Impossibility of superluminal signalling rules out causal loops in conical spacetimes

Maarten Grothus[1,2] and V. Vilasini[1,3]

[1] *Univ. Grenoble Alpes, Inria, 38000 Grenoble, France*

[2] *Univ. Grenoble Alpes, CNRS, Grenoble INP, Institut Néel, 38000 Grenoble, France*

[3] *Institute for Theoretical Physics, ETH Zürich, 8093 Zürich, Switzerland*

Contrary to popular belief, in [1] it was shown that it is theoretically possible to have operationally detectable causal loops without violating the principle of no superluminal signalling (NSS) in $(1+1)$-Minkowski spacetime. No such examples are known in $(d+1)$-Minkowski spacetimes for $d > 1$ spatial dimensions. The formalism underlying this finding is known as the *affects framework* [2], and is based on causal modelling [3, 4] which provides a-priori independent of spacetime, information-theoretic way to define causality. This allows to precisely distinguish causation and signalling (noting that these are not equivalent due to possibility of causal fine-tuning), and as a consequence, embedding these models in spacetime, it was found that NSS does not imply absence of superluminal causation (this is true in all Minkowski spacetimes, not just $d = 1$). Whether or not such superluminal causal influences compatible with NSS can be used to form operationally detectable causal loops also in $d > 1$ dimensions, has remained a key question. Our previous work [5] developed new techniques relating to (1) information-theoretic methods to infer causation from signalling, in presence of operational redundancies and fine-tuning, and (2) order-theoretic properties of spacetime geometry, in particular identifying such a property *conicality* satisfied in $d > 1$ dimensional Minkowski spacetimes but violated in $d = 1$. Building on this, the current submission resolves the open question showing that in any conical spacetime, no superluminal signalling does rule out all operationally detectable causal loops. This highlights that the relationships between information-theoretic causality and relativistic principles can generally depend on spacetime geometry.

I. BRIEF OVERVIEW OF THE AFFECTS FRAMEWORK AND CONICALITY

Causal models, signalling and causal inference The *affects framework* [2] first provides a general definition of causal models, consisting of (1) a directed graph $\mathcal{G}$ with a subset $S = \{e_X, e_Y, \dots\}$ of nodes called observed nodes (these are associated with classical random variables, RVs) and the remaining nodes being unobserved (can be described by classical, quantum or post-quantum theory), and (2) a probability distribution $P(S)$ on the observed nodes of $\mathcal{G}$ which satisfies a linking property (formally, d-separation) relative to $\mathcal{G}$. The framework then defines signalling between observed RVs of such a general causal model, via the notion of an *affects relation.* If intervening freely on a set of RVs $X \subset S$ changes the distribution of another set $Y \subset S$, then we say X affects Y and write $X \vDash Y$. In particular, $X \vDash Y$ allows one to infer that there exists $e_X \in X$ and $e_Y \in Y$ such that there is a directed path from e_X to e_Y in the causal structure $\mathcal{G}$ (or e_X is a cause of e_Y), which we will write succinctly as $X \twoheadrightarrow Y$.

The framework also defines more general forms of signalling, via *higher-order affects relations.* These take the form $X \vDash Y \mid \mathrm{do}(Z)$ (X affects Y given $\mathrm{do}(Z)$) and capture that a set of RVs X signals to the set Y, given that some interventions have been performed on another set Z. When $Z = \emptyset$, we call it a 0^{th}-order affects relation that is equivalent to $X \vDash Y$. Generally, a (higher-order) affects can contain redundancies, e.g., for $X \vDash Y \mid \mathrm{do}(Z)$ with $X = \{X_1, X_2\}$ with X_2 being a redundant RV without any causal connections. To rule out such scenarios, a notion of irreducibility has been introduced for the first [2] (Irred_1) and the third [5] (Irred_3) argument of an affects relation (i.e., the arguments X and Z for the relation $X \vDash Y \mid \mathrm{do}(Z)$), respectively. If $X \vDash Y \mid \mathrm{do}(Z)$ is Irred_1, then it allows us to infer that in any underlying causal model leading this affects relation, for all $e_X \in X$ (as opposed to just $\exists e_X$ for the case without irreducibility), there exists $e_Y \in Y$ such that $e_X \twoheadrightarrow e_Y$, which we write: for all $e_X \in X$, $e_X \twoheadrightarrow Y$. In [5], we showed that if $X \vDash Y \mid \mathrm{do}(Z)$ is Irred_3, one can infer that for all $e_Z \in Z$, $e_Z \twoheadrightarrow Y$.

Spacetime embedding, NSS and causal loops The affects framework then also defines an *embedding* of the set of RVs S of a causal model into a spacetime, modelled as a partially ordered set (*poset*) $\mathcal{T}$ where the order captures the spacetime's light cone structure. This is done by assigning each RV e_X a *location* $O(e_X) \in \mathcal{T}$, defining an ordered RV (ORV) $e_{\mathcal{X}} := (e_X, O(e_X))$. With the future light cone of an ORV defined as $\bar{\mathcal{F}}(e_{\mathcal{X}}) := \{a \in \mathcal{T} : a \succeq O(\mathcal{X})\}$, the joint future of a set $\mathcal{X}$ of ORVs is the intersection of their respective future light cones, $\bar{\mathcal{F}}_s(\mathcal{X}) = \bigcap_{e_{\mathcal{X}} \in \mathcal{X}} \bar{\mathcal{F}}(e_{\mathcal{X}})$. Having embedded information-theoretic causal models in spacetime, the relativistic principle of no superluminal signalling (NSS) is formalised as a compatibility condition between the affects relations of the causal model (capturing signalling) and the partial order of $\mathcal{T}$ (capturing light cone structure). Specifically, for each affects relation $X \vDash Y \mid \mathrm{do}(Z)$ we require $X \vDash Y \mid \mathrm{do}(Z)$ is $\mathrm{Irred}_1 \implies \bar{\mathcal{F}}_s(\mathcal{Y}) \cap \bar{\mathcal{F}}_s(\mathcal{Z}) \subseteq \bar{\mathcal{F}}_s(\mathcal{X})$. In [2] it is shown that this condition fully captures the NSS principle in spacetime, and Irred_1 plays a key role here. Further,

(a) Implications regarding the causal structure

(b) Embedding in $(1+1)$-Minkowski spacetime

Figure 1: Let $X, Y, Z, A, B \in S$. Consider $X \vDash Y, Y \vDash AB, A \vDash X, Z \vDash AB, B \vDash Z$. **(a)** No matter the underlying causal model, it can be inferred from the affects relations, that a causal loop will always be present. Here, $\twoheadrightarrow$ denotes causal relations between two RVs, while its dashed version stands for alternative causal relations, e.g. $Y \twoheadrightarrow AB$ being equivalent to $(Y \twoheadrightarrow A$ or $Y \twoheadrightarrow B)$, so at least one of the dashed arrows from Y must be present as a causal relation. **(b)** A compatible, non-degenerate embedding of these affects relations into $(1+1)$-Minkowski spacetime, where all ORVs need be located on certain light-like surfaces.

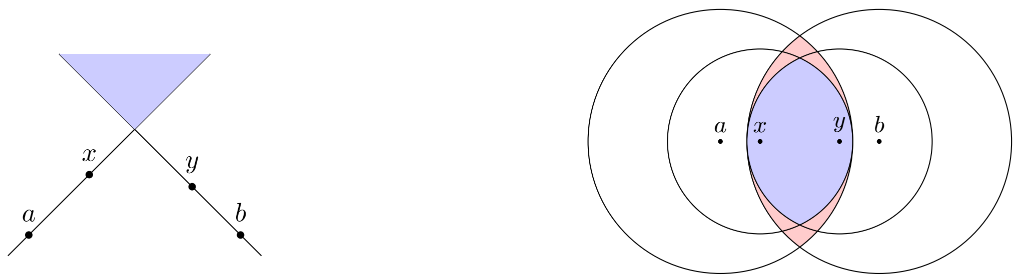

(a) $d = 1$

(b) $d = 2$

Figure 2: Representation of light cones in Minkowski spacetime for d spatial dimensions. **(a)** For $d = 1$ Minkowski spacetime does not show conicality. This is because $L_1 \coloneqq \{a, b\}$ and $L_2 \coloneqq \{x, y\}$ are distinct, yet share the same joint future. **(b)** For $d = 2$, Minkowski spacetime is conical. As can be seen from the figure which shows one particular time slice, the joint futures are distinct between L_1 and L_2 (and the red region is thus not covered by the blue one).

an embedding is considered *non-degenerate* if no two RVs are embedded into the same location.

Based on the causal inference rules mentioned above, we may (potentially) infer operationally detectable information-theoretic cycles from affects relations (i.e. signalling) alone – so-called *affects causal loops* (ACLs) [2]. The simple-most example of such a loop is given by two RVs affecting one another, i.e. both $e_X \vDash e_Y$ and $e_Y \vDash e_X$, requiring a degenerate embedding with $O(e_X) = O(e_Y)$. Several other complex classes of ACLs were defined in [2]. In Fig. 1 we illustrate a more complex example of such a causal loop which admits a non-degenerate and compatible (i.e., satisfying NSS) embedding in $(1+1)$-Minkowski spacetime.

Conicality, an order-theoretic property of spacetime geometry When embedding the RVs of a causal model into spacetime, modelled as a poset $\mathcal{T}$, the "shape" of $\mathcal{T}$ may constrain the order-theoretic properties of the embedding. One such property is *conicality* [5], capturing the requirement that the joint future region $f(L)$ (intersection of future light cones) of a finite set $L \subset \mathcal{T}$ uniquely determines the location of all points in L that contribute non-trivially to $f(L)$. In particular, this property is satisfied by $(d+1)$-Minkowski spacetime for $d > 1$ spatial dimensions, while it is violated for $d = 1$, as shown in Fig. 2. Moreover, in [5] we conjecture conicality to hold for a wide variety of spacetime configurations of this dimension without closed timelike curves.

II. SUMMARY OF MAIN CONTRIBUTIONS

Main Theorem. In conical spacetimes, all compatible embeddings of ACLs are degenerate.

Since compatibility prescisely captures the NSS principle [2, 6], this tells us that when non-degenerately embedding a causal model in any conical spacetime, the NSS principle does rule out all operationally detectable causal loops. In particular, the ACL of [2] and of Fig. 1 are compatibly and non-degenerately embedded in non-conical $(1+1)$-Minkowski spacetime. The theorem is proven in two broad steps, that also lead to useful technical results.

Step 1: Reducing higher-order to equivalent 0^{th}-order affects relations We show that for any set of possibly higher-order affects relations, there exists a set of 0^{th}-order affects relations, with the same implications for causal inference and compatibility conditions in conical spacetimes. Specifically, whenever there is a higher-order affects relation $X \vDash Y \mid \mathrm{do}(Z)$ satisfying Irred_1, there exists some $s_Z \subseteq Z$, potentially empty, such that $X \vDash Y \mid \mathrm{do}(s_Z)$ is Irred_3. Then, the combination of both relations has the same implications for causal inference and imposes the same compatibility constraints (in conical spacetimes) as a 0^{th}-order relation $X s_Z \vDash Y$ being Irred_1. In proving this result, we also advance the understanding of the interplay between Irred_1 and Irred_3, which is not yet well-understood. It follows that the original set of (possibly higher-order) affects relations implies, via causal inference, an ACL if and only if the attained set of 0^{th}-order affects relations also implies an ACL. In conical spacetimes, the original affects relations similarly satisfy compatibility (NSS) for an embedding if and only if the reduced 0^{th}-order do. Therefore, we now restrict, w.l.o.g. our attention to ACLs formed by 0^{th}-order affects relations.

Step 2: Ruling out ACLs in 0^{th}-order affects relations embedded in spacetime We sketch the proof for the case where the ACL is formed only by 0^{th}-order affects relations, illustrating the main technique intuitively by means of an example. Generally, the causal inference possible from an (0^{th}-order) affects relation $X \vDash Y$ being Irred_1 is limited, allowing one to infer only that each $e_X \in X$ is a cause of *some* element e_Y of the second argument Y. Therefore, a causal loop is implied by a set of such affects relations if and only if no matter which element in the second argument of each affects relation in such a set we "choose" to be causally influenced by the first argument, the resulting causal structure will be cyclic. Considering Fig. 1a as an example, we have some causal loop no matter whether Y causes A or B (implied by $Y \vDash AB$), and whether Z causes A or B (implied by $Z \vDash AB$). We now evaluate compatibility "along" the causal arrows (which align with the spacetime causal structure), starting from ORV $\mathcal{A}$. For a compatible and non-degenerate embedding of these affects relations in any spacetime, $A \vDash X \vDash Y$ gives $\bar{\mathcal{F}}(\mathcal{A}) \supseteq \bar{\mathcal{F}}(\mathcal{Y})$ and $Y \vDash AB$ gives $\bar{\mathcal{F}}(\mathcal{Y}) \supseteq \bar{\mathcal{F}}_s(\mathcal{AB})$. Next, we show that in conical spacetimes, compatibility further implies strictness of these subset relations on the futures, explicitly: $\bar{\mathcal{F}}(\mathcal{A}) \supsetneq \bar{\mathcal{F}}(\mathcal{Y}) \supsetneq \bar{\mathcal{F}}_s(\mathcal{AB})$. Continuing onward for $\mathcal{B}$ separately, through the same compatibility + conicality implications, we obtain $\bar{\mathcal{F}}(\mathcal{B}) \supsetneq \bar{\mathcal{F}}(\mathcal{Z}) \supsetneq \bar{\mathcal{F}}_s(\mathcal{AB})$. Moreover, conicality allows to combine these statements while preserving the strict subset relations, i.e., $\bar{\mathcal{F}}_s(\mathcal{AB}) \supsetneq \bar{\mathcal{F}}_s(\mathcal{YZ})$ and $\bar{\mathcal{F}}_s(\mathcal{YZ}) \supsetneq \bar{\mathcal{F}}_s(\mathcal{AB})$. (Note that neither of these is satisfied in the non-conical spacetime embedding shown in Fig. 1b). We therefore receive a contradiction of the form $\bar{\mathcal{F}}_s(\mathcal{AB}) \neq \bar{\mathcal{F}}_s(\mathcal{AB})$.

A generalisation and notion of fine-tuning for spacetime embeddings We also provide a generalisation of the main theorem by defining the notion of *conical embeddings*, which consider the condition of conicality only for the futures of all points $O(\mathcal{S})$ in the image of the embedding of the RVs S and the intersections of their futures. While any embedding into a conical poset $\mathcal{T}$ is trivially conical, we can have non-conical $\mathcal{T}$ within which there exists a conical embedding. This shows that causal loops without superluminal signalling require a form of fine-tuning in the spacetime embedding, captured by non-conicality of the embedding – in addition to fine-tuning in the causal model, as noted in [2]. Indeed, non-conical embeddings require the joint futures of distinct sets of spacetime events to align in a very precise "fine-tuned" manner, as we also see in Fig. 1 and Fig. 2a. Moreover, we also strengthen the observation of [2], showing that a specific form of fine-tuning in causal models, captured via *clustering* of affects relations [5] is necessary for causal loops to be embeddable without superluminal signalling.

III. DISCUSSION

Our results demonstrate definitively for the first time that the relationships between the relativistic principles of no superluminal signalling and no causal loops depends inherently on the spacetime geometry, and were established without making assumptions on the theory governing the underlying cyclic causal models. They also demonstrate that any causal loop compatible with no superluminal signalling in a spacetime requires both fine-tuning in the causal model and fine-tuning in the spacetime embedding, where a absence of (superluminal) signalling can turn into (superluminal) signalling through slight perturbations of the causal mechanisms (or spacetime locations). This extends the correspondence between the information-theoretis and spatio-temporal causal structures beyond the connections established in [5]. Some immediate open questions are discussed below.

When the underlying causal model, the cardinality/dimensions of the systems is unknown, the causal inference rules outlined here can be regarded as "complete" as argued in [7]. But it is not known whether this is the case once some of these properties are known and fixed. Can there be stronger rules for causal inference from affects relations when restricting the cardinality of the RVs, such that more sets of affects relations would certify the presence of a causal loop than before? Furthermore, the affects framework allows for cyclic causal structures in general theories but restricts them via the d-separation property. While this property holds in many scenarios, there exist cyclic models that violate d-separation [8], and alternative graph-separation properties such as σ-separation [9] and p-separation [10] have been proposed. An open question is whether our results can be extended to cyclic causal models with σ- or p-separation. Since the exact characterisation of the set of cyclic models respecting d-separation, and whether causal modelling frameworks for process matrices in quantum theory [11, 12] respect d-separation are open questions, one may consider whether spacetime-embeddable causal loops could arise from valid process matrices.

[1] V. Vilasini and R. Colbeck, Impossibility of superluminal signaling in minkowski spacetime does not rule out causal loops, Phys. Rev. Lett. **129**, 110401 (2022), arXiv:2206.12887 [quant-ph].

[2] V. Vilasini and R. Colbeck, General framework for cyclic and fine-tuned causal models and their compatibility with space-time, Phys. Rev. A **106**, 032204 (2022), arXiv:2109.12128 [quant-ph].

[3] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search* (Springer New York, NY, 1993).

[4] J. Pearl, *Causality: Models, Reasoning and Inference*, 2nd ed. (Cambridge University Press, 2009).

[5] M. Grothus and V. Vilasini, Characterizing signalling: Connections between causal inference and space-time geometry (2024), arXiv:2403.00916 [gr-qc].

[6] V. Vilasini and R. Colbeck, Theories with no superluminal signaling have greater information-processing power than theories with no superluminal causation, Phys. Rev. Res. **8**, 013070 (2026), arXiv:2402.12446 [quant-ph].

[7] M. Grothus, *Compatibility of cyclic causal structures with spacetime in general theories with free interventions*, Master's thesis, ETH Zurich (2022), arXiv:2211.03593 [quant-ph].

[8] R. M. Neal, On deducing conditional independence from d-separation in causal graphs with feedback (research note), Journal of Artificial Intelligence Research **12**, 87 (2000).

[9] P. Forré and J. M. Mooij, Markov properties for graphical models with cycles and latent variables (2017), arXiv:1710.08775 [math.ST].

[10] C. Ferradini, V. Gitton, and V. Vilasini, Cyclic functional causal models beyond unique solvability with a graph separation theorem (2025), arXiv:2502.04171 [math.ST].

[11] J. Barrett, R. Lorenz, and O. Oreshkov, Cyclic quantum causal models, Nature Commun. **12**, 885 (2021), arXiv:2002.12157 [quant-ph].

[12] C. Ferradini, V. Gitton, and V. Vilasini, Cyclic quantum causal modelling with a graph separation theorem (2025), arXiv:2502.04168 [quant-ph].

Impossibility of superluminal signalling rules out causal loops in conical spacetimes

Maarten Grothus[1, 2] and V. Vilasini[1, 3]

[1] *Univ. Grenoble Alpes, Inria, 38000 Grenoble, France*
[2] *Univ. Grenoble Alpes, CNRS, Grenoble INP, Institut Néel, 38000 Grenoble, France*
[3] *Institute for Theoretical Physics, ETH Zürich, 8093 Zürich, Switzerland*

Note: These results are part of a paper in preparation. The current draft is extracted from relevant sections of MG's recently submitted PhD thesis, and its presentation is yet to be fully optimised for a paper.

I. INTRODUCTION

When embedding causal models (including fine-tuned ones) into space-time, the absence of superluminal signalling (NSS) does not generally imply the absence of superluminal causation [1, 2]. This raises a fundamental question: how does the relativistic NSS principle constrain information-theoretic causal structures? Specifically, if space-time's causal structure is acyclic (i.e., devoid of closed timelike curves), does NSS guarantee the acyclicity of information-theoretic causal structures as well?

Surprisingly, the latter question has been shown to have a negative answer: NSS does not preclude causal loops, even in Minkowski space-time [3]. This was demonstrated by constructing a cyclic causal model embedded in 1+1-dimensional Minkowski space-time, where the existence of a causal loop could be operationally verified through interventions without violating NSS. Whether such loops can exist in Minkowski space-times with higher spatial dimensions, and the principles needed to rule them out, have been key open questions. Addressing these issues is fundamental to understanding the principles necessary for emergence of acyclic causal structures in a physical theory, which ensure a clear causal order where one event influences another without reciprocal influence.

However, answering such questions is conceptually and technically challenging. From an information-theoretic perspective, it requires methods for inferring causation from patterns of signalling among systems and operationally identifying redundant variables in signalling relations in the presence of fine-tuning, which is little explored even classically. Moreover, the causal loop identified in 1+1-dimensional Minkowski space-time [3] cannot be embedded in higher-dimensional Minkowski space-time without superluminal signalling. It was conjectured in [3] that such loops may generally be precluded in higher dimensions due to geometric properties of light cone intersections. The companion paper [1] of [3] classified several more complex classes of operationally detectable causal loops (non-exhaustively), whose space-time embeddings must be accounted for to prove this conjecture. Proving the conjecture also requires a rigorous formalisation of the geometric properties distinguishing 1+1-dimensional and higher-dimensional Minkowski space-times, focusing on order-theoretic properties of the space-time's causal structure rather than continuous manifold-like properties.

II. BACKGROUND

Contrary to popular belief, [3] establishes that the absence of superluminal signalling alone is generally insufficient to rule out operationally detectable causal loops. This raises the question whether we can impose additional conditions to limit such loops to unphysical scenarios.

Formally, this result is expressed in terms of the *affects framework* established in [1], which provided rules to infer causal relations from signalling between sets of random variables (RVs) in the context of a *causal model* [4, 5] on a set of RVs $S = \{e_X, e_Y, \ldots\}$, relating a probability distribution $P(S)$ with a directed graph $\mathcal{G}$ containing S as nodes, satisfying some Markov condition with respect to each other – in our-case, the widely used d-separation condition. In addition to S, the graph $\mathcal{G}$ may contain *unobservable nodes*, to be associated with classical, quantum or GPT systems. Generally, whenever a set of RVs $X \subset S$ signals to another set $Y \subset S$, i.e. when the value of Y can be influenced by forcing the values of X to a certain value (using an *intervention* $\mathrm{do}(X)$), we say that X affects Y and write $X \vDash Y$.

In particular, $X \vDash Y$ allows one to infer that there exists $e_X \in X$ and $e_Y \in Y$ such that e_X is a (not necessarily direct) cause of e_Y, which we will write as $X \twoheadrightarrow Y$. To isolate the effect of the causal influence of a subset $X' \subsetneq X$, we introduce the notion of a *higher-order affects relation* $X' \vDash Y \mid \mathrm{do}(X \setminus X')$: after enforcing values on $X \setminus X'$, there is still signalling from X' to Y, and hence $X' \twoheadrightarrow Y$. If for all such subsets $X' \subset X$, the respective higher-order

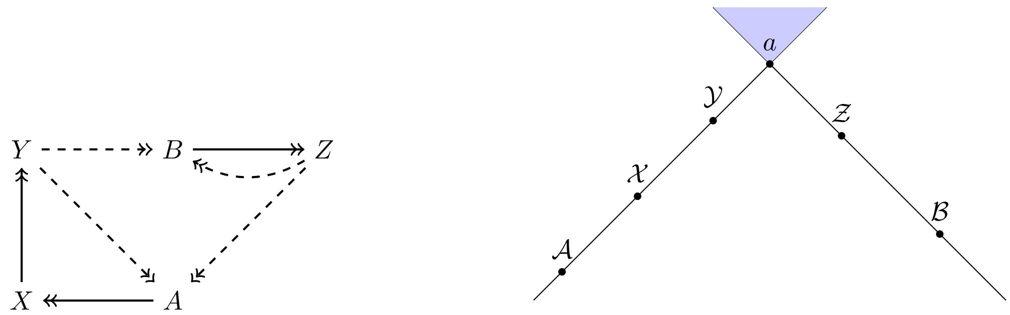

(a) Implications regarding the causal structure **(b)** Embedding in 1+1-Minkowski space-time

Figure II.1: (a) Visualisation of implications of the affects relations given in Example II.1 for the causal structure. Here, $\twoheadrightarrow$ visualises a (possibly indirect) causal relationship between two RVs, operationally inferred from a set of affects relations. Its dashed equivalent stands for alternative causal relationships implied by affects relations, e.g. $Y \twoheadrightarrow AB \iff (Y \twoheadrightarrow A \text{ or } Y \twoheadrightarrow B)$. No matter which of these alternative causal relationships is featured in the underlying causal structure, it is necessarily cyclic. **(b)** Sketch of a compatible non-degenerate embedding of Example II.1 into 1+1-Minkowski space-time. In the latter, all ORVs which are part of the causal loop are located on a light-like surface, and therefore on the boundary of the light cone of the respective earlier RVs. Indicated in blue is $\bar{\mathcal{F}}_s(\mathcal{AB})$, which is the region the causal loop can be operationally detected in, and equal to the future of a.

affects relation holds, we call $X \vDash Y$ *irreducible* in the first argument (Irred$_1$) and have $e_X \twoheadrightarrow Y$. Beyond that, we may consider generic higher-order affects relations $X \vDash Y \mid \mathrm{do}(Z)$ with X, Y, Z disjoint.[1] If such a relation is irreducible in the third argument (Irred$_3$) – according to a similar intuition as for the first one – we similarly find that $e_Z \twoheadrightarrow Y$.

In addition, the affects framework combines this information-theoretic perspective with an *embedding* of the set of RVs S into a space-time modelled as a partially ordered set (*poset*) $\mathcal{T}$, assigning each RV e_X a *location* $O(e_X) \in \mathcal{T}$. The tuple of RV and *location* is then called an ORV $e_{\mathcal{X}} := (e_X, O(e_X))$, and their causal future is defined as $\bar{\mathcal{F}}(e_{\mathcal{X}}) := \bar{J}(O(e_{\mathcal{X}})) := \{a \in \mathcal{T} : a \succeq O(\mathcal{X})\}$. An embedding is considered *compatible* if it does not allow for superluminal signalling as encoded by affects relations. For $X, Y, Z \subset S$ and for any set of ORVs $\mathcal{A}$, and for the joint ("support") future of the ORVs given by $\bar{\mathcal{F}}_s(\mathcal{A}) := \bigcap_{e^i_{\mathcal{A}} \in \mathcal{A}} \bar{\mathcal{F}}(e^i_{\mathcal{A}})$, this formalises as

$$X \vDash Y \mid \mathrm{do}(Z) \text{ is Irred}_1 \quad \Longrightarrow \quad \bar{\mathcal{F}}_s(\mathcal{Y}) \cap \bar{\mathcal{F}}_s(\mathcal{Z}) \subseteq \bar{\mathcal{F}}_s(\mathcal{X})\,. \tag{II.1}$$

Further, an embedding is considered *non-degenerate* if no two RVs are embedded into the same location.

A. Affects causal loops

Using results for causal inference as given in [1, 6], we may (potentially) infer operationally detectable information-theoretic cycles from signalling relations alone – so-called *affects causal loops* (ACLs) [1]. The simple-most example of such a loop is given by the set $\mathscr{A} = \{X \vDash Y, Y \vDash X\}$ of affects relations known to be present. While the authors of [1] differentiate at least 10 different types of causal loops of increasing complexity (with the aforementioned example being of type 1), this list is not exhaustive.[2] Hence, considering these specific classes further is only of very limited use to establish general insights on the properties of ACLs.

While it is straightforward to see that this example can not be embedded compatibly and non-trivially into physical space-time, no clear evidence exists whether we could do so for more complex ACLs. While in [3], an example of a non-degenerate compatible embedding in 1+1-Minkowski space-time is established, it canot to be embedded into higher-dimensional Minkowski space-time without implying superluminal signalling. Furthermore, it apparently requires a particular fine-tuning both in regard to the location of the random variables in space-time and with respect to their probability distributions. We give below a similar example, which allows for a bit more flexibility in the embedding (as no ORV is constrained to be located in a).

---

[1] which does not necessarily imply $XZ \vDash Y$, cf. [1], and which may again be Irred$_1$ or not

[2] Multiple examples for ACLs beyond these 10 types are given in [1, 7].

Example II.1. Let $X, Y, Z, A, B \in S$. Consider $\mathscr{A} = \{X \vDash Y, Y \vDash AB, A \vDash X, Z \vDash AB, B \vDash Z\}$. As shown in [1, Ex. B.2] and visualised in Fig. II.1a, this set of affects relations constitutes an affects causal loop (of type 7): No matter which causal relations are realised, a causal loop will always be present. Embedding this into a poset spanned by the conditions $\mathcal{A} \prec \mathcal{X} \prec \mathcal{Y} \prec a \succ \mathcal{Z} \succ \mathcal{B}$, we obtain a non-degenerate embedding. This embedding can be realised in 1+1-Minkowski space-time, as shown in Fig. II.1b, but not in the higher-dimensional case.

In Lemma II.6, we will show that this embedding can not be implemented in conical space-times, as it satisfies

$$\bar{\mathcal{F}}(\mathcal{Z}) \subsetneq \bar{\mathcal{F}}(\mathcal{B}) \quad \wedge \quad \bar{\mathcal{F}}(\mathcal{Y}) \subsetneq \bar{\mathcal{F}}(\mathcal{A}) \quad \wedge \quad \bar{\mathcal{F}}_s(\mathcal{Z}\mathcal{Y}) = \bar{\mathcal{F}}_s(\mathcal{A}\mathcal{B})\,. \tag{II.2}$$

In particular, this implies that it can not be realised in higher-dimensional Minkowski space-time.

Building upon our extensive study of the order-theoretic properties of space-time geometry and fine-tuned signalling in [8], we apply these results to show that these facts are no mere coincidences, but rather requirements for any embedding of an ACL: as we demonstrate here, it requires both a *non-conical* embedding (reviewed in Section II C) – which in Minkowski space-time, can only be realised in 1+1 dimensions – and *clustering* within its causal model, as reviewed in Section II B. Before proving this result, we will present two essential implications of these notions, providing us the groundwork required to do so.

We will continue by motivating these notions and introducing them formally, before proving this result.

B. Beyond Irreducibility: Clustering

In [6], we introduce the notion of *clustering* of affects relations:

Definition II.2. An affects relation $S_1 \vDash S_2|\mathrm{do}(S_3)$ is called clustered in the ith argument (denoted Clus_i) if there is no affects relation obtained by replacing S_i with a subset no $s_i \subsetneq S_i$.

Beyond the notion of operational irreducibility of an affects relation, clustering certifies the non-existence of any "simpler" affects relations (the affects relations it could reduce to if reducible) – and thereby, the absence of simpler affects relations resulting in equivalent implications for causal inference and compatibility. Furthermore, clustering (for any argument) is an operationally detectable signature of fine-tuning of the underlying causal mechanisms. Such possibilities for information transfer between sets of systems that is not detectable within subsets of those systems [9], appear for cryptography [10–12], quantum error correction [13, 14] and summoning quantum information in space-time [15–18].

While many types of fine-tuning are linked with clustered affects relations, not every fine-tuned causal model features clustering. In particular, [6] demonstrates this may be the case if the causal structure admits some fine-graning. However, [6] also demonstrates that causal models without clustering admit a much simpler characterisation of signalling in terms of affects relations, even if they are fine-tuned otherwise.

Corollary II.3. [6, Corollary D.7] Consider an affects relation $X \vDash Y \mid \mathrm{do}(Z), W$ in a causal model without clustering (in the first three arguments). Then there exist $e_X \in X, e_{YW} \in YW$ such that $e_X \vDash e_{YW}$.

As this imposes quite strong constraints on any compatible embeddings, it will allow us to rule out any compatible embeddings of causal loops in absence of clustering.

C. Conicality

When embedding the RVs of a causal model into space-time, modeled as a poset $\mathcal{T}$, the "shape" of $\mathcal{T}$ may generally constrain the order-theoretic properties of the embedding.

One particular such property is *conicality*, capturing the requirement that the joint future region $f(L)$ (intersection of future light cones) of a finite set $L \subset \mathcal{T}$ uniquely determines the location of all points in L that contribute non-trivially to $f(L)$. In particular, this property is satisfied by d+1-dimensional Minkowski space-time for $d > 1$, while it is violated for $d = 1$. Moreover, in [6] we conjecture conicality to hold for a wide variety of space-time configurations of this dimension without closed timelike curves.

Definition II.4 (Spanning Elements [6])**.** Let $\mathcal{T}$ be a poset and $L \subset \mathcal{T}$ be finite. Then, the set of *spanning elements*, denoted by $\mathrm{span}(L)$, is given by the union of all sets $L' \subseteq L$ that satisfy both $\bigcap_{x\in L'} \bar{J}^+(x) = \bigcap_{x\in L} \bar{J}^+(x)$ and $\nexists L'' \subsetneq L$: $\bigcap_{x\in L''} \bar{J}^+(x) = \bigcap_{x\in L} \bar{J}^+(x)$.

Alternatively, we can phrase this in terms of a set of ORVs $\mathcal{X} \subset \mathcal{S}$ on this poset. Then, the set of *spanning elements*, denoted by $\mathrm{span}(\mathcal{X})$, is given by the union of all sets $s_\mathcal{X} \subseteq \mathcal{X}$ which satisfy both $\bar{\mathcal{F}}_s(s_\mathcal{X}) = \bar{\mathcal{F}}_s(\mathcal{X})$ and $\nexists t_\mathcal{X} \subsetneq s_\mathcal{X} : \bar{\mathcal{F}}_s(t_\mathcal{X}) = \bar{\mathcal{F}}_s(\mathcal{X})$.

Definition II.5 (Conicality [6])**.** Let $\mathcal{T}$ be a poset. We say that $\mathcal{T}$ is a *conical* poset if for any two finite subsets $L_i, L_j \subset \mathcal{T}$, $f(L_i) = f(L_j) \implies \mathrm{span}(L_i) = \mathrm{span}(L_j)$ holds.

For our main theorem, the following property of conicality will be critical:

Lemma II.6. Let $\mathcal{T}$ be a poset and $\mathcal{S}$ a set of ORVs embedded non-degenerately and conically therein. Let $\mathcal{Y}_i \subset \mathcal{S}$ and $\mathcal{X}_i \in \mathcal{S} \setminus \mathcal{Y}_i$ for all i, and $\mathcal{X} = \mathrm{span}(\mathcal{X}) = \prod_i \mathcal{X}_i, \mathcal{Y} = \prod_i \mathcal{Y}_i$. Then,

$$\bar{\mathcal{F}}_s(\mathcal{Y}_i) \subsetneq \bar{\mathcal{F}}(\mathcal{X}_i) \quad \forall i \qquad \implies \qquad \bar{\mathcal{F}}_s(\mathcal{Y}) \subsetneq \bar{\mathcal{F}}_s(\mathcal{X}) \,. \tag{II.3}$$

The proof of this lemma is given in Appendix A.

This lemma encodes the crucial difference why the non-existence of compatible non-degenerate causal loops is only proven for conical embeddings: it does not hold for non-conical embeddings into (necessarily non-conical) posets, as given in Example II.1. In the conical scenario, the strict subset relation prevents portions of the light cone $\bar{\mathcal{F}}(\mathcal{X}_i)$ to coincide with portions of $\bar{\mathcal{F}}_s(\mathcal{Y})$.

III. REDUCTION TO 0^{th}-ORDER AFFECTS RELATIONS AND MAIN RESULTS

Having outlined our main ingredients, we establish theorems to rule out causal loops in two steps. First, we show that to generally rule out ACLs, it is sufficient to consider 0^{th}-order unconditional affects relations which are operationally irreducible in their interventional arguments. This significantly simplifies the structure of ACLs to potentially consider, granting a more straightforward relation between the affects relations and the causal inference and compatibility conditions they entail. Then, we demonstrate that both space-time conicality and the absence of clustering are sufficient rule out a non-degenerate embedding of ACLs.

To break down the question from the generic case to 0^{th}-order affects relations exclusively, restricting the problem from affects relations irreducible in the first argument to only those irreducible both in the first and third argument seems like an expedient approach: To do so, we make use of the main theorem of [6], for $X \vDash Y \mid \mathrm{do}(Z), W$ being $\mathrm{Irred}_{1,3}$ we achieve parallel implications for causal inference and compatibility:

Theorem III.1 ([6])**.** Let $\mathscr{A}$ be a set of (conditional) affects relations and $\mathscr{A}' \subseteq \mathscr{A}$ consist of all affects relations in $\mathscr{A}$ that are both Irred_1 and Irred_3. Then

$$(X \vDash Y \mid \mathrm{do}(Z), W) \in \mathscr{A}' \quad \implies \quad e_{XZ} \twoheadrightarrow YW \;\forall e_{XZ} \in XZ \,. \tag{III.1}$$

Consider further a compatible embedding $\mathcal{E}$ of $\mathscr{A}$ into a space-time $\mathcal{T}$. Then, **if $\mathcal{T}$ satisfies conicality or $\mathscr{A}$ is rooted in a causal model without clustering**,

$$(X \vDash Y \mid \mathrm{do}(Z), W) \in \mathscr{A}' \quad \wedge \quad \bar{\mathcal{F}}_s(\mathcal{Y}) \cap \bar{\mathcal{F}}_s(\mathcal{W}) \not\subseteq \bar{\mathcal{F}}_s(\mathcal{X}) \cap \bar{\mathcal{F}}_s(\mathcal{Z}) \tag{III.2}$$

In particular, these constraints correspond to those imposed by a 0^{th}-order affects relation $XZ \vDash YW$ being Irred_1. In this paper, we go beyond this theorem by broadening its scope to cover the implications of affects relations irreducible in the first, yet reducible in the third argument.

Lemma III.2. Consider a causal model featuring a higher-order affects relation $X \vDash Y \mid \mathrm{do}(Z), W$ satisfying Irred_1, which w.l.o.g. reduces to $X \vDash Y \mid \mathrm{do}(s_Z), W$ being Irred_3 for some $s_Z \subseteq Z$, potentially empty. If the causal model features no affects relations with Clus_3 or we have a conical embedding, $X s_Z \vDash YW$ with Irred_1 has the same implications for causal inference and compatibility (as the combination of $X \vDash Y \mid \mathrm{do}(Z), W$ Irred_1 and $X \vDash Y \mid \mathrm{do}(s_Z), W$

Irred$_3$):

$$\forall e_{XZ} \in X s_Z : \quad e_{XZ} \twoheadrightarrow YW \quad \wedge \quad \bar{\mathcal{F}}_s(\mathcal{X}) \cap \bar{\mathcal{F}}_s(s_{\mathcal{Z}}) \supseteq \bar{\mathcal{F}}_s(\mathcal{Y}) \cap \bar{\mathcal{F}}_s(\mathcal{W}) \,. \tag{III.3}$$

We prove this lemma in Appendix A. Note that this lemma does not imply that the causal model in question actually does include an affects relation $X s_Z \vDash YW$ with Irred$_1$.

The main challenge in establishing this lemma beyond Theorem III.1 lies in the fact that the interplay between irreducibliity in the first and third argument is not yet well-understood. Hence, it is not straightforward to perform this reduction, requiring a comprehensive analysis. Specifically, if an affects relation $X \vDash Y \mid \mathrm{do}(Z)$ is Irred$_1$, but Red$_3$, it is not clear that the accordingly reduced affects relation $X \vDash Y \mid \mathrm{do}(s_Z)$ should still be Irred$_1$.

Remarkably, these difficulties resolve themselves when relying on a weaker notion of irreducibility, restricted to affects relations atomic in the first argument.[3] We take a deeper look at this notion in Appendix B.

Corollary III.3. Consider a set of higher-order affects relations $\mathscr{A}$. Then, Lemma III.2 yields a set of 0^{th}-order affects relations $\mathscr{A}'$ which imply an ACL if and only if $\mathscr{A}$ implies an ACL. Furthermore, under the conditions of that lemma, both $\mathscr{A}$ and $\mathscr{A}'$ yield the same constraints for compatible embeddings into a space-time $\mathcal{T}$.

Accordingly, this corollary allows us to apply the result of Lemma III.2 to the scenario of multiple affects relations forming an ACL. At this point however, one may wonder whether there may be further causal inference rules to the ones reviewed above, which could contribute to an ACL. For individual affects relations, it is easy to see that any stronger rules would have obvious counter-examples, e.g. by including superficial nodes in the second argument of the relation. The conjunction of multiple affects relations (or beyond that, information on the cardinality of RVs, their probability distribution or other constraints on the causal structure) could however in some cases allow for additional causal inference. Nonetheless, [7, Lemma 3.12] provides arguments that in absence of such restrictions, one can generally give a causal model featuring exactly the causal influences implied by the inference rules reviewed above.[4]

Similarly to the main results of [6], we again state a pair of theorems due to clustering and conicality.

Theorem III.4. For any space-times, all compatible embeddings of ACLs rooted in a causal model without clustering are trivial.

By being more specific, this result improves on the previously known fact that fine-tuning is generally required to have any chance for embedding a causal loop into space-time [1].

We sketch the proof of this statement, which follows rather directly from Corollary II.3: due to clustering, we can ascribe each (unconditional 0^{th}-order) affects relation $X \vDash Y$ to some atomic affects relations $e_X \vDash e_Y$ with $e_X \in X, e_Y \in Y$. As for such relations, causal inference $e_X \twoheadrightarrow e_Y$ and compatibility $e_X \preceq e_Y$ are directly aligned, the embedding of any causal loop built from such relations must be trivial, as all ORVs involved must be located in a single point. The full proof is given in Appendix A.

Theorem III.5. In conical space-times, all compatible embeddings of ACLs are degenerate.

Here, we sketch the proof, illustrating the main technique employed by means of an example. Generally, the causal inference possible from an (unconditional 0^{th}-order) affects relation $X \vDash Y$ irreducible in the first argument is restricted, allowing one to infer only that each $e_X \in X$ is a cause of *some* element e_Y of the second argument Y. Therefore, a causal loop is implied by a set of such affects relations if and only if no matter which element in the second argument of each affects relation in such a set we “choose” to be causally influenced by the first argument, the resulting causal structure will be cyclic.[5] Considering Fig. II.1a as an example, we have some causal loop no matter whether Y causes A or B (implied by $Y \vDash AB$), and whether Z causes A or B (implied by $Z \vDash AB$). We now evaluate compatibility “along” the causal arrows (which due to Lemma III.2 align with the space-time causal structure), starting from ORV $\mathcal{A}$. With conicality rendering the subset relation for compatibility strict, this results

[3] Specifically, case **I** in the proof corresponds to the case where the relation is irreducible in the third argument in a weaker sense only, while case **II** represents the case where the relation is reducible in the first argument in this weaker sense.

[4] More specifically, RVs of sufficient cardinality may be fine-grained into independent (Cartesian) factors, with each affects relation only interacting with a subset of these factors. [6, 19]

[5] How to evaluate this systematically/algorithmically for a causal structure inferred from a set of affects relations is outlined in [7].

in

$$\bar{\mathcal{F}}(\mathcal{A}) \supsetneq \bar{\mathcal{F}}(\mathcal{Y}) \supsetneq \bar{\mathcal{F}}_s(\mathcal{AB}) = \bar{\mathcal{F}}_s(\mathrm{span}(\mathcal{AB}))\,, \tag{III.4}$$

where the final equation follows as by definition, generally $\bar{\mathcal{F}}_s(\cdot) = \bar{\mathcal{F}}_s(\mathrm{span}(\cdot))$. Continuing onward for $\mathcal{B}$ separately, we receive

$$\bar{\mathcal{F}}(\mathcal{B}) \supsetneq \bar{\mathcal{F}}(\mathcal{Z}) \supsetneq \bar{\mathcal{F}}_s(\mathcal{AB}) = \bar{\mathcal{F}}_s(\mathrm{span}(\mathcal{AB}))\,. \tag{III.5}$$

Combining these statements using Lemma II.6, we receive a contradiction of the form $\bar{\mathcal{F}}_s(\mathrm{span}(\mathcal{AB})) \neq \bar{\mathcal{F}}_s(\mathrm{span}(\mathcal{AB}))$, no matter the span of the sets involved. For the particular case that $\mathcal{AB} = \mathrm{span}(\mathcal{AB})$ and $\mathcal{YZ} = \mathrm{span}(\mathcal{YZ})$, we have

$$\bar{\mathcal{F}}_s(\mathcal{AB}) \supsetneq \bar{\mathcal{F}}_s(\mathcal{YZ}) \supsetneq \bar{\mathcal{F}}_s(\mathrm{span}(\mathcal{AB}))\,. \quad \lightning \tag{III.6}$$

The statement of Theorem III.5 generalises beyond conical space-times to the case of *conical embeddings*, as outlined in Definition A.2. This result is formally stated and proven as Theorem A.4. This confirms our preliminary observation that even in non-conical space-times such as 1+1-Minkowski, a fine-tuning of the embedding (cf. Remark A.3) is required to support causal loops. Additionally, this results in a nice parallel between causal models and compatible space-time embeddings: a fine-tuning of both is necessary to realise an ACL.

We conclude with an example confirming that while any embedding of an ACL must be degenerate, it is not necessarily trivial.

Example III.6. Let $A, B, C, D \in S$. Consider $\mathscr{A} = \{A \vDash BC, B \vDash AD, C \vDash AB, D \vDash AB\}$ as well as Minkowski space-time of arbitrary dimension. Then there exists a degenerate, yet non-trivial embedding given by $O(A) = O(C) \not\preceq\not\succeq O(B) = O(D)$.

IV. DISCUSSION

For space-time structure, we introduced the order-theoretic property of *conicality*, showing that it is satisfied in Minkowski space-times with $d > 1$ spatial dimensions and violated for $d = 1$. Using this, we have shown a connection (Theorem III.1) between conical space-times and causal models without a form of clustering. Both cases lead to a correspondence between the causal order constraints on the observed variables coming from compatibility with a space-time embedding and those coming from purely information-theoretic causal inference (although these two types of constraints behave differently in general). Moreover, this correspondence reveals that in conical space-times, the principle of no superluminal signalling ensures a clear temporal ordering originating from an affects relation $X \vDash Y \mid \mathrm{do}(Z), W$ that is irreducible in X and Z: between when the interventional arguments $\mathcal{X}$ and $\mathcal{Z}$ and the observational arguments $\mathcal{Y}$ and $\mathcal{W}$ are jointly accessible, with the former ordered before the latter.

Building upon these results, we demonstrated that clustering and conicality are each sufficient to certify the impossibility of affects causal loops within physical space-time – without requiring any further information regarding the causal model in question, which may involve quantum or post-quantum systems. Even beyond physical space-times, we demonstrated that any causal loop compatible with no superluminal signalling requires both a fine-tuned causal model and a fine-tuned embedding into space-time. This notion of a fine-tuned embedding is analogous to fine-tuning in operational causal models, where a non-signalling relation can turn into signalling through slight perturbations of the causal mechanisms.

A. Open questions

While we have ruled out compatible embeddings of causal loops detectable from the signalling relations in a causal model alone, some questions remain. Specifically, our (nigh-)complete ignorance of the causal model simultaneously poses the main open question left by our results. Can we still rule out compatible causal loops when restricting the cardinality of the RVs involved, e.g. to take binary values only? For this case, the completeness of the causal inference rules has not been sufficiently demonstrated. More complex causal inference rules could arise in this complex, rooted either in co-occurence of affects relations, the absence of certain affects relations, or making use not only of the signalling structure of the RVs, but also of their pre- and post-intervention correlations. Notably, the latter would constitute a class of operationally detectable causal loops beyond (pure) *affects* causal loops. We leave these questions for future work.

Furthermore, while our results apply to the wide class of d-separated causal models, the question remains whether they uphold when generalising to alternative, more lenient graph-separation conditions, which capture a wider variety of cyclic causal models. These include σ-separation [20] and p-separation [21]. Given that the applicable separation conditions for the causal modelling for (valid) process matrices in quantum theory [22, 23] are not yet well-understood, one may in particular wonder whether embeddable causal loops could arise from associated causal models. In particular, these generally assume a more general class of interventions beyond do-interventions [24–26].

Finally, we pose the question how the introduction of conicality may affect the gap in information processing power between theories with *no superluminal causation* and *no superluminal signalling* discussed in [2]. After all, we have demonstrated that already the assumption of no superluminal signalling is sufficient to rule out affects causal loops in physical space-times.

ACKNOWLEDGEMENTS

MG acknowledges financial support by l'Agence Nationale de la Recherche (ANR), project ANR-22-CE47-0012. VV acknowledges support from a government grant managed by the Agence Nationale de la Recherche under the Plan France 2030 with the reference ANR-22-PETQ-0007.

---

[1] V. Vilasini and R. Colbeck, General framework for cyclic and fine-tuned causal models and their compatibility with space-time, Phys. Rev. A **106**, 032204 (2022), arXiv:2109.12128 [quant-ph].

[2] V. Vilasini and R. Colbeck, Theories with no superluminal signaling have greater information-processing power than theories with no superluminal causation, Phys. Rev. Res. **8**, 013070 (2026), arXiv:2402.12446 [quant-ph].

[3] V. Vilasini and R. Colbeck, Impossibility of superluminal signaling in minkowski spacetime does not rule out causal loops, Phys. Rev. Lett. **129**, 110401 (2022), arXiv:2206.12887 [quant-ph].

[4] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search* (Springer New York, NY, 1993).

[5] J. Pearl, *Causality: Models, Reasoning and Inference*, 2nd ed. (Cambridge University Press, 2009).

[6] M. Grothus and V. Vilasini, Characterizing signalling: Connections between causal inference and space-time geometry (2024), arXiv:2403.00916 [gr-qc].

[7] M. Grothus, *Compatibility of cyclic causal structures with spacetime in general theories with free interventions*, Master's thesis, ETH Zurich (2022), arXiv:2211.03593 [quant-ph].

[8] M. Grothus, A. A. Abbott, A. Vanrietvelde, and C. Branciard, Routing quantum control of causal order (2025), arXiv:2507.08781 [quant-ph].

[9] J. Cotler, X. Han, X.-L. Qi, and Z. Yang, Quantum causal influence, Journal of High Energy Physics **2019**, 42 (2019), arXiv:1811.05485 [hep-th].

[10] M. Hillery, V. Bužek, and A. Berthiaume, Quantum secret sharing, Phys. Rev. A **59**, 1829 (1999), arXiv:quant-ph/9806063.

[11] D. Gottesman, Theory of quantum secret sharing, Phys. Rev. A **61**, 042311 (2000), arXiv:quant-ph/9910067.

[12] A. Karlsson, M. Koashi, and N. Imoto, Quantum entanglement for secret sharing and secret splitting, Phys. Rev. A **59**, 162 (1999).

[13] D. E. Gottesman, *Stabiliser Codes and Quantum Error Correction*, Ph.D. thesis, California Institute of Technology (1997), arXiv:quant-ph/9705052.

[14] R. Cleve, D. Gottesman, and H.-K. Lo, How to share a quantum secret, Phys. Rev. Lett. **83**, 648 (1999).

[15] A. Kent, A no-summoning theorem in relativistic quantum theory, Quantum Information Processing **12**, 1023–1032 (2012), arXiv:1101.4612 [quant-ph].

[16] A. Kent, Quantum tasks in minkowski space, Classical and Quantum Gravity **29**, 224013 (2012), arXiv:1204.4022 [quant-ph].

[17] P. Hayden and A. May, Summoning information in spacetime, or where and when can a qubit be?, Journal of Physics A: Mathematical and Theoretical **49**, 175304 (2016), arXiv:1210.0913 [quant-ph].

[18] E. Lichaire, *Causal relations in quantum information protocols: Identifying signalling and clustering with quantum systems*, Master's thesis, Inria Grenoble / Université de Montpellier (2025).

[19] V. Vilasini and R. Renner, Embedding cyclic information-theoretic structures in acyclic space-times: No-go results for indefinite causality, Phys. Rev. A **110**, 022227 (2024), arXiv:2203.11245 [quant-ph].

[20] P. Forré and J. M. Mooij, Markov properties for graphical models with cycles and latent variables (2017), arXiv:1710.08775 [math.ST].

[21] C. Ferradini, V. Gitton, and V. Vilasini, Cyclic functional causal models beyond unique solvability with a graph separation theorem (2025), arXiv:2502.04171 [math.ST].

[22] J. Barrett, R. Lorenz, and O. Oreshkov, Cyclic quantum causal models, Nature Commun. **12**, 885 (2021), arXiv:2002.12157 [quant-ph].

[23] C. Ferradini, V. Gitton, and V. Vilasini, Cyclic quantum causal modelling with a graph separation theorem (2025), arXiv:2502.04168 [quant-ph].

[24] J. Barrett, R. Lorenz, and O. Oreshkov, Quantum causal models (2019), arXiv:1906.10726 [quant-ph].

[25] R. Lorenz and S. Tull, Causal models in string diagrams (2023), arXiv:2304.07638 [cs.LO].

[26] C. Ferradini, *Connecting tensor networks to quantum causal models with applications to holography*, Master's thesis, ETH Zurich (2023).

[27] R. A. Hounnonkpe and E. Minguzzi, Globally hyperbolic spacetimes can be defined without the 'causal' condition, Classical and Quantum Gravity **36**, 197001 (2019), arXiv:1908.11701 [gr-qc].

[28] E. Witten, Light rays, singularities, and all that, Reviews of Modern Physics **92**, 045004 (2020), arXiv:1901.03928 [hep-th].

Appendix A: Proof of main results

Lemma II.6. Let $\mathcal{T}$ be a poset and $\mathcal{S}$ a set of ORVs embedded non-degenerately and conically therein. Let $\mathcal{Y}_i \subset \mathcal{S}$ and $\mathcal{X}_i \in \mathcal{S} \setminus \mathcal{Y}_i$ for all i, and $\mathcal{X} = \mathrm{span}(\mathcal{X}) = \prod_i \mathcal{X}_i, \mathcal{Y} = \prod_i \mathcal{Y}_i$. Then,

$$\bar{\mathcal{F}}_s(\mathcal{Y}_i) \subsetneq \bar{\mathcal{F}}(\mathcal{X}_i) \quad \forall i \qquad \Longrightarrow \qquad \bar{\mathcal{F}}_s(\mathcal{Y}) \subsetneq \bar{\mathcal{F}}_s(\mathcal{X})\,. \tag{II.3}$$

Proof. By conjunction, we arrive at

$$\bar{\mathcal{F}}_s(\mathcal{Y}_i) \subsetneq \bar{\mathcal{F}}(\mathcal{X}_i)\ \forall i \quad \Longrightarrow \quad \bar{\mathcal{F}}_s(\mathcal{Y}) \subsetneq \bar{\mathcal{F}}(\mathcal{X}_i)\ \forall i \quad \Longrightarrow \quad \bar{\mathcal{F}}_s(\mathcal{Y}) \subseteq \bar{\mathcal{F}}_s(\mathcal{X})\,. \tag{A.1}$$

Due to conicality of the embedding, $\bar{\mathcal{F}}_s(\mathcal{Y}) \overset{!}{=} \bar{\mathcal{F}}_s(\mathcal{X})$ would imply that the set of locations of $\mathrm{span}(\mathcal{Y})$ is equal to the set of locations of $\mathrm{span}(\mathcal{X})$. As the embedding is non-degenerate, this would imply $\mathrm{span}(\mathcal{X}) = \mathrm{span}(\mathcal{Y})$ and moreover, $\mathcal{X} = \mathrm{span}(\mathcal{Y})$. (Note that $\mathcal{X}_i$ and $\mathcal{Y}$ are not necessarily disjoint.)

However, due to $\bar{\mathcal{F}}_s(\mathcal{Y}_i) \subsetneq \bar{\mathcal{F}}(\mathcal{X}_i)$, non-degeneracy and conicality, we also have $\mathcal{X}_i \notin \mathrm{span}(\mathcal{X}_i\mathcal{Y}_i) = \mathrm{span}(\mathcal{Y}_i)$. This implies $\mathcal{X}_i \notin \mathrm{span}(\mathcal{Y})$, which poses a contradiction. $\lightning$ $\square$

Lemma A.1. For any causal model S with $X, Y, Z, W \subset S$ disjoint and $e_X \in X, e_Z \in Z$

$$e_X \not\vDash YW \mid \mathrm{do}(ZX \setminus e_Z e_X) \ \wedge \ e_X \vDash YW \mid \mathrm{do}(ZX \setminus e_X) \tag{A.2}$$

$$\Longrightarrow \quad e_Z \vDash YW \mid \mathrm{do}(ZX \setminus e_Z) \ \vee \ e_Z \vDash YW \mid \mathrm{do}(ZX \setminus e_Z e_X)\,. \tag{A.3}$$

Proof. By definition, the relations expand to

$$P_{\mathcal{G}_{\mathrm{do}(ZX \setminus e_Z e_X)}}(YW|ZX \setminus e_Z) = P_{\mathcal{G}_{\mathrm{do}(ZX \setminus e_Z)}}(YW|ZX \setminus e_Z e_X) \tag{A.4}$$

$$P_{\mathcal{G}_{\mathrm{do}(ZX \setminus e_X)}}(YW|ZX) \neq P_{\mathcal{G}_{\mathrm{do}(ZX)}}(YW|ZX \setminus e_X)\,. \tag{A.5}$$

For these equations to be compatible with one another, it is required that either $P_{\mathcal{G}_{\mathrm{do}(ZX \setminus e_Z e_X)}}(YW|ZX \setminus e_Z) \neq P_{\mathcal{G}_{\mathrm{do}(ZX \setminus e_X)}}(YW|ZX)$ or $P_{\mathcal{G}_{\mathrm{do}(ZX \setminus e_Z)}}(YW|ZX \setminus e_Z e_X) \neq P_{\mathcal{G}_{\mathrm{do}(ZX)}}(YW|ZX \setminus e_X)$, exactly corresponding to $e_Z \vDash YW \mid \mathrm{do}(ZX \setminus e_Z)$ or $e_Z \vDash YW \mid \mathrm{do}(ZX \setminus e_Z e_X)$. $\square$

This proof generalises to subsets $s_Z \subset Z$ and $s_X \subset X$. However, this generalisation will not be required in this work.

Lemma III.2. Consider a causal model featuring a higher-order affects relation $X \vDash Y \mid \mathrm{do}(Z), W$ satisfying Irred_1, which w.l.o.g. reduces to $X \vDash Y \mid \mathrm{do}(s_Z), W$ being Irred_3 for some $s_Z \subseteq Z$, potentially empty. If the causal model features no affects relations with Clus_3 or we have a conical embedding, $X s_Z \vDash YW$ with Irred_1 has the same implications for causal inference and compatibility (as the combination of $X \vDash Y \mid \mathrm{do}(Z), W$ Irred_1 and $X \vDash Y \mid \mathrm{do}(s_Z), W$ Irred_3):

$$\forall e_{XZ} \in X s_Z: \quad e_{XZ} \twoheadrightarrow YW \quad \wedge \quad \bar{\mathcal{F}}_s(\mathcal{X}) \cap \bar{\mathcal{F}}_s(s_{\mathcal{Z}}) \supseteq \bar{\mathcal{F}}_s(\mathcal{Y}) \cap \bar{\mathcal{F}}_s(\mathcal{W})\,. \tag{III.3}$$

Proof. By Lemma IV.8 of [1], $X \vDash Y \mid \mathrm{do}(Z), W$ implies $X \vDash YW \mid \mathrm{do}(Z)$ for each causal model, while preserving Irred$_1$ and Irred$_3$ Then, $X \vDash YW \mid \mathrm{do}(Z)$ can either be Red$_3$ or not. In any case there exists $X \vDash YW \mid \mathrm{do}(s_Z)$ for some $s_Z \subseteq Z$, satisfying either Irred$_3$ or having $s_Z = \emptyset$.

However, it is not necessarily clear that this relation is again irreducible in the first argument, or at least, if the relevant implications for causal inference and compatibility are maintained.

For causal models **without clustering** in the third argument of any affects relations, it is easy to see that $X \vDash YW$. The relations $e_X \vDash YW \mid \mathrm{do}(ZX \setminus e_X)$ implied for all $e_X \in X$ by irreducibility in the first argument similarly imply $e_X \vDash YW\ \forall e_X \in X$. By compatibility, this implies $\bar{\mathcal{F}}(e_\mathcal{X}) \supseteq \bar{\mathcal{F}}_s(\mathcal{YW})\ \forall e_X \in X$, which we can equivalently write as $\bar{\mathcal{F}}(\mathcal{X}) \supseteq \bar{\mathcal{F}}_s(\mathcal{YW})$. Hence, the statement follows with $s_Z = \emptyset$.

For **conical embeddings**, consider the following argument. By the original relation being Irred$_1$, we have $e_X \vDash YW \mid \mathrm{do}(ZX \setminus e_X)\ \forall e_X$. We differentiate two cases:

I) If $\forall e_Z \in Z\ \exists e_X \in X : e_X \not\vDash YW \mid \mathrm{do}(ZX \setminus e_Z e_X)$, we can use this fact together with Irred$_1$ of the original relation implying $e_X \vDash YW \mid \mathrm{do}(ZX \setminus e_X)$ to deduce additional affects relations. Following Lemma A.1, they jointly certify $e_Z \vDash YW \mid \mathrm{do}(ZX \setminus e_Z)\ \vee\ e_Z \vDash YW \mid \mathrm{do}(ZX \setminus e_Z e_X)\quad \forall e_Z \in Z$.[6] For causal inference, these alternative affects relations imply $e_Z \twoheadrightarrow YW\ \forall e_Z$. By compatibility, they imply $\bar{\mathcal{F}}(e_\mathcal{Z}) \supseteq \bar{\mathcal{F}}_s(\mathcal{Z} \setminus e_\mathcal{Z}) \cap \bar{\mathcal{F}}_s(\mathcal{X}) \cap \bar{\mathcal{F}}_s(\mathcal{YW})$, which together with $\bar{\mathcal{F}}(e_\mathcal{X}) \supseteq \bar{\mathcal{F}}_s(\mathcal{X} \setminus e_\mathcal{X}) \cap \bar{\mathcal{F}}_s(\mathcal{Z}) \cap \bar{\mathcal{F}}_s(\mathcal{YW})$ (from $e_X \vDash YW \mid \mathrm{do}(ZX \setminus e_X)$) and conicality[7] implies $\bar{\mathcal{F}}(e_{\mathcal{ZX}}) \supseteq \bar{\mathcal{F}}_s(\mathcal{YW})\ \forall e_{\mathcal{ZX}} \in \mathcal{ZX}$, recovering the same implications for causal inference and compatibility as $XZ \vDash YW$ being Irred$_1$. This amounts to the claim (for s_Z rather than Z in later iterations, see below.).

II) Otherwise, $\exists e_Z \in Z\ \ \forall e_X \in X : e_X \vDash YW \mid \mathrm{do}(ZX \setminus e_Z e_X)$[8]. By [6, Lemma F.2], for a conical embedding, this has the same implications for compatibility (and additionally, for causal inference) as $X \vDash YW \mid \mathrm{do}(Z \setminus e_Z)$ being Irred$_1$. If $Z \setminus e_Z = \emptyset$, we arrive at the claim for $s_Z = \emptyset$. If $Z \setminus e_Z \neq \emptyset$, we repeat the same argument for the relations $e_X \vDash YW \mid \mathrm{do}(ZX \setminus e_Z e_X)\ \forall e_X \in X$ in place of $e_X \vDash YW \mid \mathrm{do}(ZX \setminus e_X)\ \forall e_X \in X$, again considering for this set of relations whether case **I** or **II** applies with $Z \setminus e_Z$ in place of Z. Repeating recursively, we ultimately either end up in case **I** to recover the claim or, after $|Z|$ repetitions, reduce to $s_Z = \emptyset$ and recover it for that case. □

Theorem III.4. For any space-times, all compatible embeddings of ACLs rooted in a causal model without clustering are trivial.

Proof. Consider a set of affects relations which imply an ACL. Then by Corollary III.3, we can transform it into a set of unconditional 0^{th}-order affects relations, with the same implications for causal inference and compatibility, and continue with this transformed set.

Choose an ORV $e_\mathcal{X}$ that is part of a causal loop, as inferrable from the presence of a set of irreducible affects relations. Then, it must be part of at least one irreducible affects relation $X \vDash Y$ with $X \ni e_X$, implying $e_X \twoheadrightarrow Y$ with $Y \subset S \setminus e_X$. Due to Corollary II.3, in absence of clustering each such affects relation can be resolved such that $e_X \vDash e_Y$ for some $e_Y \in Y$, yielding the compatibility condition $e_\mathcal{X} \preceq e_\mathcal{Y}$. Iterating, if e_Y is part of the causal loop, it again admits some e_Z such that $e_Y \vDash e_Z$ and $e_\mathcal{Y} \preceq e_\mathcal{Z}$. Being restricted to unconditional 0^{th}-order affects relations atomic both in their first and second argument, the resulting ACL must be a loop of single variables $e_\mathcal{X} \preceq e_\mathcal{Y} \preceq e_\mathcal{Z} \preceq \ldots \preceq e_\mathcal{X}$ (i.e. an ACL type 2 by [1]), with all variables involved collapsing into a single point in space-time. Therefore, we obtain a trivial embedding. (See also Lemma VI.2 of [1].) □

Theorem III.5. In conical space-times, all compatible embeddings of ACLs are degenerate.

We prove a generalisation of this theorem from conical space-times to conical embeddings in arbitrary space-times, which are defined as follows:

Definition A.2 (Conical Embedding [6])**.** An embedding of a set $\mathcal{S}$ of ORVs in a space-time $\mathcal{T}$ is called *conical* if for any two finite subsets $L_i, L_j \subseteq O(\mathcal{S}) = \{O(X) | X \in S\} \subseteq \mathcal{T}$, $f(L_i) = f(L_j) \implies \mathrm{span}(L_i) = \mathrm{span}(L_j)$ holds.

Intuitively, with this definition, we consider the condition of conicality only for the futures of all points $O(\mathcal{S})$ in the image of the embedding and the intersections of their futures. Hence, any embedding into a conical poset $\mathcal{T}$ is trivially conical. However, the conicality of $\mathcal{T}$ is not necessary to obtain a conical embedding: Actually, the vast majority of embeddings into 1+1-dim. Minkowski space-time is conical as well.

---

[6] This property is analogous to imposing Irred$_3$ for $X \vDash YW \mid \mathrm{do}(Z)$ only with respect to $e_Z \in Z$ rather than all $s_Z \subseteq Z$, leading to a different notion for $|Z| > 2$.

[7] more precisely, applying an argument perfectly analogous to the proof of the main theorem of [6].

[8] For $|X| > 2$, this is not quite sufficient for irreducibility of $X \vDash YW \mid \mathrm{do}(Z \setminus e_Z)$, which is precisely the analogous statement obtained when replacing e_X with $\forall s_X \subset X$. In isolation, it implies $\bar{\mathcal{F}}(e_\mathcal{X}) \supseteq \bar{\mathcal{F}}_s(\mathcal{X} \setminus e_\mathcal{X}) \cap \bar{\mathcal{F}}_s(\mathcal{Z} \setminus e_\mathcal{Z}) \cap \bar{\mathcal{F}}_s(\mathcal{YW})\ \forall e_\mathcal{X} \in \mathcal{X}$.

Remark A.3. Specifically, such embeddings can be non-conical only if one ORV is embedded into the future light cone surface of another, constituting a kind of fine-tuning of the embedding. As a possible sufficient criterion for being located "on the light cone surface" we suggest, in purely order-theoretical terms: For $a, b \in \mathcal{T}$, b is embedded into the future light cone surface of a iff $M = \bar{J}^+(a) \cap \bar{J}^-(b) = \{x \in \mathcal{T} | a \preceq x \preceq b\}$ is totally ordered by $\prec$, i.e. $\nexists b, c \in M : \ b \npreceq \nsucceq c$. Here, M is a degenerate example of a causal diamond [27, 28]. For Minkowski space-time, this criterion is also necessary, while singularities or self-intersecting light cones could lead to violations.

Having this concept established, we return the proof of the theorem.

Theorem A.4. Any conical compatible embedding of ACLs into space-time is degenerate.

Proof. Consider a set of affects relations which imply an ACL. Then by Corollary III.3, we can transform it into a set of unconditional 0^{th}-order affects relations, with the same implications for causal inference and compatibility, and continue with this transformed set.

Choose an ORV $e_{\mathcal{X}}$ that is part of a causal loop, as inferrable from the presence of a set of irreducible affects relations. Then, it must be part of at least one irreducible affects relation $X \vDash Y$ implying $e_X \twoheadrightarrow Y$ with $Y \subset S \setminus e_X$. This yields $\bar{\mathcal{F}}_s(\text{span}(\mathcal{Y})) \subsetneq \bar{\mathcal{F}}(e_{\mathcal{X}})$ for any non-degenerate embedding due to conicality and $e_{\mathcal{X}} \notin \text{span}(\mathcal{Y})$.

Iterating, we consider $\text{span}(\mathcal{Y})$. If $\mathcal{Y}$ is to be part of a causal loop, for all $e^i_{\mathcal{Y}} \in \text{span}(\mathcal{Y})$, there is $Z_i \subset Z$ such that we can infer $e^i_Y \twoheadrightarrow Z_i$ due to an irreducible affects relation.[9] This yields

$$\bar{\mathcal{F}}_s(\text{span}(\mathcal{Z}_i)) \subsetneq \bar{\mathcal{F}}(e^i_{\mathcal{Y}}) \ \forall i \quad \Longrightarrow \quad \bar{\mathcal{F}}_s(\text{span}(\mathcal{Z})) \subsetneq \bar{\mathcal{F}}_s(\text{span}(\mathcal{Y})) \tag{A.6}$$

by Lemma II.6.

Iterating by the same argument, we can chain these strict subset relations with one another. However, the graph is finite and due to $\text{span}(\mathcal{A}) \neq \emptyset$ for any set of ORVs $\mathcal{A}$, the chain continues indefinitely – as otherwise, we would not have a loop. Hence, there exist sets $\mathcal{S}$ which appear multiple times in the chain, yielding $\bar{\mathcal{F}}_s(\text{span}(\mathcal{S})) \neq \bar{\mathcal{F}}_s(\text{span}(\mathcal{S}))$. This poses a contradiction, thereby proving the claim. ↯ □

Appendix B: A different notion of irreducibility?

Considering the proof of Lemma III.2, it is natural to wonder whether one should consider a weaker notion of irreducibility (and dually, a stronger notion of reducibility), given their more well-behaved interplay with one another. Specifically, for the first argument of an affects relation $X \vDash Y$, we would require only $e_X \vDash Y \mid \text{do}(ZX \setminus e_X)$ for all $e_X \in X$ with respect to the first argument, while for the third argument, we would require $e_Z \vDash Y \mid \text{do}(XZ \setminus e_Z) \ \vee \ e_Z \vDash Y \mid \text{do}(Z \setminus e_Z)$ for all $e_Z \in Z$.

Due to this definition, affects relations which are weakly irreducible have a cardinality of $|X| = 3$ or $|Z| = 2$ at least, to respectively be weakly irreducible in this sense without being irreducible in the usual sense. However, no example for a causal model which gives rise to an affects relation which is only weakly irreducible in some argument has been established so far.

Notably, this weaker notion is sufficient to reproduce the causal inference results we have reviewed above. Aided by the absence of clustering in the third argument or conicality of the embedding, it is also sufficient to reproduce all compatibility results [6, App. F]. In the non-conical case however, we can not rule out that replacing irreducibility in the definition of compatibility with this weaker notion may pose a stronger restriction on embeddings.

However, we do not choose to generally adopt this notion, as this weaker notion actually implies characteristics one would usually intuitively designate as reducible.

Lemma B.1. Consider an affects relation $X \vDash Y \mid \text{do}(Z)$, which is Red_1, yet satisfies $e_X \vDash Y \mid \text{do}(ZX \setminus e_X)$ for all $e_X \in X$. Then there exists $s_X \subsetneq X$ with $|s_X| \geq 2$ such that $X \setminus s_X \vDash Y \mid \text{do}(Z)$.

Proof. As $X \vDash Y \mid \text{do}(Z)$ is Red_1, there exists a non-empty subset $s_X \subsetneq X$ such that $s_X \nvDash Y | \text{do}(\tilde{s}_X Z)$, where $\tilde{s}_X := X \setminus s_X$, with $|s_X| > 2$. However, jointly with the original relation, this implies the (in-)equality chain $P_{\mathcal{G}_{\text{do}(Z)}}(Y|Z) \neq P_{\mathcal{G}_{\text{do}(XZ)}}(Y|XZ) = P_{\mathcal{G}_{\text{do}(\tilde{s}_X Z)}}(Y|\tilde{s}_X Z)$, and hence, $\tilde{s}_X \vDash Y \mid \text{do}(Z)$. □

[9] If there were $e^i_{\mathcal{Y}} \in \text{span}(\mathcal{Y})$ such that there is no such Z_i, $e_X \twoheadrightarrow e^i_Y$ would be one option to resolve $e_X \twoheadrightarrow Y$, with e^i_Y being a childless node in the causal structure. Therefore, the causal relation $e_X \twoheadrightarrow Y$ could not be part of a causal loop, would be in contradiction to our assumption.

This matches the fact shown in [6], that no affects relation which is not irreducible (in the usual sense) can be clustered.

However, one may also consider an analogously weakened property for clustering in the first or third argument – which however does not admit an as nice interpretation as the original version. Specifically, we would require only $X \setminus e_X \not\models Y \mid \mathrm{do}(Z)$ for all $e_X \in X$ for the first argument, rather than $s_X \not\models Y \mid \mathrm{do}(Z)$ for all $s_X \in Y$, and similarly $X \not\models Y \mid \mathrm{do}(Z \setminus e_Z)$ for all $e_Z \in Z$ for the third argument. Similarly, this weakened notion would be sufficient to reproduce our causal inference results for clustered affects relations. We leave a deeper study of these weaker notions for future work, in particular the interpretation and implications of this weaker notion of clustering for compatible embeddings and causal loops.