is the process matrix framework [28]. Process matrices which are incompatible with a definite acyclic (information-theoretic) causal order between the operations of the agents are regarded as *indefinite causal structures* in this framework (we do not necessarily endorse this terminology, but we will use it in order to remain consistent with the literature. It is to be noted that this refers to an indefinite order of operations in a quantum information protocol, not indefiniteness in the order or geometry relating spatio-temporal events). The theoretical framework encompasses regular quantum circuits, quantum combs [29], superpositions of direct and common cause scenarios [30] as well as more exotic processes such as *causally non-separable processes* and those that violate so-called *causal inequalities* [28, 31, 32]². These more general classes have been widely studied for the potential advantages that they may provide over regular quantum circuits in various information-theoretic tasks [33–40].

However, there are several intriguing open questions and debates regarding the physicality of such processes as well as their link to spatio-temporal concepts. A first question is regarding the set of process matrices that can be realised in accordance with standard quantum theory in Minkowski spacetime. Through bottom-up constructions, a general class of quantum circuits realising quantum controlled superpositions of the order (of agents’ operations) have been proposed [41, 42], and an important open question is whether these correspond to the largest class of physically realisable processes. On the other hand, numerous table-top experiments have been performed that claim to physically implement an indefinite causal structure (*the quantum switch* [27]) in Minkowski spacetime [23, 43–51], and their interpretation has been a subject of much debate in the community [21, 52–56]. A large class of previous theoretical and experimental results appear to suggest that there is a sense in which certain indefinite causal structure processes can be realised in Minkowski spacetime. A necessary step for clarifying these questions about quantum causality is to address the apparent paradox—how can an indefinite information-theoretic causal structure be consistent with relativistic causality in a definite and acyclic spacetime causal structure? We develop a general top-down approach for formally addressing such questions and in a manner that accommodates and links different causality notions. Such a general framework must incorporate the following desiderata, we first state them and then justify them.

1. The framework must clearly disentangle the information-theoretic and spacetime notions of causality, and characterise both under minimal but operational assumptions, which means that the former need not in general be acyclic.

² Analogous to non-separable quantum states and quantum states that violate Bell inequalities.

2. To have relevance for physical experiments, the framework must formalise what it entails for a non-acyclic information-theoretic structure to be compatible with relativistic causality in an acyclic spacetime.
3. The embedding must be general enough to model quantum systems that are not necessarily localised in spacetime.
4. The framework should permit the analysis of causal structures at different levels of detail, as the acyclicity of a causal structure can depend on the information captured by its nodes.

Both indefinite as well as cyclic causal structures correspond to non-acyclic causal structures. Interestingly, previous works [27, 28, 57–59] have demonstrated mathematical correspondences between the two, which provide very useful insights for achieving the first criterion above. However, these works focus on the details of the information-theoretic notion and do not consider spacetime causality or compatibility as required by the second criterion. In recent works involving one of the authors [2, 3], a framework applicable to quantum as well as post-quantum probabilistic theories, that meets both the first and second criteria has been developed. The present work is inspired by the general approach to information-theoretic and spatio-temporal causal structures introduced there. However, the previous framework mainly focuses on signalling between classical settings and outcomes (even if these may be generated by measuring quantum or post-quantum systems), and the criteria 3 and 4 were not considered there, which are necessary for analysing realisations of indefinite causal order processes.

There has also been much progress, in particular, coming from the research program relating to *quantum reference frames* (QRF), towards instantiating process matrices with temporal information [26, 54, 60, 61]. While these cover parts of points 1, 2 and 3, these approaches typically only describe fully unitarised quantum circuits. Moreover all systems involved, such as the reference frame and the observer whose perspective is described, are modelled as fully quantum systems. QRF approaches can thus also capture theoretical quantum gravitational realisations of process matrices where spacetime geometry is subject to quantum superpositions. However, this is a different physical regime than that of presently realisable quantum experiments, which involve classical observers (our experimental colleagues) who have access to classical reference frames (such as their wall clock) which they can utilise to order and control their actions. Such experiments effectively take place in a well-defined classical background spacetime. This is an important yet often overlooked distinction, bearing implications for the concept of causality and events, and we discuss this further in Section X.

To the best of our knowledge, previous approaches do not fully address point 2 specifically in the context of a background spacetime, while accounting for arbitrary measurements along with unitary transformations that can be physically performed by the agents involved. Point number 4 above, which has not been previously considered, is crucial for this purpose as the following classical example illustrates. Cyclic causal models are often used in classical data sciences for modelling physical scenarios with feedback [62], for instance the demand D and price P for a commodity can causally influence each other. However, we know that the situation is ultimately described by an acyclic causal structure (possibly over a larger number of nodes) where the demand D_1 at time t_1 influences the price P_2 at time $t_2 > t_1$ which in turn influences the demand D_3 at time $t_3 > t_2$ and so on. Indeed agents can verify this by intervening on D or P at some specific time t (such as through appropriate choice of market policies) and observing changes in the other variable at a later time. If we coarse-grain over the time information in this acyclic causal structure, we recover the original cyclic causal structure. This example (illustrated in Figures 1b and 1c) highlights that the cyclicity or acyclicity of a causal structure can depend on the level of detailed information encoded in its nodes. Therefore, the apparently exotic cyclic (information-theoretic) causal structure has a clear physical interpretation in such scenarios. Here we develop a framework that meets all the above criteria, which enables a similar clear and operational interpretation for spacetime realisations of so-called indefinite causal order processes.

Within our formalism we derive multiple classes of results: characterisation of purely information-theoretic causal structures and their compatibility with spacetime, no-go theorems for spacetime realisations of indefinite causal structures, examples and results establishing the applicability of our framework to a range of existing approaches to causality.

Summary of contributions We provide a summary of these contributions below, which would enable a reader interested in a subset of the results to navigate to the relevant sections and theorems. For those seeking a concise yet formal introduction to our framework, particularly the methods relevant for understanding our main no-go results for ICOs (Theorem 4, Theorem 5 and Corollary 3), along with an intuitive proof sketch for the theorems, please refer to the associated letter [63] and its supplementary materials. In Section II of this paper, we start with general abstract definitions that apply to both notions of causality and the following paragraphs summarise the contents of the remaining sections. Proofs of all results from the main text can be found in Appendix J.




FIG. 1: (a) A simple information-theoretic causal structure, over nodes X and Y embedded in Minkowski spacetime (time along vertical and space along horizontal axis). The 45 degree lines (blue) indicate light like surfaces emanating from the spacetime point where X is embedded, and we see that the information-theoretic causal influence flows within the lightcone and the two causal orders are compatible. (b) A cyclic information-theoretic causal structure over variables D and P depicting the demand and price of a commodity which can mutually influence each other. (c) A fine-grained acyclic causal structure that would physically explain (b), this captures that demand at one time influences price at a later time and so on. Notice that (b) would not be compatible with an acyclic spacetime if we associate individual spacetime locations to D and P , but we can obtain the physical fine-grained picture given in (c) if we associate spacetime regions (or sets of spacetime points) to nodes of (b), in which case we can also recover compatibility with the spacetime. This highlights that the physical interpretation of seemingly exotic cyclic information-theoretic causal structures depends on how they are linked to spatio-temporal degrees of freedom.

- A top-down approach to quantum information-theoretic causation** In Section III, we describe causal influences in a purely information-theoretic sense by defining (possibly cyclic) quantum networks that can be formed through compositions of quantum operations by means of feedback loops. We focus on the operationally relevant notion of signalling for such networks, while taking into account that signalling and causation are distinct [2, 3, 7, 59]. We introduce the concepts of *embedding* and *fine-graining* of quantum networks which enables us to formulate their *compatibility* with abstract graphs (Section III C), and to study different realisations of such networks with the possibility of analysing them at different levels of detail (Section III D).
- Compatibility between cyclic quantum networks and acyclic spacetimes** Defining spatio-temporal causal structures in terms of graphs (Section IV), we apply the concepts of embedding and compatibility to link the two notions of causality and provide an order-theoretic formalisation of *relativistic causality* (Section V). In theorem 1, we show that any (possibly cyclic) signalling structure can be compatibly embedded in an acyclic spacetime if we allow non-localisation of information in the spacetime.
- Recovering process matrices, linking indefinite and cyclic causality** After reviewing the process matrix formalism (Section VI), we show that process matrix protocols can be recovered as a special case of the general cyclic quantum networks of our framework. In particular, we recover the probability rule of the process formalism through loop composition (Lemma 2) and derive a tight correspondence between the non-definiteness of causal order in the process formalism and the cyclicity of an underlying causal structure (inferred solely through signalling relations) in our framework (Theorem 3).
- No-go result 1: non-localisation of information in spacetime** Our first main no-go theorem (Theorem 4) implies that any realisation of an indefinite causal order process which satisfies relativistic causality in a fixed acyclic spacetime must necessarily involve the non-localisation of information in spacetime, while indicating the degree to which it must be non-localised. As a corollary it implies the impossibility of such realisations where the in and output systems of agents' operations are time localised in a global reference frame (Corollary 2). These results shed light on the physical resources in spacetime which are necessary for realising such quantum processes.
- No-go result 2: recovering definite causal structure under fine-graining** Our second set of no-go results (Theorem 5 and Corollary 3) show that any realisation of an indefinite causal order process which satisfies relativistic causality in a fixed acyclic spacetime will ultimately admit, at a fine-grained level, an explanation in terms of fixed causal order process over a larger number of agents, where the (information-theoretic) causal order is definite and acyclic. This holds even for realisations where agents exchange quantum messages at

superpositions of spacetime locations. This resolves the apparent paradox regarding the two causality notions in a fixed spacetime.

- **Implications for quantum switch experiments** We consider the special case of the quantum switch (QS, an indefinite causal order process) and derive stronger versions of our two main no-go results for this case (Lemma 4 and Corollary 4), which imply that existing experimental realisations of this process in Minkowski spacetime can be explained in terms of a definite and acyclic causal order at a fine-grained level. In Section IX D, we identify a new realisation of QS in Minkowski spacetime where the in/output systems of each agent are time localised in their own reference frame, a property that was previously observed only in quantum gravitational realisations of QS. However, the conclusions of our no-go theorems also apply to such realisations, and we discuss the physical interpretation of such experiments in detail (Section IX C).
- **Operational meaning of indefinite causation in and beyond fixed spacetimes** In Section X, we connect the assumptions of our no-go theorem to the set-up assumptions of the process matrix framework which sheds light on the operational meaning of indefinite causation and its witnesses such as causal inequality violations. We also discuss the outlook that our work provides for studying different notions of events, and understanding causality and information-processing beyond the context of a fixed background spacetime.
- **Relation to other approaches to causality** In the appendices, we show that other distinct approaches to causality can also be studied in our general framework. The results of Appendix D show that the most general realisations of (cyclic) quantum networks satisfying relativistic causality in a fixed acyclic spacetime correspond to so-called *causal boxes* that were previously proposed for studying relativistic quantum cryptography [21, 64]. Appendix C describes how classical *functional causal models on cyclic graphs* (studied in classical statistics [62]) as well as *quantum Bayesian networks* [13] can be studied in our formalism, while illustrating how the concept of fine-graining introduced here applies to such causal models. Finally, Appendix I discusses the relation to recent research [2, 3], where it was shown that it is theoretically possible to have genuine causal loops between events in Minkowski spacetime, without superluminal signalling.

Finally, we refer to Appendix A for an overview of the notations used in the paper.

II. ABSTRACT CAUSAL STRUCTURES

As mentioned in the introduction, a basic feature that is common across different notions of causality is that a causal structure specifies an order relation between certain events. We will therefore use the following minimal definition of a causal structure in this paper. This abstract notion will be specified further depending on the flavour of causality (information-theoretic, spatio-temporal etc.) being considered.

Definition 1 (Causal structures). A causal structure is any directed graph \mathcal{G} where $\text{Nodes}(\mathcal{G})$ denotes the set of all nodes of \mathcal{G} or “events” of the causal structure \mathcal{G} , and $\text{Edges}(\mathcal{G})$ denotes the set of all edges of \mathcal{G} or “causal relations” of \mathcal{G} . We say that N_1 is a cause of N_2 in \mathcal{G} if there exists a directed path from the node N_1 to the node N_2 in \mathcal{G} , and in particular that N_1 is a direct cause of N_2 if there is a directed edge from N_1 to N_2 in \mathcal{G} .

To achieve the main goals laid out in the introduction, in particular to capture the fourth desideratum of being able to analyse causal structures at different levels of detail, we introduce the concept of *fine-graining of causal structures*. The following minimal definition defines fine-graining for directed graphs and applies to both information-theoretic and spacetime causal structures. For the information-theoretic case, our example with demand and price illustrated in Figures 1b and 1c paints an intuitive picture of fine-graining, the acyclic causal structure over D_1, P_2, D_3, \dots in the latter figure is a fine-graining of the cyclic causal structure over D and P in the former figure. For the spatio-temporal case, considering spacetime regions as the nodes of the causal structure, we can obtain a more fine-grained description by splitting the regions into sub-regions or individual spacetime events (cf. Section IV). Throughout the paper, we will use $\Sigma(\cdot)$ to denote the powerset of the set in its argument and we say that there is a directed path from a subset S_1 to a subset S_2 of nodes in a directed graph \mathcal{G} whenever there exists $N_1 \in S_1$ and $N_2 \in S_2$, and some nodes M_1, \dots, M_k such that \mathcal{G} contains directed edges from N_1 to M_1 , from M_l to M_{l+1} for all $l \in \{1, \dots, k-1\}$, and from M_k to N_2 .³

³ Here, the sets S_1 and S_2 can have overlap and N_1 and N_2 can be the same, in which case a directed path from N_1 to itself would correspond to a directed cycle.




FIG. 2: Fine-grainings of a directed graph (color online): A directed graph \mathcal{G} (a) and two possible fine-grainings \mathcal{G}_1^f (b) and \mathcal{G}_2^f (c) of the same graph, both of which satisfy Definition 2 but relative to different maps $\mathcal{F}_1^{\text{graph}}$ and $\mathcal{F}_2^{\text{graph}}$ (which are denoted as \mathcal{F}_1 and \mathcal{F}_2 to avoid clutter). Each node N_i in the original graph \mathcal{G} , with each $i \in \{1, 2, 3, 4\}$ depicted in a distinct color, maps to a set of nodes $\mathcal{F}_j^{\text{graph}}(N_i)$ ($j \in \{1, 2\}$) in the fine-grained graph \mathcal{G}_j^f lying within a blob of the same color. Notice that (a) is a cyclic graph while its fine-grainings are acyclic. Definition 2 is general and allows for overlaps in the image of the fine-graining map, as seen in (b). This will be relevant the spacetime causal structure, as it will allow us to consider overlapping spacetime regions. On the other hand, we will see that the more specific definition of fine-graining for information-theoretic structures will forbid such overlaps as it will map each quantum system in the coarse-grained description to a distinct set of systems in the fine-grained description.

Definition 2 (Fine-graining of a directed graph). A directed graph \mathcal{G}^f is called a fine-graining of a directed graph \mathcal{G} if there exists a map $\mathcal{F}^{\text{graph}} : \text{Nodes}(\mathcal{G}) \mapsto \Sigma(\text{Nodes}(\mathcal{G}^f))$ such that whenever there is a directed edge from N_1 to N_2 in \mathcal{G} (for $N_1, N_2 \in \text{Nodes}(\mathcal{G})$), there exists a directed path from $\mathcal{F}^{\text{graph}}(N_1)$ to $\mathcal{F}^{\text{graph}}(N_2)$ in \mathcal{G}^f .

The definition is illustrated in Fig. 2. A further property that one would intuitively expect from a concept of fine-graining is that the number of nodes in the image of $\mathcal{F}^{\text{graph}}$ is larger than the number of nodes of \mathcal{G} . While this is not required for our main results, it will be the case in all examples we illustrate in this paper (including those of Fig. 2). In the next sections, we will describe how we model general information-theoretic and spatio-temporal causal structures, and introduce more detailed definitions of fine-graining for these.

III. INFORMATION-THEORETIC CAUSAL STRUCTURES

The operational formulation of quantum theory describes quantum networks (such as circuit diagrams) formed by the composition of quantum operations. This enables information flow from an in/output system of one operation to that of another in the network, allowing for a purely information-theoretic notion of causal influence to be defined. Typically, only combinations of operations resulting in acyclic causal structures are considered such that every network singles out a direction, which we can regard as “time”. More generally, we need not restrict to acyclic information-theoretic causal structures. We often have physical scenarios with feedback where the output of a physical device is looped back and fed in to its input, and in classical statistics, the mathematics of cyclic causal structures has therefore been developed [62].

Here, we develop a formalism for composing quantum operations through feedback loops that can result in cyclic quantum networks. We will make statements about the underlying causal explanations of a network by focusing on the operational notion of signalling which can be detected through physical interventions. Finally, we introduce the concept of *fine-graining of quantum networks*, which will enable us to relate cyclic and acyclic networks in a manner that preserves the relevant information and to study different realisations of such networks.

A. Cyclic quantum networks

Notation 1. We consider completely positive linear maps (henceforth, simply referred to as CP maps or CPMs) associated with specified sets of in and output subsystems. For a CPM \mathcal{M} , we denote the sets of in and output

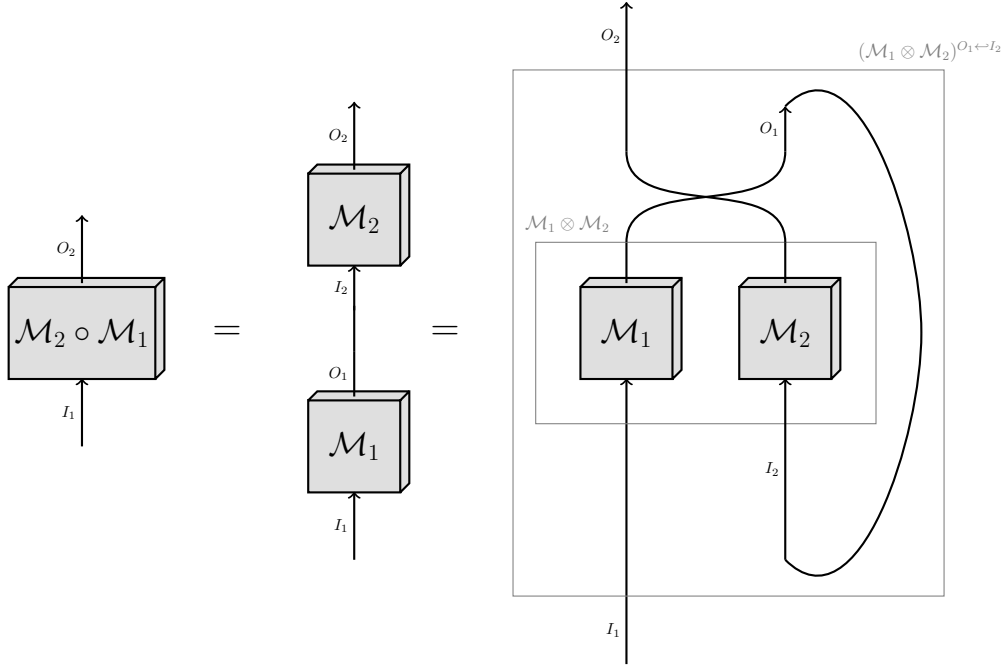


FIG. 3: Illustration of Eq. (2). Sequential composition of CPMs can be equivalently described as a parallel composition followed by loop composition.

subsystems as $\text{In}(\mathcal{M})$ and $\text{Out}(\mathcal{M})$ respectively, and only consider scenarios where these are finite sets. Whenever the map \mathcal{M} is evident from context, we simply refer to these sets as In and Out . Each element $S \in \text{In}(\mathcal{M}) \cup \text{Out}(\mathcal{M})$ is associated with a Hilbert space \mathcal{H}^S and a corresponding set $\mathcal{L}(S)$ of quantum states (or density matrices) on \mathcal{H}^S . More generally, if \mathcal{S} is a subset of in or output systems, $\mathcal{L}(\mathcal{S}) := \otimes_{S \in \mathcal{S}} \mathcal{L}(S)$ denotes the joint state space of all systems in \mathcal{S} . Moreover, for any set $\{\mathcal{M}_i\}_i$ of CPMs we denote the union of input systems as $\text{In}(\{\mathcal{M}_i\}_i) := \cup_i \text{In}(\mathcal{M}_i)$, and similarly for the output systems. We will assume that each Hilbert space \mathcal{H} of a d -dimensional quantum system has a well defined *computational basis* $\{|i\rangle\}_{i \in \{0, \dots, d-1\}}$ consisting of orthonormal vectors $|i\rangle$, where d is finite.

We consider two types of composition operations on CPMs, parallel and loop composition [21], as defined below.

Definition 3 (Parallel composition). Parallel composition of two CPMs \mathcal{M} and \mathcal{N} is a CPM $\mathcal{M} \parallel \mathcal{N} := \mathcal{M} \otimes \mathcal{N}$, with $\text{In}(\mathcal{M} \parallel \mathcal{N}) = \text{In}(\mathcal{M}) \cup \text{In}(\mathcal{N})$ and $\text{Out}(\mathcal{M} \parallel \mathcal{N}) = \text{Out}(\mathcal{M}) \cup \text{Out}(\mathcal{N})$.

Definition 4 (Loop composition [21]). Consider a CPM \mathcal{M} where $S \in \text{In}(\mathcal{M})$ and $S' \in \text{Out}(\mathcal{M})$ have isomorphic Hilbert spaces $\mathcal{H}^S \cong \mathcal{H}^{S'}$. Then we can compose the output subsystem S with the input subsystem S' through loop composition to obtain a CPM $\mathcal{M}^{S \rightarrow S'}$, with $\text{In}(\mathcal{M}^{S \rightarrow S'}) = \text{In}(\mathcal{M}) \setminus \{S'\}$ and $\text{Out}(\mathcal{M}^{S \rightarrow S'}) = \text{Out}(\mathcal{M}) \setminus \{S\}$. The action of $\mathcal{M}^{S \rightarrow S'}$ is given as follows, where we take $\rho^{\text{In} \setminus \{S'\}}$ to be an arbitrary state in $\mathcal{L}(\text{In} \setminus \{S'\})$, and treat, for simplicity (and without loss of generality), S' as the last input subsystem.

$$\mathcal{M}^{S \rightarrow S'}(\rho^{\text{In} \setminus \{S'\}}) := \sum_{i,j} |i\rangle^S \mathcal{M}(\rho^{\text{In} \setminus \{S'\}} \otimes |i\rangle\langle j|^{S'}) |j\rangle^S, \quad (1)$$

where the bases appearing in the above equation are taken by convention to be the computational basis of the relevant system.

In Appendix B we show that the set of CPMs is indeed closed under parallel and loop composition as suggested by the above definitions. Now, consider two CPMs \mathcal{M}_1 and \mathcal{M}_2 where $\text{In}(\mathcal{M}_j) = \{I_j\}$, $\text{Out}(\mathcal{M}_j) = \{O_j\}$ for $j \in \{1, 2\}$ and $\mathcal{H}^{O_1} \cong \mathcal{H}^{I_2}$. The sequential composition of \mathcal{M}_1 followed by \mathcal{M}_2 is a CPM with inputs $\text{In}(\mathcal{M}_1)$ and outputs $\text{Out}(\mathcal{M}_2)$ that acts as $\mathcal{M}_2(\mathcal{M}_1(\rho))$. This sequential composition can be equivalently obtained by first composing the two CPMs in parallel and then performing a loop composition on the single CPM obtained through the parallel composition [21] i.e.,

$$(\mathcal{M}_1 \otimes \mathcal{M}_2)^{O_1 \rightarrow I_2}(\rho) = \mathcal{M}_2(\mathcal{M}_1(\rho)) \quad \forall \rho \in \mathcal{L}(\text{In}(\mathcal{M}_1)). \quad (2)$$

This is illustrated in Fig. 3. More generally, the maps may have an arbitrary number of in and output subsystems and the sequential composition may take place by connecting one of the output subsystems of \mathcal{M}_1 to one of the input subsystems of \mathcal{M}_2 , and this can also be equivalently described as a loop composition on those subsystems in the parallel composition of the two maps. We are now ready to define what we mean by a network of CPMs.

Definition 5 (Network of CPMs). A network $\mathfrak{N} := (\mathfrak{N}^{\text{maps}}, \mathfrak{N}^{\text{comp}})$ of CPMs is specified by two elements: a set $\mathfrak{N}^{\text{maps}}$ of CPMs and a set $\mathfrak{N}^{\text{comp}}$ of connections between the CPMs in $\mathfrak{N}^{\text{maps}}$ given through loop compositions. More specifically, $\mathfrak{N}^{\text{comp}}$ is a set of the form

$$\mathfrak{N}^{\text{comp}} = \{S \leftrightarrow S' \mid S \in \text{Out}(\mathfrak{N}^{\text{maps}}), S' \in \text{In}(\mathfrak{N}^{\text{maps}}), \mathcal{H}^S \cong \mathcal{H}^{S'}\}, \quad (3)$$

where no output S is paired with more than one input S' and vice versa.

Definition 6 (Induced map of a network). The induced map \mathcal{N} of a network $\mathfrak{N} = (\mathfrak{N}^{\text{maps}}, \mathfrak{N}^{\text{comp}})$ is defined by composition of the maps in $\mathfrak{N}^{\text{maps}}$ according to $\mathfrak{N}^{\text{comp}}$ i.e.,

$$\mathcal{N} := \left(\bigotimes_{\mathcal{M} \in \mathfrak{N}^{\text{maps}}} \mathcal{M} \right)^{\{S \leftrightarrow S' \in \mathfrak{N}^{\text{comp}}\}} \quad (4)$$

This definition allows for possibly cyclic networks where there can be a directed path from an output of \mathcal{M}_1 to an input of \mathcal{M}_2 and a directed path from an output of \mathcal{M}_2 to an input of \mathcal{M}_1 (see Fig. 6a for an example). Moreover, in Appendix B, we show that given a CPM and multiple loop compositions to be performed, the loop composed CPM is well-defined independently of the order in which the compositions are performed. This implies that the induced map \mathcal{N} is itself a CPM and it is uniquely defined by the set of maps $\mathfrak{N}^{\text{maps}}$ and the set of loop compositions $\mathfrak{N}^{\text{comp}}$. The relevant set of systems in a network is then defined as follows.

Definition 7 (Systems in a network). The set of systems in a network \mathfrak{N} is defined as

$$\mathfrak{N}^{\text{sys}} := \text{In}(\mathcal{N}) \cup \text{Out}(\mathcal{N}) \cup \left(\bigcup_{S \leftrightarrow S' \in \mathfrak{N}^{\text{comp}}} S \right). \quad (5)$$

That is, $\mathfrak{N}^{\text{sys}}$ consists of all systems in the network which are not involved in a composition (which are precisely the systems in $\text{In}(\mathcal{N}) \cup \text{Out}(\mathcal{N})$), along with a single system S representing each pair S and S' of systems composed through loop composition.

Definition 8 (Sub-networks). A sub-network $\mathfrak{N}_{\text{sub}} = (\mathfrak{N}_{\text{sub}}^{\text{maps}}, \mathfrak{N}_{\text{sub}}^{\text{comp}})$ of a network $\mathfrak{N} = (\mathfrak{N}^{\text{maps}}, \mathfrak{N}^{\text{comp}})$, is itself a network where $\mathfrak{N}_{\text{sub}}^{\text{maps}} \subseteq \mathfrak{N}^{\text{maps}}$ and $\mathfrak{N}_{\text{sub}}^{\text{comp}} \subseteq \mathfrak{N}^{\text{comp}}$.

The induced map \mathcal{N}_{sub} of any sub-network is obtained by applying Eq. (4) to $\mathfrak{N}_{\text{sub}}$. Moreover, note that every network \mathfrak{N} is a sub-network of itself, and for every map $\mathcal{M} \in \mathfrak{N}^{\text{maps}}$, $(\{\mathcal{M}\}, \emptyset)$ is a sub-network of \mathfrak{N} with induced map \mathcal{M} . Of particular relevance will be networks of completely positive and trace preserving maps (CPTPMs), defined below.

Definition 9 (Network of CPTPMs). A network \mathfrak{N} of CPMs is called a network of CPTPMs if the induced map \mathcal{N}_{sub} of every sub-network $\mathfrak{N}_{\text{sub}}$ of \mathfrak{N} is a CPTPM.

Remark 1. Simply requiring each $\mathcal{M} \in \mathfrak{N}^{\text{maps}}$ to be a CPTPM does not guarantee that \mathfrak{N} is a network of CPTPMs. This is because the TP property is not always preserved by loop compositions. As we expect physical dynamics of any system under consideration, as well as of any subsystem thereof to be CP and TP, networks of CPTPMs will model examples of physical interest. As we will show later in the paper, indefinite causal order processes [27, 28] can be modelled using such networks of CPTPMs (Section VII). An interesting question for future work is to characterise the conditions on the maps and compositions of a network that would ensure that it is a network of CPTPMs. Closely related open questions are: which is the largest subset of CPTPMs that is closed under arbitrary loop compositions, and given a set of CPTPMs, which is the largest set of loop compositions which preserve the TP property?

Remark 2. A related framework is the category-theoretic approach to quantum information [65, 66], which can also describe cyclic information-theoretic causal structures from a general class of operational theories. The formalism is more abstract than ours, which focuses on quantum theory to derive more specific results for cyclic quantum causal structures. The two are however related, in particular, our formalism of cyclic quantum networks would correspond to an instance of a compact closed category. This is because loop composition as defined here can be equivalently described through pre and post-selection on maximally entangled states (i.e., in terms of post-selected teleportation, see Appendix B), which would be instances of the ‘‘cups’’ and ‘‘caps’’ of such a category.




FIG. 4: Diagrammatic representation of Definition 10 of signalling. Here \mathcal{S}_I and \mathcal{S}_O represent a subset of input and out subsystems of the CPM \mathcal{M} , and the ground symbol $\equiv\equiv\equiv$ denotes trace.

B. Signalling and causal structures of a network

Definition 10 (Signalling structure of a CPM). We say that there is a *signalling relation* from a subset $\mathcal{S}_I \subseteq \text{In}(\mathcal{M})$ of input systems to a subset $\mathcal{S}_O \subseteq \text{Out}(\mathcal{M})$ of output systems of the map \mathcal{M} and denote it as $\mathcal{S}_I \longrightarrow \mathcal{S}_O$ if there exists a CPTPM $\mathcal{N}^{\mathcal{S}_I} : \mathcal{L}(\mathcal{S}_I) \mapsto \mathcal{L}(\mathcal{S}_I)$ acting locally on \mathcal{S}_I such that the following holds.

$$\text{Tr}_{\text{Out} \setminus \mathcal{S}_O} \circ \mathcal{M} \neq \text{Tr}_{\text{Out} \setminus \mathcal{S}_O} \circ \mathcal{M} \circ \mathcal{N}^{\mathcal{S}_I}. \quad (6)$$

The set of all signalling relations of \mathcal{M} forms a directed graph $\mathcal{G}_{\mathcal{M}}^{\text{sig}}$ whose nodes are elements of the powerset (or sigma algebra) $\Sigma(\text{In} \cup \text{Out})$ of $\text{In} \cup \text{Out}$, with a directed edge \longrightarrow between two nodes whenever there is a signalling relation between those subsets. We refer to this graph as the signalling structure of \mathcal{M} .

Signalling from \mathcal{S}_I to \mathcal{S}_O captures the idea that an agent with access to the inputs can communicate information to an agent with access to the outputs \mathcal{S}_O by a suitable choice of local intervention on \mathcal{S}_I (as illustrated in Fig. 4). Notice that the definition applies to CPMs \mathcal{M} but the local operation $\mathcal{N}^{\mathcal{S}_I}$ must be a CPTPM, otherwise, we could choose $\mathcal{N}^{\mathcal{S}_I}$ to be a constant factor times a CPTPM which would make two sides of Eq. (6) trivially differ, even in maps \mathcal{M} where no physical operation could enable information flow or signalling from the inputs to the outputs. Thus the definition only holds operational meaning when the local operation $\mathcal{N}^{\mathcal{S}_I}$ are also TP. For our main results, we will only need to apply the definition to cases where \mathcal{M} is also a CPTPM, in which case, the above definition of signalling is equivalent to other natural ways in which one might formalise the ability for \mathcal{S}_I to transmit information to \mathcal{S}_O , as shown in [55, 67]. We use this formulation as our primary definition of signalling, as it encodes more explicitly how agents can communicate to each other through local interventions on nodes of the network.

Definition 11 (Signalling structure of a network of CPMs). The signalling structure $\mathcal{G}_{\mathfrak{N}}^{\text{sig}}$ of a network \mathfrak{N} is a directed graph whose nodes are elements of the powerset $\Sigma(\mathfrak{N}^{\text{sys}})$ of $\mathfrak{N}^{\text{sys}}$. We have $N_1 \longrightarrow N_2$ in $\mathcal{G}_{\mathfrak{N}}^{\text{sig}}$ (where N_1 and N_2 are disjoint subsets of $\mathfrak{N}^{\text{sys}}$) if and only if there exists a sub-network $\mathfrak{N}_{\text{sub}}$ of \mathfrak{N} whose induced map \mathcal{N}_{sub} enables signalling from N_1' to N_2 . Here, N_1' is equivalent to N_1 up to replacing every $S \in N_1$ with S' whenever $S \leftrightarrow S' \in \mathfrak{N}^{\text{comp}}$.

The above substitution of N_1 with N_1' when checking for signalling relations is to account for the identification of loop composed systems $S \leftrightarrow S' \in \mathfrak{N}^{\text{comp}}$ as a single system S . For example, consider $\mathcal{M}_1 : \mathcal{L}(I) \mapsto \mathcal{L}(S)$ and $\mathcal{M}_2 : \mathcal{L}(S') \mapsto \mathcal{L}(O)$ which are sequentially composed by connecting the output S of \mathcal{M}_1 to the input S' of \mathcal{M}_2 resulting in a network \mathfrak{N} with an induced map $\mathcal{N} : \mathcal{L}(I) \mapsto \mathcal{L}(O)$. Then $\mathfrak{N}^{\text{sys}} := \{I, O, S\}$ and we have a signalling relation $\{S\} \longrightarrow \{O\}$ in $\mathcal{G}_{\mathfrak{N}}^{\text{sig}}$, whenever $\{S'\} \longrightarrow \{O\}$ in the sub-network whose induced map is \mathcal{M}_2 .

Notation 2. Stochastic maps and deterministic functions on random variables (RV) can be viewed as special cases as CPTPMs, by encoding a RV in the computational basis of a Hilbert space. Whenever we consider purely classical examples, we will use the same labels to refer to RV as well as the quantum state space which encodes the RV.

Causal structure(s) of a network So far, we have focused on the concept of signalling, which can be operationally determined in any quantum network and we have treated causal structures only as abstract graphs. We have not




FIG. 5: (a) A classical stochastic map acting on bits (a special case of a CPTP map) where the XOR of the two output bits O_1 and O_2 equals the input bit I (b) One possible decomposition of this map where I is a cause of O_1 but not of O_2 (c) Another possible decomposition of the same map where I is not a cause of O_1 but is a cause of O_2 .

referred to the concept of information-theoretic causal structures that can be associated with a quantum network, although many such definitions exist in the literature. Most of our main results including the general no go theorems of this paper will only rely on the concept of signalling. This allows our framework and results to be generally applicable to different information-theoretic notions of causation that have been proposed in the literature, all of which agree on the definition of signalling. We will therefore not provide a full definition of information-theoretic causality here. Rather, we provide some intuition on how causality for a quantum network is typically modelled in information-theoretic approaches, through concrete physical examples. This will also help the reader appreciate the distinction between signalling and causation, even without a fully general definition of the latter.

As we have seen, the nodes of the signalling structure of a network are generally subsets of systems in the network. On the other hand, a network of maps also induces possible directed graphs on individual systems of the network, through its connectivity, which captures how information flows through systems in the network. A quantum circuit is an example of a network, and the wires in the circuit diagram tell us about the information flow. For example, in simple network involving sequential composition of quantum channels $\mathcal{M}_1 : A \mapsto B$ and $\mathcal{M}_2 : B \mapsto C$ through the B system, suggests that information can flow from A to B to C ⁴, but information cannot flow from C to A in this network. One would say that C is not a cause of A in any information-theoretic causal structure that explains this network (but A can be a cause of B for instance). In this simple example, it is clear that even without fixing the graph, we know that it must be acyclic: this network admits an explanation in terms of an acyclic information-theoretic causal structure. Notice that here, knowing that A signals to B in \mathcal{M}_1 would tell us that information does flow from A to B , hence the signalling structure tells us something about the causal structure. However, generally the signalling structure of a network (even a single CPTP map) need not fix the causal structure uniquely: we can have wires in an internal decomposition of a network which may nevertheless not lead to observable signalling at the level of the input/output behaviour of the network. The following example illustrates this fact.

Consider a CPTPM \mathcal{M} with input I and two outputs O_1 and O_2 , all of which correspond to classical bits such that \mathcal{M} is effectively a stochastic map fully described by a conditional distribution $P(O_1 O_2 | I)$ (see Fig. 5). Suppose that $P(O_1 O_2 | I)$ is such that $O_1 \oplus O_2 = I$ always holds and $P(O_1 | I)$ and $P(O_2 | I)$ are both uniform distributions, independently of I . This would imply that $\{I\}$ does not signal to $\{O_1\}$ and $\{I\}$ does not signal to $\{O_2\}$ in \mathcal{M} , however $\{I\}$ signals to $\{O_1, O_2\}$ in \mathcal{M} . Now consider the following two possible ways of implementing this map. One way is the following: \mathcal{M} generates a uniform bit K internally and sets $O_2 = K$ and $O_1 = I \oplus K$. We can then easily verify that this realises the behaviour $P(O_1 O_2 | I)$ of the channel as described above [2, 3]. Another way is to exchange the roles of O_1 and O_2 and set $O_1 = K$ and $O_2 = I \oplus K$ which also implements the same channel. If we analyse the internal connections/flow of information in the two implementations, we would say that I causally influences O_1 (since

⁴ Whether it does flow depends on the maps, e.g., if \mathcal{M}_1 is a trace and replace operation, information cannot flow from A to B

O_1 is a non-trivial function of I) but not of O_2 (which is not a function of I) in the first while I causally influences O_2 but not O_1 in the second. This example also highlights that causation does not imply signalling [2, 3, 19], in both implementations, we have that I is a cause of either O_1 or O_2 but in both cases $\{I\}$ signals to neither $\{O_1\}$ nor to $\{O_2\}$.

In summary, signalling is a property of an information-theoretic network that can be uniquely determined by an operational procedure, even when given black-box access to the network. Causal structure on the other hand, is related to particular internal decompositions of the maps in a network, and is not unique to a network. Given a particular decomposition of the maps in a network, we can infer an absence of causation from one system S to another system S' whenever there is an absence of a directed path of wires from S to S' in the decomposition. Therefore, in some situations, knowing that there exists a decomposition of a given network with the absence of directed paths between certain systems would be sufficient to conclude that the network admits an explanation in terms of an acyclic information-theoretic causal structure (even when that causal structure is not uniquely determined). As stated before, for the results of this paper, we will only require the signalling structure of the network for statements about general unknown networks. The signalling structure can nevertheless allow us to infer something about the underlying causal structure(s) of the network, and we will use this when discussing specific examples of known networks/decompositions. For further details on how information-theoretic causation can be defined, we refer the reader to the literature on quantum causality (e.g., [13, 16, 18, 19, 68, 69], section 2.5 of the thesis [52] and references therein).

Notation 3 (Signalling vs causation). We will use different types of arrows to distinguish between signalling and information-theoretic causation. In consistency with notation from a previous work [2], \mathcal{S} signals to \mathcal{S}' is denoted $\mathcal{S} \longrightarrow \mathcal{S}'$ (where \mathcal{S} and \mathcal{S}' are sets of systems) and S is a direct (information-theoretic) cause of S' will be denoted as $S \rightsquigarrow S'$ (where S and S' are individual systems). For instance, in Fig. 5b, we have $I \rightsquigarrow O_1$ but $\{I\} \not\rightarrow \{O_1\}$, which highlights the distinctions in the edges and the nodes, explicitly in the notation.

C. Compatibility with an abstract causal structure

The following general definitions will allow us to consider whether a signalling structure (operationally obtained from a possibly unknown network of maps) is compatible with different notions of causality.

Definition 12 (Embedding systems in a causal structure). An embedding \mathcal{E} of a set \mathcal{S} of systems in a causal structure \mathcal{G} is a mapping $\mathcal{E} : \mathcal{S} \mapsto \text{Nodes}(\mathcal{G})$ from elements of \mathcal{S} to nodes of \mathcal{G} . For each system $S \in \mathcal{S}$, we will use $S^{\mathcal{E}(S)}$ to denote the system embedded on the node $\mathcal{E}(S)$ of \mathcal{G} , and refer to $S^{\mathcal{E}(S)}$ as the \mathcal{G} -embedded system.

Definition 13 (Embedding networks of CPMs in a causal structure). An embedding of a network \mathfrak{N} of CPMs in a causal structure \mathcal{G} with respect to an embedding \mathcal{E} of the systems $\mathcal{S} := \mathfrak{N}^{\text{sys}}$ in \mathcal{G} is a network $\mathfrak{N}_{\mathcal{G}, \mathcal{E}}$ that is equivalent to \mathfrak{N} up to a relabelling of each system $S \in \mathcal{S}$ to the corresponding \mathcal{G} -embedded system $S^{\mathcal{E}(S)}$. Further, for every $S \rightsquigarrow S' \in \mathfrak{N}^{\text{comp}}$, we require that the embedding satisfies $\mathcal{E}(S) = \mathcal{E}(S')$.

Definition 14 (Compatibility of a signalling structure with a causal structure). We say that a signalling structure \mathcal{G}^{sig} over a set \mathcal{S} of systems is compatible with an embedding \mathcal{E} in a causal structure \mathcal{G} if the following holds.

$$\mathcal{S}_1 \longrightarrow \mathcal{S}_2 \text{ in } \mathcal{G}^{\text{sig}}$$

$$\Downarrow$$

$$\exists S_1 \in \mathcal{S}_1, S_2 \in \mathcal{S}_2 \text{ such that there is a directed path from } \mathcal{E}(S_1) \text{ to } \mathcal{E}(S_2) \text{ in } \mathcal{G}.$$

Once we embed a network in a causal structure by associating systems in the network to nodes in the causal structure, then compatibility allows us to infer a causal influence from a node associated with some $S_1 \in \mathcal{S}_1$ to a node associated with some $S_2 \in \mathcal{S}_2$ whenever \mathcal{S}_1 signals to \mathcal{S}_2 , or equivalently, compatibility ensures that in the absence of any causal influences from the nodes associated with \mathcal{S}_1 to those associated with \mathcal{S}_2 , we cannot have signalling from \mathcal{S}_1 to \mathcal{S}_2 . For example, in the case of information-theoretic causal structures of CPTP maps, a decomposition such as Fig. 5b where there is no causal influence (in this context, no internal wires) from an input system I to an output system O_2 implies (as we can easily check) that I does not signal to O_1 in that decomposition. In the spatio-temporal case, if we have two systems S_I and S_O of a network embedded at space-like separated locations in a spacetime, there are no directed paths between the corresponding locations in the spatio-temporal causal structure and compatibility would require that S_I does not signal to S_O in the network. In Section IV, we will apply these concepts to generally define relativistic causality in a spacetime.

Remark 3. A compatibility condition between cyclic information-theoretic causal models and acyclic graphs, for the purpose of providing a general formulation of relativistic causality principles, was also proposed in recent works [2, 3] (in the context of general theories, not necessarily restricted to quantum theory). The relationships and differences between the current formalism and that of [2, 3] are discussed in Appendix I.

D. Fine-graining of quantum networks

The general formalism laid out thus far allows for networks which can be associated with cyclic information-theoretic causal structures. For instance, imagine a network consisting of two maps $\mathcal{M}_1 : A \mapsto B$ and $\mathcal{M}_2 : B \mapsto A$ with $\mathcal{H}^A \cong \mathcal{H}^B$, such that we compose the in/output systems with the same label to obtain a causal structure with $A \rightsquigarrow B$ and $B \rightsquigarrow A$ as illustrated in Fig. 6a and Fig. 6c⁵. Upon a closer inspection, it may turn out that we can encode the information about this network in a new network where the information about A is now encoded in two systems A_1 and A_2 such that the causal structure over these “fine-grained” systems is acyclic, with $A_1 \rightsquigarrow B$ and $B \rightsquigarrow A_2$ as shown in Fig. 6d.

In Appendix C, we provide illustrative examples of the concept of fine-graining of cyclic networks and their causal structures. The examples highlight (at least) two distinct physical ways in which fine-graining becomes relevant: namely *fine-graining by splitting into subsystems* and *fine-graining through uncertainty in location (relative to an abstract causal structure)*. The former case corresponds to a scenario where a cyclic network involving a 4-dimensional system A , when equivalently representing A in terms of two 2-dimensional subsystems A_1 and A_2 , turns out to be an acyclic network over the subsystems and can therefore be explained by an acyclic fine-grained causal structure. The latter models a situation where the information content of A is encoded either into a system A_1 or a system A_2 , where A_1 and A_2 are some systems (of any dimension) embedded in an abstract causal structure. This would correspond to a scenario where the information content of A is no longer localised in the abstract causal structure, this kind of non-localisation of information can also be due to purely classical uncertainty as illustrated by Example 3 in Appendix C, or it may be quantum, as will be the case in spacetime realisations of certain processes such as the *quantum switch* [27], which we will discuss later in the paper (Section IX). In the following, we define the concept of fine-graining more rigorously and generally for arbitrary cyclic networks.

Definition 15 (Fine-graining of systems). We say that a set \mathcal{S}^f of systems is a fine-graining of a set \mathcal{S} of systems if there exists an injective map $\mathcal{F} : \mathcal{S} \mapsto \Sigma(\mathcal{S}^f)$ that maps each system $S \in \mathcal{S}$ to a subset of systems $\mathcal{F}(S) \subseteq \mathcal{S}^f$ such that $\mathcal{F}(S) \cap \mathcal{F}(S') = \emptyset$ whenever $S \neq S'$. We refer to such a map \mathcal{F} as a systems fine-graining.

Definition 16 (Fine-graining of a CPM). Consider two CPMs \mathcal{M} and \mathcal{M}^f whose in and output systems are related by a systems fine-graining $\mathcal{F} : \text{In}(\mathcal{M}) \cup \text{Out}(\mathcal{M}) \mapsto \Sigma(\text{In}(\mathcal{M}^f) \cup \text{Out}(\mathcal{M}^f))$ such that $\mathcal{F}(S') \subseteq \text{In}(\mathcal{M}^f)$ for all $S' \in \text{In}(\mathcal{M})$ and $\mathcal{F}(S) \subseteq \text{Out}(\mathcal{M}^f)$ for all $S \in \text{Out}(\mathcal{M})$. Then the CPM \mathcal{M}^f is called a fine-graining of the CPM \mathcal{M} with respect to the system fine-graining \mathcal{F} if there exist a pair of CPTPMs $\text{Enc} : \mathcal{L}(\text{In}(\mathcal{M})) \mapsto \mathcal{L}(\text{In}(\mathcal{M}^f))$ and $\text{Dec} : \mathcal{L}(\text{Out}(\mathcal{M}^f))^* \mapsto \mathcal{L}(\text{Out}(\mathcal{M}))$, which we refer to as the *encoder* and the *decoder* respectively, such that the following properties are satisfied. Here, $\mathcal{L}(\text{Out}(\mathcal{M}^f))^*$ is the subspace of $\mathcal{L}(\text{Out}(\mathcal{M}^f))$ defined by the image of $\mathcal{M}^f \circ \text{Enc}$.

1. **Recovery of \mathcal{M} from \mathcal{M}^f :** The CPM \mathcal{M} can be recovered from the CPM \mathcal{M}^f through the encoding and decoding procedure i.e.,

$$\mathcal{M} = \text{Dec} \circ \mathcal{M}^f \circ \text{Enc} \quad (7)$$

2. **Preservation of signalling relations:** For every directed edge $\mathcal{S}_1 \longrightarrow \mathcal{S}_2$ in the signalling structure $\mathcal{G}_{\mathcal{M}}^{\text{sig}}$ of \mathcal{M} , there exists a corresponding directed edge $\mathcal{F}(\mathcal{S}_1) \longrightarrow \mathcal{F}(\mathcal{S}_2)$ in the signalling structure $\mathcal{G}_{\mathcal{M}^f}^{\text{sig}}$ of \mathcal{M}^f , where $\mathcal{F}(\mathcal{S}_i) := \bigcup_{S \in \mathcal{S}_i} \mathcal{F}(S)$, for $i \in \{1, 2\}$.

See Fig. 7 for an illustration of Eq. (7). The above definition only requires that any signalling properties of the coarse-grained (original) map \mathcal{M} must be replicated in the fine-grained map \mathcal{M}^f , but not the converse. Indeed, the fine-grained description can have additional structure that can reveal itself at the level of the individual fine-grained systems, that may not be reproduced when grouping them together as a single coarse-grained system. This is illustrated in the following simple example.

⁵ We would have this causal structure whenever both maps allow signalling from their single input to single output system, which is taken to be the case in this example




FIG. 6: (a) is a cyclic network which is described in Example 2 (in Appendix C). The network can be uniquely associated with a cyclic causal structure (c). The network (a) can be fine-grained to an acyclic network shown in (b) which can be uniquely associated with the acyclic causal structure of (d). The causal structure (d) is then a fine-graining of the causal structure of (c). In this example, all systems are classical and the maps \mathcal{M}_1 and \mathcal{M}_2 implement the deterministic functions f_A and f_B which are described in the text of Example 2. More generally, quantum networks with cyclic causal structures, and the associated fine-grained networks can also be modelled in our formalism.

Example 1 (“Coarse-graining” does not preserve signalling relations). Let \mathcal{M} correspond to a classical stochastic map from an input bit I to an output bit O that discards the input I and deterministically prepares $O = 0$ as the output. Then we immediately have $\{I\} \not\rightarrow \{O\}$ in \mathcal{M} . We now construct a fine-graining of \mathcal{M} where $\mathcal{F}(I) \rightarrow \mathcal{F}(O)$ holds. Consider the map \mathcal{M}^f with the input systems I_1, I_2 and output systems O_1, O_2 (all of which correspond to classical bits) acting as $\mathcal{M}^f : (I_1, I_2) \mapsto (O_1 = I_1 \oplus I_2, O_2 = I_1 \oplus I_2)$ (where \oplus denotes modulo-2 addition). Consider the systems fine-graining $\mathcal{F}(I) = \{I_1, I_2\}$ and $\mathcal{F}(O) = \{O_1, O_2\}$, an encoder that maps the input value $I = 0$ randomly to one of the values $(I_1, I_2) \in \{(0, 0), (1, 1)\}$ and $I = 1$ to one of the values in $(I_1, I_2) \in \{(0, 1), (1, 0)\}$ ⁶, and a decoder that maps the values $(O_1, O_2) \in \{(0, 0), (1, 1)\}$ to $O = 0$ and $(O_1, O_2) \in \{(0, 1), (1, 0)\}$ to $O = 1$ i.e., $O = O_1 \oplus O_2$. In other words, the coarse grained variables encode the parity of the corresponding two fine-grained variables. We can see that \mathcal{M}^f is indeed a fine-graining of \mathcal{M} relative to these encoders and decoders. It satisfies Eq. (7) since \mathcal{M} maps every $I \in \{0, 1\}$ to $O = 0$ while \mathcal{M}^f maps every input to outputs that satisfy $O_1 = O_2$ which will always be decoded to

⁶ Note that the encoder has a single input I and two outputs I_1, I_2 such that $I_1 \oplus I_2 = I$ i.e., it is precisely the map depicted in Fig. 5 with I_1 and I_2 playing the role of O_1 and O_2 .




FIG. 7: Illustration of the first condition of Definition 16, which requires that \mathcal{M} can be recovered from \mathcal{M}^f through an encoding-decoding scheme. We have illustrated the simplified case where \mathcal{M} has a single input system I and a single output system O , but Definition 16 also applies to \mathcal{M} with arbitrary number of in and output subsystems. We can regard \mathcal{M}^f as a fine-graining of \mathcal{M} if this condition is satisfied, and the signalling relations of \mathcal{M} are preserved in \mathcal{M}^f , for the given map \mathcal{M} , the latter condition translates to the requirement that $\{I\} \rightarrow \{O\}$ in \mathcal{M} implies $\{I_1, \dots, I_n\} \rightarrow \{O_1, \dots, O_m\}$ in \mathcal{M}^f .

$O = 0$. It trivially satisfies the second condition of Definition 16 since \mathcal{M} has no non-trivial signalling relations to be preserved. However, we can see that $\{I_1, I_2\} \rightarrow \{O_1, O_2\}$ in \mathcal{M}^f since $\mathcal{M}^f(I_1 = 0, I_2 = 0) \neq \mathcal{M}^f(I_1 = 0, I_2 = 1)$.

Definition 17 (Fine-graining of a network of CPMs). Let \mathfrak{N} and \mathfrak{N}^f be two networks of CPMs such that there exists a systems fine-graining $\mathcal{F} : \mathfrak{N}^{\text{sys}} \mapsto \Sigma(\mathfrak{N}^{f,\text{sys}})$. Then the network \mathfrak{N}^f is called a fine-graining of the network \mathfrak{N} with respect to \mathcal{F} if for every sub-network $\mathfrak{N}_{\text{sub}}$ of \mathfrak{N} , there exists a unique sub-network $\mathfrak{N}_{\text{sub}}^f$ of \mathfrak{N}^f such that the induced CPM $\mathcal{N}_{\text{sub}}^f$ is a fine-graining of the induced CPM \mathcal{N}_{sub} , with respect to the systems fine-graining \mathcal{F} .

With these definitions, the following lemma immediately follows.

Lemma 1. *Fine-graining a network preserves its signalling relations] Let the network \mathfrak{N}^f be a fine-graining of the network \mathfrak{N} with respect to a systems fine-graining \mathcal{F} . Then for every directed edge $\mathcal{S}_1 \rightarrow \mathcal{S}_2$ in the signalling structure $\mathcal{G}_{\mathfrak{N}}^{\text{sig}}$ of \mathfrak{N} , there exists a corresponding directed edge $\mathcal{F}(\mathcal{S}_1) \rightarrow \mathcal{F}(\mathcal{S}_2)$ in the signalling structure $\mathcal{G}_{\mathfrak{N}^f}^{\text{sig}}$ of \mathfrak{N}^f , where $\mathcal{F}(\mathcal{S}_i) := \cup_{S \in \mathcal{S}_i} \mathcal{F}(S)$ for $i \in \{1, 2\}$.*

This lemma follows from the above definitions because the signalling structure of a network \mathfrak{N} is obtained from the signalling relations present in the CPMs induced by each sub-network $\mathfrak{N}_{\text{sub}}$ of \mathfrak{N} . Since the CPM induced by each sub-network $\mathfrak{N}_{\text{sub}}^f$ of the fine-grained network \mathfrak{N}^f must itself be a fine-graining of the corresponding induced CPM of the original network \mathfrak{N} (by Definition 17), it follows from Definition 17 that the signalling relations of each sub-network in \mathfrak{N} is preserved in a corresponding sub-network of \mathfrak{N}^f and therefore that the signalling relations of \mathfrak{N} are preserved in \mathfrak{N}^f .

IV. SPATIO-TEMPORAL CAUSAL STRUCTURES

We model a spacetime in general and order-theoretic terms, without assuming any manifold topology or symmetries, but focusing on its causal structure.

Definition 18 (Fixed acyclic spacetime). A fixed acyclic spacetime corresponds to a partially ordered set \mathcal{T} associated with the order relation \leq . $P \leq Q$ is denoted as $P < Q$ whenever $P, Q \in \mathcal{T}$ are distinct elements. $P < Q$, $P > Q$ and $P \not\star Q$ represent P being in the past of, future of and neither in the past nor future of Q respectively, with respect to this order relation.

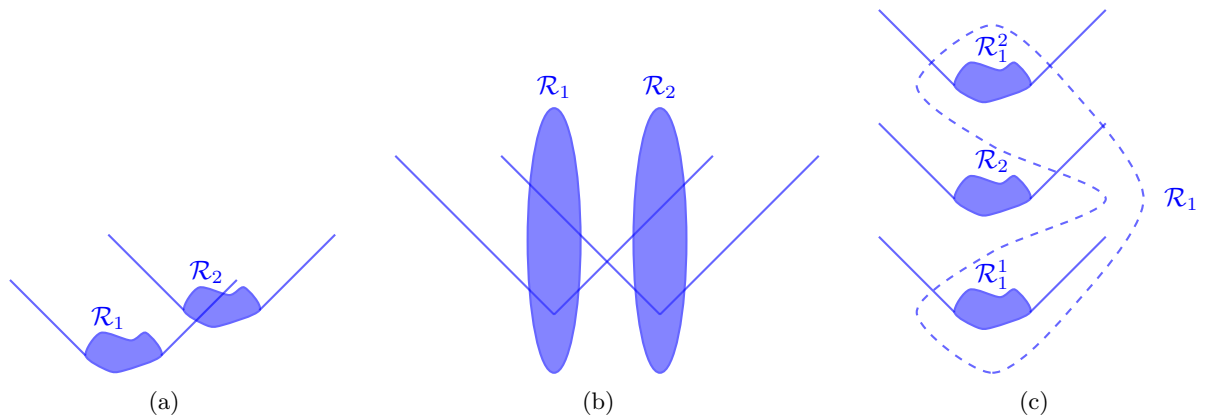


FIG. 8: Examples of region causal structures in Minkowski spacetime, where the 45 degree lines denote forward light-like surfaces. The region causal structure depicted in (a) is acyclic with $\mathcal{R}_1 \xrightarrow{R} \mathcal{R}_2$ (but $\mathcal{R}_2 \not\xrightarrow{R} \mathcal{R}_1$) while the region causal structures depicted in (b) and (c) are cyclic with $\mathcal{R}_1 \xrightarrow{R} \mathcal{R}_2$ and $\mathcal{R}_2 \xrightarrow{R} \mathcal{R}_1$. In (c), it is evident that the region \mathcal{R}_1 is composed of two disjoint regions \mathcal{R}_1^1 and \mathcal{R}_1^2 such that the region causal structure over \mathcal{R}_1^1 , \mathcal{R}_1^2 and \mathcal{R}_2 is acyclic with $\mathcal{R}_1^1 \xrightarrow{R} \mathcal{R}_2 \xrightarrow{R} \mathcal{R}_1^2$. Since the spacetime is acyclic, it is easy to see that (b) can also be partitioned into a set of sub-regions such that the resulting region causal structure is acyclic.

In the remainder of the paper, we will often refer to a fixed acyclic spacetime simply as a spacetime. The partial order is an abstraction of the light-cone structure. In particular, any globally hyperbolic spacetime with a well-defined arrow of time can be described this way. The partial order on spacetime events induces an order relation on spacetime regions as follows.

Definition 19 (Order relation on spacetime regions). Let $\mathcal{R}_1, \mathcal{R}_2 \subseteq \mathcal{T}$ be two distinct subsets of locations (or “regions”) in a spacetime \mathcal{T} . We say that $\mathcal{R}_1 \xrightarrow{R} \mathcal{R}_2$ if there exists at least one pair of distinct locations $P_1 \in \mathcal{R}_1$ and $P_2 \in \mathcal{R}_2$ such that $P_1 < P_2$. More generally, we will refer to a directed graph $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ as a region causal structure of \mathcal{T} if $\text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg}}) \subseteq \Sigma(\mathcal{T})$ and its edges are given by the order relation \xrightarrow{R} induced by the partial order of \mathcal{T} .

Note that \xrightarrow{R} is neither a pre-order nor a partial order relation as it is non-transitive and we can in general have $\mathcal{R} \xrightarrow{R} \mathcal{R}'$ as well as $\mathcal{R} \xrightarrow{R} \mathcal{R}'$ for two spacetime regions (see Fig. 8 for illustrative examples of region causal structures).

Region partitions Choosing a partition of each spacetime region $\mathcal{R} \in \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg}})$ into mutually disjoint sub-regions $\mathcal{R} = \cup_i \mathcal{R}_i$ ($\mathcal{R}_i \cap \mathcal{R}_{i'} = \emptyset$ for all $i \neq i'$) defines a fine-graining $\mathcal{G}_{\mathcal{T}}^{\text{reg},f}$ of the graph $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ (Definition 2), where $\mathcal{G}_{\mathcal{T}}^{\text{reg},f}$ is generated by considering each node $\mathcal{R} \in \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg}})$ and including every sub-region \mathcal{R}_i appearing in the partition of \mathcal{R} as a node of $\mathcal{G}_{\mathcal{T}}^{\text{reg},f}$. The order relation (defining the edges of $\mathcal{G}_{\mathcal{T}}^{\text{reg},f}$) is then induced by the spacetime order relation $<$ as per Definition 19.⁷ While the spacetime \mathcal{T} and regions $\mathcal{R} \subseteq \mathcal{T}$ can in general be infinite sets, we only consider finite partitions where a given region is split into finitely many sub-regions. This is because for any physical protocol, agents’ actions occur only within a finite set of well-defined spacetime regions. In fact, it is often useful to take the spacetime \mathcal{T} to be a finite poset, which represents the finite set of spacetime events which are relevant to the physical protocol that one wishes to model on the spacetime. In such cases, we can partition each region $\mathcal{R} \in \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg}})$ in a region causal structure of \mathcal{T} in terms of the individual spacetime locations in \mathcal{T} that comprise it $\mathcal{R} = \cup_{P \in \mathcal{R}} P$. Consequently, the spacetime \mathcal{T} itself, when represented as a directed acyclic graph corresponds to a fine-graining of any region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ of the spacetime, which we will refer to as the *maximal fine-graining* of $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$. Note that while $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ can in general be cyclic, its maximal fine-graining, with nodes corresponding to elements of \mathcal{T} , would always be acyclic since \mathcal{T} is a partially ordered set. Our results do not rely on the existence of such a maximal fine-graining, but it will often be useful to refer to it in examples and explanations.

⁷ That $\mathcal{G}_{\mathcal{T}}^{\text{reg},f}$ is a fine-graining of the graph $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ follows from Definition 19 which implies that whenever $\mathcal{R} = \cup_i \mathcal{R}_i \xrightarrow{R} \mathcal{R}' = \cup_j \mathcal{R}'_j$, there exist \mathcal{R}_i and \mathcal{R}'_j such that $\mathcal{R}_i \xrightarrow{R} \mathcal{R}'_j$.

V. LINKING INFORMATION-THEORETIC AND SPATIO-TEMPORAL CAUSAL STRUCTURES

Until this point, the spacetime structure is simply an abstract (acyclic) causal structure, which is devoid of any operational meaning before we embed physical systems in it. An agent may assign physical meaning to a spacetime point P when it can be associated with some operational event such as “I received message from Bob at the spacetime location P ”. In our formalism, the in and output systems of our networks of maps model such operational events. Physically realising a network of maps in a spacetime, involves associating spacetime regions with the in and output systems of the network. We formalise these ideas below, by applying the general concepts that we have already defined in the previous sections.

A. Realisations of cyclic quantum networks in fixed acyclic spacetimes

We have already defined what it means to embed systems and networks of CPMs in an abstract causal structure, in Definition 12 and Definition 13. The following two definitions of spacetime embeddings naturally follow from applying these general definitions to region causal structures of a spacetime.

Definition 20 (Spacetime embedding of systems). A spacetime embedding $\mathcal{E} : \mathcal{S} \mapsto \Sigma(\mathcal{T})$ of a set \mathcal{S} of systems in a fixed acyclic spacetime \mathcal{T} assigns a spacetime region $\mathcal{E}(S) \in \Sigma(\mathcal{T})$ to each $S \in \mathcal{S}$. The image of \mathcal{E} defines a region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ of \mathcal{T} where $\mathcal{E}(S)$ for every $S \in \mathcal{S}$ corresponds to a distinct node of $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$. Then \mathcal{E} corresponds to an embedding of \mathcal{S} in the directed graph $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ (according to Definition 12), and can equivalently be viewed as a mapping $\mathcal{E} : \mathcal{S} \mapsto \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg}})$.

It is possible to have $\mathcal{E}(S) = \mathcal{E}(S')$ for distinct systems $S, S' \in \mathcal{S}$, however, these will correspond to distinct nodes in the region causal structure implied by the image of \mathcal{E} . This is helpful for modelling situations where distinct agents act within the same spacetime region on distinct set of systems e.g., S may be Alice’s system and S' may be Bob’s system and Alice and Bob may be situated within the same lab.

Definition 21 (Spacetime embedding of a network of CPMs). A spacetime embedding $\mathcal{E} : \mathcal{S} \mapsto \Sigma(\mathcal{T})$ of the systems $\mathcal{S} := \mathfrak{N}^{\text{sys}}$ of a network \mathfrak{N} induces an embedding $\mathfrak{N}_{\mathcal{G}_{\mathcal{T}}^{\text{reg}}, \mathcal{E}}$ of \mathfrak{N} (cf. Definition 13) in the region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ defined by the image of \mathcal{E} . We will denote $\mathfrak{N}_{\mathcal{G}_{\mathcal{T}}^{\text{reg}}, \mathcal{E}}$ as $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}$ for short.

Spacetime embedded networks and fine-graining Let $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}$ be a spacetime embedding of a network \mathfrak{N} and $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ be the region causal structure defined by the image of the embedding \mathcal{E} . Recall from Section IV that by choosing a partition for each region $\mathcal{R} \in \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg}})$ into disjoint sub-regions $\mathcal{R} = \bigcup_i \mathcal{R}_i$, we can define a fine-graining $\mathcal{G}_{\mathcal{T}}^{\text{reg}, f}$ of $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$. This induces a systems fine-graining $\mathcal{F}(S) = \{S^{\mathcal{R}_i}\}_i$, of the systems $S \in \mathfrak{N}^{\text{sys}}$ (or equivalently, the corresponding $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ -embedded systems) into $\mathcal{G}_{\mathcal{T}}^{\text{reg}, f}$ -embedded systems, where $S^{\mathcal{R}_i}$ labels the system embedded in the region $\mathcal{R}_i \in \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg}, f})$. We can then consider a fine-graining $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}^f$ of the network \mathfrak{N} (or equivalently, the spacetime embedded network $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}$ cf. Definition 21) with respect to each such system fine-graining \mathcal{F} induced by a choice of partition on regions in $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$. In particular, when we maximally fine-grain each region in the image of \mathcal{E} in terms of the individual spacetime points, $\mathcal{E}(S) = \{P \in \mathcal{T}\}_{P \in \mathcal{E}(S)}$, each system S is correspondingly fine-grained to a set of systems $\mathcal{F}_{\text{max}}(S) = \{S^P\}_{P \in \mathcal{E}(S)}$, such that each spacetime location $P \in \mathcal{E}(S)$ becomes associated with a Hilbert space \mathcal{H}^{S^P} .

Definition 22 (Fixed spacetime realisation of a network of CPMs). A realisation of a network \mathfrak{N} in a fixed acyclic spacetime \mathcal{T} with respect to an embedding $\mathcal{E} : \mathfrak{N}^{\text{sys}} \mapsto \Sigma(\mathcal{T})$ is specified by the spacetime embedded network $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}$ along with at least one fine-graining $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}^f$ of $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}$ which is associated with an acyclic region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}, f}$.

If the image of the spacetime embedding \mathcal{E} itself corresponds to an acyclic region causal structure, then the spacetime realisation is fully specified by the spacetime embedded network $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}$ (which is equivalent to the original network \mathfrak{N} , up to the labelling of each system S with the spacetime region it is embedded in). If this is not the case, we can partition the regions in the image of \mathcal{E} until we obtain an acyclic region causal structure and specify a fine-graining $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}^f$ of $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}$ relative to that partitioning of regions, in order to fully specify the spacetime realisation. In particular, any fine-graining $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}^f$ of the network defined relative to the maximal systems fine-graining \mathcal{F}_{max} will always have this property. However, it is also possible to have other non-maximal partitions of the regions that still lead to an acyclic region causal structure, and it would be sufficient to only consider these in the specification of a spacetime realisation.

Remark 4. Notice that the above definition only refers to the acyclicity of the fine-grained region causal structure and does not place any constraints on the causal or signalling structure of the fine-grained network $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}^f$. This is

a minimal requirement which captures the essence of regarding a background spacetime as being acyclic: we can in-principle describe any protocol immersed in an acyclic spacetime in terms of what happens (or does not happen) within a set of spacetime regions which have a well-defined ordering.

B. Relativistic causality

For networks embedded/realised in spacetime as discussed in the previous sub-section, both information-theoretic and spatio-temporal causal orders come into play, and both are defined over the same set of nodes: the spacetime embedded systems which carry both system labels and spacetime region labels. The information-theoretic order is given by the edges \rightsquigarrow , the properties of which we can infer from the signalling order \longrightarrow defined over subsets of spacetime embedded systems. The latter is given by the order relation \xrightarrow{R} on regions. Relativistic causality principles impose constraints on the relations between these informational and spatio-temporal types of orders, which have so far been treated independently in our framework. In particular, we consider the relativistic causality principle of no signalling outside the future light-cone. The following definitions provide a natural sufficient condition for capturing this principle in our framework, as a compatibility condition between \longrightarrow and \xrightarrow{R} . In general there are distinct relativistic causality principles which can be considered, for instance, depending on whether they are formulated using the causal order \rightsquigarrow or signal order \longrightarrow and whether we consider necessary or sufficient conditions [2, 3, 70] (see Appendix I for further details). In this paper, we will focus on the following definitions for capturing relativistic causality, which follows as a special case of Definition 14 of compatibility between a signalling structure and an abstract causal structure, when the latter is a region causal structure of a spacetime.

Definition 23 (Relativistic causality for signalling structures). Let \mathcal{G}^{sig} be a signalling structure over a set \mathcal{S} of systems and \mathcal{E} be an embedding of the systems $S \in \mathcal{S}$ in a region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ associated with a spacetime \mathcal{T} . Then, we say that the signalling structure \mathcal{G}^{sig} satisfies relativistic causality with respect to the embedding \mathcal{E} in the spacetime \mathcal{T} if \mathcal{G}^{sig} is compatible with the causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ (cf. Definition 14).

Then we have the following theorem.

Theorem 1. *[Embedding arbitrary, cyclic signalling relations in spacetime] For every signalling structure \mathcal{G}^{sig} , there exists a fixed acyclic spacetime \mathcal{T} and an embedding \mathcal{E} of \mathcal{G}^{sig} in a region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ of \mathcal{T} that respects relativistic causality.*

Definition 24 (Relativistic causality for spacetime realisations of networks). A spacetime realisation of a network \mathfrak{N} of CPMs is said to satisfy relativistic causality relative to an embedding \mathcal{E} in a spacetime \mathcal{T} only if the signalling structure of each network in the specification of the spacetime realisation satisfies the relativistic causality condition of Definition 23. We will then refer to such a spacetime realisation as a *physical spacetime realisation*.

The above is a minimal definition of what constitutes a physical realisation of a network in a spacetime, it ensures that no interventions that an agent could perform could allow them to signal outside the future lightcone of the spacetime. In the next sections, we apply our general framework and results to address questions relating to the physical spacetime realisations of indefinite causal order quantum processes.

VI. INDEFINITE CAUSAL STRUCTURES IN THE PROCESS MATRIX FRAMEWORK

The concept of indefinite causal structures was independently proposed within different information-theoretic frameworks, to model scenarios where agents can perform quantum operations that are not applied in a definite and acyclic causal order. The causaloid framework of Lucien Hardy [24], the formalism of higher-order computation and maps [27], and the process matrix framework [28] are prominent examples. Here, we consider the process matrix framework (which is known to be mathematically equivalent to higher-order maps), since the question of physicality of indefinite causal structures has most commonly been posed and discussed in this formalism. The framework describes multi-agent information-processing scenarios without assuming a global, acyclic causal order between the operations of the agents, but while assuming the local validity of quantum theory within the lab of each agent.

Local behaviour: quantum instruments. Each agent A acts within their respective local laboratory, which is associated an input Hilbert space \mathcal{H}^{A^I} of dimension d_{A^I} and an output Hilbert space \mathcal{H}^{A^O} of dimension d_{A^O} . A^I and A^O will be used as shorthand to represent the set of all linear operators over \mathcal{H}^{A^I} and \mathcal{H}^{A^O} respectively. The operations that can be performed within the lab of each agent A are described by a set of *quantum instruments*, one for each setting a which is a classical variable specifying the choice of operation performed on A^I . For each setting a , a quantum instrument is a set of CP maps, $\mathcal{J}_a^A = \{\mathcal{M}_{x|a}^A\}_{x=1}^m$ with $\mathcal{M}_{x|a}^A : A^I \mapsto A^O$ and $\sum_{x=1}^m \mathcal{M}_{x|a}^A$ being a CPTP map [28, 31], where x parametrizes the possible outcomes of the operation (for some set $\{1, \dots, m\}$ of possible outcome values). Quantum instruments can be equivalently represented by the set of Choi-Jamiołkowski states $\{M_{x|a}^{A^I A^O} = [\mathcal{I} \otimes \mathcal{M}_{x|a}^A(|\mathbb{1}\rangle\langle\mathbb{1}|)]^T\}_{x=1}^m$, where $|\mathbb{1}\rangle := \sum_j |j\rangle^{A^I} |j\rangle^{A^I}$ and T denotes matrix transposition with respect to the chosen orthonormal basis $|j\rangle^{A^I}$ of \mathcal{H}^{A^I} .

Global behaviour: process matrix. If the scenario consists of N agents $\{A_1, \dots, A_N\}$, the probability $P(x_1, \dots, x_N | a_1, \dots, a_N)$ that the N agents observe the outcomes (x_1, \dots, x_N) for a choice of settings (a_1, \dots, a_N) can be expressed using the Choi-Jamiołkowski representation of the local operations as follows [28, 31],

$$P(x_1, \dots, x_N | a_1, \dots, a_N) = P\left(\mathcal{M}_{x_1|a_1}^{A_1}, \dots, \mathcal{M}_{x_N|a_N}^{A_N}\right) = \text{Tr}\left[\left(M_{x_1|a_1}^{A_1^I A_1^O} \otimes \dots \otimes M_{x_N|a_N}^{A_N^I A_N^O}\right) W\right], \quad (8)$$

for a Hermitian operator $W \in A_1^I \otimes A_1^O \otimes \dots \otimes A_N^I \otimes A_N^O$, known as the *process matrix*. The set of valid process matrices is characterised by the set of all such Hermitian operators W that yield positive normalised probabilities for all possible CP maps $\{\mathcal{M}_{x_k|a_k}^{A_k}\}_{k=1}^N$. This is required to hold also for local operations that can additionally act on ancillary quantum systems, where the ancillas between multiple labs may be entangled [28].

The global behaviour can equivalently be described as a *higher-order map* which maps each set of local operations to a probability distribution [27, 71]. This higher-order map can itself be represented as a CPTP map \mathcal{W} from the input systems A_1^O, \dots, A_N^O (corresponding the outputs of the agents) to the output systems A_1^I, \dots, A_N^I (corresponding to the inputs of the agents), as illustrated in Fig. 9. The process matrix $W = \mathcal{I} \otimes \mathcal{W}(|\mathbb{1}\rangle\langle\mathbb{1}|)$ is then the Choi matrix of the CPTPM \mathcal{W} , where $|\mathbb{1}\rangle$ corresponds to the unnormalised maximally entangled state over two copies of the input Hilbert space of \mathcal{W} [31].⁸ In the rest of this paper, we will refer to the CPTPM \mathcal{W} as the process map, and this will be our object of interest.

Different classes of processes Several different classes of processes have been proposed and studied in the process matrix framework. Here, we review these concepts for the bipartite case and refer the reader to [32, 71] for the general definitions. Firstly, we have the class of *fixed order processes*, which are processes that are compatible with a definite acyclic causal order. In particular, this includes standard causally ordered quantum circuits. For instance a process $W^{A < B}$ is said to be compatible with the fixed order Alice before Bob if Alice's outcome x is independent of Bob's setting b i.e., $P(x|ab) = P(x|a)$. Similarly, $W^{B < A}$ is a fixed order process that is compatible with the order $B < A$ whenever $P(y|ab) = P(y|b)$. Note that a process where neither agent can communicate to the other is compatible with both orders according to this definition. The general definition of a fixed order process is reviewed in Appendix E.

More generally, we can have probabilistic mixtures of fixed order processes, which leads to the set of causally separable processes as well as causally non-separable processes that cannot be cast in this form. These concepts are analogous to separable and non-separable quantum states. Moreover, in analogy to Bell local and Bell non-local quantum states and correlations, we also have the concepts of causal and non-causal processes and correlations. These are reviewed below for the bipartite case, the general definitions can be found in [31, 32].

Definition 25 (Causally separable and non-separable processes [31]). A bipartite process matrix W is said to be *causally separable* iff it decomposes as

$$W = qW^{A < B} + (1 - q)W^{B < A}, \quad (9)$$

for some $q \in [0, 1]$, where $W^{A < B}$ and $W^{B < A}$ are process matrices compatible with the fixed ordering between agents indicated in the respective superscripts. W is said to be *causally non-separable* otherwise.

⁸ Note that the Choi matrix of \mathcal{W} and that of the local operations defined in the previous paragraphs differs by a transpose, this is a choice of convention made in the process matrix framework, that makes the notation and calculations more convenient.

Definition 26 (Causal and non-causal processes and correlations [28]). A bipartite process matrix W is said to be *causal* iff for all choices of local operations, the joint probability $P(xy|ab)$ generated by W (for outcomes x and y and settings a and b of agents A and B respectively) decomposes as follows for some $q \in [0, 1]$

$$P(xy|ab) = qP^{A<B}(xy|ab) + (1 - q)P^{B<A}(xy|ab), \quad (10)$$

where $P^{A<B}$ is a probability distribution compatible with the causal order $A < B$ by disallowing signalling from B to A i.e., $P^{A<B}(x|ab) = P^{A<B}(x|a)$, and similarly for $P^{B<A}$. The process W is called *non-causal* otherwise. Similarly, distributions $P(xy|ab)$ are said to be causal/non-causal depending on whether they can be decomposed as above.

Both causally non-separable processes and non-causal processes are regarded as having indefinite causal orders in this framework, as they are incompatible with an explanation in terms of a fixed acyclic order of the operations, or any convex mixture of such fixed orders. Moreover, non-fixed order but causally separable processes can be regarded as representing classically indefinite orders.

Set-up assumptions The process matrix framework rests on the assumption that each agent acts within a closed laboratory, that interacts with the environment only to let in one input system A^I , apply an operation on it and then let out one output system A^O , with the local events unfolding in that order (even though no global order between the operational events of different agents is assumed)⁹ and the lab being closed to in/output otherwise. These correspond to the closed labs (CL) and local order (LO) assumptions, and they are meant to ensure that each agent acts once and only once. Moreover, free choice (FC) is also assumed: that the agents can freely chose the local operation performed in their lab, that is the setting modelling the choice of operation in each lab has no other causes relevant to the scenario.

Causal inequalities [28] derives a linear inequality constraint on the joint probabilities $P(xy|ab)$ under four assumptions, this is referred to as a *causal inequality* and is shown to be a necessary constraint on causal distributions. The first three assumptions are the set-up assumptions LO, CL and FC of the process matrix framework. The fourth is an additional assumption referred to as *causal structure* (CS) which states that the input and output events of the agents are localised in a fixed partial order such as a background spacetime that prohibits signalling outside the future. Therefore, it is argued that if the correlations violate a causal inequality in a scenario where CL, LO and FC are satisfied, this would be certify a violation of the assumption CS which relates to the existence of a definite causal structure. Hence causal inequalities are commonly regarded as device-independent witnesses of the indefiniteness of causal order. In Section XC, we will analyse the CS assumption in more detail in light of our results to draw more fine-grained insights on the interpretation of causal inequalities and their violation.

VII. INDEFINITE CAUSAL STRUCTURES AS INSTANCES OF CYCLIC QUANTUM NETWORKS

In this section, we show that indefinite causal structures as described in the process matrix formalism can be modelled as instances of the cyclic quantum networks allowed in our framework, and derive some connecting results.

Extended local maps To capture the set of quantum instruments associated with an agent in the process framework as a single CPTP map, we introduce additional in and output systems corresponding to the classical setting and outcome of the agent. Observe that the local operation \mathcal{M}^A of an agent A can be equivalently modelled as a CPTP map from the input systems A^I and A^s to the output systems A^O and A^o , where A^I and A^O are the quantum in/output systems we have seen before, A^s and A^o model the local in/output systems carrying the classical setting a (that specifies the choice of operation to be applied on the input system A^I) and a possible measurement outcome x of the agent. We can encode the classical setting and outcome, a and x as quantum states $|a\rangle$ and $|x\rangle$ in the computational basis. If the map does not implement a measurement but implements a transformation from A^I to A^O depending on a setting choice a on A^s , the output state on A^o will correspond to the deterministic outcome represented by a fixed state $|\perp\rangle$, and the output system A^o can simply be ignored in this case. An outcome is non-trivial if it does not equal \perp . We refer to the map $\mathcal{M}^{A_k} : A_k^I \otimes A_k^s \mapsto A_k^O \otimes A_k^o$ for each agent A_k as the *extended local map* or *extended*

⁹ Note that no spacetime information is explicitly considered but even in the absence of information about the absolute time of occurrence of these operational events, it is in principle possible to ensure that they are ordered in a certain way.

local operation of that agent. From the extended map, we can then recover the action the CP maps $\mathcal{M}_{x|a}^A$ (defined in the process framework) on an input state ρ^{A^I} as follows.

$$\mathcal{M}_{x|a}^A(\rho^{A^I}) := \text{Tr}_{A^O} \left[\left(|x\rangle\langle x|^{A^O} \otimes \mathcal{I}^{A^O} \right) \left[\mathcal{M}^A(|a\rangle\langle a|^{A^S} \otimes \rho^{A^I}) \right] \right]. \quad (11)$$

Notice that $\mathcal{M}_a^A(\cdot) := \sum_x \mathcal{M}_{x|a}^A(\cdot) = \text{Tr}_{A^O} \left[\mathcal{M}^A(|a\rangle\langle a|^{A^S} \otimes (\cdot)) \right]$ is then a CPTP map from A^I to A^O , since we sum over a complete set of outcomes, $\sum_x |x\rangle\langle x| = \mathbb{1}$. We will refer to \mathcal{M}_a^A as a *fixed local map* as opposed to the extended local map \mathcal{M}^A , since it corresponds to the local map of agent A associated with plugging in a fixed choice of setting a .

The process network and joint probabilities The action of the process matrix on the local operations can be equivalently understood as a loop composition between the process map \mathcal{W} and the extended local maps $\{\mathcal{M}^{A_k}\}_{k=1}^N$, as shown in Figure 9. Let us denote the in/outputs of the local maps as A_k^I, A_k^O and corresponding systems of \mathcal{W} as A_k^I, A_k^O (note that the prime is associated with systems that are inputs of the corresponding map). Then for an N partite process matrix W , this action generates the network $\mathfrak{N}_{\mathcal{W},N}$ where

$$\begin{aligned} \mathfrak{N}_{\mathcal{W},N}^{\text{maps}} &:= \mathcal{W} \cup \{\mathcal{M}^{A_k}\}_{k=1}^N, \\ \mathfrak{N}_{\mathcal{W},N}^{\text{comp}} &:= \{A_k^I \hookrightarrow A_k^I, A_k^O \hookrightarrow A_k^O\}_{k=1}^N. \end{aligned} \quad (12)$$

We refer to the network $\mathfrak{N}_{\mathcal{W},N}$ as the *process network* associated with the process matrix W (or process map \mathcal{W}). The induced map of this network will be denoted as $\mathcal{N}_{\mathcal{W},N} := \mathcal{P}_{\mathcal{W}}$ as it encodes the probabilities $P(x_1, \dots, x_N | a_1, \dots, a_N)$ and the number of parties N is usually clear from context. Notice that the induced map $\mathcal{P}_{\mathcal{W}}$ obtained in this way has only classical input systems ($\{A_k^S\}_k$ carrying the setting choices of all agents) and classical output systems ($\{A_k^O\}_k$ carrying the outcomes of the agents).

If we input a choice $\{a_k\}_k$ for settings to the map $\mathcal{P}_{\mathcal{W}}$ and post-select on a set $\{x_k\}_k$ of outcomes on the corresponding output, we get the joint probability of obtaining those outcomes given those setting choices as the success probability of that post-selection. Explicitly, given an N -partite process map \mathcal{W} and a set $\{\mathcal{M}^{A_k}\}_{k=1}^N$ of extended local maps, the joint probability of obtaining a set $\{x_k\}_k$ of outcomes given a choice $\{a_k\}_k$ of the settings is obtained as follows.

$$P(x_1, \dots, x_N | a_1, \dots, a_N) = \frac{\text{Tr} \left[\Pi_{x_1} \otimes \dots \otimes \Pi_{x_N} \left(\mathcal{P}_{\mathcal{W}}(|a_1\rangle\langle a_1| \otimes \dots \otimes |a_N\rangle\langle a_N|) \right) \right]}{\sum_{x_1, \dots, x_N} \text{Tr} \left[\Pi_{x_1} \otimes \dots \otimes \Pi_{x_N} \left(\mathcal{P}_{\mathcal{W}}(|a_1\rangle\langle a_1| \otimes \dots \otimes |a_N\rangle\langle a_N|) \right) \right]}, \quad (13)$$

where the projector $\Pi_{x_k} = |x_k\rangle\langle x_k|$ projects the state on the system A_k^O to $|x_k\rangle$. The following result shows that this reproduces the probability rule of the process matrix framework.

Lemma 2. *[Probabilities from loop composition] For every process map \mathcal{W} , the joint probability distribution obtained through the induced map $\mathcal{P}_{\mathcal{W}}$ of the network $\mathfrak{N}_{\mathcal{W},N}$ as in Equation (13) is equivalent to that obtained in the process matrix framework through Equation (8). Moreover, W being a valid process matrix guarantees that the denominator of Equation (13) equals unity.*

Notice that for a general map \mathcal{W} , the denominator is needed to ensure that we get normalised probabilities, since a map such as $\mathcal{P}_{\mathcal{W}}$ formed by loop composition of CPTP maps could in general be trace decreasing and considering the numerator alone may not result in a valid normalised distribution. The above result shows that in the special case where the Choi matrix of \mathcal{W} is a valid process matrix, the numerator alone defines a normalised probability distribution.

Sub-networks and reduced processes The network $\mathfrak{N}_{\mathcal{W},N}$ involves composing the N -partite process map \mathcal{W} together with the local operations of all N agents, the sub-networks of $\mathfrak{N}_{\mathcal{W},N}$ include networks $\mathfrak{N}_{\mathcal{W},l < N}$ where \mathcal{W} is composed with the local operations of l out of N agents, with $l < N$. Without loss of generality, we take these to be the first l agents, $\{A_1, \dots, A_l\}$, then the maps and compositions in the network $\mathfrak{N}_{\mathcal{W},l < N}$ are defined exactly as in Eq. (12) but with the index k going from 1 to l (instead of N).




FIG. 9: The process matrix W can be modelled as a higher-order transformation which, given a set of CP (but possibly trace decreasing) maps $\mathcal{M}_{x|a}^A$ and $\mathcal{M}_{y|b}^B$ (one for each agent), maps them to a probability distribution $P(xy|ab)$ given by Equation (8). This is illustrated in the top left. This can be equivalently viewed in terms of a network formed by the loop composition of the process map W with the extended local operations of the agents as shown on the top right, which models their classical setting and outcome as additional in and output wires. The induced map of the network formed in this way is denoted as \mathcal{P}_W (bottom middle), it is a map with classical in and outputs that encodes the joint probability distribution $P(xy|ab)$ as shown in Lemma 2.

In the following lemma, we relate the sub-networks $\mathfrak{N}_{W,l < N}$ with the notion of a reduced process matrix [31]. Given an N -partite process matrix W and a CPTP map $\mathcal{M}_{a_j}^{A_j}$ (associated with a specific setting choice a_j) for the j^{th} agent with Choi matrix $M_{a_j}^{A_j^I A_j^O}$, the reduced process matrix [31] for the remaining $N - 1$ agents is given as

$$\overline{W}(M_{a_j}^{A_j^I A_j^O}) := \text{Tr}_{A_j^I A_j^O} \left[\left(\mathbb{1}^{A_1^I A_1^O} \otimes \dots \otimes M_{a_j}^{A_j^I A_j^O} \otimes \dots \otimes \mathbb{1}^{A_N^I A_N^O} \right) \cdot W \right], \quad (14)$$

where $\mathbb{1}^{A_i^I A_i^O}$ is an identity matrix on $A_i^I \otimes A_i^O$. Analogously, given a fixed local operation (which is a CPTPM) for each of the first $l < N$ agents, the reduced process $\overline{W}(M_{a_1}^{A_1^I A_1^O}, \dots, M_{a_l}^{A_l^I A_l^O})$ associated with the remaining $N - l$ can be defined.

Lemma 3. [Sub-networks and reduced processes] For any N -partite process map W , the sub-network $\mathfrak{N}_{W,l < N}$ of the process network formed by composing W with some set of fixed local operations $\{\mathcal{M}_{a_k}^{A_k}\}_{k=1}^l$ for the first $l < N$ agents

is a CPTP map whose Choi matrix is the reduced process matrix $\overline{W}(M_{a_1}^{A_1^I A_1^O}, \dots, M_{a_i}^{A_i^I A_i^O})$.

These results lead to the following useful theorem, that every valid process network is a network of CPTP maps.

Theorem 2. *The network $\mathfrak{N}_{\mathcal{W}, N}$ formed the composition of any valid process map \mathcal{W} with a set of extended local operations $\{\mathcal{M}^{A_k}\}_{k=1}^N$ is a network of CPTP maps.*

These results enable us to apply our framework to process matrices and infer their properties using the observable signalling relations. Recall that the signalling relations between sets of quantum systems in a network is defined through the induced CPTPMs of its sub-networks. On the other hand, we can also consider signalling at the level of the probabilities, by considering whether the outcome of a set of agents is correlated with a freely chosen setting of another agent. Say $N = 2$ with the two agents being A and B . One notion of signalling from A to B is at the level of the probabilities, which would say $P(y|ab) \neq P(y|b)$. This corresponds to signalling at between the classical systems A^s and B^o of the network. On the other hand, we also can define signalling from A to B at the level of the quantum systems of the process map \mathcal{W} (which corresponds to a sub-network of the process network) by considering whether $\{A^O\}$ signals to $\{B^I\}$ in \mathcal{W} . Does $P(y|ab) \neq P(y|b)$ imply $\{A^O\}$ signals to $\{B^I\}$ in \mathcal{W} and vice-versa? In Appendix F, we show that this is indeed the case whenever W is a valid process matrix. More generally we prove that for any N -partite process, there is an equivalence between signalling at the level of the probabilities over classical settings and outcomes and signalling at the level of the corresponding quantum systems in the network. An upcoming work [72] shows with a counter example, that when \mathcal{W} is a CPTP map whose Choi matrix is not a valid process matrix, $P(y|ab) \neq P(y|b)$ does not in general imply $\{A^O\}$ signals to $\{B^I\}$ in \mathcal{W} . Therefore \mathcal{W} being associated with a valid process matrix is sufficient for this equivalence to hold in the corresponding (possibly cyclic) network. Whether it is also a necessary condition, is left as an interesting open question for future work.

Indefinite causality and cyclicity of compatible causal structures To conclude this section, we present the following theorem that connects the non-fixed order processes to cyclicity of causal structures compatible with the corresponding network, whose proof interestingly relies on the correspondence between the two notions of signalling mentioned above.

Theorem 3. *A process matrix W is not a fixed order process if and only if the signalling structure of the corresponding network $\mathfrak{N}_{\mathcal{W}, N}$ certifies the cyclicity of any causal structure \mathcal{G} that it is compatible with.*

Here, by certifying the cyclicity of \mathcal{G} , we mean that compatibility of the signalling structure with \mathcal{G} will imply the presence of at least one directed cycle in \mathcal{G} . This theorem has implications both for causal inference in cyclic quantum causal models, as well as for physical processes that can be implemented, consistently with relativistic causality in a spacetime. These different implications can be derived based on what \mathcal{G} represents, for causal inference it would be the information-theoretic causal structure of the network and for relativistic causality, it would be the causal structure defined through the light-cones of a spacetime. The next section focuses on the implications of this result for spacetime realisations of processes.

VIII. NO-GO RESULTS FOR SPACETIME REALISATIONS OF QUANTUM PROCESSES

Applying our general framework, we now establish and discuss a number of no-go results for spacetime realisations of indefinite causal order processes.

A. Localisation of information in spacetime

The following no-go theorem follows from the general result we have shown in Theorem 3. It tells us that three, assumptions of physical interest and relevance cannot be simultaneously satisfied. The first assumption requires that the process that we wish to realise in a spacetime is not of fixed order, which is necessary for it to be considered an indefinite causal structure in the process matrix framework. The second assumption ensures that the realisation is physically meaningful, i.e., satisfies relativistic causality in the given spacetime embedding of the process network. The third assumption concerns the embedding, and captures that the systems of the process network are sufficiently localised in the spacetime, an assumption that is usually satisfied when embedding typical abstract quantum information protocols in spacetime.

Theorem 4. *[No-go theorem for spacetime realisations of processes] No realisation (Definition 22) of an N -partite process network $\mathfrak{N}_{\mathcal{W}, N}$ in a fixed spacetime structure \mathcal{T} (Definition 18) with respect to a spacetime embedding \mathcal{E} can simultaneously satisfy the following three assumptions.*

1. W is not a fixed order process (Definition 28).
2. The spacetime realisation satisfies relativistic causality (cf. Definition 24).
3. The region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ induced by the embedding \mathcal{E} (Definition 21) is acyclic.

A special case of an embedding that respects condition 3 is one which assigns to each system S of the network $\mathfrak{N}_{\mathcal{W},N}$ a single spacetime location $\mathcal{E}(S) \in \mathcal{T}$, since the region causal in this case is a subgraph of \mathcal{T} which will be acyclic since \mathcal{T} is a poset. This gives the following corollary regarding the impossibility of perfectly localising systems in the spacetime.

Corollary 1. [Spacetime localisation] *For every non-fixed order process \mathcal{W} , it is impossible to implement the corresponding network $\mathfrak{N}_{\mathcal{W},N}$ in a fixed spacetime consistently with relativistic causality, through an embedding that localises all the in/output systems in the spacetime.*

All the above statements hold irrespective of the choice of reference frame used to describe the spacetime, since they only depend on the order relation between spacetime points which is an agent/frame independent notion according to Definition 18 (as is the case also in special relativity theory). We can also obtain a frame-dependent statement by considering the spatial and temporal coordinates of all the spacetime locations from the perspective of a single agent. For this, we must first add some more structure to our definition of spacetime to include details about spacetime co-ordinates.

Definition 27 (spacetime co-ordinates and time localisation). Given a spacetime structure \mathcal{T} (cf. Definition 18), each agent A can express every location $P \in \mathcal{T}$ in terms of a spatial co-ordinate $\vec{r}_P^A \in \mathbb{R}^n$ and temporal co-ordinate $t_P^A \in \mathbb{R}$ as (\vec{r}_P^A, t_P^A) . Moreover, we require that whenever $P < Q$ for some $P, Q \in \mathcal{T}$, $t_P^A < t_Q^A$ for all agents A i.e., agents always agree on the order relations even if they may disagree on the co-ordinate assignment. Then, we say that a spacetime region $\mathcal{R} \subseteq \mathcal{T}$ is time localised with respect to an agent A if $t_P^A = t_{P'}^A$ for any $P, P' \in \mathcal{R}$.

We then obtain the following corollary.

Corollary 2. [Time localisation in a global frame] *No realisation of an N -partite process network $\mathfrak{N}_{\mathcal{W},N}$ in a fixed spacetime structure \mathcal{T} with respect to a spacetime embedding \mathcal{E} can simultaneously satisfy the following three assumptions.*

1. W is not a fixed order process (cf. Definition 28).
2. The spacetime realisation satisfies relativistic causality (cf. Definition 24).
3. The region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ induced by \mathcal{E} is such that there exists an agent A from whose perspective every region $\mathcal{R} \in \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg}})$ is time-localised (cf. Definition 27).

The above theorem requires all regions to be time-localised when described in the single global frame of A . A weaker requirement would be to require each spacetime region associated to an in/output system in the process network is time-localised from the perspective of the agent associated with that in/output system. However, the theorem no longer holds when replacing the third condition with this weaker requirement. To illustrate this, in Section IX D, we propose a protocol for realising the quantum switch (a causally non-separable process) in Minkowski spacetime where the two agents in relative inertial motion perceive their respective in/output systems to be time-localised in their own frame.

Remark 5. A purely information-theoretic notion of *time-delocalised subsystems* was previously introduced [54], and the realisation of processes in terms of such subsystems has been studied. While the broad message of Corollary 2 accords with the intuition suggested by these works, that physical realisations of indefinite causal structure processes are linked to the non-localisation of information in time, the definitions and assumptions backing the mathematical model of this information non-localisation are different in our framework as compared to this previous work. In particular, as we show in the next section, even though information contained in the systems of the original process are not necessarily localised in (space)time (in the sense defined here, see Theorem 4 and Corollary 1), spacetime realisations of any process satisfying relativistic causality in a fixed spacetime will ultimately be compatible with a definite acyclic fine-grained causal structure. This leads to different conclusions and implications for experimental realisations of the quantum switch, than those drawn in the framework of [54], as we discuss in Section IX C.

B. Existence of a fine-grained explanation in terms of a fixed order process

Theorem 4 permits the realisation of networks associated with non-fixed order processes consistently with relativistic causality, as long as the associated region causal structure is cyclic. The results of this section concern the fine-grained description of such realisations, showing that this description will ultimately be associated with that of a process with a fixed acyclic causal order over a larger number of operations. This generalises the intuition from the the classical demand and price example (Fig. 1) to cyclic quantum causal structures described by process matrices.

To simplify the presentation of this theorem, we incorporate a natural and physically motivated property, which is satisfied in experiments where the devices implementing the operations have a fixed (non-zero) processing time for converting in to outputs. More generally, this can be formalised as a correspondence between the spacetime regions assigned to each agent's in and output quantum systems. We say that two regions \mathcal{R}_1 and \mathcal{R}_2 which are the nodes of some region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ have a correspondence from \mathcal{R}_1 to \mathcal{R}_2 if for any fine-graining $\mathcal{G}_{\mathcal{T}}^{\text{reg},f}$ of $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$, through which \mathcal{R}_1 maps to a set of sub-regions $\{\mathcal{P}_1, \dots, \mathcal{P}_n\} \subseteq \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg},f})$, \mathcal{R}_2 maps to a corresponding set of sub-regions $\{\mathcal{Q}_1, \dots, \mathcal{Q}_n\} \subseteq \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg},f})$ such that $\mathcal{P}^i \xrightarrow{R} \mathcal{Q}^i$ for all $i \in \{1, \dots, n\}$. Then we say that the regions \mathcal{R}_1 and \mathcal{R}_2 have cardinality n relative to this fine-graining, and denote it as $|\mathcal{R}_1| = |\mathcal{R}_2| = n$, when the fine-graining is clear from context. For any given spacetime realisation of the process network $\mathfrak{N}_{\mathcal{W},N}$, $|\mathcal{R}|$ will denote the cardinality of \mathcal{R} relative to the fine-graining $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}^f$ in the specification of the realisation (Definition 22). We will require that the inputs A_k^s and A_k^I of each agent A_k are assigned a spacetime region $\mathcal{R}^{A_k^I}$ and the outputs A_k^O and A_k^o are assigned a spacetime region $\mathcal{R}^{A_k^O}$ such that there is a correspondence from the input to the output region. A physical example of such a correspondence would be a situation where Alice may receive inputs in a region \mathcal{R}^{A^I} consisting of the spacetime locations $(\vec{x}, t_1), \dots, (\vec{x}, t_n)$ and produce outputs in a region \mathcal{R}^{A^O} consisting of the locations $(\vec{x}, t_1 + \delta), \dots, (\vec{x}, t_n + \delta)$ for some fixed $\delta > 0$. For the purpose of the following results, we consider spacetime realisations of processes where such an input-output correspondence is satisfied. Note that this still permits general scenarios where agents' receive/send quantum messages at a superposition of different times (as we will see, for the quantum switch in Section IX C).

Theorem 5. *Any spacetime realisation of a process network $\mathfrak{N}_{\mathcal{W},N}$ in a spacetime \mathcal{T} relative to an embedding \mathcal{E} that satisfies relativistic causality will admit an explanation in terms of a fine-grained network $\mathfrak{N}_{\mathcal{W}^f,M}$ which will be described by a fixed order and therefore causally separable process W^f associated with $M \geq N$ (possibly communicating) agents. Here M is given in terms of the cardinality of the associated input spacetime regions $M = \sum_{k=1}^N |\mathcal{R}^{A_k^I}|$.*

We clarify what we mean by ‘‘possibly communicating agents’’ in the above theorem. In a usual process matrix involving at least two agents A_1 and A_2 , the operations performed by A_1 and A_2 are taken to be in product form $\mathcal{M}^{A_1} \otimes \mathcal{M}^{A_2}$ with any communication between the agents going through the process. In general, plugging in operations on $A_1^I \otimes A_2^I$ which are not of this form into a process matrix can result in non-normalised probabilities. In the above fine-grained description, we allow for situations where non-product operations can be plugged into the fine-grained process. This models situations where A_1 and A_2 can be communicating, as long as the communication respects relativistic causality. We show in the proof of the theorem that the fine-grained process W^f will necessarily be such that it will lead to valid probabilities even when acting on these more general set of operations. Allowing for such operations at the level of the information-theoretic description is similar in spirit to multi-round processes introduced in [73], a formalism that generalises the process matrix framework to scenarios where agents can act in multiple rounds while carrying along a local memory from previous rounds. With these results, we immediately obtain the following corollary.

Corollary 3. *No fixed spacetime realisation of a process network $\mathfrak{N}_{\mathcal{W},N}$ in a spacetime \mathcal{T} relative to an embedding \mathcal{E} can satisfy the following conditions simultaneously*

1. *The spacetime realisation satisfies relativistic causality (cf. Definition 24)*
2. *The spacetime realisation of the network does not admit any explanation in terms of a causally separable process that is consistent with the perspectives of all the agents involved.*

In the above, perspectives of agents refer to their description of the spacetime realisation relative to their classical reference frames which they use for specifying their co-ordinate system in the fixed spacetime, which satisfies the natural properties detailed in Definition 27. Further discussions on the physical implications of these results and the operational meaning of the fine-grained process can be found in Section IX C. There, we apply these results to the special case of the quantum switch process and its experimental realisations and also discuss related arguments in the previous literature which were made for this case.

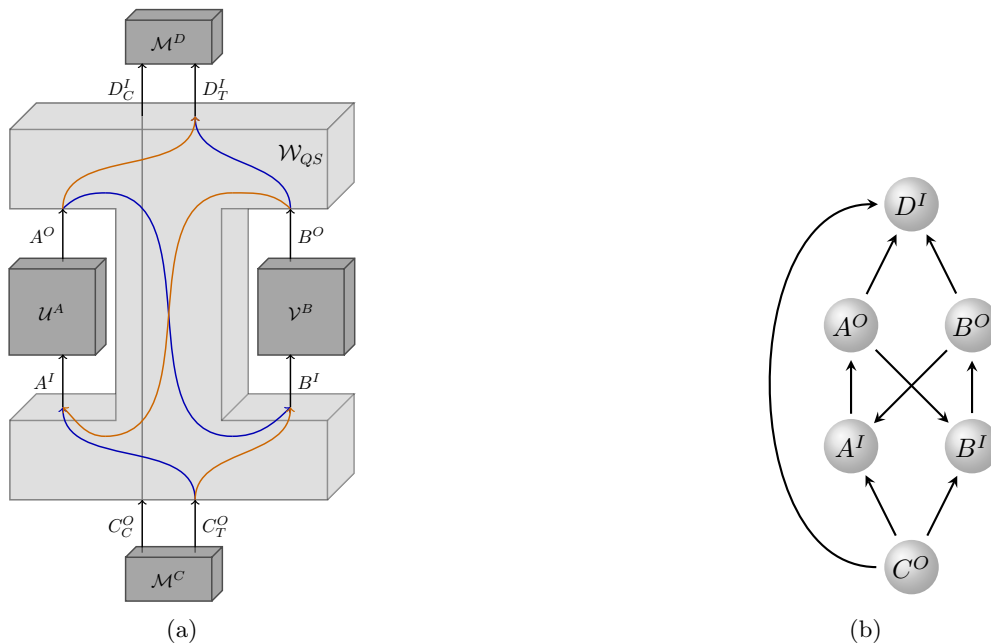


FIG. 10: (a) The quantum switch can be described as a 4-partite process matrix, where agent C is in the past of all others and their local operation prepares the control and target systems in the initial state given on the right hand side of Eq. (15), with C_C^O and C_T^O playing the role of C and T . The agents A and B apply their local operations $\mathcal{U}^A : A^I \mapsto A^O$ and $\mathcal{V}^B : B^I \mapsto B^O$ respectively, where the process map \mathcal{W}_{QS} sends the target system to A or B 's lab in an order that depends on the control. And finally, an agent D in the future of all others receives the target from A or B and control directly from C (gray path) which will be in the final state given in the left hand side of Eq. (15), with D_C^I and D_T^I playing the role of C and T . D can then perform operations/measurements on this state. The process matrix, W for the quantum switch represents a controlled superposition of the orders $C < A < B < D$ (blue path) and $C < B < A < D$ (orange path) of operations on the target system. (b) The causal structure associated with the network formed by the composition of the process map \mathcal{W}_{QS} with non-trivial local operations \mathcal{U}^A and \mathcal{V}^B is cyclic. \mathcal{W}_{QS} is known to be a unitary, thus the causal structure depicted here coincides with the signalling structure over these systems [59]. Here we have grouped together C 's outputs as C_O and D 's inputs as D^I .

IX. CAUSALITY IN THE QUANTUM SWITCH

To illustrate our general formalism and results, we apply it to the example of the quantum switch (QS), which is a particularly popular indefinite causal order process that has been claimed to be experimentally realised in Minkowski spacetime [43, 45, 74]. QS corresponds to a unitary, causally non-separable process, and we can therefore make stronger statements for this case than Theorem 4 which applies to general processes. This clarifies the physical interpretation of such experiments which have been long debated, while consistently unifying the information-theoretic and spatio-temporal pictures on the causal structures involved. Further, we propose a new realisation of QS in Minkowski spacetime, with agents in relative inertial motion, where the agents cannot distinguish the order in which they act even if they were to measure the local time at which they perform their operation. We explicitly show that this too admits the same physical interpretation, in terms of a definite and acyclic fine-grained causal structure, once we account for both spatial and temporal degrees of freedom.

A. The QS process

The quantum switch (QS) [27] was originally defined as a higher-order transformation which takes as input two qudit quantum channels (for simplicity, take them to be unitaries \mathcal{U}^A and \mathcal{V}^B), and outputs a quantum channel that applies \mathcal{U}^A and \mathcal{V}^B on a target system in a quantum controlled superposition of orders. Explicitly, the final channel obtained from the action of the higher-order transformation on \mathcal{U}^A and \mathcal{V}^B implements the following transformation on an initial state of a control qubit C and a target qudit T given on the left (where α and β are arbitrary normalised




FIG. 11: Classical and quantum spacetime realisations of the quantum switch: (a) This is a schematic of a linear optical experimental implementation of QS that can be realised in Minkowski spacetime, reproduced from [37]. Here, the control qubit is encoded in the polarisation of a photon and the target qubit is encoded in a different degree of freedom of the same photon e.g., angular momentum modes. A horizontally polarized photon is transmitted by all the polarizing beam splitters (PBSs) and takes the path where the unitary \mathcal{U}^A is applied first and then \mathcal{V}^B while a vertically polarized photon is reflected by all PBSs and follows the path where \mathcal{V}^B is applied before \mathcal{U}^A . The unitaries \mathcal{U}^A and \mathcal{V}^B act only on the target degrees of freedom. This was experimentally realised in several works, such as [43, 74]. (b) This is a theoretical proposal [26] for implementing the QS process using a quantum superposition of gravitating masses, which results in a superposition of spacetime structures. Here, Alice and Bob are each in possession of their own clock C_A and C_B which are initially synchronised. A gravitating mass is prepared in a quantum superposition of macroscopically distinguishable spatial configurations depending on the state of a control qubit. If the control is in the state $|0\rangle$, the mass is placed closer to Bob such that the clock C_B ticks slower than C_A due to gravitational time dilation enabling Alice to send the target system to Bob at a proper time $t_A = 3$, such that it is received by Bob at his proper time $t_B = 3$. This mass configuration is labelled as $\kappa_{A<B}$. If the control is in the state $|1\rangle$, the mass is placed closer to Alice, enabling Bob to send the target system to Alice at $t_B = 3$ such that it is received by Alice at $t_A = 3$. This mass configuration is labelled as $\kappa_{B<A}$. Here, irrespective of the control, both agents receive the target system in their lab at the same local time. The co-ordinate axes are in the frame of a distant observer, whose time intervals (small gray dots) are unaffected by the mass configurations.

amplitudes and $|\psi\rangle^T$ is an arbitrary qudit state).

$$(\alpha|0\rangle^C + \beta|1\rangle^C) \otimes |\psi\rangle^T \mapsto \alpha|0\rangle^C \otimes \mathcal{V}^B \mathcal{U}^A |\psi\rangle^T + \beta|1\rangle^C \otimes \mathcal{U}^A \mathcal{V}^B |\psi\rangle^T \quad (15)$$

It is important that each of \mathcal{U}^A and \mathcal{V}^B are applied no more than once, else, the above transformation can be easily simulated using causally ordered quantum circuits, using multiple queries to the operations [27]. In the process matrix framework, QS can be modelled as a 4-partite process W associated with a cyclic causal structure, as shown in Fig. 10. The closed labs assumption ensures that each agent in the theoretical process matrix description acts exactly once. Further details on the process matrix description of QS are reviewed in Appendix E.

Realisations of QS in different physical regimes Figure 11 provides an overview of two proposed realisations of the quantum switch transformation, where it has been argued that there is a sense in which the operations are queried no more than once each. The former corresponds to table-top optical setups in Minkowski spacetime. The latter corresponds to a theoretical proposal for realising the quantum switch through a quantum superposition of gravitating masses [26], and involves a quantum indefinite spacetime structure.

B. No go result for QS

We now derive a stronger version of our general no-go result, Theorem 4 for the particular case of QS. In the following, we will call a fixed local operation $\mathcal{M}_{a_k}^{A_k} : A_k^I \mapsto A_k^O$ of some agent A_k non-trivial if $\{A_k^I\} \rightarrow \{A_k^O\}$ in $\mathcal{M}_{a_k}^{A_k}$. For this result, we will consider the network $\mathfrak{N}_{\mathcal{W}_{QS}, U, V}$ which is formed by the composition of the process map \mathcal{W}_{QS} of the quantum switch (which has In = $\{C_C^O, C_T^O, A^O, B^O\}$ and Out = $\{D_C^I, D_T^I, A^I, B^I\}$) with the non-trivial and fixed local operations $\mathcal{U}^A : A^I \mapsto A^O$ and $\mathcal{V}^B : B^I \mapsto B^O$ of Alice and Bob, obtained by loop composing the systems with the same labels. This is distinct from the process network considered in the previous sections, as the local maps \mathcal{U}^A




FIG. 12: Fine-graining of QS associated with its spacetime realisation: (a) A schematic of the fine-grained network associated with typical spacetime realisations of QS. It is associated with a 6-partite process \mathcal{W}_{QS}^f which is a fixed order process but implements, for arbitrary unitaries \mathcal{U}^A and \mathcal{V}^B , the QS transformation of Eq. (15) from the control and target received on C^O to corresponding systems in D^I when Alice and Bob (who now act at two distinct times t_1 and $t_2 > t_1$) apply the same unitary at both times. More detailed descriptions and diagrams of this network, which describe the maps in the internal decomposition of \mathcal{W}_{QS}^f can be found in Appendix G. (b) The causal structure of the fine-grained network, which is definite and acyclic.

and \mathcal{V}^B here correspond to a fixed choice of setting and are not the same as extended local maps which include all possible setting choices. In the rest of the paper, whenever we refer to the process network of QS, we mean $\mathfrak{N}_{\mathcal{W}_{QS}, U, V}$ and not the network with the extended maps.

Lemma 4. [No-go result for the quantum switch] Let $\mathfrak{N}_{\mathcal{W}_{QS}, U, V}$ be the network describing the action of \mathcal{W}_{QS} on any two non-trivial local operations $\mathcal{U}^A : A^I \mapsto A^O$ and $\mathcal{V}^B : B^I \mapsto B^O$, where \mathcal{W}_{QS} is the process map associated with the process matrix W_{QS} of the quantum switch. Then any realisation of this network in a fixed spacetime \mathcal{T} with respect to an embedding \mathcal{E} cannot simultaneously satisfy both of the following assumptions

1. The spacetime realisation satisfies relativistic causality (cf. Definition 24).
2. The region causal structure over the nodes $\mathcal{E}(S)$ for $S \in \{A^I, A^O, B^I, B^O\}$ is acyclic.

The first assumption of Theorem 4 is automatically satisfied here as \mathcal{W}_{QS} is a causally non-separable and hence, is not a fixed order process. Moreover, this is a stronger statement than that of Theorem 4 applied to this process because Theorem 4 deals with the extended local maps (which include all possible setting choices) while the above statement follows for any fixed choice of settings for A and B associated with non-trivial local operations. For a

general process, the analogous statement with fixed operations may not hold.¹⁰

C. Implications for experimental realisations

We first state a corollary of Lemma 4 and Corollary 3 which applies to experimental realisations of QS in a fixed spacetime. We then describe a fine-grained network associated with such a realisation, and discuss in detail, the physical interpretation of QS experiments in Minkowski spacetime offered by our results.

Corollary 4. *Any experiment realising the action of the 4-partite process map \mathcal{W}_{QS} on non-trivial local operations \mathcal{U}^A and \mathcal{V}^B without violating relativistic causality in a fixed spacetime will admit an explanation in terms of a fine-grained network with the following properties*

1. *It is a network formed by an M -partite process map \mathcal{W}_{QS}^f with $M > 4$ together with operations of the M parties.*
2. *\mathcal{W}_{QS}^f is a fixed order process compatible with an acyclic causal structure over the M parties.*
3. *The fine-grained description is consistent with the perspectives of all agents involved.*

These results prove the existence of a fine-grained description with the desired properties. For illustration, we present in detail a particular fine-graining of QS that describes its spacetime realisations while satisfying all these properties. In this case, we can also analyse the causal structure (and not just the signalling structure). The fine-grained network and causal structure are illustrated in Fig. 12. This fine-grained network in fact coincides with a description of QS originally proposed in the causal box framework [21]. By connecting both process matrices and causal boxes to our more general framework, our work provides a mapping between the process matrix and causal box descriptions of QS (even though one admits an indefinite and other a definite information-theoretic causal structure) and reveals that these can be understood as coarse and fine-grained descriptions that are related through appropriate encoders and decoders. We describe the encoding and decoding in Appendix G. Here we outline the main features of this fine-grained network and associated fixed-order process \mathcal{W}_{QS}^f .

The network $\mathfrak{N}_{\mathcal{W}_{QS}, U, V}$ associated with the original process, involves the set $\{C_C^O, C_T^O, A^I, A^O, B^I, B^O, D_C^I, D_T^I\}$ of systems. Suppose it is realised in Minkowski spacetime according to the following embedding.¹¹

$$\begin{aligned}
 \mathcal{E}(C_C^O) = \mathcal{E}(C_T^O) &:= \{L^C\} \\
 \mathcal{E}(A^I) &:= \{P_1, P_2\}, \\
 \mathcal{E}(A^O) &:= \{Q_1, Q_2\}, \\
 \mathcal{E}(B^I) &:= \{R_1, R_2\}, \\
 \mathcal{E}(B^O) &:= \{T_1, T_2\}, \\
 \mathcal{E}(D_C^I) = \mathcal{E}(D_T^I) &:= \{L^D\}.
 \end{aligned} \tag{16}$$

Where the spacetime points are ordered as follows

$$\begin{aligned}
 L^C &< P_1 < Q_1 < R_2 < T_2 < L^D, \\
 L^C &< R_1 < T_1 < P_2 < Q_2 < L^D.
 \end{aligned} \tag{17}$$

In the fine-grained process, C and D 's systems have isomorphic states spaces to the original ones while each of A and B 's systems split into two systems, one for each spacetime location. For instance, A^I in the coarse-grained process maps to $\{A^{I, P_1}, A^{I, P_2}\}$ under the associated systems fine-graining map. To avoid clutter, we simply label these as

¹⁰ For general processes, whether or not agent A_i can signal to an agent A_j can depend on the choice of local operations of A_j as well as those of the remaining agents. In the quantum switch, signalling from A to B (whenever possible, i.e., the initial state $\alpha|0\rangle^C + \beta|1\rangle^C$ of the control has $\alpha \neq 0$) can be verified by A by suitable local choices of operations alone, independently of B 's operation, and the symmetric statement holds for B .

¹¹ In this example, the embedding is specified relative to the maximal fine-graining, i.e., in terms of individual spacetime locations comprising the region.

$\{A_1^I, A_2^I\}$ and similarly for the remaining systems of A and B . Then the fine-grained process \mathcal{W}_{QS}^f thus obtained, has $\text{In} = \{C_C^O, C_T^O, A_1^O, A_2^O, B_1^O, B_2^O\}$ and $\text{Out} = \{D_C^I, D_T^I, A_1^I, A_2^I, B_1^I, B_2^I\}$. The local operation \mathcal{U}^A (taking it to be a unitary for simplicity) of Alice gets fine-grained to $\mathcal{U}_{1,2}^{A,f} = \mathcal{U}_1^{A,f} \otimes \mathcal{U}_2^{A,f} : A_1^I \otimes A_2^I \mapsto A_1^O \otimes A_2^O$ and similarly for Bob. In the fine-grained network, the presence of a d -dimensional target system at a spacetime location is modelled as a single d -dimensional message at that location and the absence of the target at a location as a vacuum state $|\Omega\rangle$ (zero message state) at the location. The fine-grained systems of Alice and Bob are therefore each $(d+1)$ -dimensional when the original target is d -dimensional, as they additionally include the 1-dimensional vacuum state $|\Omega\rangle$.

For the sake of illustration, suppose that the spacetime locations are such that the inputs A_1^I and B_1^I are associated with a time co-ordinate t_1 while A_2^I and B_2^I with a time co-ordinate $t_2 > t_1$ in some global classical reference frame. Then, when the control is $|0\rangle$, the target qudit is sent to Alice at time t_1 and then Bob at time t_2 (and the vacuum state to Bob at time t_1 and Alice at time t_2) and when the control is $|1\rangle$, the target is sent to Bob at time t_1 and Alice at time t_2 (and the vacuum state to Alice at time t_1 and Bob at time t_2). The fine-grained local operations at each time are such that they leave the vacuum state invariant and apply a qudit unitary \mathcal{U}_i^A whenever the input is non-vacuum, where $i \in \{1, 2\}$. That is,

$$\begin{aligned} \mathcal{U}_i^{A,f} |\Omega\rangle^{A_i} &= |\Omega\rangle^{A_i^O} \\ \mathcal{U}_i^{A,f} |\psi\rangle^{A_i} &= \mathcal{U}_i^A |\psi\rangle^{A_i^O} \end{aligned} \quad (18)$$

In particular, composing the fine-grained process map \mathcal{W}_{QS}^f with the corresponding local operations $\mathcal{U}_{1,2}^{A,f}$ and $\mathcal{V}_{1,2}^{B,f}$ implements the following transformation (see Appendix G for further details) on the control and target (where C and T on the left are associated with the corresponding output systems of Charlie, and on the right, to corresponding inputs of Danny).

$$(\alpha|0\rangle^C + \beta|1\rangle^C) \otimes |\psi\rangle^T \mapsto \alpha|0\rangle^C \otimes \mathcal{V}_2^B \mathcal{U}_1^A |\psi\rangle^T + \beta|1\rangle^C \otimes \mathcal{U}_2^A \mathcal{V}_1^B |\psi\rangle^T \quad (19)$$

Setting $\mathcal{U}_1^A = \mathcal{U}_2^A$ and $\mathcal{V}_1^B = \mathcal{V}_2^B$, this recovers the QS transformation of Eq. (15). However, accounting for all possible interventions that the agents can physically perform in the spacetime (which includes time-dependent unitaries), we see the fine-grained, 6-partite process of Fig. 12 which has a definite acyclic causal structure accurately captures these possibilities. Indeed, in the coarse grained description, the process matrix W_{QS} there is no such degree of freedom available to Alice and Bob through which they can implement a different unitary depending on the order in which they act but this is possible in the spacetime realisation. In this fine-graining, the spacetime location at which Alice/Bob receive a non-vacuum state is perfectly correlated with the control state and therefore the order in which the two agents act. There are QS realisations in Minkowski spacetime [45, 49] where such perfect correlations don not exist. Our results nevertheless apply here, as discussed in Remark 6 in the appendix.

Role of vacuum states Similar constructions using vacuum states, for describing experimental realisations of QS have been proposed in the literature [21, 22, 52, 53, 55]. The associated fine-graining of QS has the property that each agent acts exactly once on a d -dimensional non-vacuum state. One can construct other valid fine-grainings where agents act on multiple d -dimensional non-vacuum messages. Arguably however, we could not regard such fine-grainings as corresponding to “faithful realisations” of the process as they would not keep with the spirit of the closed labs assumption of the process matrix framework (Section VI) which requires that each agent acts exactly once on the target qudit. The general results of this paper do not refer to or depend on the existence of such vacuum states, they prove the existence of an acyclic fine-grained causal structure for physical spacetime realisations of any process irrespective of whether the realisation can be considered “faithful” in this sense, therefore generalising previous arguments for QS to all processes and also to a larger class of spacetime realisations. Modelling the set-up assumptions of the process framework at the fine-grained level is a subject of an upcoming work, based on [75]. This enables further no-go results regarding the set of physically and faithfully realisable processes in a fixed spacetime (see also [77]).

Physical interpretation of the experiments Whether QS experiments in Minkowski spacetime implement or simulate an indefinite causal structure has been a subject of long-standing debate. Instead of directly focusing on this question, we addressed a fundamental and overlooked question that underpins the discussion: how can an indefinite information-theoretic causal structure (such as that of QS) be compatible with relativistic causality in a definite spatio-temporal causal structure (such as Minkowski spacetime)? We have found that in such experiments, at the

fine-grained level, both causal structures can be regarded as being definite and acyclic. This supports the conclusion that these experiments do not implement a fundamentally indefinite causal structure, but they are useful for simulating these exotic causal structures.

Our work also highlights that such experiments are nevertheless interesting and worthy of further investigation, as they involve the non-localisation of information in spacetime. A better characterisation of the underlying experimental resources, which can distinguish between quantum vs classical forms of non-localisation, would enhance our understanding of spatio-temporal quantum correlations which can play an important role in relativistic quantum information processing. Being explainable by an definite information-theoretic causal structure would only make it easier to analyse potentially new applications of such experiments in a fixed spacetime while carrying forth many intuitions from standard quantum information formalism. To this effect, in Appendix D we show that all physical spacetime realisations of networks in our framework can equivalently be viewed as causal boxes [21], a previously known formalism for modelling quantum messages exchanged at superpositions of spacetime locations, and which has applications in relativistic quantum cryptography [64].

That being said, there are other previous works which present arguments for the opposite conclusion, of regarding these experiments as genuine implementations of indefinite information-theoretic causal structures. A central piece in our results is our formalisation of relativistic causality, which we believe also highlights a core difference between these two opposite sides of the debate. We discuss this aspect further in the next paragraphs.

Relativistic causality and agents' interventions Relativistic causality in spacetime implies that no physical operation an agent can physically perform in one region can transmit information to another region beyond the first region's future light-cone. It is therefore important to take into account the set of all possible physical operations that an agent can in-principle perform in a given scenario. Let us focus on this aspect, and compare the original process matrix description and the experimental realisations of QS. In process matrix, agents can intervene in their closed laboratories, while the process models the uncontrollable external environment. In contrast, experimental QS setups occur within a single laboratory where experimenters can, in principle, intervene and control any part of the experiment. Our fine-grained description of the spacetime realization ensures that even with these interventions, superluminal signaling remains impossible. It is worth noting that the counterfactual possibility of different interventions is also central to information-theoretic approaches to causality based on causal inference [1].

This distinction between considering *all physically possible interventions* vs *actually performed interventions* highlights an important difference between the two sides of the debate. In the fine-grained description of QS, if the two "Alices" and "Bobs" are constrained to perform identical operations at their respective spacetime locations, the action of the fine-grained network representing a fixed order, causally separable process can become indistinguishable from that of the coarse-grained network which represents a causally non-separable process (also see the paragraph following Eq. (19)). More broadly, when starting with an acyclic information-theoretic causal structure and limiting interventions so they can not individually act on certain nodes, it can appear as if the nodes where agents' operations take place no longer have a definite order. A similar effect also arises in the classical examples of Fig. 6 and Fig. 14, where by restricting interventions to always act jointly on A_1 and A_2 without distinguishing them as separate nodes, the locus of agents' operations would define the nodes $\{A_1, A_2\}$ and B between which there is bidirectional causation and signalling, and hence no definite ordering. Analogously, in the spacetime perspective, we could have $P_1 < Q < P_2$ but when focusing on the region $\mathcal{R} = \{P_1, P_2\}$ as a whole, we cannot assign a definite ordering to \mathcal{R} and Q . In different physical regimes beyond fixed spacetimes, the fundamental structure of allowed physical interventions might enable realizations of indefinite causal structures, where both fine and coarse-grained descriptions agree on the lack of a definite causal order. Thus, it would be intriguing to establish mathematical connections between our framework and other formalisms contributing to this debate, as this would provide insights also for physical regimes beyond classical spacetime structures and classical reference frames (see Section XE for further discussion).

D. A realisation of QS in Minkowski spacetime with systems localised in time

We propose a new realisation of QS in Minkowski spacetime with the property that the in and output events are time localised for both Alice and Bob in their respective frames. The associated protocol requires Alice and Bob to be in relative motion with respect to each other, at a constant velocity. Previous quantum switch protocols in a fixed spacetime typically involve spatial localisation and time non-localisation, in contrast, ours will involve spatial non-localisation and time localisation (with respect to the local reference frames). It also demonstrates that Corollary 2 no longer holds when only requiring time localisation with respect to each local frame, as opposed to a single global reference frame.




FIG. 13: Quantum switch in Minkowski spacetime with time localised systems: Alice labels her spacetime locations with respect to the co-ordinate system associated with the black axes (orthogonal) and Bob with respect to the gray axes (non-orthogonal), which is related to Alice's frame by a Lorentz transformation. Eq. (16) specifies how the spacetime points in the figure are associated with the in/output systems of the agents A , B , C and D . When the control is in the state zero, the target follows the blue trajectory through the spacetime (left side in figure, first line of Eq. (17)), going to Alice first and then to Bob and when the control is in the state $|1\rangle$, the target follows the orange trajectory through spacetime (right side in figure, second line of Eq. (17)) going to Bob first and then to Alice. Irrespective of the order, Alice's in and output events are localised at times t^{A^I} and t^{A^O} in her frame and Bob's in and output events are localised at time t^{B^I} and t^{B^O} in his frame. However, the target degree of freedom is nevertheless not localised in the spacetime as the protocol implements a quantum controlled superposition of the two trajectories, to Alice and Bob this appears as non-localised in space but localised in time.

The network $\mathfrak{N}_{\mathcal{W}_{QS,U,V}}$ is realised in Minkowski spacetime according to the same embedding as Eq. (16) and Eq. (17), and thus corresponds to the same fine-grained network as that of Fig. 12 at a purely information-theoretic level. The only distinction is that we no longer require A_1^I and B_1^I to be localised at time t_1 and A_2^I , B_2^I at time $t_2 > t_1$ as in Fig. 12. Instead, this realisation has the property that P_1 and P_2 (associated with Alice's inputs) have the same time co-ordinate t^{A^I} in while Q_1 and Q_2 (associated with Alice's outputs) have the same time co-ordinate $t^{A^O} > t^{A^I}$ in Alice's frame, and similarly R_1 and R_2 (associated with Bob's inputs) have the time co-ordinate t^{B^I} while T_1 and T_2 (associated with Bob's outputs) have the time co-ordinate $t^{B^O} > t^{B^I}$ in Bob's frame. Thus the regions associated with Alice and Bob's systems are time localised regions (cf. Definition 27). Further details of the protocol can be found in Appendix G and the intuition behind the proposed realisation and its main features are illustrated in Fig. 13.

X. TOWARDS DEMYSTIFYING INDEFINITE CAUSATION

The concept of indefinite causal structures (in the sense of indefinite order of information-theoretic operations) has introduced valuable mathematical tools for studying quantum processes. However, its physical interpretation has remained debated. Consensus is lacking regarding their relationship to spatio-temporal causality notions, the operational meaning of indefinite causality witnesses like causal inequalities, and on their implications for the quantum

nature of causal structures. In order to facilitate formal discussions and consensus on these matters, we have developed a general theoretical framework that formalizes different causality notions, connecting them through operational formulations of fundamental principles, under relatively minimal physical and mathematical assumptions. More generally, our graph theoretic-formulation of physical considerations makes it possible to analyse how other models of spacetime must structurally deviate from a fixed acyclic spacetime in order to evade the no-go theorems we have derived. We discuss the concrete implications of our work for interpreting indefinite causality within and beyond fixed background spacetimes and outline different directions for future research relating to this discussion. In Section XI, we also outline future research directions of more pragmatic relevance.

A. Different notions of events

Causality is closely related to the concept of events, with each type of causal structure implying a notion of events associated with its the nodes. In spatio-temporal causal structures, events correspond to spacetime points or regions, while in information-theoretic structures, events are linked to the in/output systems of operational procedures, reflecting where agents can intervene. Both notions are pertinent in physical experiments. Instantiating the abstract process matrix description in a specific physical regime, like a background spacetime, can provide agents access to more degrees of freedom not present in original coarse-grained process description. This can potentially expand the set of information-theoretic events that are physically relevant.

This is exemplified in our analysis of the quantum switch where Alice, in the process matrix, can only apply \mathcal{U}^A to her input system A^I independently of whether she acts before or after Bob, whereas in the spacetime realization we have analysed, she can apply different unitaries \mathcal{U}_1^A or \mathcal{U}_2^A based on the time order (associated with the background spacetime structure, cf. Eq. (19)). She must restrict these unitaries to be identical in order to realize the QS transformation of Eq. (15). Here Alice is associated with a single pair of in/output events in the process matrix, and two pairs in/output events in the spacetime realization. The former corresponds to a *causally non-separable process* with one operational event for Alice and one for Bob, while the latter picture is described by a *fixed-order, causally separable process* involving more operational events. This emphasizes the importance of distinguishing whether (a) the original indefinite causal structure's behaviour is replicated by restricting the agents' operations in a fine-grained fixed-order processes which physically allows for a larger set of operations, or (b) even when considering all physical interventions that agents can perform in a given realisation of a process in some regime, no deviations from the original causally non-separable process can be observed. Our results indicate that in fixed and acyclic background spacetimes, case (a) always holds. The possibility of realising the transformation of Eq. (19) in quantum gravitational realisations of QS, and the associated notion(s) of events in such scenarios has been discussed in [78]. More generally, formalizing our framework's requirements in regimes beyond fixed spacetime and exploring whether (a) or (b) is satisfied in different realizations is an intriguing avenue for future research, especially given our demonstration that (b) is impossible in a fixed spacetime.

B. Relating the process matrix assumptions to those of our no-go theorems

As discussed in Section VI, in [28], causal inequalities are shown to be necessary conditions on correlations generated by protocols satisfying the assumptions: free choice (FC), local order (LO), closed labs (CL) and that events are localised in a global causal structure (CS). Thus, a violation of causal inequalities under FC, LO and CL would indicate a violation of (CS), which is often interpreted as certifying the indefiniteness of the causal structure.

However, a more careful look reveals that (CS) is essentially two assumptions: (CS1) there exists a global partial order in which the process is realised such that signalling is only possible from past to future with respect to this partial order and (CS2) the input/output events of every agent in the process are localised in this partial order. Within our framework, a fixed spacetime inherently defines this partial order, and (CS1) is automatically satisfied for any process realisation which respects our relativistic causality condition (Definition 24) in a fixed spacetime. (CS2) then becomes an additional constraint on the spacetime embedding, requiring that each system S is embedded at a single spacetime point $\mathcal{E}(S) \in \mathcal{T}$ (rather than a non-trivial region). This constraint implies the acyclicity of the region causal structure. Then it immediately follows from Theorem 4 and Corollary 1 that a process realised according to (CS)=(CS1)+(CS2) in any partial order must be a fixed order process and hence can not lead to correlations violating causal inequalities when the agents perform freely chosen operations. Interestingly our definitions of *spacetime realisation* and *relativistic causality* automatically account for the (LO) and (CL) assumptions once the localisation condition of (CS2) is satisfied. Specifically, (LO) follows because the signalling relations of the process network always include $\{A^I\} \longrightarrow \{A^O\}$ for

every agent A , and relativistic causality together with (CS2) ensure that $\mathcal{E}(A^I) < \mathcal{E}(A^O)$ has the required local input-output ordering relative to the embedding partial order. (CL) is satisfied since each agent is associated with a single in/output pair of events (both in the informational and spatio-temporal sense) and hence acts only once. This extends the statement of [28], showing that formalising (CS) as done in our framework, (CS) alone (/(CS)+(FC)) is sufficient to imply that the process being realised is a fixed order process (./does not generate correlations that violate causal inequalities).

C. Operational meaning of causal inequalities

The previous sub-section discussed the implications of imposing (CS) = (CS1) + (CS2). By formalising these assumptions independently, our framework enables us to explore scenarios where (CS1) holds but not (CS2). This would correspond to fixed spacetime realisations of processes associated with a cyclic region causal structure, violating 3. in Theorem 4 while satisfying 1. and 2. One can then ask: can a process realisation satisfying (CS1), (FC), (LO) and (CL) violate causal inequalities? Before answering this question, it is essential to note that, even if such a violation of causal inequalities were possible, it would not imply the indefiniteness of the causal structure, as Theorem 5 implies that (CS1) alone guarantees a fine-grained explanation in terms of a fixed causal order process that all agents agree on. A violation of causal inequalities under the assumptions (CS1), (FC), (LO) and (CL) (if it were possible) would however provide a device-independent certification of the non-localization of information (i.e., violation of (CS2)) in a fixed and acyclic background causal structure.

Having clarified the operational meaning of a potential violation of causal inequalities in a fixed spacetime, we will investigate in a follow-up work whether such violations are indeed possible. This requires formally defining the (LO) and (CL) assumptions for spacetime realizations where (CS1) holds but not (CS2), where an agent may receive a superposition of vacuum and non-vacuum states at a given spacetime location. To the best of our knowledge, such a formalization is missing in the previous literature, but has been discussed in recent works involving one of the authors (the unpublished abstract [75], pre-print [76] and the master’s thesis [77]) and will be explicitly addressed in forthcoming works based on this. Moreover the results of [77] suggest that under the conditions of (CS1) (or equivalently, satisfying relativistic causality in a fixed spacetime) + (LO) + (CL), the only realizable processes are effectively those corresponding to quantum circuits with quantum-controlled superposition of orders [41], which include QS, and do not violate causal inequalities when combined with local operations satisfying (FC). Interestingly, through a different formulation of (CL) and (LO) and of non-localisation of information in time, [79] reaches the conclusion that causal inequality violations are physically realisable. There is no mathematical contradiction between these two results; the (CL) and (LO) assumptions of [79] operate at the level of the coarse-grained causal structure of the process matrix (where vacuum states do not feature), and this does not imply our formulation of these assumptions which operate at the fine-grained level of the spacetime realization (which includes vacuum and non-vacuum states). In the latter case, one wishes to ensure that each agent A acts on a single non-vacuum message of dimension d_{A^I} (dimension of A ’s input system in the original process matrix description) in the spacetime realisation, which can be physically verified using coherent counters for each agent (this is the case for the fine-graining of QS presented in Section IX C and Appendix G). In our view, such a formulation of these constraints at the level of the fine-grained spacetime realisation is necessary for realising a process in a “faithful” or “loophole free” manner. Establishing mathematical connections between the different formulations of (LO) and (CL), at the coarse and fine-grained levels, would inform a rigorous discussion in the community about what would constitute a “loophole free” violation of causal inequalities.

Finally, we note an intriguing recent work [80] that derives a different kind of causal inequality than those discussed above, by including an additional “relativistic causality” assumption. Such inequalities can in fact be violated by an extended quantum switch protocol. In particular, this involves a quantum switch between Alice₁ and Alice₂ and a Bob who is space-like separated to the two Alices, who measures a system which is entangled with the control of the quantum switch. The relativistic causality assumption relates to a non-signalling condition between Bob and the two Alices. In this case as well, if a loophole free violation of such an inequality were observed in a fixed spacetime implementation of the quantum switch, then it would not certify any causal indefiniteness (at the fine-grained level) as implied by our results, but can instead be seen as a certification of the violation of (CS2) or witnessing non-classicality in the spacetime localisation of systems. This is further discussed in our associated letter [63].

D. Cyclic causal structures and insights from causal inference

The lack of a definite acyclic causal structure need not imply that the causal structure is indefinite (as commonly interpreted in the indefinite causality literature), but can also mean that the causal structure is definite but cyclic.

Indeed, connections between indefinite and cyclic causal structures have been previously observed [27, 28, 57–59]. While these results provide significant insights on the information-theoretic causal structure of process matrices, they do not directly address the questions regarding physical realisability of quantum processes in a spacetime or the relation between such cyclic causal structures and relativistic causality in acyclic spacetimes, which we have addressed here. Building on these previous insights, our work shows a tight correspondence (cf. Theorem 3) between non-fixed order process matrices and cyclic causal structures as characterised by operationally observable signalling relations.

Interestingly, cyclic information-theoretic causal structures can capture both physical scenarios with feedback [2, 21, 62] as well model physics in the presence of exotic closed timelike curves (CTCs) [3, 58, 81–83]. The physical distinction between these two situations come from how the cyclic causal structure is linked with spatio-temporal degrees of freedom [2, 3], e.g., whether the nodes of the information-theoretic structure are embedded into spacetime regions or spacetime events (as discussed in Fig. 1). Linking indefinite causal structures of the process framework and cyclic quantum networks has allowed us to provide a clear physical interpretation to their realisations in an acyclic spacetime, which is analogous to the interpretation of cyclic causal models used in classical data sciences to model physical processes with feedback (which are not regarded as exotic CTCs). This interpretation is possible due to our relativistic causality condition which is a necessary constraint on the signalling relations of spacetime realised network needed to ensure that there are *no superluminal (information-theoretic) causal influences*. Interestingly, in recent works [2, 3] involving one of the authors, it is shown that using a weaker relativistic causality principle which is necessary for ruling out *superluminal signalling* (as opposed to causation, which is distinct from signalling), it is not possible to rule out genuine causal loops from being embedded in Minkowski spacetime. Moreover, there can be a certifiable gap between the information processing possibilities in theories constrained by these distinct relativistic causality principles [70]. In Appendix I, we further discuss the relation between the distinct relativistic causality/compatibility conditions of [2, 3] and the present work, as well as their physical implications for causal loops. Characterising the general relation between such classes of causal loops found in the causal modelling framework of [2, 3] and process matrices remains an open problem.

More broadly, on the topic of cyclic causal structures, there are at least two interesting and inter-related research directions for future research. Firstly, the study of fine-graining of cyclic causal networks as initiated here can also provide a pathway for investigating the simulation of CTCs in a fixed acyclic spacetime by unravelling them into acyclic structures. Such simulations have been previously considered using experimental post-selection [82]. In general, we can simulate cyclic causal structures using acyclic ones even without invoking post-selection (the quantum switch network being an example, see Section IX). Which are the set of cyclic networks that admit acyclic fine-grainings without post-selection? How is this set related to the set of networks of CPTPMs which describe process matrices?

The second direction relates to causal inference in the presence of cycles. It is known that several causal inference results which hold in classical and quantum acyclic models (such as the *d-separation* theorem [1, 13, 59]) fail in cyclic models, already in the classical case, and there is active research in characterising the domain on applicability of the *d-separation* theorem in classical cyclic models [62]. A particular consequence of the *d-separation* theorem is that if two nodes in a causal graph associated with classical variables, say a and y have no directed paths between each other, then we must have $P(ya) = P(y)P(a)$ in the resulting correlations. This does not hold in general cyclic models, but our results of Appendix F imply that this would hold for the subset of cyclic networks corresponding to valid process matrices. An upcoming work [72] further develops the causal modelling aspects of cyclic quantum networks, providing insights on graph separation properties such as *d-separation*. The precise subset of cyclic quantum networks that respect the *d-separation* theorem, and their relation to the set of valid process networks remain open problems. Addressing them would simultaneously shed light on open questions in classical causal inference and the causal decompositions of process matrices, and would also provide useful graph-theoretic tools for addressing the first question of this paragraph relating to spacetime realisations.

E. Beyond classical spacetimes and classical reference frames

We have modelled a fixed spacetime structure as a global partial order, which is the same in the all of reference frames (cf. Definition 27). This is the case, for instance in special relativity, where reference frames are classical, the light cone structure of the spacetime is invariant under Lorenz transformations which relate such classical frames. Are there any potential features of information-theoretic causal structures which are impossible to achieve in a fixed spacetime but may be possible in scenarios that deviate from this model of fixed spacetime? We discuss this further in light of our no-go results.

Firstly, our model of a fixed spacetime imposes an agent independent notion of spacetime localisation/non-localisation. That is, if a system S is localised at a spacetime point $P \in \mathcal{T}$, all agents agree that S is localised at P , and if S is in a superposition of being at $P_1, P_2 \in \mathcal{T}$, all agents agree that S is not localised and is non-localised between P_1 and P_2 (even if they may describe the points using different co-ordinates). This property is indeed satisfied in standard special and general relativity where classical reference frames, sharing a common origin, are used. This is also the case in most existing quantum experiments (such as the quantum switch experiments discussed in Section IX C). However, this property can fail (even classically) if different agents only have access to reference frames that are mutually unsynchronised, such that they cannot unambiguously order events. This highlights that the underlying assumptions about reference frames and about what information/degrees of freedom are accessible to agents play a crucial role in the identification and ordering of events (already classically).

More fundamentally, we may no longer expect an absolute notion of localisation to hold when spacetime co-ordinates are described using quantum reference frames (QRFs) since superposition and entanglement can become frame dependent in such settings, even using perfect QRFs [84, 85]. As we have seen in Figure 11, realisations of the quantum switch in quantum gravitational spacetimes have been considered, and analysed using quantum clocks (temporal QRFs) [26, 60]. Here it is possible that Alice sees her operation (and systems) to be localised in her proper time while she sees Bob's operation in a superposition of occurring earlier and later than her (and similarly for Bob). However, our quantum switch protocol in Minkowski spacetime (Section IX D) illustrates that this property of agent-dependent time localisation is not specific to quantum gravitational or QRF based realisations, but can be achieved in a fixed spacetime. What is impossible in our model of fixed spacetime is agent-dependent spacetime localisation.

To characterise whether agent-dependent spacetime localisation is possible in quantum gravitational realisations of processes, further analysis is needed which accounts for spatial and temporal information, and sheds light on how spacetime distances could be measured in these scenarios. For example, in the quantum gravitational realisation of QS, if the agents measure spatial distances with respect to the gravitating mass, then by construction, their spatial co-ordinate would depend on the branch of the superposition they are in (cf. Figure 11) and they would not be spatially localised. On the other hand, if each agent localises their input event at the origin of their spatial co-ordinates, then each agent can possibly describe their own in/output events as being spatially and temporally localised while describing the in/output events of other agents as not being localised, in the gravitational QS. In a similar spirit, we note that a protocol to distinguish optical realizations of the quantum switch in Minkowski spacetime from theoretical quantum gravitational realizations was previously proposed [53]. The protocol involves an additional agent, the Friend F to whom Alice and Bob send a photon every time they act on a non-trivial target system, and F performs a non-demolition measurement to test a property relating to the time of arrival of the photon. In Appendix H, we show that this protocol can be generalised to arbitrary fixed spacetime realisations by formulating the main property in order-theoretic terms, and showing it is impossible to achieve in any fixed spacetime. On the other hand, [53] suggests there can exist theoretical quantum gravitational models that achieve the property.

Our results for fixed spacetime together with this discussion highlights the following point: whether spatio-temporal co-ordinates are described in terms of classical or quantum RFs, whether the RFs used by different agents are synchronised etc. play an important role in the fine-grained description of these realisations as well as the physical interpretation of the associated causal structure. We envisage that extending our framework to allow for agent-dependent partial orders and fine-grained causal structures can enable a description of both notions of causality and their interplay in scenarios beyond the current model of fixed spacetime, where spacetime localization becomes subjective and where the spacetime may exhibit certain quantum effects. Generalizing our framework and testing the applicability of our no-go results in quantum gravitational realizations of quantum processes and QRF-based descriptions of spacetime coordinates presents another inter-disciplinary avenue for future research, which can have applications for understanding proposed table-top quantum gravity experiments [86, 87].

XI. CONCLUSIONS AND OUTLOOK

We have developed a general theoretical framework for studying quantum and relativistic causality, which satisfies the 4 desiderata laid out in the introduction. Applying the formalism to indefinite causal order processes, we derived several no-go results regarding their physical realisations in a background acyclic spacetime. In the context of the long-standing debate [21, 52–56] regarding the interpretation of experimental realisations [43, 45, 74] of the quantum switch in Minkowski spacetime, this work furnishes formal theorems in support of the argument that these experiments do not implement a fundamentally indefinite causal structure. Rather, they are useful for simulating the behaviour

of such exotic structures (see Section IX C) using definite and acyclic causal structures (in the spatio-temporal *and* information theoretic sense).

Our work also highlights that such experiments still remain an interesting subject of further study, as they involve the non-localisation of information in space and in time. The result of Appendix D, that any such quantum experiment in a spacetime can be described as *causal boxes* [21], enable their analysis in terms of a well-defined causal and compositional structure, and using fine-grained tools which account for information-theoretic and spatio-temporal aspects. Causal boxes also have natural applications for studying composable security in practical relativistic quantum cryptographic protocols [64]. Obtaining a fine-grained characterisation of the spatio-temporal quantum resources and correlations involved in these experiments, certifying their non-classicality and developing their applications for relativistic quantum information processing still remain important future directions that can be investigated by building on this work and the related literature (e.g., [22, 23]).

More generally, apart from process matrices and causal boxes, our formalism also provides a platform for studying and relating other information-theoretic approaches to causality: such as functional causal models on cyclic graphs, classical and quantum Bayesian networks, and quantum split-node causal models (see Appendix C). There is much scope for further developing the information-theoretic aspects of our framework for specific purposes. As discussed in Section X D, there are intriguing open questions regarding graph separation properties in cyclic quantum causal models which are related to ongoing research directions in classical causal inference [62]. Moreover, the concept of fine-graining introduced here, and our associated results, also highlight interesting parallels between information-theoretic and spatio-temporal notions of causality.

Our compatibility condition, utilized in formulating relativistic causality, has broader scope. It allows to investigate the compatibility between detectable information flow in quantum networks with a partial order or acyclic graph, where the former is formulated independently of spacetime and without assuming a unique causal direction (allowing for cycles). This general approach to both information-theoretic and spacetime structures suggests avenues for extending techniques from the quantum causality literature towards the exploration of how spacetime structure may emerge from properties of quantum information, which is an active research area within the quantum gravity community (see, for example, [88–91]). Finally, there have been recent works on generalising higher-order transformations (such as process matrices) in a quantum-theory independent manner [92], as well as work on linking information-theoretic and spatio-temporal causality in a theory-independent manner [2, 3]. Building on these insights, it would be interesting to consider whether our approach would generalise to information-theoretic structures in post-quantum generalised probabilistic theories (GPTs), and which GPTs would still be constrained by similar no-go theorems.

In the natural process of science, we may often need to update our preconceived notions in light of new experiments. Bell’s theorem [5] has set an unprecedented example in highlighting the power of no-go theorems in this scientific process— establishing what is impossible to achieve within certain physical regimes tells us how physics in new regimes challenges our prior intuitions and how we can exploit these new physical phenomena for useful practical tasks. A theory of quantum gravity may further challenge our current understanding of causality in stronger ways than Bell’s theorem, but we cannot fully anticipate how. The approach of disentangling different causality notions and carefully reconnecting them, as developed in [2, 3] and the present paper, can help better prepare for such challenges. We have shown how this approach provides concrete theorems for reinterpreting existing experiments and also to analyse how causality in more exotic physical regimes (such as quantum gravity) may differ. As we have outlined throughout this text, there are several fascinating questions that still remain open and areas where consensus is lacking regarding the interpretation of quantum and relativistic causal structures. We hope that our work, by formalising connections between different notions and approaches to causality, facilitates the process of building a consensus and also addressing some of these important open questions that span different research directions.

Note Shortly after the first arXiv version of this manuscript was posted, an independent work [55] appeared which also proposes a related concept of fine-graining for quantum causal structures, specifically in the context of the quantum switch. While our framework and definition of fine-graining have mathematical and conceptual differences from those of this work, the conclusions drawn about experimental realisations of the quantum switch are similar.

Acknowledgements VV thanks Lin-Qing Chen, Augustin Vanrietvelde, Āslav Brukner, Āmin Baumeler, Esteban Castro-Ruiz, Hlér Kristjánsson, Lía del Rio, Nick Ormrod, Lorenzo Maccone and Ognjan Oreshkov for interesting discussions on quantum and relativistic causality. VV’s research has been supported by an ETH Postdoctoral Fellowship. VV and RR acknowledge support from the ETH Zurich Quantum Center, the Swiss National Science Foundation via project No. 200021_188541 and the QuantERA programme via project No. 20QT21_187724.

Appendix A: Overview of symbols and notation

Here, we provide an overview of the notations used in the paper, with a particular focus on the graph-theoretic notation as the framework involves various types of order relations.

General notation

- \mathcal{H}^S : Hilbert space of a system S
- $\mathcal{L}(\mathcal{H}^S)$: Density operators on \mathcal{H}^S or state space of S
- \mathcal{I}^S : The identity channel on system S
- $\mathbb{1}^S$: Identity matrix on system \mathcal{H}^S used for denoting states
- \oplus : Addition modulo 2
- $\Sigma(\mathcal{S})$: Powerset of a set \mathcal{S}
- $P(X)$: Probability distribution of a random variable X
- $P(x)$: Probability of an event x (such as a measurement outcome)

Graphs, arrows and order relations

- \mathcal{G} : Directed graph
- \hookrightarrow : $S \hookrightarrow S'$ denotes that an output system S is loop composed to an input system S' of some CPMs in a network of CPMs
- \longrightarrow : Order relation on sets of systems which indicates (information-theoretic) *signalling* between them
- \rightsquigarrow : Order relation on systems which indicates (information-theoretic) *causal influence* between them
- \mathcal{T} : Partially ordered set used to denote spacetime structure
- $<, \leq, >, \geq$: Partial order relation of a spacetime \mathcal{T} , with $P < Q$ (equivalently $Q > P$) denoting that P and Q are distinct with P being strictly in the causal past of Q while $P \leq Q$ (equivalently $Q \geq P$) implies $P < Q$ or $P = Q$
- $\not\prec$: Used to denote that two spacetime locations are unordered relative to the partial order of the spacetime
- \xrightarrow{R} : Order relation associated with a region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ of a spacetime \mathcal{T}
- \xrightarrow{G} : Edges of a generic directed graph \mathcal{G}

Appendix B: Further details of loop composition

Lemma 5. [Closedness under composition] *CPMs are closed under arbitrary parallel and loop composition.*

Proof. Note that parallel composition $\mathcal{M} \otimes \mathcal{N}$ of two CPMs can be equivalently written as the following sequential composition

$$\mathcal{M} \otimes \mathcal{N} = (\mathcal{I}^{\text{Out}(\mathcal{M})} \otimes \mathcal{N}) \circ (\mathcal{M} \otimes \mathcal{I}^{\text{In}(\mathcal{N})}). \quad (\text{B1})$$

Therefore $\mathcal{M} \otimes \mathcal{N} \otimes \mathcal{I}^A$ for any system A , is equivalent to $(\mathcal{I}^{\text{Out}(\mathcal{M})} \otimes \mathcal{N} \otimes \mathcal{I}^A) \circ (\mathcal{M} \otimes \mathcal{I}^{\text{In}(\mathcal{N})} \otimes \mathcal{I}^A)$. Complete positivity of the parallel composition $\mathcal{M} \otimes \mathcal{N}$ requires that $\mathcal{M} \otimes \mathcal{N} \otimes \mathcal{I}^A$ maps positive states to positive states. This follows from the fact that $(\mathcal{M} \otimes \mathcal{I}^{\text{In}(\mathcal{N})} \otimes \mathcal{I}^A)$ maps positive states to positive states (by the CP property of \mathcal{M}) and $(\mathcal{I}^{\text{Out}(\mathcal{M})} \otimes \mathcal{N} \otimes \mathcal{I}^A)$ also maps positive states to positive states (by the CP property of \mathcal{N}).

To see that loop composition preserves complete positivity, we show that loop composition connecting an output subsystem S to an input subsystem S' of the CPM \mathcal{M} can be equivalently represented as follows: a preparation of a maximally entangled state $|\Phi^+\rangle^{S'R} := \sum_i |ii\rangle^{S'R}$ on the input S' and an isomorphic ancilla R with $\mathcal{H}^{S'} \cong \mathcal{H}^R$, followed

by action of \mathcal{M} , and finally a post-selection of the same maximally entangled state $|\Phi^+\rangle^{SR} := \sum_i |ii\rangle^{SR}$ on the output S along with the ancilla R .¹² Explicitly, this results in an overall map acting on input systems $\text{In}\{S'\}$ and having output systems $\text{Out}\{S\}$, which acts on arbitrary input states $\rho^{\text{In}\{S'\}}$ as follows.

$$\begin{aligned} & \text{Tr}_{SR} \left[|\Phi^+\rangle \langle \Phi^+|^{SR} \left(\mathcal{M} \otimes \mathcal{I}^R (\rho^{\text{In}\{S'\}} \otimes |\Phi^+\rangle \langle \Phi^+|^{S'R}) \right) |\Phi^+\rangle \langle \Phi^+|^{SR} \right] \\ &= \langle \Phi^+|^{SR} \left(\mathcal{M} \otimes \mathcal{I}^R (\rho^{\text{In}\{S'\}} \otimes |\Phi^+\rangle \langle \Phi^+|^{S'R}) \right) |\Phi^+\rangle^{SR}. \end{aligned} \quad (\text{B2})$$

It is clear that the above map acts as a CPM on the input ρ as it is formed from the sequential composition of CPMs (preparation of $|\Phi^+\rangle$ on a disjoint set of subsystems, application of \mathcal{M} and \mathcal{I} , projection on $|\Phi^+\rangle$ and tracing out, are all CPMs). In the previous paragraph, we have already argued why the sequential composition of the CPMs in Eq. (B1) is itself a CPM, it is easy to see that this argument holds for arbitrary CPMs and sequential compositions. Finally, we show that the map defined above is in fact the same as the loop composed map $\mathcal{M}^{S \rightarrow S'}$, which will complete the proof and establish that loop compositions map CPMs to CPMs.

$$\begin{aligned} & \langle \Phi^+|^{SR} \left(\mathcal{M} \otimes \mathcal{I}^R (\rho^{\text{In}\{S'\}} \otimes |\Phi^+\rangle \langle \Phi^+|^{S'R}) \right) |\Phi^+\rangle^{SR} \\ &= \sum_{i,j,k,l} \langle ii|^{SR} \left(\mathcal{M} \otimes \mathcal{I}^R (\rho^{\text{In}\{S'\}} \otimes |jj\rangle \langle kk|^{S'R}) \right) |ll\rangle^{SR} \\ &= \sum_{i,j,k,l} \langle i|^{S'} \left(\mathcal{M} (\rho^{\text{In}\{S'\}} \otimes |j\rangle \langle k|^{S'}) \right) |l\rangle^S \langle i|^R \mathcal{I}^R |j\rangle \langle k|^R |l\rangle^R \\ &= \sum_{i,k} \langle i|^{S'} \left(\mathcal{M} (\rho^{\text{In}\{S'\}} \otimes |i\rangle \langle k|^{S'}) \right) |k\rangle^S \\ &= \mathcal{M}^{S \rightarrow S'} (\rho^{\text{In}\{S'\}}). \end{aligned} \quad (\text{B3})$$

□

Lemma 6. [Composition order independence] Consider a CPM \mathcal{M} such that $S', R' \in \text{In}(\mathcal{M})$ and $S, R \in \text{Out}(\mathcal{M})$ with $\mathcal{H}^S \cong \mathcal{H}^{S'}$ and $\mathcal{H}^R \cong \mathcal{H}^{R'}$. Then we have the following

$$\mathcal{M}^{\{S \rightarrow S', R \rightarrow R'\}} := \left(\mathcal{M}^{S \rightarrow S'} \right)^{R \rightarrow R'} = \left(\mathcal{M}^{R \rightarrow R'} \right)^{S \rightarrow S'}. \quad (\text{B4})$$

That is, the order of loop compositions do not matter.

Proof. For simplicity and without loss of generality, we take R' to be the first input subsystem and S' to be the last input subsystem of \mathcal{M} . Then we have for arbitrary input state $\rho^{\text{In}\{R', S'\}}$

$$\left(\mathcal{M}^{S \rightarrow S'} \right)^{R \rightarrow R'} (\rho^{\text{In}\{R', S'\}}) = \sum_{k,\ell} \sum_{i,j} \langle k|^R \langle i|^S \left(\mathcal{M} (|k\rangle \langle \ell|^{R'} \otimes \rho^{\text{In}\{R', S'\}} \otimes |i\rangle \langle j|^{S'}) \right) |j\rangle^S |\ell\rangle^R \quad (\text{B5})$$

$$\left(\mathcal{M}^{R \rightarrow R'} \right)^{S \rightarrow S'} (\rho^{\text{In}\{R', S'\}}) = \sum_{i,j} \sum_{k,\ell} \langle k|^R \langle i|^S \left(\mathcal{M} (|k\rangle \langle \ell|^{R'} \otimes \rho^{\text{In}\{R', S'\}} \otimes |i\rangle \langle j|^{S'}) \right) |j\rangle^S |\ell\rangle^R \quad (\text{B6})$$

These two expressions only differ in the order of the summation, but this order does not matter (these are finite sums) and we can equivalently write $\sum_{i,j} \sum_{k,\ell} = \sum_{k,\ell} \sum_{i,j} = \sum_{i,j,k,\ell}$ and we have the following, which completes the proof.

$$\begin{aligned} & \left(\mathcal{M}^{S \rightarrow S'} \right)^{R \rightarrow R'} (\rho^{\text{In}\{R', S'\}}) = \left(\mathcal{M}^{R \rightarrow R'} \right)^{S \rightarrow S'} (\rho^{\text{In}\{R', S'\}}) \\ &= \sum_{i,j,k,\ell} \langle k|^R \langle i|^S \left(\mathcal{M} (|k\rangle \langle \ell|^{R'} \otimes \rho^{\text{In}\{R', S'\}} \otimes |i\rangle \langle j|^{S'}) \right) |j\rangle^S |\ell\rangle^R \\ &:= \mathcal{M}^{\{S \rightarrow S', R \rightarrow R'\}} (\rho^{\text{In}\{R', S'\}}). \end{aligned} \quad (\text{B7})$$

¹² The fact that we use non-normalised maximally entangled states does not affect the arguments regarding the CP property, since the overall normalisation constant that would result from using normalised states in the pre and post selection would be $\frac{1}{d_S}$, which is positive.

We note that the composition order independence under multiple loop compositions was first shown in [21] for so-called causal boxes which are CPMs with spacetime labels, satisfying relativistic notions of causality in the spacetime. The relation between our framework and causal boxes is discussed in Appendix D. The above proof closely mirrors the original proof, but we have provided it here for completeness. We note that the original proof also applies to the infinite dimensional case (as causal boxes can act on infinite dimensional Fock spaces as input), where the loop composition is defined in a different manner (in terms of the infinite dimensional Choi-Jamiolkowski representation) but the definition reduces to the finite dimensional definition provided here. For this paper, the finite dimensional case will suffice. \square

Applying this lemma recursively, we obtain the following corollary.

Corollary 5. *Let \mathcal{M} be a CPM, let $\{S_k\}_{k \in \{0,1,\dots,K\}}$ be any K distinct elements of $\text{Out}(\mathcal{M})$ and let $\{S'_k\}_{k \in \{0,1,\dots,K\}}$ be any K distinct elements of $\text{In}(\mathcal{M})$ such that we have $\mathcal{H}^{S_k} \cong \mathcal{H}^{S'_k}$ for all $k \in \{0,1,\dots,K\}$. Then, denoting $\mathfrak{M} = \{S_k \leftrightarrow S'_k\}_{k \in \{0,1,\dots,K\}}$, we have, for any two permutations π and σ of the set $\{0,1,\dots,K\}$,*

$$\begin{aligned} \mathcal{M}^{\{S_k \leftrightarrow S'_k \in \mathfrak{M}\}} &:= \left(\left(\left(\mathcal{M}^{S_{\pi(1)} \leftrightarrow S'_{\pi(1)}} \right)^{S_{\pi(2)} \leftrightarrow S'_{\pi(2)}} \right) \dots \right)^{S_{\pi(K)} \leftrightarrow S'_{\pi(K)}} \\ &= \left(\left(\left(\mathcal{M}^{S_{\sigma(1)} \leftrightarrow S'_{\sigma(1)}} \right)^{S_{\sigma(2)} \leftrightarrow S'_{\sigma(2)}} \right) \dots \right)^{S_{\sigma(K)} \leftrightarrow S'_{\sigma(K)}} \end{aligned} \quad (\text{B8})$$

Appendix C: Illustrative examples of fine-graining and link to causal models

The causal modelling paradigm [1] provides an information-theoretic approach for connecting observable correlations and effects of interventions to underlying causal explanations. It was originally developed for scenarios that can be described by classical variables, and has found useful applications across (classical) data sciences. In recent years, there has been great progress in generalising this approach to define quantum and non-classical causal models [7, 12–20]. We briefly outline some prominent causal modelling frameworks and how their main objects of study can be seen as special cases of the quantum information-theoretic networks that can be described in our formalism.

In a classical deterministic (or functional) causal models, the nodes of the causal structure (a directed graph, as before) are random variables and each variable X with a non-empty set $\text{par}(X)$ of parents is obtained by applying a deterministic function $f_X : \text{par}(X) \mapsto X$ to the parental variables. For parentless variables X , a probability distribution $P(X)$ is specified in the model. More generally, the formalism of classical Bayesian networks allows for an arbitrary stochastic map (or equivalently, a classical channel) from the parental set of each variable to the variable. For the case of acyclic causal structures, this has also been generalised to quantum and post-quantum probabilistic theories by considering the information-theoretic channels of the theory [13]. Every quantum Bayesian network (which includes functional causal models and classical Bayesian networks as special cases) uniquely identifies an acyclic quantum circuit constructed by the parallel and sequential composition of quantum channels (CPTPMs) which correspond to networks of CPTPMs (since CPTPMs are closed under parallel and sequential composition). Therefore our framework can describe these as special case. Moreover, [19, 59] introduces frameworks for split-node quantum causal models for acyclic and cyclic graphs. Here each node N is not a system, but rather a “slot” associated with an input system N_I and output system N_O between which a quantum operation $\mathcal{M}_N : N_I \mapsto N_O$ can be plugged in. Indefinite causal order process matrices can be described as cyclic causal models in this formalism. Objects arising in such split-node causal models can also be studied in our framework analogously to process matrices (cf. Section VII), by defining an extended local map in each slot.

Further details on the general mapping between these approaches will appear in future work [72]. Here, we focus on intuitive examples that illustrate how our framework can capture cyclic functional models, and how our concept of fine-graining apply to them. We will discuss two types of fine-grainings that can be considered. The procedure would be analogous for the quantum cases, as illustrated for the cyclic causal structure of the quantum switch in Section IX and Appendix G.

Example 2 (Fine-graining by splitting into smaller subsystems). Consider the directed cyclic graph of Fig. 6c and the following functional causal model on this graph where the nodes A , B , and D correspond to random variables with cardinality 4, 2 and 2 respectively. Let $P(D)$ be some probability distribution on D and the functions of the

model be $f_A : (D, B) \mapsto A$ which maps $(D, B) = (0, 0), (0, 1), (1, 0), (1, 1)$ to $A = 0, 1, 2, 3$ respectively, and $f_B : A \mapsto B$ sets $B = 0$ whenever $A \in \{0, 1\}$ and $B = 1$ whenever $A \in \{2, 3\}$. This causal model can equivalently be described as a cyclic network (Fig. 6a), where we have two CPTPMs $\mathcal{M}_1 : \mathcal{L}(D) \otimes \mathcal{L}(B') \mapsto \mathcal{L}(A)$ and $\mathcal{M}_2 : \mathcal{L}(A') \mapsto \mathcal{L}(B)$ which are connected by loop compositions $A \mapsto A'$ and $B \mapsto B'$, along with a preparation on D , which can also be seen as a CPTPM. The in/output systems of \mathcal{M}_1 and \mathcal{M}_2 encode the RVs in the computational basis and these CPTPMs implement the functions f_A and f_B respectively (in the computational basis).

Now, notice that we can split A into two bits A_1, A_2 by identifying $A = 0, 1, 2, 3$ with $(A_1, A_2) = (0, 0), (0, 1), (1, 0), (1, 1)$ through an isomorphism. Then D specifies the first bit A_1 while B specifies the second bit A_2 , and in addition B is itself the first bit A_1 . Thus we can equivalently describe the situation through the functional relations $A_1 = D, B = A_1$ and $A_2 = B$, with the same distribution $P(D)$ over D . This specifies a causal model on the acyclic causal structure Fig. 6d. Similarly, this defines an acyclic network as shown in Fig. 6b which is a fine-graining of the cyclic network of Fig. 6a associated with the systems fine-graining $\mathcal{F}(A) = \{A_1, A_2\}$, $\mathcal{F}(B) = \{B\}$ and $\mathcal{F}(D) = \{D\}$ and with a fine-grained map \mathcal{M}_i^f associated with each \mathcal{M}_i of the original network. The relevant encoders and decoders simply act as identities on the B and D systems and map between the 4-dimensional system A and the two 2-dimensional systems A_1 and A_2 through the isomorphism mentioned above.¹³

Our definitions of fine-graining do not place any constraints on how the dimensions of a system S correspond to the dimensions of the set of fine-grained systems $\mathcal{F}(S)$. In the above example, we fine-grained the network by splitting the 4-dimensional system A into two smaller subsystems A_1 and A_2 of 2-dimensions each. We now present a simple example of another type of fine-graining which can be associated with the uncertainty in location of a message in a background causal structure, where the dimensions of the coarse-grained system does not correspond to a product of the dimensions of the corresponding fine-grained systems.

Suppose that Alice and Bob share a classical channel and a common source of randomness Λ and they execute the following protocol. Whenever $\Lambda = 0$, Alice independently generates a bit A and sends it to Bob who outputs a bit $B = A$ and whenever $\Lambda = 1$, Bob independently generates a bit B and sends it to Alice who outputs $A = B$. Whenever $\Lambda = 0$, Alice acts before Bob and whenever $\Lambda = 1$, Alice acts after Bob and whenever Λ is unknown, there appears to be no definite acyclic order between Alice and Bob's operations. This is however a physically plausible protocol, which we can interpret as follows in terms of a definite acyclic causal structure, by breaking down Alice's operations (and associated system A) in two steps (associated with A_1 and A_2) and distinguishing when she exchanges a bit valued message with Bob and when she exchanges no messages with Bob (equivalently modelled as a trivial message $\Omega \neq 0, 1$ exchanged with Bob). Then, when $\Lambda = 0$, A_1 contains a bit valued message which is sent from Alice to Bob while $A_2 = \Omega$ and when $\Lambda = 1$, $A_1 = \Omega$ while A_2 equals the bit valued message that Alice receives from Bob. Thus the bit valued message which was earlier denoted by A and related to B through a non-definite causal order is now a bit that is "non-localised" in a well defined acyclic causal structure over a greater number of nodes A_1, A_2 and B with A_1 preceding B preceding A_2 . The following example formalises this protocol as a cyclic causal model (associated with a cyclic network in our framework) which can be fine-grained to an acyclic causal model (associated with an acyclic network). For this, consider the following two acyclic causal models associated with two different causal orders between the variables A and B .

- CM₀: Graph $R_A \rightsquigarrow A \rightsquigarrow B$ with a model where $P(R_A)$ is arbitrary, $A = R_A, B = A$.
- CM₁: Graph $R_B \rightsquigarrow B \rightsquigarrow A$ with a model where $P(R_B)$ is arbitrary, $B = R_B, A = B$.

The cyclic causal model of the following example can be seen as realising a probabilistic mixture of these two acyclic models.

Example 3 (Fine-graining through uncertainty in location). Consider the cyclic causal structure of Figure 14c and the following causal model over this causal structure. The parentless nodes R_A, R_B and Λ are distributed according arbitrary distributions $P(R_A), P(R_B)$ and $P(\Lambda)$ where $P(\Lambda)$ is not a deterministic distribution. The non-parentless nodes are related by the functional dependences: $A = (\Lambda \oplus 1) \cdot R_A \oplus \Lambda \cdot B := f_A(\Lambda, R_A, B)$ and

¹³ By construction of this example, the only difference between the coarse- and fine-grained networks is in the isomorphism relating the states of the 4-dimensional to those of the two 2-dimensional systems. Therefore it is immediate that the encoders and decoders mentioned in the text would satisfy both properties of Definition 16 for relating the induced map of every sub-network of the original network to the corresponding induced map of the fine-grained network. That is, the acyclic network defined here is indeed a fine-graining of the original cyclic network, as per Definition 17.




FIG. 14: (a) and (b) illustrate the original and fine-grained networks of Example 3, while (c) and (d) illustrate the corresponding causal structures respectively. The maps \mathcal{M}_1 and \mathcal{M}_2 implement the deterministic functions f_A and f_B described in the text of Example 3.

$B = (\Lambda \oplus 1) \cdot A \oplus \Lambda \cdot R_B := f_B(R_B, \Lambda, A)$, where \oplus denotes modulo-2 addition and \cdot denotes multiplication (or logical AND). This causal model can equivalently be described as a cyclic network in our framework, as shown in Fig. 14a where the CPTP maps \mathcal{M}_1 and \mathcal{M}_2 implement the functions f_A and f_B respectively on the corresponding systems (which encode the associated variables). Notice that depending on whether $\Lambda = 0$ or $\Lambda = 1$, we effectively have the acyclic causal model CM_0 or CM_1 , and in a general (since $P(\Lambda)$ is non-deterministic), a probabilistic mixture of the two.

We now construct a fine-graining of this cyclic causal model, which will be associated with the acyclic causal structure of Figure 14d. Consider a causal model on this acyclic graph, where the parentless nodes R_A , R_B and Λ have the same distributions as before. The non-parentless nodes are B , A_1 and A_2 . B is binary as before, while A_1 and A_2 can now take an additional value associated with a symbol Ω , thus making them trinary. The functional dependencies are: $A_1 = (\Lambda \oplus 1) \cdot R_A \oplus \Lambda \cdot \Omega := f_{A_1}(\Lambda, R_A)$, $B = (\Lambda \oplus 1) \cdot A_1 \oplus \Lambda \cdot R_B := f_B(R_B, \Lambda, A_1)$, $A_2 = (\Lambda \oplus 1) \cdot \Omega \oplus \Lambda \cdot B := f_{A_2}(\Lambda \cdot B)$, where $\Omega \oplus 0 = \Omega$, $0 \cdot \Omega = 0$ and $1 \cdot \Omega = \Omega$. As illustrated in Fig. 14b, this defines a fine-graining of the original network, where the fine-graining \mathcal{M}_1^f of the CPM \mathcal{M}_1 encodes the functions f_{A_1} and f_{A_2}

and the fine-graining \mathcal{M}_2^f of the CPM \mathcal{M}_2 encodes the function f_B (along with a trace on A_2). A_1 and A_2 correspond to qutrits encoding the corresponding variables in the basis $\{|0\rangle, |1\rangle, |\Omega\rangle\}$. Here we have, whenever $\Lambda = 0$, the acyclic model CM_0 but with A_1 playing the role of A and $A_2 = \Omega$ and whenever $\Lambda = 1$, the model CM_1 with A_2 taking the place of A but $A_1 = \Omega$.

Let us describe explicitly, the encoders and decoders involved in the above example. Consider the encoding-decoding scheme that relates $\mathcal{M}_1 : \Lambda \otimes R_A \otimes B \mapsto A$ to $\mathcal{M}_1^f : \Lambda \otimes R_A \otimes B \mapsto A_1 \otimes A_2$. The encoder Enc_1 in this case is a CPTPM from the inputs of \mathcal{M}_1 to the inputs of \mathcal{M}_1^f , since these are the same systems, the encoder is trivial (identity on all systems). The decoder Dec_1 is a map from A_1 and A_2 to A which must be CPTP on the image of $\mathcal{M}_1^f \circ \text{Enc}_1$. Notice that this image only has support on the following values of A_1 and A_2 : $\{(0, \Omega), (1, \Omega), (\Omega, 0), (\Omega, 1)\}$. The decoder performs the mapping $\{(i, \Omega), (\Omega, i)\} \mapsto i$ for every $i \in \{0, 1\}$. This can be physically realised through a measurement of A_1 and A_2 in the basis $\{|0\rangle, |1\rangle, |\Omega\rangle\}$ followed by reporting only the value which lies in $\{0, 1\}$. It is then easy to check that $\mathcal{M}_1 = \text{Dec}_1 \circ \mathcal{M}_1^f \circ \text{Enc}_1$ holds, and also that all signalling relations of \mathcal{M}_1 (each of $\{R_A\}$, $\{\Lambda\}$ and $\{B\}$ signal to $\{A\}$) are preserved in \mathcal{M}_1^f (each of $\{R_A\}$, $\{\Lambda\}$ and $\{B\}$ signal to $\{A_1, A_2\}$). Now consider the encoding-decoding scheme that relates $\mathcal{M}_2 : \Lambda \otimes R_B \otimes A \mapsto B$ to $\mathcal{M}_2^f : \Lambda \otimes R_B \otimes A_1 \otimes A_2 \mapsto B$. In this case, the decoder Dec_2 is trivial. The encoder $\text{Enc}_2 : \Lambda \otimes R_B \otimes A \mapsto \Lambda \otimes R_B \otimes A_1 \otimes A_2$ acts as identity on R_B and whenever $\Lambda = 0$, maps A to $A_1 = A$ and $A_2 = \Omega$ and whenever $\Lambda = 1$, maps A to $A_1 = \Omega$ and $A_2 = A$. It is then easy to verify that $\mathcal{M}_2 = \text{Dec}_2 \circ \mathcal{M}_2^f \circ \text{Enc}_2$ holds and all signalling relations of \mathcal{M}_2 are preserved in \mathcal{M}_2^f .

To show that the network of Fig. 14b is indeed a fine-graining of the network of Fig. 14a according to our Definition 17, we still need to establish that the above encoding and decoding for the individual maps implies the existence of encoding and decoding schemes between every sub-network in the coarse and fine-grained networks. The remaining sub-networks of our cyclic network are: the network itself (whose induced map is simply a constant number, since it has no in/outputs), the set of sub-networks formed by the sequential composition $\mathcal{M}_2 \circ \mathcal{M}_1$ (through the A system) composed with any subset of the preparations on Λ , R_A and R_B and another set of sub-networks formed by the sequential composition $\mathcal{M}_1 \circ \mathcal{M}_2$ (through the system B) composed with any subset of the preparations. The encoders and decoders for the latter two sets of sub-networks are obtained from those constructed above, by simply applying the appropriate sequential composition to both sides of $\mathcal{M}_1 = \text{Dec}_1 \circ \mathcal{M}_1^f \circ \text{Enc}_1$ and $\mathcal{M}_2 = \text{Dec}_2 \circ \mathcal{M}_2^f \circ \text{Enc}_2$ along with composing the same set of preparations (noting that Λ , R_A and R_B remain unchanged in the two networks). In particular, the sub-network obtained by applying the sequential composition $\mathcal{M}_2 \circ \mathcal{M}_1$ on all preparations has an induced map with a single in and output system, both isomorphic to B . The original cyclic network can be recovered by performing a loop composition on B in this sub-network. The same is true in the fine-grained scenario, the sequential composition $\mathcal{M}_2^f \circ \mathcal{M}_1^f$ acting on the initial preparations yields an induced map having a single in and output isomorphic to B and recovers the fine-grained network through loop composition on this in-output pair. Since the coarse and fine-grained sub-networks act identically in this case (as we have argued above), the loop compositions on B in the two cases will yield identical results, and the same induced map for the original network (Fig. 14a) and fine-grained network (Fig. 14b). This shows that the latter network is indeed a fine-graining of the former relative to this encoding-decoding scheme.

These examples highlight that, even in purely classical scenarios, the causal structure that we associate with a protocol, as well its cyclic/acyclic nature depends on the information captured by the nodes of the causal structure and the level of detail to which we specify this information. Example 3 illustrates a cyclic functional causal model that is inspired by the classical switch process [27] which implements a classical mixture of the orders $A < B$ and $B < A$ in the process framework (cf. Section VI). The quantum generalisation of the process, the quantum switch can also be formalised as a cyclic network in our framework, which admits an acyclic fine-graining that has a similar physical interpretation (see Section IX and Appendix G for details).

Appendix D: Physical spacetime realisations are causal boxes

In this appendix, we connect spacetime realisations of quantum networks as defined in our top-down formalism to the previously known framework of *causal boxes* [21]. The main result of this section indicates that the most general physical protocols that can be physically realised in a spacetime are those described by causal boxes. Our present formalism and proofs apply to any protocol involving a finite number of finite-dimensional quantum systems, but there is scope for generalising them to the infinite dimensional case in the future (noting that the causal box framework can indeed model infinite dimensional systems). This result provides a top-down justification for the causal box framework, at least for the case of such finite protocols. For such protocols, we can effectively treat the spacetime

poset \mathcal{T} as being finite, without loss of generality. Only a finite set of spacetime locations (or sufficiently localised regions) will be relevant for describing such protocols and we can take \mathcal{T} to only include those locations as all our statements only concern spacetime locations that can embed a system of the protocol. We will not review the causal box framework here but simply outline why it is a powerful approach with several useful features.

Causal boxes describe quantum information processing protocols in an acyclic background spacetime where agents may send/receive messages at a superposition of different spacetime locations, in a manner consistent with relativistic causality. They generalise the standard quantum information formalism to explicitly include spacetime information and also to treat space and time on a more equal footing, in the spirit of special relativity. The standard formalism defines quantum states on fixed time slices (possibly non-localised in space) while causal box formalism allows us to consider states that are not localised in both space and time. Moreover, they preserve the desired mathematical properties of the usual formalism: they are closed under composition, admit a circuit decomposition into a well-defined sequence of isometries that are compatible with the partial order of the spacetime (the sequence representation), as well as admitting Stinespring and Choi matrices and thus being purifiable to unitaries. More generally, the framework is defined on a full infinite-dimensional Fock space and can capture superpositions of different number of messages exchanged between physical devices/agents. The formalism has useful applications for relativistic quantum information processing, in particular for studying composable cryptographic security also for practical relativistic quantum schemes [64].

Causal boxes are defined relative to a fixed partially order set (which models the spacetime), and they satisfy a certain causality condition relative to this poset. The following theorem connects our physical spacetime realisations of networks in our framework to causal boxes (more precisely these realisations map to the subset of *finite* causal boxes), showing that the causality condition of the causal box formalism can be derived from our relativistic causality condition (or more abstractly, compatibility between signalling structure of CPTPMs and the partial order).

Theorem 6. *Let \mathfrak{N} be a network of CPTPMs whose signalling structure is compatible with an acyclic causal structure \mathcal{G} . Then the induced map of every sub-network of \mathfrak{N} is a causal box relative the partial order induced by \mathcal{G} , therefore such a network \mathfrak{N} is itself completely described by the causal box formalism.*

Proof. Consider any acyclic causal structure \mathcal{G} . Being a directed acyclic graph, it induces a partial order, where we say that $N_1 < N_2$ relative to this partial order if and only if there exists a directed path from N_1 to N_2 in \mathcal{G} . Denote the induced poset as $\mathcal{T}^{\mathcal{G}}$. Let $\mathcal{E}(S) \in \text{Nodes}(\mathcal{G})$ define an embedding of the systems $S \in \mathfrak{N}^{sys}$ in \mathcal{G} and equivalently, in the partial order $\mathcal{T}^{\mathcal{G}}$ i.e., $\mathcal{E}(S) \in \mathcal{T}^{\mathcal{G}}$. We can then associate a “region” $\mathcal{R}^{\mathcal{S}}$ or subset of locations in $\mathcal{T}^{\mathcal{G}}$ with every subset of $\mathcal{S} \subseteq \mathfrak{N}^{sys}$, with $\mathcal{R}^{\mathcal{S}} := \cup_{S \in \mathcal{S}} \mathcal{E}(S)$. Then define

$$\text{Past}(\mathcal{R}^{\mathcal{S}}) := \{P \in \mathcal{T}^{\mathcal{G}} : \exists Q \in \mathcal{R}^{\mathcal{S}}, P < Q\}, \quad (\text{D1})$$

where $<$ is the partial order relation associated with $\mathcal{T}^{\mathcal{G}}$. $\text{Past}(\mathcal{R}^{\mathcal{S}})$ corresponds to all points that are in the past of the region $\mathcal{R}^{\mathcal{S}}$ relative to this partial order.

Let us now focus on the induced map \mathcal{N} of \mathfrak{N} with In and Out denoting the set of in and output systems of \mathcal{N} . Consider $\mathcal{S}_O \subseteq \text{Out}$, and a corresponding $\mathcal{S}_I \subseteq \text{In}$ defined as follows

$$\mathcal{S}_I := \{S \in \text{In} : \mathcal{E}(S) \in \text{Past}(\mathcal{R}^{\mathcal{S}_O})\}. \quad (\text{D2})$$

That is, \mathcal{S}_I consists of all inputs of \mathcal{N} which are embedded at locations in the past of the region $\mathcal{R}^{\mathcal{S}_O}$ associated to \mathcal{S}_O . Then, recalling that any poset \mathcal{T} imposes an order relation \xrightarrow{R} on its region $\Sigma(\mathcal{T})$ (cf. Definition 19), $\mathcal{R}^{\text{In} \setminus \mathcal{S}_I} \not\xrightarrow{B} \mathcal{R}^{\mathcal{S}_O}$ follows from the above definition of \mathcal{S}_I . Imposing compatibility between the signalling structure of \mathfrak{N} and the causal structure \mathcal{G} implies that we must have $\text{In} \setminus \mathcal{S}_I$ does not signal to \mathcal{S}_O in \mathfrak{N} . Applying Definition 10 to the CPTPM \mathcal{N} , this non-signalling relation implies that for all states $\rho \in \mathcal{L}(\text{In} \setminus \mathcal{S}_I)$,

$$\text{Tr}_{\text{Out} \setminus \mathcal{S}_O} \circ \mathcal{N} = \text{Tr}_{\text{Out} \setminus \mathcal{S}_O} \circ \mathcal{N} \circ \text{Tr}_{\text{In} \setminus \mathcal{S}_I}^{\rho}, \quad (\text{D3})$$

where $\text{Tr}_{\text{In} \setminus \mathcal{S}_I}^{\rho}$ is a particular choice of local operation on $\text{In} \setminus \mathcal{S}_I$ which discards the incoming state on the subsystems $\text{In} \setminus \mathcal{S}_I$ and replaces it with the fixed state ρ . As the embedding \mathcal{E} associates each system to a node in \mathcal{G} , we can equivalently represent $\text{Tr}_{\mathcal{S}}$ as $\text{Tr}_{\mathcal{R}^{\mathcal{S}}}$ (where $\mathcal{R}^{\mathcal{S}}$ is the region or subset of nodes in \mathcal{G} , associated to the set of systems \mathcal{S}), to capture the fact that it traces out systems embedded in the region $\mathcal{R}^{\mathcal{S}}$. Then defining a causality function $\chi : \mathcal{R} \mapsto \text{Past}(\mathcal{R})$ for every $\mathcal{R} \subseteq \mathcal{T}^{\mathcal{G}}$, and denoting $\text{Tr}_{\text{Out} \setminus \mathcal{S}_O} \circ \mathcal{N} = \mathcal{N}^{\mathcal{R}^{\mathcal{S}_O}}$ to capture that this map only produces outputs in the region $\mathcal{R}^{\mathcal{S}_O}$ (when embedded in \mathcal{G} according to \mathcal{E}) Eq. (D3) is equivalent to

$$\mathcal{N}^{\mathcal{R}^{S_O}} = \mathcal{N}^{\mathcal{R}^{S_O}} \circ \text{Tr}_{\mathcal{T}^{\mathcal{G}} \setminus \chi(\mathcal{R}^{S_O})}^{\rho}. \quad (\text{D4})$$

Fixing ρ to be the vacuum state $|\Omega\rangle\langle\Omega|$ (0-message space of the causal box Fock space), this is precisely the causality condition for regarding \mathcal{N} as a finite causal box, relative to the causality function χ [21]. It tells us that the outputs of a causal box in a region \mathcal{R} , are independent of the input systems of the causal box that are not in the past of \mathcal{R} , relative to the partial order. We can see that χ is indeed a valid finite causality function. This requires us to show that for every $P \in \mathcal{T}^{\mathcal{G}}$, there exists a finite $n \in \mathbb{N}$ such that $P \notin \chi^n(\mathcal{T}^{\mathcal{G}})$, which ensures that every point in the partial order can be reached in a finite number of steps, through the causality function [21]. To show this, notice that the function χ defined above has the property that $\chi(\mathcal{R}) \subset \mathcal{R}$ for all $\mathcal{R} \subseteq \mathcal{T}^{\mathcal{G}}$. This is because $\mathcal{T}^{\mathcal{G}}$ being a partial order ensures that every region \mathcal{R} contains at least one “latest” point $L \in \mathcal{R}$ such that $L \not\prec P$ for all $P \in \mathcal{R}$. Since $\mathcal{T}^{\mathcal{G}}$ is finite, every point P will eventually be excluded from $\chi^n(\mathcal{T}^{\mathcal{G}})$ for some finite number n applications of χ . This establishes that the induced map \mathcal{N} of \mathfrak{N} is a causal box with respect to the partial order induced by \mathcal{G} . The same argument as above can be extended to the induced map of every sub-network of \mathfrak{N} to show that they are all causal boxes relative to \mathcal{G} (and using the same embedding $\mathcal{E}(S) \in \text{Nodes}(\mathcal{G})$ for $S \in \mathfrak{N}^{sys}$). \square

Corollary 6. *[Physical spacetime realisations are causal boxes] Consider a spacetime realisation of a network \mathfrak{N} of CPTPMs that satisfies relativistic causality in a spacetime \mathcal{T} with respect to an embedding \mathcal{E} . Let $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}^f$ be the fine-graining in the specification of the realisation, which is associated with an acyclic region causal structure $\mathcal{G}_{\mathfrak{N}}^{\text{reg}, f}$. Then the induced maps of all sub-networks of $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}^f$ are causal boxes relative to the order relation defined by $\mathcal{G}_{\mathfrak{N}}^{\text{reg}, f}$. Therefore all physical spacetime realisations of networks in our framework can be described within the causal box formalism.*

Proof. This follows from the proof of the above theorem, by taking $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}^f$ to play the role of the network \mathfrak{N} used in that proof, and the acyclic causal structure $\mathcal{G}_{\mathfrak{N}}^{\text{reg}, f}$ to play the role of \mathcal{G} , while noting that relativistic causality (cf. Definition 24) in particular requires the signalling structure of $\mathfrak{N}_{\mathcal{T}, \mathcal{E}}^f$ to be compatible with the causal structure $\mathcal{G}_{\mathfrak{N}}^{\text{reg}, f}$. \square

Appendix E: Further details of the process matrix framework

In Section VI, we described different classes of processes for the bipartite case. In particular, the set of fixed order processes are relevant for our main results and we review the full multi-partite definition for this class of processes below. This is adapted from definitions originally provided in [31, 32], where the only adaptation is to distinguish the in and output systems A^I and A^O of each agent A rather than grouping them together as a single element.

Definition 28 (Fixed order processes). An N -partite process matrix W is said to be a fixed order process, if there exists a partial order $\mathcal{K}(\mathcal{S}_{I/O})$ on the set $\mathcal{S}_{I/O} = \{A_1^I, A_1^O, \dots, A_N^I, A_N^O\}$ of the input and output systems of the N agents and associated with the binary relations $<_{\mathcal{K}}$ (first element precedes the second), $>_{\mathcal{K}}$ (first element succeeds the second) and $\not\prec_{\mathcal{K}}$ (the elements are unordered) such that the following conditions are satisfied

1. For any $i \in \{1, \dots, N\}$, $A_i^I <_{\mathcal{K}} A_i^O$.
2. For any agent A_i and a subset A^S of the remaining agents, such that $A_i^O \not\prec_{\mathcal{K}} A_S^I, \forall A_S \in A_S$ (which is denoted in short as $A_i^O \not\prec_{\mathcal{K}} A_S^I$) with respect to the partial order $\mathcal{K}(\mathcal{S}_{I/O})$, the joint probability distribution $P(x_1, \dots, x_N | a_1, \dots, a_N)$ (cf. Equation (8)) obtained from W for any choice of local measurements of the N agents does not allow the outcome of any of the agents in A_S to depend on the setting of agent A_i . That is, taking x_S to denote the set of outcomes $\{x_S\}_{S \in A_S}$ of the agents in A_S , we have the following whenever $A_i^O \not\prec_{\mathcal{K}} A_S^I$ with respect to $\mathcal{K}(\mathcal{S}_{I/O})$

$$\begin{aligned} P(x_S | a_1, \dots, a_{i-1}, a_i, a_{i+1}, \dots, a_N) &= \sum_{x_j \notin x_S} P(x_1, \dots, x_N | a_1, \dots, a_N) \\ &= P(x_S | a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_N). \end{aligned} \quad (\text{E1})$$

We now review the process matrix description of the quantum switch discussed in Section IX. The quantum switch can be modelled as a 4-partite process involving the agents A , B , C and D . Where the agent C has a trivial (i.e., 1-dimensional) input system and D has a trivial output system, which we can ignore in the description. C 's output decomposes as $C^O = C_C^O \otimes C_T^O$ and D 's input as $D^I = D_C^I \otimes D_T^I$ corresponding to the control and target systems respectively. Thus C and D act on the control and target, with C acting in the global past of all other agents (i.e., cannot be signalling to by others) and D in the global future of all other agents (i.e., cannot signal to the others). A and B only act on the target qudit. Therefore the dimensions of the non-trivial systems are $d_{C_C^O} = d_{D_C^I} = 2$ and $d_{C_T^O} = d_{D_T^I} = d_{A^I} = d_{A^O} = d_{B^I} = d_{B^O} = d$ where $d \geq 2$. The corresponding process matrix is pure (or equivalently, the process map \mathcal{W}_{QS} is unitary) i.e., $W^{QS} = |W^{QS}\rangle\langle W^{QS}|$ where

$$|W^{QS}\rangle = |\mathbb{1}\rangle^{C_T^O A^I} |\mathbb{1}\rangle^{A^O B^I} |\mathbb{1}\rangle^{B^O D_T^I} |00\rangle^{C_C^O D_C^I} + |\mathbb{1}\rangle^{C_T^O B^I} |\mathbb{1}\rangle^{B^O A^I} |\mathbb{1}\rangle^{A^O D_T^I} |11\rangle^{C_C^O D_C^I} \quad (\text{E2})$$

The situation is illustrated in Figure 10. If the lab C prepares an initial state $(\alpha|0\rangle + \beta|1\rangle)^{C_C^O} \otimes |\psi\rangle^{C_T^O}$ (for arbitrary normalised amplitudes α, β and qudit state $|\psi\rangle$) and labs A and B perform the respective operations $\mathcal{U}^A : A^I \mapsto A^O$ and $\mathcal{V}^B : B^I \mapsto B^O$ (taken to be unitaries), the final state arriving at lab D is given as follows and recovers the final state of Eq. (15) on the control and target.

$$\begin{aligned} & \left((\alpha\langle 0| + \beta\langle 1|)^{C_C^O} \otimes \langle \psi^* |^{C_T^O} \otimes \langle \mathcal{U}^{A^*} |^{A^I A^O} \otimes \langle \mathcal{V}^{B^*} |^{B^I B^O} \right) \cdot |W^{QS}\rangle \\ & = \alpha |0\rangle^{D_C^I} \otimes (\mathcal{V}^B \mathcal{U}^A |\psi\rangle)^{D_T^I} + \beta |1\rangle^{D_C^I} \otimes (\mathcal{U}^A \mathcal{V}^B |\psi\rangle)^{D_T^I}, \end{aligned} \quad (\text{E3})$$

where $\langle \psi^* |$ denotes the complex conjugate of $\langle \psi | = |\psi\rangle^\dagger$ in the computational basis $\{|0\rangle, |1\rangle\}$, such that $\langle \psi^* | i \rangle = \langle i | \psi \rangle$, $i \in \{0, 1\}$. $|\mathcal{U}^{A^*}\rangle^{A^I A^O} = (\mathcal{I} \otimes \mathcal{U}^{A^*}) |\mathbb{1}\rangle^{A^I A^I}$ and similarly for \mathcal{V}^B , where $*$ denotes the complex conjugate in the chosen orthonormal basis.

The process matrix W^{QS} is known to be causally non-separable (i.e., cannot be decomposed as in Equation (9)) but nevertheless causal (i.e., always produces probabilities that decompose as per Equation (10)) [31]. The process map $\mathcal{W}_{QS} : C_C^O \otimes C_T^O \otimes A^O \otimes B^O \mapsto D_C^I \otimes D_T^I \otimes A^I \otimes B^I$ is related to W_{QS} through the Choi isomorphism $W_{QS} = \mathcal{I}^{\text{In}} \otimes \mathcal{W}_{QS}(|\mathbb{1}\rangle\langle\mathbb{1}|^{\text{In}})$, and is obtained from W_{QS} through the inverse Choi isomorphism. Here the superscript In corresponds to the total input system $C_C^O C_T^O A^O B^O$ of the process map. We can also obtain the action of the process map more directly, by considering how it acts on the computational basis $\{|ijkl\rangle^{\text{In}}\}_{ijkl}$ with $i \in \{0, 1\}$ and $j, k, l \in \{0, \dots, d-1\}$, which provides a complete orthonormal basis for its inputs. The process map is unitary, and when the control is $|0\rangle^{C_C^O}$, \mathcal{W}_{QS} behaves as a fixed order process with A before B . In this case it forwards the target from, C_T^O to A^I , then from A^O to B^I and from B^O to D_T^I (as can be read off from the first term of Eq. (E2)). When the control is $|1\rangle^{C_C^O}$, it behaves as a fixed order process with B before A , forwarding C_T^O to B^I , B^O to A^I and A^O to D_T^I and in both cases (as can be read off from the first term of Eq. (E2)), and the control being forwarded to D_C^I . Therefore we have, for all $j, k, l \in \{0, \dots, d-1\}$,

$$\begin{aligned} \mathcal{W}_{QS} : & |0\rangle^{C_C^O} |j\rangle^{C_T^O} |k\rangle^{A^O} |l\rangle^{B^O} \mapsto |0\rangle^{D_C^I} |l\rangle^{D_T^I} |j\rangle^{A^I} |k\rangle^{B^I} \\ \mathcal{W}_{QS} : & |1\rangle^{C_C^O} |j\rangle^{C_T^O} |k\rangle^{A^O} |l\rangle^{B^O} \mapsto |1\rangle^{D_C^I} |k\rangle^{D_T^I} |l\rangle^{A^I} |j\rangle^{B^I} \end{aligned} \quad (\text{E4})$$

Appendix F: Equivalence of notions of signalling in process networks

In this section we define two notions of signalling which are relevant in process networks, one is formulated at the level of the probabilities of the classical settings and outcomes and the other is formulated at the level of the quantum in and output systems of the process map and local operations. We then show an equivalence between these two notions which holds whenever the process map represents a valid process matrix.

Definition 29 (Probabilistic signalling). Let W be an N -partite process matrix associated with the probability distribution $P(x_1, \dots, x_N | a_1, \dots, a_N)$. We say that an agent A_i probabilistically signals to a set A_S of the remaining agents if and only if the setting a_i of A_i is correlated with outcomes $x_S := \{x_j\}_{j \in S}$ of the subset of the remaining $N-1$ agents, associated with the indices $S \subseteq \{1, \dots, i-1, i+1, \dots, N\}$ i.e.,

$$P(x_S | a_1, \dots, a_{i-1}, a_i, a_{i+1}, \dots, a_N) \neq P(x_S | a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_N), \quad (\text{F1})$$

where $P(x_S | a_1, \dots, a_{i-1}, a_i, a_{i+1}, \dots, a_N) = \sum_{x_j : j \notin S} P(x_1, \dots, x_N | a_1, \dots, a_N)$.

Definition 30 (Quantum signalling). Let W be an N -partite process matrix and $\mathfrak{N}_{\mathcal{W},N}$ be the corresponding network in our framework. Let A_i be an agent and A_S be any subset of the remaining agents, and $\mathfrak{N}_{\mathcal{W},i\cup\overline{S}}$ be the sub-network of $\mathfrak{N}_{\mathcal{W},N}$ obtained from composing \mathcal{W} with the extended local operations of all agents except those in $A_i \cup A_S$. We say that an agent A_i quantumly signals to a set A_S of the remaining agents if and only if $\{A_i^O\}$ signals to $\{A_j^I\}_{j \in S}$ in the induced map $\mathcal{N}_{\mathcal{W},i\cup\overline{S}}$.

Theorem 7. [Equivalence of two notions of signalling] *Whenever W is a valid N -partite process matrix, A_i probabilistically signals to A_S if and only if A_i quantumly signals to A_S .*

Proof. We first prove the equivalence for bipartite processes W over agents A (Alice) and B (Bob), and later show how this implies the general N -partite result.

Probabilistic signalling implies quantum signalling Here, we wish to prove that $P(y|ab) \neq P(y|a)$ implies that $\{A^O\}$ signals to $\{B^I\}$ in the process map \mathcal{W} . We will prove the equivalent contrapositive statement which is that $\{A^O\}$ does not signal to $\{B^I\}$ in \mathcal{W} implies $P(y|ab) = P(y|a)$.

First we show that $\{A^O\}$ does not signal to $\{B^I\}$ in \mathcal{W} implies $\{A^O\}$ does not signal to $\{A^I, B^I\}$ in \mathcal{W} . This follows from a result about the decomposition of process matrices proven in [28] which shows that any bipartite process matrix $W \in A^I \otimes A^O \otimes B^I \otimes B^O$ can be expressed in the Hilbert-Schmidt basis in the following general form, where a Hilbert-Schmidt basis for linear operators on a Hilbert space \mathcal{H}^X is given by a set of matrices $\{\sigma_\mu^X\}_{\mu=0}^{d_X^2-1}$ with $\sigma_0^X = \mathbb{1}^X$, $\text{Tr} \sigma_\mu^X \sigma_\nu^X = d_X \delta_{\mu\nu}$ and $\text{Tr} \sigma_j^X = 0$ for all $j \in \{1, \dots, d_X^2 - 1\}$.

$$\begin{aligned}
W &= \frac{1}{d_{A^I} d_{B^I}} \left(\mathbb{1}^{A^I A^O B^I B^O} + \sigma^{B \leq A} + \sigma^{A \leq B} + \sigma^{A \not\leq B} \right), \\
\sigma^{B \leq A} &:= \sum_{i,j>0} c_{ij} \sigma_i^{A^I} \otimes \mathbb{1}^{A^O} \otimes \mathbb{1}^{B^I} \otimes \sigma_j^{B^O} + \sum_{i,j,k>0} d_{ijk} \sigma_i^{A^I} \otimes \mathbb{1}^{A^O} \otimes \sigma_j^{B^I} \otimes \sigma_k^{B^O} \\
\sigma^{A \leq B} &:= \sum_{i,j>0} e_{ij} \mathbb{1}^{A^I} \otimes \sigma_i^{A^O} \otimes \sigma_j^{B^I} \otimes \mathbb{1}^{B^O} + \sum_{i,j,k>0} f_{ijk} \sigma_i^{A^I} \otimes \sigma_j^{A^O} \otimes \sigma_k^{B^I} \otimes \mathbb{1}^{B^O} \\
\sigma^{A \not\leq B} &:= \sum_{i>0} \mathcal{V}_i \sigma_i^{A^I} \otimes \mathbb{1}^{A^O} \otimes \mathbb{1}^{B^I} \otimes \mathbb{1}^{B^O} + \sum_{i>0} x_i \mathbb{1}^{A^I} \otimes \mathbb{1}^{A^O} \otimes \sigma_i^{B^I} \otimes \mathbb{1}^{B^O} \\
&\quad + \sum_{i,j>0} g_{ij} \sigma_i^{A^I} \otimes \mathbb{1}^{A^O} \otimes \sigma_j^{B^I} \otimes \mathbb{1}^{B^O},
\end{aligned} \tag{F2}$$

where $c_{ij}, d_{ijk}, e_{ij}, f_{ijk}, \mathcal{V}_i, x_i, g_{ij} \in \mathbb{R}$. We then use the fact that the non-signalling relations $\{A^O\} \not\rightarrow \{B^I\}$ and $\{A^O\} \not\rightarrow \{A^I, B^I\}$ in the process map \mathcal{W} can be equivalently expressed as the following conditions on its Choi matrix (which is by construction, the process matrix W),

$$\begin{aligned}
\text{Tr}_{A^I} W &= \text{Tr}_{A^I A^O} W \otimes \frac{\mathbb{1}^{A^O}}{d_{A^O}} \\
W &= \text{Tr}_{A^O} W \otimes \frac{\mathbb{1}^{A^O}}{d_{A^O}}.
\end{aligned} \tag{F3}$$

Then we can immediately see for each term of W , except the term $\sigma^{A \leq B}$, both conditions of the above equation are satisfied when replacing W with the term under consideration. Thus a violation of the second condition by a given W could only come from the presence of terms of the form of $\sigma^{A \leq B}$. We see that the term $\sigma^{A \leq B}$ always violates both conditions, whenever there are any non-zero e_{ij} or f_{ijk} coefficients. This is because $\sigma_j^{A^O} \neq \mathbb{1}^{A^O}$ by construction (since the j values start from 1, and for all such $\sigma_j^{A^O}$, the trace is 0 which is distinct from the trace of the identity). Thus, tracing and replacing the A^O system will change the term, and violate the second condition, and by the same argument we can see that it will also change the term even if we have already traced out A^I , hence the first condition will also be violated.

Thus we have shown that in each term of a valid W , if there is a violation of the second condition of Eq. (F3), then there will also be a violation of the first condition. This implies (by linearity) that for all valid W 's this implication holds, and consequently that $\{A^O\}$ signals to $\{A^I, B^I\}$ implies $\{A^O\}$ signals to $\{B^I\}$. Moreover, $\{A^O\}$ signals to $\{A^I, B^I\}$ which is equivalent to the second condition of Eq. (F3) is also shown to be equivalent to $P(y|ab) = P(y)$ in [32]. Therefore, $P(y|ab) \neq P(y|b)$ implies that $\{A^O\}$ signals to $\{A^I, B^I\}$ in \mathcal{W} , which in turn implies that $\{A^O\}$ signals to $\{B^I\}$ in \mathcal{W} and altogether, this establishes that A probabilistically signals to B implies A quantumly signals to B .

Quantum signalling implies probabilistic signalling Suppose that we have the quantum signalling relation $\{A^O\}$ signals to $\{B^I\}$ in \mathcal{W} . For any CPTP map \mathcal{W} , this is equivalent [67] to the condition that there exists a state σ^{B^O} on B^O and distinct states ρ^{A^O} and $\tilde{\rho}^{A^O}$ on A^O such that

$$\text{Tr}_{A^I} \circ \mathcal{W}(\rho^{A^O} \otimes \sigma^{B^O}) \neq \text{Tr}_{A^I} \circ \mathcal{W}(\tilde{\rho}^{A^O} \otimes \sigma^{B^O}).$$

In other words, the state on B^I when \mathcal{W} acts on $\rho^{A^O} \otimes \sigma^{B^O}$ is distinct from the state on B^I when it acts on $\tilde{\rho}^{A^O} \otimes \sigma^{B^O}$ and hence there exists a measurement $\{\mathcal{M}_{y|b}^B\}_y$ that B can perform to distinguish these states with a non-zero probability. Then, we can define A 's local operation \mathcal{M}^A to be such that whenever the setting $a = 0$ is input on A^s , it discards the input on A^I and prepares the state ρ^{A^O} to send to \mathcal{W} and whenever the setting $a = 1$ is input on A^s , it discards the input on A^I and prepares the state $\tilde{\rho}^{A^O}$ to send to \mathcal{W} . This in turn implies that $P(y|a=0, b) \neq P(y|a=1, b)$ or $P(y|ab) \neq P(y|b)$, and establishes the claim.

Generalisation to the multi-partite case For the multi-partite case, we consider the signalling relation A_i signals to the set A_S in terms of the two notions of signalling. We can map this to the proof for the bipartite case which we have shown above, by treating A_i as A in the above proof, and the set A_S of agents together as B . In this case, instead of signalling in \mathcal{W} , we would consider signalling in the induced map $\mathcal{N}_{\mathcal{W}, \overline{i\cup S}}$ of the sub-network $\mathfrak{N}_{\mathcal{W}, \overline{i\cup S}}$ (note that in the bipartite case, this sub-network reduces to the network (\mathcal{W}, \emptyset) whose induced map is \mathcal{W}). Then the proof of the statement that probabilistic implies quantum signalling carries forth unchanged, we can again prove the equivalent contrapositive statement using the same arguments, showing that $\{A_i^O\}$ does not signal to $\{A_j^I\}_{j \in S}$ in $\mathcal{N}_{\mathcal{W}, \overline{i\cup S}}$ (i.e., no quantum signalling) implies that $\{A_i^O\}$ does not signal to $\{A_j^I\}_{j \in i\cup S}$ in $\mathcal{N}_{\mathcal{W}, \overline{i\cup S}}$ which implies that A_i does not probabilistically signal to A_S .

For the statement that quantum signalling implies probabilistic signalling, the only aspect that needs additional consideration is that, when we treat Bob as a collection of more than one agent A_S in the multi-partite case, then he can only apply extended local maps that act independently on each input A_j^I for $j \in S$ and not arbitrary joint operations. Bob's role, as the receiver of the signal is to be able to distinguish any two distinct states that arise on the subsystems $\{A_j^I\}_{j \in S}$ which are the outputs of the CPTP map $\mathcal{N}_{\mathcal{W}, \overline{i\cup S}}$, which arise as a result of different local operations that Alice may perform on the input A_i^O of the same map. Quantum theory satisfies the property that any two states of a multi-partite system can be distinguished through local measurements performed on the subsystems i.e., there will exist a local measurement on each subsystem such that the outcome probabilities for the two states differ when performing those local measurements simultaneously on all the systems. This is known as the *local distinguishability* or *local discriminability* property of quantum theory [93, 94]. This property ensures that all the arguments used above for proving the equivalence in the bipartite case also carry over to the general multi-partite case. \square

Appendix G: Fine-grained acyclic network describing a spacetime realisation of QS

Here we describe a spacetime realisation of the QS network $\mathfrak{N}_{\mathcal{W}_{QS}, U, V}$ that satisfies relativistic causality, and the corresponding fine-grained network which will be associated with a definite and acyclic (information-theoretic) causal structure (as shown in Theorem 5 and Corollary 3). The realisation is associated with the spacetime embedding given in Eq. (16) and Eq. (17) of the main text and the fine-grained description will correspond to a *causal box* [21], as expected from our results of Appendix D. We first detail the fine-grained causal box description, and then show that this is indeed a fine-graining of the original network $\mathfrak{N}_{\mathcal{W}_{QS}, U, V}$ as per Definition 17.

1. Causal box description of the fine-grained network

Consider the quantum switch process map \mathcal{W}_{QS} described in Appendix E, which is known to be a unitary, along with local unitaries \mathcal{U}^A and \mathcal{V}^B of Alice and Bob. Recall from the main text that the corresponding fine-grained process map \mathcal{W}_{QS}^f associated with the spacetime embedding of Eq. (16) and Eq. (17) has the in and output systems, $\text{In} = \{C_C^O, C_T^O, A_1^O, A_2^O, B_1^O, B_2^O\}$ and $\text{Out} = \{D_C^I, D_T^I, A_1^I, A_2^I, B_1^I, B_2^I\}$, where $d_{C_C^O} = d_{D_C^I} = 2$, $d_{C_T^O} = d_{D_T^I} = d$ ($d \geq 2$) and $d_{A_1^I} = d_{A_2^I} = d_{A_1^O} = d_{A_2^O} = d_{B_1^I} = d_{B_2^I} = d_{B_1^O} = d_{B_2^O} = d + 1$. The $d + 1$ dimensional systems are associated with the Hilbert space $\mathbb{C}^d \oplus |\Omega\rangle$, where $|\Omega\rangle$ is the vacuum state. The fine-grained local operations $\mathcal{U}_{1,2}^{A,f} : A_1^I \otimes A_2^I \mapsto A_1^O \otimes A_2^O$




FIG. 15: An equivalent representation of the fine-grained network of Fig. 12: This representation makes explicit that \mathcal{W}_{QS}^f itself is a CPTPM which is composed with the fine-grained local operations through feedback loops. The blue path and orange wires represent the paths of the non-vacuum target when the control is $|0\rangle$ and $|1\rangle$ respectively and correspond to the blue and orange paths in Fig. 12. Although there are feedback loops involved, \mathcal{W}_{QS}^f is a causal box and decomposes into a sequence of CPTPMs as shown in Fig. 16. This means that every causal influence between the systems involved in this network, is that such whenever S is a cause of S' , then S is embedded in the past of S' according to the spacetime embedding associated with this realisation (Eq. (16) and Eq. (17)). As the spacetime forms a partial order, this implies that the causal structure of this network is acyclic as shown in Fig. 12.

and $\mathcal{V}_{1,2}^{B,f} : B_1^I \otimes B_2^I \mapsto B_1^O \mapsto B_2^O$ are given as

$$\mathcal{U}_{1,2}^{A,f} = \mathcal{U}_1^{A,f} \otimes \mathcal{U}_2^{A,f}, \quad \mathcal{V}_{1,2}^{B,f} = \mathcal{V}_1^{B,f} \otimes \mathcal{V}_2^{B,f}, \quad (\text{G1})$$

where $\mathcal{U}_i^{A,f}$ and $\mathcal{V}_i^{B,f}$ act on vacuum and non-vacuum states as in Eq. (18) (leaving the vacuum $|\Omega\rangle$ invariant and applying a qudit unitary \mathcal{U}_i^A or \mathcal{V}_i^B on a non-vacuum input). We repeat this action here for convenience.

$$\begin{aligned} \mathcal{U}_i^{A,f} |\Omega\rangle^{A_i} &= |\Omega\rangle^{A_i^O} \\ \mathcal{U}_i^{A,f} |\psi\rangle^{A_i} &= \mathcal{U}_i^A |\psi\rangle^{A_i^O} \end{aligned} \quad (\text{G2})$$

The resultant fine-grained network is identical to that of Fig. 12 at the information-theoretic level (maps and compositions) and can be equivalently represented as shown in Fig. 15.

We now proceed to describe the internal decomposition of \mathcal{W}_{QS}^f . This is identical to the description of the quantum switch originally given in the causal box framework [21] and we repeat it here in our notation. \mathcal{W}_{QS}^f can be decomposed

into three internal CPTP maps, \mathcal{QS}_1^f , \mathcal{QS}_2^f and \mathcal{QS}_3^f , which are illustrated in Fig. 16. When \mathcal{W}_{QS}^f is composed with the fine-grained local operations $\mathcal{U}_{1,2}^{A,f}$ and $\mathcal{V}_{1,2}^{B,f}$ described above, we have the following well-defined sequence of operations in the network.

1. \mathcal{QS}_1^f Receives the initial state of control and target $(\alpha|0\rangle + \beta|1\rangle)^{C_C^O} \otimes |\psi\rangle^{C_T^O}$ from C , stores the control in an internal memory Q , and forwards $|\psi\rangle$ to A_1^I and $|\Omega\rangle$ to B_1^I when the control is in state $|0\rangle$, and $|\Omega\rangle$ to A_1^I and $|\psi\rangle$ to B_1^I when the control is in the state $|1\rangle$.
2. $\mathcal{U}_1^{A,f}$ and $\mathcal{V}_1^{B,f}$ are applied on the inputs A_1^I and B_1^I and the corresponding outputs A_1^O and B_1^O are fed back to \mathcal{W}_{QS}^f .
3. \mathcal{QS}_2^f Forwards the state in A_1^O to B_2^I and the state in B_1^O to A_2^I .
4. $\mathcal{U}_2^{A,f}$ and $\mathcal{V}_2^{B,f}$ are applied on the inputs A_2^I and B_2^I and the corresponding outputs A_2^O and B_2^O are fed back to \mathcal{W}_{QS}^f .
5. \mathcal{QS}_3^f consults the control qubit stored in Q , if it is in the state $|0\rangle$, forwards the state in the output B_2^O to D_T^I and if the control is in the state $|1\rangle$, forwards the output A_2^O to D_T^I , and in both cases forwards the control in Q to D_C^I .

One can easily verify (see [21] for details) that this implements the transformation of Eq. (19) from $C_C^O \otimes C_T^O$ to $D_C^I \otimes D_T^I$ whenever the local maps act on the vacuum as given in Eq. (G2) and in the case where $\mathcal{U}_1^{A,f} = \mathcal{U}_2^{A,f}$ and $\mathcal{V}_1^{B,f} = \mathcal{V}_2^{B,f}$ this implements the quantum switch transformation of Eq. (15). Moreover, as shown in [21], one can additionally include counters with the local operations, initialised to $|0\rangle$ and which increment by 1 whenever the local operation is applied on a non-vacuum state. At the end of the protocol, both Alice and Bob's counters would then record 1, which verifies that in this network, each of Alice and Bob's operations act exactly once on a non-vacuum d -dimensional state.

Moreover, this sequence decomposition ensures that the fine-grained causal structure of the network is acyclic as shown in Fig. 12, with a causal influence $S \rightsquigarrow S'$ only if the fine-grained systems S and S' are associated with spacetime locations P^S and $P^{S'}$ (through the embedding \mathcal{E} of Eq. (16)) where $P^S < P^{S'}$. Consequently, the signalling relations of the fine-grained network are compatible with the partial order of the spacetime. We can view the fine-grained map \mathcal{W}_{QS}^f as a process map associated with the agents A_1, A_2, B_1, B_2, C and D . Notice that the in and output systems of these 6 agents are localised in the spacetime according to the embedding Eq. (16), and satisfy relativistic causality with respect to this embedding. This implies, by Corollary 1 that \mathcal{W}_{QS}^f must correspond to a fixed order process. The fine-grained causal structure of Fig. 12 along with the embedding Eq. (16) make it clear that the order involves C and D being in the global past and future respectively of other agents, A_1 before B_2 and B_1 before A_2 (with A_1 and A_2, B_1 and B_2 not necessarily ordered among themselves).

Here, we have only considered the 0-message (vacuum) and 1-message spaces of d dimensions, for Alice and Bob. More generally, all fine-grained systems can be modelled as Fock spaces which allow for multiple messages of a given dimension and the above fine-grained operations can be consistently extended to the full Fock space as shown in [21]. In follow up work based on the master's thesis [77], we show that such causal box descriptions associated with an acyclic causal structure, and their extensions to the full Fock space can also be obtained for a much general class of processes called QCQCs [41], which realise arbitrary quantum controlled superpositions of orders.

2. Encoders and decoders relating the fine and coarse grained networks

We now describe the encoder and decoder that allow us to recover the coarse-grained network $\mathfrak{N}_{\mathcal{W}_{QS}, U, V}$ from the fine-grained causal box description of the previous sub-section, and show that these respect the conditions of Definition 17. Having described the coarse and fine-grained systems of the network, we know that the encoder and decoder relating \mathcal{W}_{QS} and \mathcal{W}_{QS}^f must be maps of the following form.

$$\begin{aligned} \text{Enc}_{QS} : C_C^O \otimes C_T^O \otimes A^O \otimes B^O &\mapsto C_C^O \otimes C_T^O \otimes A_1^O \otimes A_2^O \otimes B_1^O \otimes B_2^O, \\ \text{Dec}_{QS} : D_C^I \otimes D_T^I \otimes A_1^I \otimes A_2^I \otimes B_2^I \otimes B_1^I &\mapsto D_C^I \otimes D_T^I \otimes A^I \otimes B^I. \end{aligned} \tag{G3}$$




FIG. 16: Sequence representation of the causal box \mathcal{W}_{QS}^f of Figure 15

A first requirement of Definition 17 is that Enc_{QS} must be CPTP on the whole input space while Dec_{QS} must be CPTP on the image of $\mathcal{W}_{QS}^f \circ \text{Enc}_{QS}$. We completely specify the action of Enc_{QS} by its action on a complete orthonormal basis of $C_C^O \otimes C_T^O \otimes A^O \otimes B^O$, which is given by the computational basis $\{|ijkl\rangle\}_{i \in \{0,1\}, j,k,l \in \{0,\dots,d-1\}}$. This is given as follows, for any $j, k, l \in \{0, \dots, d-1\}$

$$\begin{aligned} \text{Enc}_{QS} : |0\rangle^{C_C^O} |j\rangle^{C_T^O} |k\rangle^{A^O} |l\rangle^{B^O} &\mapsto |0\rangle^{C_C^O} |j\rangle^{C_T^O} |k\rangle^{A_1^O} |\Omega\rangle^{A_2^O} |\Omega\rangle^{B_1^O} |l\rangle^{B_2^O}, \\ \text{Enc}_{QS} : |1\rangle^{C_C^O} |j\rangle^{C_T^O} |k\rangle^{A^O} |l\rangle^{B^O} &\mapsto |0\rangle^{C_C^O} |j\rangle^{C_T^O} |\Omega\rangle^{A_1^O} |k\rangle^{A_2^O} |l\rangle^{B_1^O} |\Omega\rangle^{B_2^O}. \end{aligned} \quad (\text{G4})$$

We can see that such an encoder is in fact unitarity (and therefore a CPTPM). We could physically realise it as follows, including an ancilla E_A for Alice and ancilla E_B for Bob, both initialised in the state $|\Omega\rangle$, then performing a controlled swap which sends A^O to A_1^O and E_A to A_2^O when $|0\rangle^{C_C^O}$ and A^O to A_2^O and E_A to A_1^O when $|1\rangle^{C_C^O}$, and another analogous controlled swap for Bob. Then we can see that the image of $\mathcal{W}_{QS}^f \circ \text{Enc}_{QS}$ (using that \mathcal{W}_{QS}^f can be decomposed as in Fig. 16) is always such that it lies within the subspace defined by the set of projectors $\{\Pi_{0jkl}^f, \Pi_{1jkl}^f\}_{jkl \in \{0,\dots,d-1\}}$ associated with the following states

$$\begin{aligned} |\Psi_{0jkl}^f\rangle^{\text{Out}} &:= |0\rangle^{D_C^I} |j\rangle^{D_T^I} |k\rangle^{A_1^I} |\Omega\rangle^{A_2^I} |\Omega\rangle^{B_1^I} |l\rangle^{B_2^I}, \\ |\Psi_{1jkl}^f\rangle^{\text{Out}} &:= |1\rangle^{D_C^I} |j\rangle^{D_T^I} |\Omega\rangle^{A_1^I} |k\rangle^{A_2^I} |l\rangle^{B_1^I} |\Omega\rangle^{B_2^I}. \end{aligned} \quad (\text{G5})$$

Then consider a decoder that acts as follows for all i, j, k

$$\begin{aligned} \text{Dec}_{QS} : |\Psi_{0jkl}^f\rangle^{\text{Out}} &\mapsto |0\rangle^{D_C^I} |j\rangle^{D_T^I} |k\rangle^{A^I} |l\rangle^{B^I}, \\ \text{Dec}_{QS} : |\Psi_{1jkl}^f\rangle^{\text{Out}} &\mapsto |1\rangle^{D_C^I} |j\rangle^{D_T^I} |k\rangle^{A^I} |l\rangle^{B^I}. \end{aligned} \quad (\text{G6})$$

Such a decoder will be unitary (and therefore CPTP) on the image of $\mathcal{W}_{QS}^f \circ \text{Enc}_{QS}$. We could physically realise this by sending the state in A_1^I or A_2^I to A^I depending on D_C^I and tracing out the other state, and similarly for Bob. Note that when applied on the image of $\mathcal{W}_{QS}^f \circ \text{Enc}_{QS}$, the state being traced out will always be the vacuum $|\Omega\rangle$ and will decouple from the rest such that the decoder is unitary on the subspace and maps pure states to pure states as described.

The next step is to show that $\mathcal{W}_{QS} = \text{Dec}_{QS} \circ \mathcal{W}_{QS}^f \circ \text{Enc}_{QS}$ holds for the encoder and decoder defined above. \mathcal{W}_{QS} is unitary and its action is as given in Eq. (E4). From the causal box description of \mathcal{W}_{QS}^f given in the previous sub-section, we can see that it is unitary on the image of Enc_{QS} and acts as follows (noting from Eq. (G4) that the image of the encoder is a subspace of $C_C^O \otimes C_T^O \otimes A_1^O \otimes A_2^O \otimes B_1^O \otimes B_2^O$ that is spanned by states of the form given

below)

$$\begin{aligned}\mathcal{W}_{QS}^f &: |0\rangle^{C_C^O} |j\rangle^{C_T^O} |k\rangle^{A_1^O} |\Omega\rangle^{A_2^O} |\Omega\rangle^{B_1^O} |l\rangle^{B_2^O} \mapsto |0\rangle^{D_C^I} |l\rangle^{D_T^I} |j\rangle^{A_1^I} |\Omega\rangle^{A_2^I} |\Omega\rangle^{B_1^I} |k\rangle^{B_2^I}, \\ \mathcal{W}_{QS}^f &: |1\rangle^{C_C^O} |j\rangle^{C_T^O} |\Omega\rangle^{A_1^O} |k\rangle^{A_2^O} |l\rangle^{B_1^O} |\Omega\rangle^{B_2^O} \mapsto |1\rangle^{D_C^I} |k\rangle^{D_T^I} |\Omega\rangle^{A_1^I} |l\rangle^{A_2^I} |j\rangle^{B_1^I} |\Omega\rangle^{B_2^I}.\end{aligned}\tag{G7}$$

Notice that the final states above correspond to the states $|\Psi_{0ljk}^f\rangle$ and $|\Psi_{1klj}^f\rangle$ of Eq. (G5). Then it is immediate from combining the description of the encoder (Eq. (G4)), the description of the fine-grained process map (Eq. (G7)) and that of the decoder (Eq. (G6)), that we recover the action of the coarse-grained process (Eq. (E4)) i.e., $\mathcal{W}_{QS} = \text{Dec}_{QS} \circ \mathcal{W}_{QS}^f \circ \text{Enc}_{QS}$ holds. Moreover, it is easy to see that all the signalling relations are preserved, in particular the signalling relations $\{A^O\} \rightarrow \{B^I\}$ and $\{B^O\} \rightarrow \{A^I\}$ between the main agents in \mathcal{W}_{QS} is preserved in the corresponding signalling relations $\{A_1^O, A_2^O\} \rightarrow \{B_1^I, B_2^I\}$ and $\{B_1^O, B_2^O\} \rightarrow \{A_1^I, A_2^I\}$ in \mathcal{W}_{QS}^f .

Recall that in order to regard the causal box realisation of the process network $\mathfrak{N}_{\mathcal{W}_{QS}, U, V}$ as its fine-graining, we must compare all corresponding sub-networks of the two descriptions and find corresponding encoders and decoders to translate between their induced maps. \mathcal{W}_{QS} that we have analysed above is the induced map of one such sub-network of $\mathfrak{N}_{\mathcal{W}_{QS}, U, V}$. Other sub-networks of $\mathfrak{N}_{\mathcal{W}_{QS}, U, V}$ are associated with the individual local maps $\mathcal{U}^A, \mathcal{V}^B$ (without any composition) as well as the networks formed by composing \mathcal{W}_{QS} with one or both of \mathcal{U}^A and \mathcal{V}^B . The latter type of networks either correspond to fully connecting a local operation to \mathcal{W}_{QS} by composing the inputs and outputs of the operation with \mathcal{W}_{QS} , or sequentially composing it to \mathcal{W}_{QS} (in any order) which would correspond to connecting only its input system or only its output system to the corresponding system of \mathcal{W}_{QS} with the same label. In all of these cases, we can construct the corresponding encoders and decoders in an entirely analogous manner as we did for \mathcal{W}_{QS} in the above paragraphs, and we therefore do not repeat the construction for each case. For instance, consider the sub-network of $\mathfrak{N}_{\mathcal{W}_{QS}, U, V}$ obtained by composing \mathcal{V}^B but not \mathcal{U}^A with \mathcal{W}_{QS} . The induced map then maps $C_C^O \otimes C_T^O \otimes A^O$ to $D_C^I \otimes D_T^I \otimes A^I$ and the corresponding induced map in the fine-grained network maps $C_C^O \otimes C_T^O \otimes A_1^O \otimes A_2^O$ to $D_C^I \otimes D_T^I \otimes A_1^O \otimes A_2^O$. The encoder and decoder in the case would be constructed in the same manner as Eq. (G4) and Eq. (G6) to ensure that the non-vacuum and vacuum states end up in the right fine-grained system for Alice depending on the control. They would also satisfy all the required properties by the same arguments as before.

The only remaining sub-networks are those corresponding to the individual local operations \mathcal{U}^A and \mathcal{V}^B . Consider the fine-grained local maps of Eq. (G1), which act on the vacuum state according to Eq. (18). Further suppose that one the fine-grained operations \mathcal{U}_i^A acting on the non-vacuum subspace, say for $i = 1$, is identical to the coarse-grained operations \mathcal{U}^A , and similarly for Bob i.e., $\mathcal{U}_1^A = \mathcal{U}^A$ and $\mathcal{V}_1^B = \mathcal{V}^B$. Here, the encoders and decoders are maps of the following form.

$$\begin{aligned}\text{Enc}_A &: A^I \mapsto A_1^I \otimes A_2^I, & \text{Enc}_B &: B^I \mapsto B_1^I \otimes B_2^I, \\ \text{Dec}_A &: A_1^O \otimes A_2^O \mapsto A^O, & \text{Dec}_B &: B_1^O \otimes B_2^O \mapsto B^O.\end{aligned}\tag{G8}$$

Then it is sufficient to take encoders are simply identity channels from the coarse-grained system (e.g., A^I) to the first fine-grained system (e.g., A_1^I) and prepare an arbitrary, independent state on the second fine-grained system. Similarly, it is sufficient to consider decoders that act as identity channels from the first fine-grained system to the coarse-grained system and trace out the second fine-grained system. Then it is clear that the necessary properties to regard the map $\mathcal{U}_{1,2}^{A,f}$ as a fine-graining of \mathcal{U}^A and the map $\mathcal{V}_{1,2}^{B,f}$ as a fine-graining of \mathcal{V}^B are trivially satisfied.

From the above, it follows that the causal box description of the previous sub-section is indeed a fine-graining of the process network (Definition 17) and preserves all of its signalling relations, even though it admits an acyclic causal structure (unlike the original coarse-grained process network which has a cyclic causal structure).

Remark 6 (Local distinguishability of the order). The fixed spacetime realisation of QS described in Figure 12 and in this section is such that the spacetime location at which Alice or Bob's operation is applied (on a non-vacuum state) is perfectly correlated with the control. Therefore, Alice and Bob can perfectly and locally distinguish the two orders by measuring the time of arrival (or more generally, spacetime location of arrival) of a non-vacuum state to their lab. This would of course collapse the superposition of orders. Alternatively, they can use the location of arrival as a control parameter based on which they apply different unitary operations depending on the order (cf. Eq. (19)). In [45, 49], a different type of experimental realisation of QS in Minkowski spacetime has been proposed, which has the interesting property that local measurements by Alice and Bob of the (space)time location at which they receive a non-vacuum state would not reveal significant information about the control since there is a large uncertainty in this spacetime

location even when the operations are applied in a fixed order (i.e., where the control is in one of the computational basis states). Our results are general and apply to both these type of realisations, and independently of the agents' ability to perfectly distinguish the orders locally. For the realisations of [45, 49], the fine-grained description would not correspond to that of Figure 12 but one where Alice and Bob act on a larger number of fine-grained systems (associated with more than two spacetime locations) such that the location at which they receive a non-vacuum state would not tell them much about the order. This means that we would need to consider a larger number of agents than 6 in order to describe the experiment using a fixed order processes (cf. Theorem 5), and this also means that we may not be able to perfectly implement the transformation of Eq. (19) where the unitary applied by each agent is perfectly correlated with the order in which they act. However, this does not affect the conclusions of Theorem 5 that the fine-grained description of such an experiment would correspond to a fixed order process that is compatible with a definite and acyclic causal structure.

Appendix H: Comparing fixed spacetime and quantum gravitational realisations

In [53], a protocol has been proposed for distinguishing between the physically realised optical realisations and theoretically proposed gravitational realisations of QS in a manner that does not disturb the coherence between the different branches of the superposition. For the definite spacetime case, [53] focuses on a QS implementation in Minkowski spacetime described with respect to a single global frame where the in/output events of Alice and Bob are spatially localised but not temporally localised. Here we show that their core argument can be generalised to arbitrary fixed spacetime realisations.

The main property that the protocol of [53] seeks to distinguish is whether the realisation involves a superposition of orders between two spacetime events (one for Alice and one for Bob) or whether at least one agent is associated with more than one spacetime event. The protocol introduces an additional agent F , hereby known as ‘‘Friend’’ to whom Alice and Bob send out photons in addition to performing their usual operations in the QS scenario. F is assumed to be spatially localised¹⁴ and can measure information regarding the times of arrival of the photons arriving from Alice and Bob to decide whether or not the local operation in each lab was a single spacetime event. In the optical implementation of QS, at least one of Alice and Bob must act (on a non-vacuum state) at an earlier time or a distinct later time depending coherently on the control qubit (while the other agent may act at a fixed spacetime location). In this case, at least one agent is associated with two spacetime events which translates into F receiving at least one of the agents' photons arriving to them at a coherent superposition of different times, in the optical realisation in Minkowski spacetime. On the other hand, [53] proposes (theoretical) quantum gravitational realisations where the superposition of spacetime metrics can be used to ensure that any photon from Alice always arrives to F at the same time t^A and any photon from Bob always arrives to F at the same time t^B . A non-demolition measurement is then performed by F to distinguish these two scenarios without collapsing the superposition of orders (see [53] for details).

We can model this protocol by considering a new 5-partite process map \mathcal{W}_{QS}^{+F} obtained from the 4-partite quantum switch map \mathcal{W}_{QS} by including the agent F with input systems F_A^I and F_B^I and a trivial output space. We can give additional outputs A_F^O and B_F^O to Alice and Bob for the photons being sent to F and the process vector $|W_{QS}^{+F}\rangle$ is then identical to $|W_{QS}\rangle$ of Equation (E2), but with the additional factor $|\mathbb{1}\rangle_{A_F^O F_A^I} |\mathbb{1}\rangle_{B_F^O F_B^I}$ in both terms of the superposition, representing the identity channels from A_F^O to F_A^I and B_F^O to F_B^I . [53] argue that in the Minkowski spacetime realisation, assuming that Alice and Bob communicate to F using photons (i.e., light speed communication), then it cannot be the case that F (assumed to be spatially localised) receives Alice's photon at time t^A and Bob's at time t^B independently of the order in which they act. This is because such a realisation of QS involves at least one agent (say Alice) acting at multiple spacetime events, and the light-like surfaces emanating from the two events would intersect F 's spatially localised worldline at two distinct times. That is, it is impossible for the systems F_A^I and F_B^I to both be localised in the spacetime if we assume light speed communication from A and B to F , in a realisation satisfying relativistic causality. Notice that without this assumption, it is in-principle possible to adjust the communication speed according to the order in which Alice acts such that her communication reaches F at the same time (e.g., if Alice acts earlier, she uses free space communication and if she act later, she uses a fibre optic cable

¹⁴ In a classical background spacetime, which is what we consider here, the notion of spatial localisation is well defined. To model a quantum spacetime, [53] considers a region of the manifold within which there can be a superposition of metrics while there is a well defined classical geometry outside this region where localisation is well-defined as in classical spacetimes. The friend F is assumed to be spatially localised in such a classical region of the spacetime.

with a suitable refractive index). Here, we suggest a way to generalise the arguments of [53] to arbitrary spacetime realisations, without assuming light speed communication.

For this, consider an arbitrary but definite and acyclic spacetime \mathcal{T} , and that F_A^I is embedded at a single spacetime location $P^{F_A^I}$ and F_B^I at the location $P^{F_B^I}$, taking $P^{F_B^I} < P^{F_A^I}$ (the argument for $P^{F_A^I} < P^{F_B^I}$ is analogous). Denoting $\text{Past}(P) := \{Q \in \mathcal{T} | Q < P\}$ to be the past of a spacetime point P , we have $\text{Past}(P^{F_B^I}) \subset \text{Past}(P^{F_A^I})$. Then, relativistic causality, along with the fact that Alice and Bob must in general, signal to F in the protocol implies that $\mathcal{R}^{A^O} \xrightarrow{R} P^{F_A^I}$ and $\mathcal{R}^{B^O} \xrightarrow{R} P^{F_B^I}$. Now if we can ensure that $P^{F_B^I}$ is outside the lightcone of all points in \mathcal{R}^{A^O} i.e., that it is impossible for Alice to have communicated to F such that Alice's message could arrive to F at the location $P^{F_B^I}$ in any relativistically causal protocol, then we would have $\mathcal{R}^{A^O} \subseteq \text{Past}(P^{F_A^I}) \setminus \text{Past}(P^{F_B^I})$. Together with $\mathcal{R}^{B^O} \subseteq \text{Past}(P^{F_B^I})$, this would imply that $\mathcal{R}^{A^O} \not\subseteq \mathcal{R}^{B^O}$. However, since we have $\{A^O\} \longrightarrow \{B^O\}$ in the QS protocol, this would be a violation of relativistic causality.

This shows that the following properties cannot be simultaneously achieved in any fixed spacetime realisation of [53]'s modified QS protocol that satisfies relativistic causality: (1) the spatially localised agent F receives the communication from Alice at a single spacetime location $P^{F_A^I} = (\vec{r}^F, t^{F_A^I})$ and from Bob at a distinct, single location $P^{F_B^I} = (\vec{r}^F, t^{F_B^I})$ and (2) for at least one agent, say A , the spacetime location associated with the other agent, $P^{F_B^I}$ is such that $P^{F_B^I}$ is outside the future light-cone of A 's output locations i.e., $\mathcal{R}^{A^O} \not\subseteq P^{F_B^I}$. Notice that (2) is a weaker condition than requiring A and B to communicate to F only at light speed. If consider Minkowski spacetime and we restrict communication from A and B to F to be light speed communication, then we can recover the argument of [53] as a special case. In this case, assuming (1) and relativistic causality, would imply that every point $Q \in \mathcal{R}^{A^O}$ must lie on the past light-like surface of $P^{F_A^I}$. Taking $P^{F_B^I} < P^{F_A^I}$ w.l.o.g, and recalling that $P^{F_B^I}$ and $P^{F_A^I}$ have the same spatial co-ordinate \vec{r}^F , it is clear that the past light-like surface of $P^{F_A^I}$ cannot have any overlap with the past light-cone of $P^{F_B^I}$. Therefore $\mathcal{R}^{A^O} \subseteq \text{Past}(P^{F_A^I}) \setminus \text{Past}(P^{F_B^I})$ and consequently $\mathcal{R}^{A^O} \not\subseteq \mathcal{R}^{B^O}$ holds. However, in a gravitational implementation where the spacetime locations are described with respect to local quantum reference frames [60], it might nevertheless be possible to satisfy all these conditions in a physical implementation of QS, as suggested by the gravitational realisations proposed in [53]. It would be interesting to study, in the future, how (2) can be imposed and justified in quantum gravitational realisations.

Appendix I: Relation to causal loops without superluminal signalling

In recent works [2, 3] involving one of the authors, a formalism for describing cyclic and non-classical information-theoretic causal models as well as their compatibility with spacetime has been developed. While the present paper is inspired by the spirit of this work, namely to disentangle the two causality notions and relate them through compatibility conditions (which capture relativistic causality constraints), there are some important differences between the two formalisms as they were developed for describing different types of scenarios. For instance, [2, 3] can model certain beyond-quantum theories and correlations known as *jamming non-local theories/correlations* [95], however it focused on interventions on observed classical nodes in information-theoretic causal models rather than interventions directly on the underlying non-classical systems. On the other hand, the present formalism focuses on quantum information-theoretic networks where agents can intervene on quantum systems, which enables us to explicitly model indefinite causal structures and process matrices. In the following, we discuss the relation between the two formalisms, and in particular, clarify the link between our results about the acyclic fine-grained causal structure of spacetime realised networks and the surprising result of [3] regarding the possibility of causal loops without superluminal signalling in Minkowski spacetime.

As highlighted in Section III B and in previous works [2, 3, 7, 59], information-theoretic causation and signalling are distinct notions. It is possible to have causal influence (\rightsquigarrow) without signalling (\longrightarrow), which implies that *no superluminal causation* and *no superluminal signalling* correspond to two distinct relativistic causality principles. Let us illustrate this more concretely with the CPTPM $\mathcal{M} : I \mapsto O_1 \otimes O_2$ from Fig. 5, where $\{I\} \longrightarrow \{O_1, O_2\}$ but $\{I\} \not\rightarrow \{O_1\}$ and $\{I\} \not\rightarrow \{O_2\}$. Suppose that the in/output systems are embedded at distinct spacetime locations P^S , $S \in \{I, O_1, O_2\}$. Then in any decomposition of \mathcal{M} , I must be a cause of at least one of O_1 or O_2 , the decompositions




FIG. 17: (a) The directed cyclic graph over the nodes A, B, C, Λ represents the information-theoretic causal structure of the causal loop of [3]. The loop corresponds to a classical functional causal model (reviewed in Appendix C) on this causal structure where all nodes are binary variables, the parentless node Λ is associated with the uniform distribution $P(\Lambda)$, and the remaining nodes depend on their parent nodes through the functional relations: $A = \Lambda \oplus B$, $C = \Lambda$ and $B = A \oplus C$. Here only A, B and C are accessible to agents, Λ is a latent node and cannot be intervened upon. The signalling relations of the causal model are such that B signals jointly to A and C but not individually to A or C . The causal model is embedded in (1+1)-Minkowski spacetime by assigning a single spacetime event to each of its nodes as shown, the 45 degree straight lines (blue) denote light-like surfaces. In the given embedding, it is shown in [3] that this causal loop does not lead to superluminal signalling, since the signalling from B to $\{A, C\}$ can only be verified in the blue shaded region which is contained in the future of B . Notice however that the causal influence from B to A is superluminal in this embedding. (b) Equivalent representation of the causal loop of [3] as a cyclic network in our framework. Notice that the map \mathcal{M}_1 acts identically to the map \mathcal{M}_1 of Fig. 5b, with B, A and C playing the role of I, O_1 and O_2 respectively. If this network is embedded in spacetime, our relativistic causality condition (Definition 23) would require at least one of A or C to be embedded in the future light-cone of B (since $\{B\}$ signals to $\{A, C\}$), the spacetime embedding of (a) for this network therefore violates the relativistic causality condition of the present paper, although it respects the strictly weaker relativistic causality condition of [2, 3] and does not lead to superluminal signalling as discussed in the text (and shown in [3]).

illustrated in Fig. 5 have this property.¹⁵ Therefore, a necessary condition for ensuring that information-theoretic causal influences flow from past to future in the spacetime (i.e., there is no superluminal causation) is that at least one of $P^I < P^{O_1}$ or $P^I < P^{O_2}$ must hold. This is what our relativistic causality condition of Definition 23 imposes. On the other hand, noting that the signalling $\{I\} \rightarrow \{O_1, O_2\}$ can only be verified when O_1 and O_2 are jointly accessed, imposing that the joint future (intersection of the future light-cones) of P^{O_1} and P^{O_2} is contained within the future of I would be sufficient for ensuring no superluminal signalling in this scenario [2].¹⁶ This condition is strictly weaker than requiring $P^I < P^{O_1}$ or $P^I < P^{O_2}$ to hold, as either of these would imply that the joint future of the outputs would be contained in the inputs but the converse is not true: the joint future condition can be satisfied even when P^I, P^{O_1} and P^{O_2} are space-like separated [95]. The compatibility condition of [2, 3] which is necessary and sufficient for ensuring no superluminal signalling in that framework, imposes these weaker conditions.

¹⁵ For if it were not (i.e., we had a decomposition where there is no internal wire connection I to either of O_1 or O_2 , then $\{I\}$ signals to $\{O_1, O_2\}$ would be impossible.

¹⁶ Provided that the systems I, O_1 and O_2 are the only ones that are accessible to agents (systems involved in the internal decomposition of \mathcal{M} being outside their access).

The difference between the two compatibility conditions, which correspond to two distinct (but related) relativistic causality principles as outlined above, leads to different characterisations of the set of information processing protocols which can be realised in a fixed spacetime. Interestingly, only demanding no superluminal signalling (the weaker condition) does not rule out the possibility of genuine causal loops being embedded in Minkowski spacetime as shown in [3]. In the language of our formalism, such genuine loops would correspond to those where the maximally fine-grained description is also cyclic i.e., the causal loop of [3] is a counter-example which shows that Theorem 5 would not hold under the weaker compatibility condition. [3] does not suggest that such loops are physical in any way, rather, the main message of [3] is that the principle of no superluminal signalling is not sufficient for ruling out causal loops in Minkowski spacetime, as commonly believed. In light of this, it remains important to understand the physical principles that can rule them out. Our work suggests that the stronger relativistic causality constraint (which is necessary for ruling out superluminal causation) is sufficient for ruling out such genuine causal loops in a fixed spacetime, at least in quantum theory. Whether this is also a necessary condition for ruling out such loops (even in quantum and classical theories) in a fixed spacetime, and whether the sufficiency result generalises to post-quantum theories remain open questions. In Fig. 17 we illustrate the information-theoretic causal model of the causal loop proposed in [3], the cyclic network by which we can model it within the present framework and the spacetime embedding proposed in [3] whereby the causal loop does not lead to superluminal signalling. While we have discussed some of the relationships and distinctions between the two formalisms, characterising the precise relationship(s) between the classes of causal loops proposed in [2, 3] and the cyclic quantum networks proposed here, specifically the subset that describes indefinite causal structure processes, remains an interesting open question.

We stress that although our compatibility condition is stronger than the condition of [2, 3] and is necessary for ensuring the absence of superluminal causation, it is formulated only at the level of the operational signalling structure of the network and without referring to the underlying causal decompositions of the maps involved. This is because the signalling, although distinct from causation, does allow us to operationally infer the presence of certain causal influences. For the questions of interest in the present work, namely for understanding the physical and experimental realisations of indefinite causal structures in spacetime, the stronger compatibility condition that we have used is the more relevant one, as we are not yet aware of the existence of any physical mechanisms for superluminal causation in Minkowski spacetime. However, one way to evade the consequences of our main no-go theorems would be through a theoretical model which incorporates such superluminal causal mechanisms (which, could in-principle be possible without violating the principle of superluminal signalling). However, such causal explanations are typically disfavoured as they would need to be fine-tuned [4, 7].

Appendix J: Proofs of all results from the main text

Theorem 1. *[Embedding arbitrary, cyclic signalling relations in spacetime] For every signalling structure \mathcal{G}^{sig} , there exists a fixed acyclic spacetime \mathcal{T} and an embedding \mathcal{E} of \mathcal{G}^{sig} in a region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ of \mathcal{T} that respects relativistic causality.*

Proof. In our framework, a signalling structure $\mathcal{G}_{\mathfrak{N}}^{\text{sig}}$ of some network \mathfrak{N} (which will simply be denoted as \mathcal{G}^{sig} in the rest of the proof) is in general a directed graph where the nodes are subsets of in/output systems in the network, i.e., $\text{Nodes}(\mathcal{G}^{\text{sig}}) \subseteq \Sigma(\mathfrak{N}^{\text{sys}})$. An embedding \mathcal{E} of \mathcal{G}^{sig} in a spacetime \mathcal{T} corresponds to an assignment of spacetime regions to each system in $\mathfrak{N}^{\text{sys}}$. This immediately implies an embedding for all systems in $\Sigma(\mathfrak{N}^{\text{sys}})$: for any subset \mathcal{S} of $\mathfrak{N}^{\text{sys}}$ the spacetime region assigned to \mathcal{S} through the embedding \mathcal{E} is simply the union $\mathcal{R}^{\mathcal{S}} = \bigcup_{S \in \mathcal{S}} \mathcal{R}^S$ of all the spacetime regions assigned to the individual elements $S \in \mathcal{S}$ under the embedding \mathcal{E} . Thus, in order to establish the theorem statement, we first find an embedding for systems in $\mathfrak{N}^{\text{sys}}$, such that the signalling relations in \mathcal{G}^{sig} over the nodes corresponding to systems in $\mathfrak{N}^{\text{sys}}$ respects relativistic causality. We will later see that this immediately implies an embedding of all the nodes of \mathcal{G}^{sig} such that relativistic causality is still preserved. Therefore, for the purpose of the next few paragraphs, we will treat \mathcal{G}^{sig} as a directed graph \mathcal{G} over the nodes $\mathfrak{N}^{\text{sys}}$ with directed edges \longrightarrow , and generalise the result to the powerset $\Sigma(\mathfrak{N}^{\text{sys}})$ at the end.

We first need to set out some nomenclature. For every node N of a directed graph \mathcal{G} , let the set $\text{Par}(N) := \{N' \in \text{Nodes}(\mathcal{G}) : N' \longrightarrow N \in \mathcal{G}\}$ denote the set of all parents of the node N , the set $\text{Ch}(N) := \{N' \in \text{Nodes}(\mathcal{G}) : N \longrightarrow N' \in \mathcal{G}\}$ denote the set of all children of the node N , the set $\text{Anc}(N) := \{N' \in \text{Nodes}(\mathcal{G}) : \exists \text{ directed path } N' \longrightarrow \dots \longrightarrow N \in \mathcal{G}\}$ denote the set of all ancestors of N in \mathcal{G} and the set $\text{Desc}(N) := \{N' \in \text{Nodes}(\mathcal{G}) : \exists \text{ directed path } N \longrightarrow \dots \longrightarrow N' \in \mathcal{G}\}$ denote the set of all descendants of N in \mathcal{G} . Then the set of all nodes in \mathcal{G} that are involved in at least one cycle is defined as

$$\text{Cycle}(\mathcal{G}) = \{N \in \text{Nodes}(\mathcal{G}) : N \in \text{Anc}(N)\}.$$

That is, every node N that is its own ancestor belongs to the set $\text{Cycle}(\mathcal{G})$. Note that since $N_1 \in \text{Anc}(N_2)$ is equivalent to $N_2 \in \text{Desc}(N_1)$, we could have equivalently defined this set as $\text{Cycle}(\mathcal{G}) = \{N \in \text{Nodes}(\mathcal{G}) : N \in \text{Desc}(N)\}$. In particular this implies that $N \in \text{Anc}(N) \cap \text{Desc}(N)$ (or N is contained in its own strongly connected component [62]), for every $N \in \text{Cycle}(\mathcal{G})$. Moreover, \mathcal{G} is a directed acyclic graph if and only if $\text{Cycle}(\mathcal{G}) = \emptyset$.

The rest of the proof proceeds as follow. We will show that we can fine-grain (cf. Definition 2) any signalling structure \mathcal{G} into a directed acyclic graph \mathcal{G}' by “splitting” nodes in $\text{Cycle}(\mathcal{G})$, such that when the split nodes in \mathcal{G}' are recombined, we get back the original graph \mathcal{G} . Once we have a directed acyclic graph \mathcal{G}' , we can always embed it in a partially ordered spacetime \mathcal{T} though an embedding \mathcal{E}' that respects relativistic causality. We can then coarse-grain the embedding \mathcal{E}' of \mathcal{G}' to an embedding \mathcal{E} of the original structure \mathcal{G} that also respects relativistic causality, which will establish the theorem statement.

Construct a new directed graph \mathcal{G}' from the original directed graph \mathcal{G} as follows. If $\text{Cycle}(\mathcal{G}) = \emptyset$, set $\mathcal{G}' = \mathcal{G}$. If $\text{Cycle}(\mathcal{G}) \neq \emptyset$, then split every node $N \in \text{Cycle}(\mathcal{G})$ into two nodes N_1 and N_2 such that $\text{Par}(N_1) = \text{Par}(N) \setminus \text{Cycle}(\mathcal{G})$ and $\text{Ch}(N_1) = \text{Ch}(N)$ and $\text{Par}(N_2) = \text{Par}(N)$ and $\text{Ch}(N_2) = \text{Ch}(N) \setminus \text{Cycle}(\mathcal{G})$ i.e., N_1 contains no incoming arrows from nodes in $\text{Cycle}(\mathcal{G})$ but all the same outgoing arrows as N does in \mathcal{G} while N_2 contains no outgoing arrows to nodes in $\text{Cycle}(\mathcal{G})$ but all the same incoming arrows as N does in \mathcal{G} . Nodes that do not belong to $\text{Cycle}(\mathcal{G})$, and other edges not featuring in the above construction remain unaffected. This fully defines \mathcal{G}' . \mathcal{G}' constructed in this way is a fine-graining of \mathcal{G} (cf. Definition 2). More precisely, when we recombine N_1 and N_2 back into a single node N , for each pair of split nodes and without altering the edge structure, we recover the original graph \mathcal{G} since $\text{Par}(N_1 \cup N_2) = \text{Par}(N)$ and $\text{Ch}(N_1 \cup N_2) = \text{Ch}(N)$ for every $N \in \text{Cycle}(\mathcal{G})$, and the nodes $N \notin \text{Cycle}(\mathcal{G})$ were not split or altered in going from \mathcal{G} to \mathcal{G}' . This implies that we can define a graph fine-graining $\mathcal{F}^{\text{graph}} : \text{Nodes}(\mathcal{G}) \mapsto \Sigma(\text{Nodes}(\mathcal{G}'))$ such that $\mathcal{F}^{\text{graph}}(N) = \{N_1, N_2\}$ for all $N \in \text{Cycle}(\mathcal{G})$ and $\mathcal{F}^{\text{graph}}(N) = N$ otherwise, with respect to which \mathcal{G}' is a fine-graining of \mathcal{G} .

Now, we can show that \mathcal{G}' is a directed acyclic graph. For this, observe that a node of a graph can belong to the cycle of the graph only if it has a parent and a child that belong to the cycle of the graph. This is because $N \in \text{Cycle}(\mathcal{G})$ implies that there is a directed path from N to itself in \mathcal{G} which implies a directed path from N to at least one $N' \in \text{Par}(N)$ and this directed path together with $N' \rightarrow N$ tells us that $N' \in \text{Cycle}(\mathcal{G})$. Similarly, $N \in \text{Cycle}(\mathcal{G})$ implies a directed path from at least one $N'' \in \text{Ch}(N)$ to N and this directed path together with $N \rightarrow N''$ implies that $N'' \in \text{Cycle}(\mathcal{G})$. Next, notice that if $N \notin \text{Cycle}(\mathcal{G})$, then by construction, we have $N \notin \text{Cycle}(\mathcal{G}')$. Moreover, for each $N \in \text{Cycle}(\mathcal{G})$, we have a pair of nodes N_1 and N_2 in \mathcal{G}' where all parents of N_1 and all children of N_2 are nodes of the form $N \notin \text{Cycle}(\mathcal{G})$ (and hence $N \notin \text{Cycle}(\mathcal{G}')$). Then, it follows that no node N in \mathcal{G}' can have both a parent and a child in $\text{Cycle}(\mathcal{G}')$, and therefore that N itself cannot belong to $\text{Cycle}(\mathcal{G}')$. In other words, this establishes that $\text{Cycle}(\mathcal{G}') = \emptyset$ or that \mathcal{G}' is acyclic.

Since \mathcal{G}' is a directed acyclic graph, there exists an embedding $\mathcal{E}' : \text{Nodes}(\mathcal{G}') \mapsto \mathcal{T}$ of \mathcal{G}' in a partially ordered set \mathcal{T} (associated with the order relation $<$) such that $N_i \rightarrow N_j$ in $\mathcal{G} \Rightarrow \mathcal{E}'(N_i) < \mathcal{E}'(N_j)$. By virtue of being a partial order, \mathcal{T} satisfies our minimal definition of spacetime structure, according to Definition 18. Then the required embedding \mathcal{E} of \mathcal{G} in the spacetime \mathcal{T} simply associates two spacetime locations with each node $N \in \text{Cycle}(\mathcal{G})$, the two locations being precisely those assigned by \mathcal{E}' to each of the split nodes i.e., $\mathcal{E}(N) := \{\mathcal{E}'(N_1), \mathcal{E}'(N_2)\}$. For all nodes $N \notin \text{Cycle}(\mathcal{G})$, $\mathcal{E}(N) = \mathcal{E}'(N)$ noting that these nodes never got split. Then it is clear that the embedding \mathcal{E} of \mathcal{G} respects relativistic causality whenever the embedding \mathcal{E}' of \mathcal{G}' respects relativistic causality (which it does by construction).

We now describe how the proof generalises to case where $\text{Nodes}(\mathcal{G}) = \Sigma(\mathfrak{N}^{sys})$. For this, note that our above proof covers all cases where \mathcal{G} has the property that $\mathcal{S}_1 \rightarrow \mathcal{S}_2$ for two subsets $\mathcal{S}_1, \mathcal{S}_2$ of \mathfrak{N}^{sys} , then there exists $S_1 \in \mathcal{S}_1$ and $S_2 \in \mathcal{S}_2$ such that $S_1 \rightarrow S_2$. However, suppose that we have a signalling relation $\mathcal{S}_1 \rightarrow \mathcal{S}_2$ in \mathcal{G} such that there is no signalling relation between individual elements of these two sets (this is indeed possible, see for instance Fig. 5). The relativistic causality condition implied by this signalling relation on the corresponding spacetime embedding is that $\mathcal{R}^{\mathcal{S}_1} \xrightarrow{R} \mathcal{R}^{\mathcal{S}_2}$ (cf. Definition 23). Since $\mathcal{R}^{\mathcal{S}_1} = \bigcup_{S_1 \in \mathcal{S}_1} \mathcal{R}^{S_1}$ (and similarly for \mathcal{S}_2), this is equivalent to saying that there exists $S_1 \in \mathcal{S}_1$ and $S_2 \in \mathcal{S}_2$ such that the corresponding spacetime regions satisfy $\mathcal{R}^{S_1} \xrightarrow{R} \mathcal{R}^{S_2}$. In other words, the relativistic causality constraints on the spacetime embedding of \mathcal{G} are the same irrespective of whether or not \mathcal{G} satisfies the aforementioned property. Thus the above proof also applies to establish the theorem statement for signalling structures \mathcal{G} not satisfying this property and hence applies to all signalling structures.

Finally, we note that according to Definition 18 any partially ordered set corresponds to a spacetime. This rather minimal definition allows us to derive general results that only depend on the order relation between spacetime

points and does not require the spacetime to have any further symmetries, or a smooth differentiable structure. However, under this minimal definition, one might regard two different partially ordered sets \mathcal{T} and \mathcal{T}' as two different “spacetimes”. On the other hand, if we consider the more standard method of modelling spacetime as a differentiable manifold \mathcal{M} , as done in relativistic physics, we could sample different sets of points on the same manifold to generate different partially ordered sets¹⁷ \mathcal{T} and \mathcal{T}' from the same spacetime. If we model spacetime as a globally hyperbolic manifold that ensures the absence of closed timelike curves, then the statement of the present theorem would instead become “For every signalling structure \mathcal{G}^{sig} and every globally hyperbolic manifold \mathcal{M} , there exists an embedding \mathcal{E} of \mathcal{G}^{sig} in a region causal structure $\mathcal{G}_{\mathcal{M}}^R$ of \mathcal{M} that respects relativistic causality, where each node of $\mathcal{G}_{\mathcal{M}}^R$ is a finite set of points in \mathcal{M} .” This can be shown as follows. If \mathcal{G} is a directed acyclic graph, then it can be embedded in any globally hyperbolic manifold \mathcal{M} through an embedding $\mathcal{E} : \text{Nodes}(\mathcal{G}) \mapsto \mathcal{M}$ that assigns a point in \mathcal{M} to each node of \mathcal{G} (see [53] for an explicit construction of such an embedding for the acyclic case). This is because the graph has a finite number of nodes and we can always sample a suitable set of points in the manifold having the required order relations. One can apply this embedding to the acyclic graph \mathcal{G}' constructed in the proof above, this would define the embedding $\mathcal{E}' : \text{Nodes}(\mathcal{G}') \mapsto \mathcal{M}$. The rest of the proof will be the same as the above case for partially ordered sets \mathcal{T} . \square

Lemma 2. [Probabilities from loop composition] *For every process map \mathcal{W} , the joint probability distribution obtained through the induced map $\mathcal{P}_{\mathcal{W}}$ of the network $\mathfrak{N}_{\mathcal{W},N}$ as in Equation (13) is equivalent to that obtained in the process matrix framework through Equation (8). Moreover, \mathcal{W} being a valid process matrix guarantees that the denominator of Equation (13) equals unity.*

Proof. Recall from Eq. (12) that the induced map of the network $\mathfrak{N}_{\mathcal{W},N}$ which we denote by $\mathcal{P}_{\mathcal{W}}$ is given as follows

$$\mathcal{P}_{\mathcal{W}} := \left(\mathcal{W} \bigotimes_{k=1}^N \mathcal{M}^{A_k} \right)^{\{A_k^I \mapsto A_k'^I, A_k^O \mapsto A_k'^O\}_{k=1}^N}, \quad (\text{J1})$$

where \mathcal{W} is an N -partite process map with inputs $\{A_k'^O\}_k$ and outputs $\{A_k^I\}_k$, while \mathcal{M}^{A_k} is the extended local map of the agent A_k , associated with a quantum input $A_k'^I$ and classical input A_k^s (for the setting a_k) and quantum output A_k^O and a classical output A_k^o (for the outcome x_k). Here the primed and corresponding unprimed systems are isomorphic to each other, which allows them to be loop composed together. Then the map $\mathcal{P}_{\mathcal{W}}$ obtained this way has N classical inputs $\{A_k^s\}_k$ and N classical outputs $\{A_k^o\}_k$. The above equation can be read as a parallel composition of the process map with the N local maps followed by $2N$ loop compositions.

In the following, for brevity, we detail the proof for the bipartite case. However, the proof readily generalises to the N agent case. In the bipartite case, taking the agents to be A and B with local settings associated with input systems A^s, B^s and outcomes associated with output systems A^o, B^o , the parallel composition yields the map $\mathcal{M}^A \otimes \mathcal{M}^B \otimes \mathcal{W}$ with input systems $\{A'^I, A^s, B'^I, B^s, A'^O, B'^O\}$ and output systems $\{A^O, A^o, B^O, B^o, A^I, B^I\}$. Applying the loop formula of Eq. (1) to compose the corresponding prime and unprimed systems, we obtain the following expression for the induced map $\mathcal{P}_{\mathcal{W}}$ acting on an input state $|a\rangle\langle a|^{A^s} \otimes |b\rangle\langle b|^{B^s}$.

$$\begin{aligned} & \mathcal{P}_{\mathcal{W}}(|a\rangle\langle a|^{A^s} \otimes |b\rangle\langle b|^{B^s}) \\ &= \sum_{k\dots r} \langle kmoq|^{A^I A^O B^I B^O} \left(\mathcal{M}_a^A(|k\rangle\langle l|^{A'^I}) \otimes \mathcal{M}_b^B(|o\rangle\langle p|^{B'^I}) \otimes \mathcal{W}(|mq\rangle\langle nr|^{A'^O B'^O}) \right) |lnpr\rangle^{A^I A^O B^I B^O} \\ &= \sum_{k\dots r} \langle knor|^{A^I A^O B^I B^O} \left([\mathcal{M}_a^A(|k\rangle\langle l|)]^T \otimes [\mathcal{M}_b^B(|o\rangle\langle p|)]^T \otimes \mathcal{W}(|mq\rangle\langle nr|) \right) |lmpq\rangle^{A^I A^O B^I B^O}, \end{aligned}$$

where we have used the notation $\mathcal{M}_a^A(|k\rangle\langle l|^{A'^I}) := \mathcal{M}^A(|a\rangle\langle a|^{A^s} \otimes |k\rangle\langle l|^{A'^I})$ (and similarly for B 's operation) and suppressed the system labels where they are evident from context. Denoting the factor $\left([\mathcal{M}_a^A(|k\rangle\langle l|^{A'^I})]^T \otimes \right.$

¹⁷ That is, if the manifold is globally hyperbolic, in more exotic spacetimes with closed timelike curves, we can also obtain pre-ordered sets from sampling suitable points.

$[\mathcal{M}_b^B(|o\rangle\langle p|^{B^I})]^T \otimes \mathcal{W}(|m\rangle\langle n|^{A^O} \otimes |q\rangle\langle r|^{B^O})$ by (...), introducing factors of the identity $\mathcal{I} = \sum_j |j\rangle\langle j|$, and then rearranging the resulting inner products we have

$$\begin{aligned}
& \mathcal{P}_{\mathcal{W}}(|a\rangle\langle a|^{A^s} \otimes |b\rangle\langle b|^{B^s}) \\
&= \sum_{ijst} \sum_{k\dots r} \langle k|^{A^I} \langle n|j\rangle\langle j|^{A^O} \langle o|^{B^I} \langle r|s\rangle\langle s|^{B^O} \left(\dots \right) |i\rangle^{A^I} \langle i|l\rangle^{A^I} |m\rangle^{A^O} |t\rangle^{B^I} \langle t|p\rangle^{B^I} |q\rangle^{B^O} \\
&= \sum_{ijst} \sum_{k\dots r} \langle i|^{A^I} |l\rangle\langle k|^{A^I} \langle j|^{A^O} \langle t|^{B^I} |p\rangle\langle o|^{B^I} \langle s|^{B^O} \left(\dots \right) |i\rangle^{A^I} |m\rangle\langle n|^{A^O} |j\rangle^{A^O} |t\rangle^{B^I} |q\rangle\langle r|^{B^O} |s\rangle^{B^O} \\
&= \text{Tr}_{A^I A^O B^I B^O} \left[\sum_{k\dots r} |l\rangle\langle k|^{A^I} \otimes |p\rangle\langle o|^{B^I} \otimes \left(\dots \right) \otimes |m\rangle\langle n|^{A^O} \otimes |q\rangle\langle r|^{B^O} \right] \\
&= \text{Tr}_{A^I A^O B^I B^O} \left[\left(\sum_{kl} |l\rangle\langle k|^{A^I} \otimes [\mathcal{M}_a^A(|k\rangle\langle l|^{A^I})]^T \right) \otimes \left(\sum_{op} |p\rangle\langle o|^{B^I} \otimes [\mathcal{M}_b^B(|o\rangle\langle p|^{B^I})]^T \right) \right. \\
&\quad \left. \otimes \left(\sum_{mnr} |mq\rangle\langle nr|^{A^O B^O} \otimes \mathcal{W}(|mq\rangle\langle nr|^{A^O B^O}) \right) \right]
\end{aligned}$$

Now, we wish to calculate the probability that the output of $\mathcal{P}_{\mathcal{W}}(|a\rangle\langle a|^{A^s} \otimes |b\rangle\langle b|^{B^s})$ is $|x\rangle^{A^O} \otimes |y\rangle^{B^O}$ i.e., the outcomes x and y are obtained by Alice and Bob upon measuring the settings a and b . This is given by Eq. (13). We now show that the numerator of this expression i.e., $\text{Tr} \left[\left(|x\rangle\langle x|^{A^O} \otimes |y\rangle\langle y|^{B^O} \right) \left(\mathcal{P}_{\mathcal{W}}(|a\rangle\langle a|^{A^s} \otimes |b\rangle\langle b|^{B^s}) \right) \right]$ equals the process matrix probability rule of Eq. (8). Using the expression for $\mathcal{P}_{\mathcal{W}}(|a\rangle\langle a|^{A^s} \otimes |b\rangle\langle b|^{B^s})$ derived above, along with Eq. (11) to absorb the outcome projectors into the definition of the maps $\mathcal{M}_{x|a}^A$ and $\mathcal{M}_{y|b}^B$, we immediately obtain the following which coincides with Eq. (8).

$$P(xy|ab) = \text{Tr} \left[\left(M_{x|a}^{A^I A^O} \otimes M_{y|b}^{B^I B^O} \right) W \right], \quad (\text{J2})$$

where $W = \mathcal{I} \otimes \mathcal{W} |\mathbb{1}\rangle\langle\mathbb{1}|$ is the process matrix and $M_{x|a}^{A^I A^O} = \left[\mathcal{I} \otimes \mathcal{M}_{x|a}^A \left(|\mathbb{1}\rangle\langle\mathbb{1}| \right) \right]^T$ is the Choi matrix of the local map $\mathcal{M}_{x|a}^A$ (and similarly for B 's operation) as defined in Section VI. Whenever W is valid process matrix, this distribution is normalised, which implies that $\sum_x \text{Tr} \left[\left(M_{x|a}^{A^I A^O} \otimes M_{y|b}^{B^I B^O} \right) W \right] = 1$, which in turn implies that the denominator of Eq. (13) equals unity. \square

Lemma 3. [Sub-networks and reduced processes] For any N -partite process map \mathcal{W} , the sub-network $\mathfrak{N}_{\mathcal{W}, l < N}$ of the process network formed by composing \mathcal{W} with some set of fixed local operations $\{\mathcal{M}_{a_k}^{A_k}\}_{k=1}^l$ for the first $l < N$ agents is a CPTP map whose Choi matrix is the reduced process matrix $\overline{W}(M_{a_1}^{A_1^I A_1^O}, \dots, M_{a_l}^{A_l^I A_l^O})$.

Proof. The proof method is very similar to that of Lemma 2 but we provide it here for completeness and follow the same notation as the previous proof. Again, we restrict to the bipartite case for simplicity but the proof easily generalises to the multi-partite case. Consider a bipartite process map \mathcal{W} associated with the agents $A_1 \equiv A$ and $A_2 \equiv B$, with input systems A^O, B^O and output systems A^I, B^I . Let $\mathcal{M}_a^A : A^I \mapsto A^O$ be a local operation of the agent A associated with a particular setting a , and where all the primed systems are isomorphic to the corresponding unprimed ones (we only add the primes on A 's systems as we will consider only their composition in the next steps). The sub-network $\mathfrak{N}_{\mathcal{W}, 1 < 2}$ (which we denote here as $\mathfrak{N}_{\mathcal{W}, A}$) formed by composing \mathcal{W} only with the operation of A , through $A^I \mapsto A^I, A^O \mapsto A^O$ will have an induced map $\mathcal{N}_{\mathcal{W}, A}$ which is associated with the input B^O and output B^I . We now show that for every choice of setting a , the Choi matrix of the induced sub-network map $\mathcal{N}_{\mathcal{W}, A}$ is the reduced process matrix $\overline{W}(M_a^{A^I A^O})$ over B (cf. Equation (14)). This will complete the proof as the same arguments, under exchange of A and B apply for showing that the Choi matrix of the subnetwork $\mathcal{N}_{\mathcal{W}, B}$ is the reduced process matrix $\overline{W}(M_b^{B^I B^O})$ on A , and these two are the only non-trivial reduced processes here. The Choi matrix of $\mathcal{N}_{\mathcal{W}, A}$ is given as

$$\sum_{m,n} |m\rangle\langle n|^{B^O} \otimes \mathcal{N}_{\mathcal{W}, A}(|m\rangle\langle n|^{B^O}).$$

Using the composition operation, we can write $\mathcal{N}_{\mathcal{W},A}(|m\rangle\langle n|^{B^O})$ as

$$\begin{aligned}\mathcal{N}_{\mathcal{W},A}(|m\rangle\langle n|^{B^O}) &= \sum_{i,j,k,l} \langle i|^{A^I} \langle k|^{A^O} \left(\mathcal{M}_a^A \otimes \mathcal{W} \right) \left(|i\rangle\langle j|^{A^I} \otimes |km\rangle\langle ln|^{A^O B^O} \right) |j\rangle^{A^I} |l\rangle^{A^O} \\ &= \sum_{i,j,k,l} \langle k|^{A^O} \mathcal{M}_a^A \left(|i\rangle\langle j|^{A^I} \right) |l\rangle^{A^O} \otimes \langle i|^{A^I} \mathcal{W} \left(|km\rangle\langle ln|^{A^O B^O} \right) |j\rangle^{A^I}\end{aligned}$$

Plugging this back into the Choi matrix, inserting factors of the identity and rearranging, we can write the Choi matrix of $\mathcal{N}_{\mathcal{W},A}$ as

$$\begin{aligned}& \sum_{i,j,k,l,m,n} |m\rangle\langle n|^{B^O} \otimes \langle k|^{A^O} \mathcal{M}_a^A \left(|i\rangle\langle j|^{A^I} \right) |l\rangle^{A^O} \otimes \langle i|^{A^I} \mathcal{W} \left(|km\rangle\langle ln|^{A^O B^O} \right) |j\rangle^{A^I} \\ &= \sum_{i,j,k,l,m,n} |m\rangle\langle n|^{B^O} \otimes \langle l|^{A^O} \left[\mathcal{M}_a^A \left(|i\rangle\langle j|^{A^I} \right) \right]^T |k\rangle^{A^O} \otimes \langle i|^{A^I} \mathcal{W} \left(|km\rangle\langle ln|^{A^O B^O} \right) |j\rangle^{A^I} \\ &= \sum_{i,j,k,l,m,n,p,q} |m\rangle\langle n|^{B^O} \otimes \langle l|^{A^O} |p\rangle\langle p|^{A^O} \left[\mathcal{M}_a^A \left(|i\rangle\langle j| \right) \right]^T |k\rangle^{A^O} \otimes \langle i|^{A^I} |q\rangle\langle q|^{A^I} \mathcal{W} \left(|km\rangle\langle ln| \right) |j\rangle^{A^I} \\ &= \sum_{i,j,k,l,m,n,p,q} |m\rangle\langle n|^{B^O} \otimes \langle p|^{A^O} \left[\mathcal{M}_a^A \left(|i\rangle\langle j| \right) \right]^T |k\rangle^{A^O} \langle l|p\rangle^{A^O} \otimes \langle q|^{A^I} \mathcal{W} \left(|km\rangle\langle ln| \right) |j\rangle^{A^I} \langle i|q\rangle^{A^I} \\ &= \sum_{p,q} |pq\rangle^{A^O A^I} \left(\mathbb{1}^{B^I B^O} \otimes \sum_{i,j} |j\rangle\langle i| \otimes \left[\mathcal{M}_a^A \left(|i\rangle\langle j| \right) \right]^T \right) \left(\sum_{k,l,m,n} |km\rangle\langle ln| \otimes \mathcal{W} \left(|km\rangle\langle ln| \right) \right) |pq\rangle^{A^O A^I} \\ &= \text{Tr}_{A^I A^O} \left(\left(\mathbb{1}^{B^I B^O} \otimes M_a^{A^I A^O} \right) \cdot W \right) = \overline{W} \left(M_a^{A^I A^O} \right),\end{aligned}$$

where we have used the definition of the Choi matrices of the local operations and the process map (see Section VI) in the last line, along with that of the reduced process matrix (Equation (14)). Moreover, we have identified the primed and unprimed versions of the systems A^I and A^O as they get composed to form effectively a single system. Moreover the reduced process matrix $\overline{W} \left(M_a^{A^I A^O} \right)$ being a valid process matrix implies that the map $\mathcal{N}_{\mathcal{W},A}$ of which it is the Choi matrix is a CPTP map [31]. This completes the proof. \square

Theorem 2. *The network $\mathfrak{N}_{\mathcal{W},N}$ formed the composition of any valid process map \mathcal{W} with a set of extended local operations $\{\mathcal{M}^{A^k}\}_{k=1}^N$ is a network of CPTP maps.*

Proof. The sub-networks of $\mathcal{N}_{\mathcal{W},N}$ are formed by taking subsets of its maps and compositions (cf. Eq. (12)). Subsets of $\mathcal{N}_{\mathcal{W},N}^{\text{maps}}$ are of two types: either the subset includes \mathcal{W} or it does not. In the latter case, none of the compositions of $\mathcal{N}_{\mathcal{W},N}^{\text{comp}}$ can be applied, as all compositions involve \mathcal{W} . Thus such sub-networks are always guaranteed to be CPTP since they correspond to the parallel composition (i.e., tensor product) of CPTP maps, and we only need to consider sub-networks formed by composing \mathcal{W} with some subset of the extended local operations through some subset of the compositions in $\mathcal{N}_{\mathcal{W},N}^{\text{comp}}$. W.l.o.g, consider the subset of the first $l < N$ extended local operations. The compositions associated with these operations and \mathcal{W} are $\{A_k^I \hookrightarrow A_k'^I, A_k^O \hookrightarrow A_k'^O\}_{k=1}^l$. If we were to perform all these compositions, we obtain the sub-network $\mathfrak{N}_{\mathcal{W},l < N}$, whose induced map is CPTP as it corresponds to a reduced process matrix associated with the remaining $N - l$ agents, as we have shown in Lemma 3.

Suppose we perform only a subset of the compositions in $\{A_k^I \hookrightarrow A_k'^I, A_k^O \hookrightarrow A_k'^O\}_{k=1}^l$. Consider all the agents A_k for which both compositions $\{A_k^I \hookrightarrow A_k'^I, A_k^O \hookrightarrow A_k'^O\}$ are included in this subset. Again, w.l.o.g, we can take these to be the first $p < l$ agents. As the order in which we perform the compositions does not matter (cf. Lemma 6), we can perform all the operations in the subset under consideration by first performing the compositions $\{A_k^I \hookrightarrow A_k'^I, A_k^O \hookrightarrow A_k'^O\}_{k=1}^p$ and then the rest. By Lemma 3 it again follows that performing the compositions $\{A_k^I \hookrightarrow A_k'^I, A_k^O \hookrightarrow A_k'^O\}_{k=1}^p$ results in a CPTP map $\mathcal{N}_{\mathcal{W},p < N}$. Now by construction, the remaining compositions, include at most one of $A_k^I \hookrightarrow A_k'^I$ or $A_k^O \hookrightarrow A_k'^O$ for each agent in $\{A_{p+1}, \dots, A_l\}$. It is easy to see that each of these compositions would correspond to a regular sequential composition of the CPTP map $\mathcal{N}_{\mathcal{W},p < N}$ with a CPTP map \mathcal{M}^{A^k} for an agent A_k in $\{A_{p+1}, \dots, A_l\}$. The sequential composition of two CPTP maps (in any given order) is always a CPTP map. This shows that the remaining compositions also yield a CPTP map. We have covered all possible sub-networks in the above and

shown their induced maps to be CPTP. This establishes that our original network is a network of CPTP maps, as required. \square

Theorem 3. *A process matrix W is not a fixed order process if and only if the signalling structure of the corresponding network $\mathfrak{N}_{\mathcal{W},N}$ certifies the cyclicity of any causal structure \mathcal{G} that it is compatible with.*

Proof. Observe that in the network $\mathfrak{N}_{\mathcal{W},N}$, we have $\{A_k^I\}$ signals to $\{A_k^O\}$ for every agent A_k since the extended local map for each agent encodes all possible operations between their these in and output systems and in particular the identity channel which clearly signals from the input to output. This implies that when we compatibly embed (Definition 14) the systems of $\mathfrak{N}_{\mathcal{W},N}$ in a causal structure \mathcal{G} with directed edges $\xrightarrow{\mathcal{G}}$ (thereby allowing us to treat each system as a node of a directed graph with these edges), then we must have $A_k^I \xrightarrow{\mathcal{G}} A_k^O$ for each $k \in \{1, \dots, N\}$. By Definition 28 of fixed ordered processes and Definition 29, we can see that a process W is a fixed order process if and only if there exists a partial order on the in/output systems of all agents where A_i^O precedes A_j^I in the partial order for some $j \in \mathcal{S}$ whenever an agent A_i probabilistically signals to a set of agents $A_{\mathcal{S}}$, and A_i^I precedes A_k^O in this partial order for every agent A_k . By Theorem 7, A_i probabilistically signals to $A_{\mathcal{S}}$ is equivalent to $\{A_i^O\}$ signals to $\{A_j^I\}_{j \in \mathcal{S}}$ in the network. Recall that compatibility of a signalling relation $\{A_i^O\}$ signals to $\{A_j^I\}_{j \in \mathcal{S}}$ with an embedding in a causal structure \mathcal{G} implies that there exists $j \in \mathcal{S}$ such that there is a directed path $A_i^O \xrightarrow{\mathcal{G}} \dots \xrightarrow{\mathcal{G}} A_j^I$ with respect to the order induced by \mathcal{G} on the embedded systems (cf. Definition 14). Then the above arguments, taken together imply that W is a fixed order process if and only if there exists a directed acyclic graph \mathcal{G} with which the signalling relations of the process network $\mathfrak{N}_{\mathcal{W},N}$ are compatible. Equivalently W is not a fixed order process if and only if any directed graph \mathcal{G} with which the signalling relations of $\mathfrak{N}_{\mathcal{W},N}$ are compatible is a directed cyclic graph i.e., the signalling relations of $\mathfrak{N}_{\mathcal{W},N}$ being compatible with an embedding in some \mathcal{G} certify the cyclicity of \mathcal{G} whenever W is not a fixed order process. \square

Theorem 4. *[No-go theorem for spacetime realisations of processes] No realisation (Definition 22) of an N -partite process network $\mathfrak{N}_{\mathcal{W},N}$ in a fixed spacetime structure \mathcal{T} (Definition 18) with respect to a spacetime embedding \mathcal{E} can simultaneously satisfy the following three assumptions.*

1. W is not a fixed order process (Definition 28).
2. The spacetime realisation satisfies relativistic causality (cf. Definition 24).
3. The region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ induced by the embedding \mathcal{E} (Definition 21) is acyclic.

Proof. Recall that a spacetime realisation of a network $\mathfrak{N}_{\mathcal{W},N}$ is specified by a spacetime embedding $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}$ of the network (obtained by associating a spacetime region $\mathcal{E}(S) \subseteq \mathcal{T}$ for each system S in the network) along with a fine-graining $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}^f$. For this theorem, the fine-graining and its properties will be irrelevant, we only need to focus on the spacetime embedded network $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}$ to prove the result. Observe that the relativistic causality condition of Definition 24 implies, in particular, that the signalling structure of $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}$ must be compatible with the region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ induced by the embedding \mathcal{E} . Moreover, the signalling structure of $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}$ is, by construction, the same as the signalling structure of $\mathfrak{N}_{\mathcal{W},N}$, since the only difference in the two networks is in the relabelling of each system S to the corresponding spacetime embedded system $S^{\mathcal{E}(S)}$. This implies that the signalling structure of $\mathfrak{N}_{\mathcal{W},N}$ must be compatible with the causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ if condition 2 of the theorem must hold. Now if we impose condition 1, that the associated process W is not a fixed order process, it immediately follows from Theorem 3 that $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ must contain a directed cycle and therefore that it is not an acyclic graph, which would violate condition 3. This proves the theorem. \square

Corollary 2. *[Time localisation in a global frame] No realisation of an N -partite process network $\mathfrak{N}_{\mathcal{W},N}$ in a fixed spacetime structure \mathcal{T} with respect to a spacetime embedding \mathcal{E} can simultaneously satisfy the following three assumptions.*

1. W is not a fixed order process (cf. Definition 28).
2. The spacetime realisation satisfies relativistic causality (cf. Definition 24).
3. The region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ induced by \mathcal{E} is such that there exists an agent A from whose perspective every region $\mathcal{R} \in \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg}})$ is time-localised (cf. Definition 27).

Proof. This corollary follows from noting that the first two assumptions are identical to the first two assumptions of Theorem 4, while the third assumption here implies the third assumption of the theorem. More explicitly, assumption 3 here requires that for each system S , the corresponding spacetime region $\mathcal{R}^S \subseteq \mathcal{T}$ assigned to S by the embedding \mathcal{E} is such that all spacetime points in \mathcal{R}^S have the same time coordinate, say t^S in a global reference frame. Then, Definitions 18 and 19 tell us that $\mathcal{R}^{S_1} \xrightarrow{R} \mathcal{R}^{S_2}$ implies $t^{S_1} < t^{S_2}$. Since we can never have a sequence of times in some global reference frame such that $t^{S_1} < t^{S_2} < \dots < t^{S_n} < t^{S_1}$, it follows that the set of spacetime regions satisfying assumption 3 of this corollary can never contain a sequence of regions such that $\mathcal{R}^{S_1} \xrightarrow{R} \mathcal{R}^{S_2} \xrightarrow{R} \dots \xrightarrow{R} \mathcal{R}^{S_n} \xrightarrow{R} \mathcal{R}^{S_1}$, i.e., the regions satisfy assumption 3 of Theorem 4. \square

Theorem 5. *Any spacetime realisation of a process network $\mathfrak{N}_{\mathcal{W},N}$ in a spacetime \mathcal{T} relative to an embedding \mathcal{E} that satisfies relativistic causality will admit an explanation in terms of a fine-grained network $\mathfrak{N}_{\mathcal{W}^f,M}$ which will be described by a fixed order and therefore causally separable process W^f associated with $M \geq N$ (possibly communicating) agents. Here M is given in terms of the cardinality of the associated input spacetime regions $M = \sum_{k=1}^N |\mathcal{R}^{A_k}|$.*

Proof. We prove the bipartite case here for simplicity, but the proof method readily generalises to the N -partite case. Let W be a bipartite process matrix associated with the agents A and B and $\mathfrak{N}_{\mathcal{W},N}$ be the associated network with $N = 2$. Consider a spacetime embedding $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}$ of $\mathfrak{N}_{\mathcal{W},N}$ in a fixed acyclic spacetime \mathcal{T} . If the region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ induced by the embedding \mathcal{E} is acyclic, then this spacetime embedding fully specifies a spacetime realisation of the network (Definition 22). In this case, if we impose the first condition of the theorem, namely that relativistic causality condition of Definition 24 is satisfied, then it follows that the signalling structure of $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}$ is compatible with the causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ (which is acyclic in the case under consideration). This signalling structure is the same as the signalling structure of $\mathfrak{N}_{\mathcal{W},N}$ (since the spacetime embedded network only differs from the original network in the fact that each system S is additionally labelled by the corresponding spacetime region $\mathcal{E}(S)$), which implies that the signalling structure of $\mathfrak{N}_{\mathcal{W},N}$ is compatible with an acyclic causal structure, which is only possible if W is a fixed order process (cf. Theorem 3).

Therefore we consider the case where $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ is cyclic. In this case, we require a fine-graining $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}^f$ of $\mathfrak{N}_{\mathcal{W},N}$, associated with an acyclic region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg},f}$ in order to fully specify the spacetime realisation. Moreover, $\mathcal{G}_{\mathcal{T}}^{\text{reg},f}$ would be a fine-graining of the region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$ induced by the original embedding \mathcal{E} . Invoking the input-output correspondence between the regions in the embedding, we can denote by $\mathcal{R}^{A^I} \in \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg}})$ the region assigned to A 's input systems A^s and A^I by \mathcal{E} and by $\mathcal{R}^{A^O} \in \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg}})$ the region assigned to A 's output systems A^o and A^O , and similarly for B . Further, since $\mathcal{G}_{\mathcal{T}}^{\text{reg},f}$ is a fine-graining of $\mathcal{G}_{\mathcal{T}}^{\text{reg}}$, there is a fine-graining map that maps $\mathcal{R}^{A^I} \in \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg}})$ to a set of nodes $\{\mathcal{P}_1^I, \dots, \mathcal{P}_n^I\} \subseteq \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg},f})$ which partition it, $\mathcal{R}^{A^I} = \bigcup_{j=1}^n \mathcal{P}_j^I$. The input-output correspondence ensures that $\mathcal{R}^{A^O} \in \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg}})$ is mapped to a corresponding set of nodes $\{\mathcal{P}_1^O, \dots, \mathcal{P}_n^O\} \subseteq \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg},f})$ under the fine-graining, which partition it as $\mathcal{R}^{A^O} = \bigcup_{j=1}^n \mathcal{P}_j^O$ and where $\mathcal{P}_i^I \xrightarrow{R} \mathcal{P}_i^O$ for each $i \in \{1, \dots, n\}$. Similarly, we have the set of regions $\{\mathcal{Q}_1^I, \dots, \mathcal{Q}_m^I\} \subseteq \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg},f})$ associated with Bob's input B^I under the fine-graining, and a corresponding set of regions $\{\mathcal{Q}_1^O, \dots, \mathcal{Q}_m^O\} \subseteq \text{Nodes}(\mathcal{G}_{\mathcal{T}}^{\text{reg},f})$ associated with B^O , where $\mathcal{Q}_j^I \xrightarrow{R} \mathcal{Q}_j^O$ for each $j \in \{1, \dots, m\}$, with $\mathcal{R}^{B^I} = \bigcup_{j=1}^m \mathcal{Q}_j^I$ and $\mathcal{R}^{B^O} = \bigcup_{j=1}^m \mathcal{Q}_j^O$.

Then in the fine-grained network $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}^f$, we have n input and n output systems associated with Alice (an input system for each \mathcal{P}_i^I and an output system for each \mathcal{P}_i^O) and m input and m output systems associated with Bob (an input system for each \mathcal{Q}_j^I and an output system for each \mathcal{Q}_j^O). Moreover, since $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}^f$ is a fine-graining of $\mathfrak{N}_{\mathcal{W},N}$, each sub-network of the former must be a fine-graining of a corresponding sub-network of the latter (Definition 17). As the process map \mathcal{W} itself corresponds to a sub-network of $\mathfrak{N}_{\mathcal{W},N}$, the fine-grained network $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}^f$ must be associated with a fine-grained process map $\mathcal{W}_{\mathcal{T},\mathcal{E}}^f$ whose inputs are the n output systems of Alice, along with the m output systems of Bob, and whose outputs are the n input systems of Alice and the m input systems of Bob. We then split Alice into n and Bob into m agents, associated with the corresponding in/output systems in order to model $\mathcal{W}_{\mathcal{T},\mathcal{E}}^f$ as an $n+m$ -partite process map. In the following we show that whenever the spacetime realisation satisfies the relativistic causality condition, $\mathcal{W}_{\mathcal{T},\mathcal{E}}^f$ is indeed a valid $n+m$ -partite process map which corresponds to a fixed order process.

In Corollary 6 we have shown that spacetime realisations of any network that respect the relativistic causality condition correspond to so-called *causal boxes* which are CTPMs that are closed under arbitrary compositions.

Imposing relativistic causality for the spacetime realisation of the original network $\mathfrak{N}_{\mathcal{W},N}$, this result implies that the induced map of each sub-network of $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}^f$ must be a causal box and in particular that $\mathcal{W}_{\mathcal{T},\mathcal{E}}^f$ is a causal box. This implies that composing $\mathcal{W}_{\mathcal{T},\mathcal{E}}^f$ together with any choice of extended local operations (which are themselves CPTP maps) between each input-output pair from the $n+m$ inputs associated with $\{\mathcal{P}_1^I, \dots, \mathcal{P}_n^I, \mathcal{Q}_1^I, \dots, \mathcal{Q}_m^I\}$ and $n+m$ outputs associated with $\{\mathcal{P}_1^O, \dots, \mathcal{P}_n^O, \mathcal{Q}_1^O, \dots, \mathcal{Q}_m^O\}$ can only lead to valid normalised probabilities. This composition is also a valid composition of causal boxes as the input regions are ordered before the corresponding output regions relative to the partial order induced by the acyclic region causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg},f}$. From this argument, we conclude that $\mathcal{W}_{\mathcal{T},\mathcal{E}}^f$ is a valid $n+m$ -partite process map whenever the given spacetime realisation of $\mathfrak{N}_{\mathcal{W},N}$ satisfies the relativistic causality condition. Then it immediately follows from Theorem 3, just as in the first paragraph of this proof, that $\mathcal{W}_{\mathcal{T},\mathcal{E}}^f$ must be a fixed-order process as its signalling structure is compatible with an acyclic causal structure $\mathcal{G}_{\mathcal{T}}^{\text{reg},f}$ (which is implied by the relativistic causality condition).

Finally, we note that $\mathfrak{N}_{\mathcal{W},N,\mathcal{T},\mathcal{E}}^f$ being a fine-graining of $\mathfrak{N}_{\mathcal{W},N}$ also implies that the extended local maps \mathcal{M}^A and \mathcal{M}^B of Alice and Bob in the original network get fine-grained to extended local maps $\mathcal{M}^{A,f}$ and $\mathcal{M}^{B,f}$ acting on n and m in/output systems respectively. As we have not placed any constraints on the fine-grained local maps, these need not factorise as $\mathcal{M}^{A,f} = \otimes_{i=1}^n \mathcal{M}_i^{A,f}$ with each $\mathcal{M}_i^{A,f}$ acting between the in/output systems associated with \mathcal{P}_i^I and $\mathcal{P}_i^O \xleftarrow{R} \mathcal{P}_i^I$. Thus the $n+m$ agents in the fine-grained process can perform non-product local operations by communicating to each other. The fact that the spacetime realisation satisfying relativistic causality can be described as a causal box guarantees that the process $\mathcal{W}_{\mathcal{T},\mathcal{E}}^f$ will lead to valid normalised probabilities even when composed with such non-product local operations. This completes the proof.

As an additional remark, we note that above argument using causal boxes also guarantees that each fine-grained local operation will admit a decomposition into a sequential composition of CPTP maps (the sequence representation), whenever relativistic causality is imposed. This means that the local operations correspond to quantum combs (these are defined, for instance in [29]) and allows us to recover the product structure of the local operations through a simple trick originally proposed in [73]. One can always map a process and set of non-product local operations modelled as quantum combs (on which the action of the process leads to normalised probabilities) to a new process and a set of product local operations such that the composition is identical in the two cases, yielding the same correlations. This can be done by absorbing all communication channels in the non-product local operations into the definition of the new process. \square

Lemma 4. [No-go result for the quantum switch] *Let $\mathfrak{N}_{\mathcal{W}_{QS},U,V}$ be the network describing the action of \mathcal{W}_{QS} on any two non-trivial local operations $\mathcal{U}^A : A^I \mapsto A^O$ and $\mathcal{V}^B : B^I \mapsto B^O$, where \mathcal{W}_{QS} is the process map associated with the process matrix W_{QS} of the quantum switch. Then any realisation of this network in a fixed spacetime \mathcal{T} with respect to an embedding \mathcal{E} cannot simultaneously satisfy both of the following assumptions*

1. *The spacetime realisation satisfies relativistic causality (cf. Definition 24).*
2. *The region causal structure over the nodes $\mathcal{E}(S)$ for $S \in \{A^I, A^O, B^I, B^O\}$ is acyclic.*

Proof. We first show that in \mathcal{W}_{QS} , when the initial state of the control and target systems is $|\psi_C\rangle := \alpha|0\rangle + \beta|1\rangle$ and $|\psi_T\rangle$ respectively where α and β are both non-zero amplitudes and $|\psi_T\rangle$ is an arbitrary qudit state, then we have $\{A^O\} \longrightarrow \{B^I\}$ irrespective of the state input on B^O and $\{B^O\} \longrightarrow \{A^I\}$ irrespective of the state input on A^O , in \mathcal{W} . This implies that both signalling relations can be realised in the network $\mathfrak{N}_{\mathcal{W}_{QS},U,V}$ for any fixed choice of non-trivial operations \mathcal{U}^A and \mathcal{V}^B without needing to consider the extended local operations. This is in contrast to the general case of Theorem 4 where the signalling relation being realised may depend on the choice of local operations, such that not all signalling relations allowed by \mathcal{W} are realised when it is composed with a given fixed set of local operations.

Consider the action of the process map $\mathcal{W}_{QS} : C_C^O \otimes C_T^O \otimes A^O \otimes B^O \mapsto D_C^I \otimes D_T^I \otimes A^I \otimes B^I$ on the input state $|\Psi\rangle^{C_C^O C_T^O A^O B^O} := (\alpha|0\rangle + \beta|1\rangle)^{C_C^O} \otimes |\psi_T\rangle^{C_T^O} \otimes |\psi_A\rangle^{A^O} \otimes |\psi_B\rangle^{B^O}$, where $|\psi_A\rangle$ and $|\psi_B\rangle$ are arbitrary qudit states. We see that \mathcal{W}_{QS} maps the initial state $|\Psi\rangle^{C_C^O C_T^O A^O B^O}$, to the final state

$$\alpha|0\rangle^{D_C^I} |\psi_B\rangle^{D_T^I} |\psi_T\rangle^{A^I} |\psi_A\rangle^{B^I} + \beta|1\rangle^{D_C^I} |\psi_A\rangle^{D_T^I} |\psi_B\rangle^{A^I} |\psi_T\rangle^{B^I} := |\Phi\rangle^{D_C^I D_T^I A^I B^I}.$$

This gives,

$$\begin{aligned} \text{Tr}_{D_C^I D_T^I A^I} [\mathcal{W}_{QS}(|\Psi\rangle\langle\Psi|^{C_C^O C_T^O A^O B^O})] &= |\alpha|^2 |\psi_A\rangle\langle\psi_A|^{B^I} + |\beta|^2 |\psi_T\rangle\langle\psi_T|^{B^I} \\ \text{Tr}_{D_C^I D_T^I B^I} [\mathcal{W}_{QS}(|\Psi\rangle\langle\Psi|^{C_C^O C_T^O A^O B^O})] &= |\alpha|^2 |\psi_T\rangle\langle\psi_T|^{A^I} + |\beta|^2 |\psi_B\rangle\langle\psi_B|^{A^I} \end{aligned} \tag{J3}$$

Notice that $\text{Tr}_{D_C^I D_T^I A^I} [\mathcal{W}_{QS}(|\Psi\rangle\langle\Psi|)]$ which is the output on B^I does depend on ψ_A , the input on A^O but not depend on ψ_B , the input on B^O . Similarly $\text{Tr}_{D_C^I D_T^I B^I} [\mathcal{W}_{QS}(|\Psi\rangle\langle\Psi|)]$ depends on ψ_B but does not depend on ψ_A . This implies that given the knowledge of the initial control and target states, $\alpha|0\rangle + \beta|1\rangle$ and $|\psi_T\rangle$, Alice and Bob can signal to each other by suitable choices of $|\psi_A\rangle$ and $|\psi_B\rangle$ on their respective output systems A^O and B^O , irrespective of the local operation of the other agent. The above proof easily generalises to arbitrary input states $\rho^{C_C^O C_T^O A^O B^O} := |\psi_C\rangle\langle\psi_C|^{C_C^O} \otimes |\psi_T\rangle\langle\psi_T|^{C_T^O} \otimes \rho^{A^O B^O}$, where the input state $\rho^{A^O B^O}$ on A^O and B^O may be an entangled state, $\text{Tr}_{D_C^I D_T^I A^I} [\mathcal{W}_{QS}(|\Psi\rangle\langle\Psi|)]$ depends only on the marginal of the initial state over A^O which is unaffected by local operations on B^O .

This establishes that the network $\mathfrak{N}_{\mathcal{W}_{QS}, U, V}$ gives rise to a directed cycle of signalling relations $\{A^O\} \longrightarrow \{B^I\} \longrightarrow \{B^O\} \longrightarrow \{A^I\} \longrightarrow \{A^O\}$ for any non-trivial \mathcal{U}^A and \mathcal{V}^B (where $\{A^I\} \longrightarrow \{A^O\}$ and $\{B^I\} \longrightarrow \{B^O\}$ come from the non-triviality of these operations). Then implies (by relativistic causality) that the spacetime regions must satisfy $\mathcal{P}^{A^O} \xrightarrow{R} \mathcal{P}^{B^I} \xrightarrow{R} \mathcal{P}^{B^O} \xrightarrow{R} \mathcal{P}^{A^I} \xrightarrow{R} \mathcal{P}^{A^O}$, which violates assumption 3. This completes the proof. \square

-
- [1] Judea Pearl. Causality: Models, reasoning, and inference. *Second edition, Cambridge University Press*, 2009.
 - [2] V. Vilasini and Roger Colbeck. General framework for cyclic and fine-tuned causal models and their compatibility with spacetime. *Phys. Rev. A*, 106:032204, Sep 2022. <https://link.aps.org/doi/10.1103/PhysRevA.106.032204>.
 - [3] V. Vilasini and Roger Colbeck. Impossibility of superluminal signaling in Minkowski spacetime does not rule out causal loops. *Phys. Rev. Lett.*, 129:110401, Sep 2022. <https://link.aps.org/doi/10.1103/PhysRevLett.129.110401>.
 - [4] V. Vilasini and Roger Colbeck. A causal modelling analysis of Bell scenarios in spacetime: implications of jamming non-local correlations for relativistic causality principles, 2023. <https://arxiv.org/abs/2311.18465>.
 - [5] J.S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1964.
 - [6] Howard M. Wiseman and Eric G. Cavalcanti Causarum Investigatio and the Two Bell's Theorems of John Bell *Quantum [Un]Speakables II: Half a Century of Bell's Theorem* (Springer, Switzerland, 2017), edited by Reinhold Bertlmann and Anton Zeilinger. pp. 119-142. https://link.springer.com/chapter/10.1007/978-3-319-38987-5_6
 - [7] Christopher J Wood and Robert 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(3):033002, 2015. <https://iopscience.iop.org/article/10.1088/1367-2630/17/3/033002>.
 - [8] Robert R. Tucci. Quantum Bayesian nets. *International Journal of Modern Physics B*, 09(03):295–337, Jan 1995. <http://dx.doi.org/10.1142/S0217979295000148>.
 - [9] M. S. Leifer. Quantum dynamics as an analog of conditional probability. *Physical Review A*, 74(4), Oct 2006. <http://dx.doi.org/10.1103/PhysRevA.74.042310>.
 - [10] Kathryn B. Laskey. Quantum causal networks, 2007. <https://arxiv.org/abs/0710.1200>.
 - [11] M.S. Leifer and D. Poulin. Quantum graphical models and belief propagation. *Annals of Physics*, 323(8):1899–1946, Aug 2008. <http://dx.doi.org/10.1016/j.aop.2007.10.001>.
 - [12] M. S. Leifer and Robert W. Spekkens. Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference. *Physical Review A*, 88(5):052130, nov 2013. <https://link.aps.org/doi/10.1103/PhysRevA.88.052130>.
 - [13] Joe Henson, Raymond Lal, and Matthew Pusey. Theory-independent limits on correlations from generalized Bayesian networks. *New Journal of Physics*, 16(11):113043, nov 2014.
 - [14] Jacques Pienaar and Časlav Brukner. A graph-separation theorem for quantum causal models. *New Journal of Physics*, 17(7):73020, 2015. <https://iopscience.iop.org/article/10.1088/1367-2630/17/7/073020>.
 - [15] Katja Ried, Megan Agnew, Lydia Vermeyden, Dominik Janzing, Robert W. Spekkens, and Kevin J. Resch. A quantum advantage for inferring causal structure. *Nature Physics*, 11(5):414–420, Mar 2015. <http://dx.doi.org/10.1038/nphys3266>.
 - [16] Fabio Costa and Sally Shrapnel. Quantum causal modelling. *New Journal of Physics*, 18(6):63032, 2016. <https://iopscience.iop.org/article/10.1088/1367-2630/18/6/063032>.
 - [17] Tobias Fritz. Beyond Bell's Theorem II: Scenarios with arbitrary causal structure. *Communications in Mathematical Physics*, 341(2):391–434, Nov 2015. <http://dx.doi.org/10.1007/s00220-015-2495-5>.
 - [18] John-Mark A. Allen, Jonathan Barrett, Dominic C. Horsman, Ciarán M. Lee, and Robert W. Spekkens. Quantum Common Causes and Quantum Causal Models. *Physical Review X*, 7(3):031021, jul 2017. <http://link.aps.org/doi/10.1103/PhysRevX.7.031021>.
 - [19] Jonathan Barrett, Robin Lorenz, and Ognjan Oreshkov. Quantum causal models, 2020. <https://arxiv.org/abs/1906.10726>.
 - [20] Jacques Pienaar. Quantum causal models via quantum Bayesianism. *Physical Review A*, 101(1), Jan 2020. <http://dx.doi.org/10.1103/PhysRevA.101.012104>.

- [21] Christopher Portmann, Christian Matt, Ueli Maurer, Renato Renner, and Bjorn Tackmann. Causal Boxes: Quantum Information-Processing Systems Closed under Composition. *IEEE Transactions on Information Theory*, 63(5):3277–3305, 2017. <http://ieeexplore.ieee.org/document/7867830/>.
- [22] Giulio Chiribella and Hlér Kristjánsson. Quantum Shannon theory with superpositions of trajectories. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 475(2225):20180903, may 2019. <https://doi.org/10.1098%2Frspa.2018.0903>.
- [23] Giulia Rubino, Lee A. Rozema, Daniel Ebler, Hlér Kristjánsson, Sina Salek, Philippe Allard Guérin, Alastair A. Abbott, Cyril Branciard, Časlav Brukner, Giulio Chiribella, and et al. Experimental quantum communication enhancement by superposing trajectories. *Physical Review Research*, 3(1), Jan 2021. <http://dx.doi.org/10.1103/PhysRevResearch.3.013093>.
- [24] Lucien Hardy. Probability Theories with Dynamic Causal Structure: A New Framework for Quantum Gravity, sep 2005. <http://arxiv.org/abs/gr-qc/0509120>.
- [25] Lucien Hardy. Towards quantum gravity: a framework for probabilistic theories with non-fixed causal structure. *Journal of Physics A: Mathematical and Theoretical*, 40(12):3081–3099, Mar 2007. <http://dx.doi.org/10.1088/1751-8113/40/12/S12>.
- [26] Magdalena Zych, Fabio Costa, Igor Pikovski, and Časlav Brukner. Bell’s theorem for temporal order. *Nature Communications*, 10(1):3772, 2019. <https://doi.org/10.1038/s41467-019-11579-x>.
- [27] Giulio Chiribella, Giacomo Mauro D’Ariano, Paolo Perinotti, and Benoit Valiron. Quantum computations without definite causal structure. *Physical Review A*, 88(2):022318, aug 2013. <https://link.aps.org/doi/10.1103/PhysRevA.88.022318>.
- [28] Ognjan Oreshkov, Fabio Costa, and Časlav Brukner. Quantum correlations with no causal order. *Nature Communications*, 3:1092, oct 2012. <https://www.nature.com/articles/ncomms2076>.
- [29] Giulio Chiribella, Giacomo Mauro D’Ariano, and Paolo Perinotti. Theoretical framework for quantum networks. *Physical Review A*, 80(2), Aug 2009. <http://dx.doi.org/10.1103/PhysRevA.80.022339>.
- [30] Adrien Feix and Časlav Brukner. Quantum superpositions of ‘common-cause’ and ‘direct-cause’ causal structures. *New Journal of Physics*, 19 123028, 2017. <http://dx.doi.org/10.1088/1367-2630/aa9b1a>.
- [31] Mateus Araújo, Cyril Branciard, Fabio Costa, Adrien Feix, Christina Giarmatzi, and Časlav Brukner. Witnessing causal nonseparability. *New Journal of Physics*, 17(10):102001, 2015. <https://iopscience.iop.org/article/10.1088/1367-2630/17/10/102001>.
- [32] Ognjan Oreshkov and Christina Giarmatzi. Causal and causally separable processes. *New Journal of Physics*, 18(9):93020, 2016. <https://iopscience.iop.org/article/10.1088/1367-2630/18/9/093020>.
- [33] Philippe Allard Guérin, Adrien Feix, Mateus Araújo, and Časlav Brukner. Exponential communication complexity advantage from quantum superposition of the direction of communication. *Phys. Rev. Lett.*, 117:100502, Sep 2016. <https://link.aps.org/doi/10.1103/PhysRevLett.117.100502>.
- [34] Giulio Chiribella, Manik Banik, Some Sankar Bhattacharya, Tamal Guha, Mir Alimuddin, Arup Roy, Sutapa Saha, Sristy Agrawal, and Guruprasad Kar. Indefinite causal order enables perfect quantum communication with zero capacity channels. *New Journal of Physics*, 23(3):033039, Mar 2021. <http://dx.doi.org/10.1088/1367-2630/abe7a0>.
- [35] Giulio Chiribella. Perfect discrimination of no-signalling channels via quantum superposition of causal structures. *Physical Review A*, 86(4), Oct 2012. <http://dx.doi.org/10.1103/PhysRevA.86.040301>.
- [36] Xiaobin Zhao, Yuxiang Yang, and Giulio Chiribella. Quantum metrology with indefinite causal order. *Physical Review Letters*, 124(19), May 2020. <http://dx.doi.org/10.1103/PhysRevLett.124.190503>.
- [37] Mateus Araújo, Adrien Feix, Fabio Costa, and Časlav Brukner. Quantum circuits cannot control unknown operations. *New Journal of Physics*, 16(9):93026, 2014. <https://iopscience.iop.org/article/10.1088/1367-2630/16/9/093026>.
- [38] Tamal Guha, Mir Alimuddin, and Preeti Parashar. Thermodynamic advancement in the causally inseparable occurrence of thermal maps. *Physical Review A*, 102(3), Sep 2020. <http://dx.doi.org/10.1103/PhysRevA.102.032215>.
- [39] David Felce and Vlatko Vedral. Quantum refrigeration with indefinite causal order. *Physical Review Letters*, 125(7), Aug 2020. <http://dx.doi.org/10.1103/PhysRevLett.125.070603>.
- [40] Tamal Guha, Saptarshi Roy, Kyrylo Simonov, and Zoltán Zimborás. Activation of thermal states by quantum switch-driven thermalization and its limits, 2022. <https://arxiv.org/abs/2208.04034>.
- [41] Julian Wechs, Hippolyte Dourdent, Alastair A. Abbott, and Cyril Branciard. Quantum circuits with classical versus quantum control of causal order, 2021. <https://arxiv.org/abs/2101.08796>.
- [42] Tom Purves and Anthony J. Short. Quantum theory cannot violate a causal inequality, 2021. <https://arxiv.org/abs/2101.09107>.
- [43] Lorenzo M Procopio, Amir Moqanaki, Mateus Araújo, Fabio Costa, Irati Alonso Calafell, Emma G Dowd, Deny R Hamel, Lee A Rozema, Časlav Brukner, and Philip Walther. Experimental superposition of orders of quantum gates. *Nature Communications*, 6:7913, August 2015. <https://www.nature.com/articles/ncomms8913>.
- [44] Giulia Rubino, Lee A. Rozema, Adrien Feix, Mateus Araújo, Jonas M. Zeuner, Lorenzo M. Procopio, Časlav Brukner, and Philip Walther. Experimental verification of an indefinite causal order. *Science Advances*, 3(3), 2017. <http://advances.sciencemag.org/content/3/3/e1602589>.
- [45] K. Goswami, C. Giarmatzi, M. Kewming, F. Costa, C. Branciard, J. Romero, and A.G. White. Indefinite causal order in a quantum switch. *Physical Review Letters*, 121(9), Aug 2018. <http://dx.doi.org/10.1103/PhysRevLett.121.090503>.
- [46] Kejin Wei, Nora Tischler, Si-Ran Zhao, Yu-Huai Li, Juan Miguel Arrazola, Yang Liu, Weijun Zhang, Hao Li, Lixing You, Zhen Wang, and et al. Experimental quantum switching for exponentially superior quantum communication complexity. *Physical Review Letters*, 122(12), Mar 2019. <http://dx.doi.org/10.1103/PhysRevLett.122.120504>.
- [47] C. T. Marco Ho, Fabio Costa, Christina Giarmatzi, and Timothy C. Ralph. Violation of a causal inequality in a spacetime

- with definite causal order, 2019. <https://arxiv.org/abs/1804.05498>.
- [48] Yu Guo, Xiao-Min Hu, Zhi-Bo Hou, Huan Cao, Jin-Ming Cui, Bi-Heng Liu, Yun-Feng Huang, Chuan-Feng Li, Guang-Can Guo, and Giulio Chiribella. Experimental transmission of quantum information using a superposition of causal orders. *Physical Review Letters*, 124(3), Jan 2020. <http://dx.doi.org/10.1103/PhysRevLett.124.030502>.
- [49] K. Goswami and J. Romero. Experiments on quantum causality. *AVS Quantum Science*, 2(3):037101, Oct 2020. <http://dx.doi.org/10.1116/5.0010747>.
- [50] Márcio M. Taddei, Jaime Cariñe, Daniel Martínez, Tania García, Nayda Guerrero, Alastair A. Abbott, Mateus Araújo, Cyril Branciard, Esteban S. Gómez, Stephen P. Walborn, and et al. Computational advantage from the quantum superposition of multiple temporal orders of photonic gates. *PRX Quantum*, 2(1), Feb 2021. <http://dx.doi.org/10.1103/PRXQuantum.2.010320>.
- [51] David Felce, Nicetu Tibau Vidal, Vlatko Vedral, and Eduardo O. Dias. Indefinite causal orders from superpositions in time, 2021. <https://arxiv.org/abs/2107.08076>.
- [52] V. Vilasini. Causality in quantum theory (and beyond). Masters Thesis, ETH Zürich, 2017. https://foundations.ethz.ch/wp-content/uploads/2019/07/vilasini_master_thesis-v2.pdf.
- [53] Nikola Paunković and Marko Vojinović. Causal orders, quantum circuits and spacetime: distinguishing between definite and superposed causal orders. *Quantum*, 4:275, May 2020. <http://dx.doi.org/10.22331/q-2020-05-28-275>.
- [54] Ognjan Oreshkov. Time-delocalized quantum subsystems and operations: on the existence of processes with indefinite causal structure in quantum mechanics. *Quantum*, 3:206, December 2019. <https://doi.org/10.22331/q-2019-12-02-206>.
- [55] Nick Ormrod, Augustin Vanrietvelde, and Jonathan Barrett. Causal structure in the presence of sectorial constraints, with application to the quantum switch. arXiv:2204.10273, 2022. <https://arxiv.org/abs/2204.10273>.
- [56] Viktoria Kabel, Anne-Catherine de la Hammette, Luca Apadula, Carlo Cepollaro, Henrique Gomes, Jeremy Butterfield, Časlav Brukner. Identification is Pointless: Quantum Reference Frames, Localisation of Events, and the Quantum Hole Argument, 2024. <https://arxiv.org/abs/2402.10267>
- [57] Amin Baumeler and Stefan Wolf. The space of logically consistent classical processes without causal order. *New Journal of Physics*, 18(1):013036, jan 2016.
- [58] Mateus Araújo, Philippe Allard Guérin, and Amin Baumeler. Quantum computation with indefinite causal structures. *Physical Review A*, 96(5), Nov 2017. <http://dx.doi.org/10.1103/PhysRevA.96.052315>.
- [59] Jonathan Barrett, Robin Lorenz, and Ognjan Oreshkov. Cyclic quantum causal models. *Nature Communications*, 12(1), Feb 2021. <http://dx.doi.org/10.1038/s41467-020-20456-x>.
- [60] Esteban Castro-Ruiz, Flaminia Giacomini, Alessio Belenchia, and Časlav Brukner. Quantum clocks and the temporal localisability of events in the presence of gravitating quantum systems. *Nature Communications*, 11(1), May 2020. <http://dx.doi.org/10.1038/s41467-020-16013-1>.
- [61] Veronika Baumann, Marius Krumm, Philippe Allard Guérin, and Časlav Brukner. Page-Wootters formulation of indefinite causal order, 2021. <https://arxiv.org/abs/2105.02304>.
- [62] Stephan Bongers, Patrick Forré, Jonas Peters, and Joris M. Mooij. Foundations of structural causal models with cycles and latent variables. *The Annals of Statistics*, 49(5):2885 – 2915, 2021. <https://doi.org/10.1214/21-AOS2064>.
- [63] V. Vilasini and Renato Renner. Fundamental limits for realizing quantum processes in spacetime. *Phys. Rev. Lett.*, 133:080201, 2024. <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.133.080201>.
- [64] V Vilasini, Christopher Portmann, and Lídia del Río. Composable security in relativistic quantum cryptography. *New Journal of Physics*, 21(4):043057, apr 2019. <https://iopscience.iop.org/article/10.1088/1367-2630/ab0e3b/meta>.
- [65] Samson Abramsky and Bob Coecke. A categorical semantics of quantum protocols. Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science, 2004, pp. 415-425. <https://ieeexplore.ieee.org/document/1319636>.
- [66] Samson Abramsky and Bob Coecke. Categorical quantum mechanics, 2008. <https://arxiv.org/abs/0808.1023>.
- [67] Benjamin Schumacher and Michael D. Westmoreland. Locality and information transfer in quantum operations. *Quantum Information Processing*, 4(1):13–34, 2005. <https://doi.org/10.1007/s1128-004-3193-y>.
- [68] Robin Lorenz and Jonathan Barrett. Causal and compositional structure of unitary transformations. *Quantum*, 5:511, July 2021. <https://doi.org/10.22331/q-2021-07-28-511>.
- [69] Paolo Perinotti. Causal influence in operational probabilistic theories. *Quantum*, 5:515, aug 2021. <https://doi.org/10.22331/q-2021-08-03-515>.
- [70] V. Vilasini and Roger Colbeck. Information-processing in theories constrained by no superluminal causation vs no superluminal signalling, 2024. <https://arxiv.org/abs/2402.12446>.
- [71] Mateus Araújo, Adrien Feix, Miguel Navascués, and Časlav Brukner. A purification postulate for quantum mechanics with indefinite causal order. *Quantum*, 1:10, nov 2016.
- [72] Carla Anais Ferradini, V. Vilasini, and Victor Gitton. A causal modelling framework for classical and quantum cyclic causal structures. In-preparation based on CAF’s semester project supervised by VV and VG. Presented as poster at *Quantum Physics and Logic, Solstice of Foundations*, 2022.
- [73] Timothé e Hoffreumon and Ognjan Oreshkov. The multi-round process matrix. *Quantum*, 5:384, jan 2021. <https://doi.org/10.22331/q-2021-01-20-384>.
- [74] Giulia Rubino, Lee A. Rozema, Francesco Massa, Mateus Araújo, Magdalena Zych, Časlav Brukner, and Philip Walther. Experimental entanglement of temporal orders. In *Quantum Information and Measurement (QIM) V: Quantum Technologies*, page S3B.3. Optical Society of America, 2019. <http://www.osapublishing.org/abstract.cfm?URI=QIM-2019-S3B.3>.
- [75] V. Vilasini, Lídia del Río, and Renato Renner. Causality in definite and indefinite spacetimes, 2020. Extended abstract from QPL 2020, <https://wdi.centralesupelec.fr/users/valiron/qplmfps/papers/qs01t3.pdf>, talk version at <https://wdi.centralesupelec.fr/users/valiron/qplmfps/papers/qs01t3.pdf>.

[//www.youtube.com/watch?v=iBCntE0zxcg8](https://www.youtube.com/watch?v=iBCntE0zxcg8).

- [76] M. Salzger and V. Vilasini. Mapping indefinite causal order processes to composable quantum protocols in a spacetime, 2024. <https://arxiv.org/abs/2404.05319>
- [77] M. Salzger and V. Vilasini. Connecting indefinite causal order processes to composable quantum protocols in a spacetime. Paper in-preparation based on MS's *master's thesis* supervised by VV. Presented as talk at [Quantum Physics and Logic](#), 2022.
- [78] Natália S. Móller, Bruna Sahdo, and Nelson Yokomizo. Gravitational quantum switch on a superposition of spherical shells, 2023. <https://arxiv.org/abs/2306.10984>.
- [79] Julian Wechs, Cyril Branciard, and Ognjan Oreshkov. Existence of processes violating causal inequalities on time-delocalised subsystems, 2022. <https://arxiv.org/abs/2201.11832>.
- [80] Tein van der Lugt, Jonathan Barrett, and Giulio Chiribella. Device-independent certification of indefinite causal order in the quantum switch *Nature Communications*, 14, 5811, 2023. <https://doi.org/10.1038/s41467-023-40162-8>.
- [81] Seth Lloyd, Lorenzo Maccone, Raul Garcia-Patron, Vittorio Giovannetti, and Yutaka Shikano. Quantum mechanics of time travel through post-selected teleportation. *Physical Review D*, 84(2), jul 2011. <https://doi.org/10.1103/PhysRevD.84.025007>.
- [82] Seth Lloyd, Lorenzo Maccone, Raul Garcia-Patron, Vittorio Giovannetti, Yutaka Shikano, Stefano Pirandola, Lee A. Rozema, Ardavan Darabi, Yasaman Soudagar, Lynden K. Shalm, and Aephraim M. Steinberg. Closed timelike curves via postselection: Theory and experimental test of consistency. *Physical Review Letters*, 106(4), jan 2011. <https://doi.org/10.1103/PhysRevLett.106.040403>.
- [83] Germain Tobar and Fabio Costa. Reversible dynamics with closed time-like curves and freedom of choice. *Classical and Quantum Gravity*, 37(20):205011, sep 2020. <https://doi.org/10.1088/2F1361-6382/2Faba4bc>.
- [84] Flaminia Giacomini, Esteban Castro-Ruiz, and Časlav Brukner. Quantum mechanics and the covariance of physical laws in quantum reference frames. *Nature Communications*, 10(1):494, 2019. <https://doi.org/10.1038/s41467-018-08155-0>.
- [85] Philipp A. Höhn, Alexander R. H. Smith, and Maximilian P. E. Lock. Trinity of relational quantum dynamics. *Physical Review D*, 104(6), sep 2021. <https://doi.org/10.1103/PhysRevD.104.066001>.
- [86] Sougato Bose, Anupam Mazumdar, Gavin W. Morley, Hendrik Ulbricht, Marko Toroš, Mauro Paternostro, Andrew A. Geraci, Peter F. Barker, M. S. Kim, and Gerard Milburn. Spin entanglement witness for quantum gravity. *Physical Review Letters*, 119(24), dec 2017. <https://doi.org/10.1103/PhysRevLett.119.240401>.
- [87] C. Marletto and V. Vedral. Gravitationally induced entanglement between two massive particles is sufficient evidence of quantum effects in gravity. *Physical Review Letters*, 119(24), dec 2017. <https://doi.org/10.1103/PhysRevLett.119.240402>.
- [88] Mark Van Raamsdonk. Building up spacetime with quantum entanglement. *General Relativity and Gravitation*, 42(10):2323–2329, 2010. <https://www.worldscientific.com/doi/abs/10.1142/S0218271810018529>.
- [89] J. Maldacena and L. Susskind. Cool horizons for entangled black holes. *Fortschritte der Physik*, 61(9):781–811, 2013. <http://dx.doi.org/10.1002/prop.201300020>.
- [90] Alexander Jahn and Jens Eisert. Holographic tensor network models and quantum error correction: a topical review. *Quantum Science and Technology*, 6(3):033002, 2021. <http://dx.doi.org/10.1088/2058-9565/ac0293>.
- [91] Achim Kempf. Replacing the notion of spacetime distance by the notion of correlation. *Frontiers in Physics*, 9, 2021. <http://dx.doi.org/10.3389/fphy.2021.655857>.
- [92] Matt Wilson, Giulio Chiribella, and Aleks Kissinger. Quantum supermaps are characterized by locality, 2023. <https://arxiv.org/abs/2205.09844>.
- [93] Giulio Chiribella, Giacomo Mauro D’Ariano, and Paolo Perinotti. Probabilistic theories with purification. *Physical Review A*, 81(6), jun 2010. <https://doi.org/10.1103/PhysRevA.81.062348>.
- [94] Giulio Chiribella, Giacomo Mauro D’Ariano, and Paolo Perinotti. Informational derivation of quantum theory. *Physical Review A*, 84(1), jul 2011. <https://doi.org/10.1103/PhysRevA.84.012311>.
- [95] Jacob Grunhaus, Sandu Popescu, and Daniel Rohrlich. Jamming nonlocal quantum correlations. *Physical Review A*, 53(6):3781–3784, jun 1996. <https://doi.org/10.1103/PhysRevA.53.3781>.