Impossibility of superluminal signalling in Minkowski space-time does not rule out causal loops

V. Vilasini^{1,2} and Roger Colbeck²

¹Institute for Theoretical Physics, ETH Zurich, 8093 Zürich, Switzerland ²Department of Mathematics, University of York, Heslington, York YO10 5DD. (Dated: 6th September, 2022)

Causality is fundamental to science, but it appears in several different forms. One is relativistic causality, which is tied to a space-time structure and forbids signalling outside the future. A second is an operational notion of causation that considers the flow of information between physical systems and interventions on them. In In [Vilasini and Colbeck, Phys. Rev. A. xx (2022)], we propose a framework for characterising when a causal model can coexist with relativistic principles such as no superluminal signalling, while allowing for cyclic and non-classical causal influences and the possibility of causation without signalling. In a theory without superluminal causation, both superluminal signalling and causal loops are not possible in Minkowski space-time. Here we demonstrate that if we only forbid superluminal signalling, superluminal causation remains possible and show the mathematical possibility of causal loops that can be embedded in a Minkowski space-time without leading to superluminal signalling. The existence of such loops in the given space-time could in principle be operationally verified using interventions. This establishes that the physical principle of no superluminal signalling is not by itself sufficient to rule out causal loops between Minkowski space-time events. Interestingly, the conditions required to rule out causal loops in a space-time depend on the dimension. Whether such loops are possible in three spatial dimensions remains an important open question.

Introduction.—Understanding cause-effect relations is central to the scientific method, yet there are several inequivalent notions of causality. Often, it is defined with respect to a background space-time structure, after which causal structure and space-time structure are treated synonymously. An alternative is to define causality operationally and independently of space-time. One way to do this is through causal models, which are based on intervening on physical systems and analysing the resulting correlations [1, 2]. This is the approach we take here. Causal models have been extensively applied to situations involving classical variables, being used for instance for medical testing [3, 4], economic predictions [1, 5], and machine learning [6–8].

Bell's theorem [9] demonstrates that classical causal models cannot explain quantum correlations within the causal structure that is naturally associated with a Bell experiment [10]. This has fuelled several approaches for providing causal explanations to quantum and more general non-classical correlations. One approach is to develop causal models [10–25] that go beyond classical random variables and allow quantum or even post-quantum systems [26, 27] to be causes. Other approaches suggest modifying the causal structure itself without necessarily considering non-classical causes, such as allowing for additional causal influences that go outside the future light cone (e.g., non-local hidden variable theories [28]) or against the direction of time (retro-causality [29]). Such causal influences must remain hidden in order to prevent superluminal signalling at the observed level. More radical approaches lean towards giving up the standard understanding of causation as being acyclic, and replacing it with a suitable notion of logical consistency [30-32]. While these alternatives correspond to different descriptions of the underlying causal model, they are all compatible with the impossibility of superluminal signalling in Minkowski space-time [33].

In our associated paper [34], we have developed a general framework for causation that can describe nonclassical and cyclic causal influences as well as causation without signalling. We do so by keeping the direction of causation and the order of events in space-time distinct. The former is modelled operationally and we characterize when the two are *compatible* with each other i.e., when we can assign space-time locations to the random variables in the causal model without leading to signalling outside the space-time's future (we call this assignment an embedding). Here we ask: does the ability to compatibly embed a causal model in an acyclic space-time (such as Minkowski space-time) imply that the operational predictions of the causal model can be reproduced within an acyclic causal structure? If so, it would not be necessary to consider cyclic causation.

Within relativistic physics, causal influences are taken to flow within the light cone, making both causal loops and superluminal signalling impossible in Minkowski space-time. Here we relax this assumption, and require only that observable signalling is limited by the light cone structure. In scenarios in which there is causation without correlation (i.e., a *fine-tuned* influence), answering the above question is more challenging since finetuned influences can act outside the lightcone without leading to signalling. Our framework allows treatment of both correlations and interventions on physical systems in general scenarios with cyclic causal influences. It can model fine-tuned causal influences as well as latent quantum and post-quantum causes, and be used to characterise conditions under which a causal model can be compatibly embedded in a space-time. Here, we apply this framework to demonstrate the mathematical possibility of causal loops between events in (1+1)-Minkowski space-time where these loops can be operationally detected without superluminal signalling, providing an explicit example. We also show that the observable predictions of this loop cannot be reproduced in any acyclic causal model, answering the aforementioned question in the negative. Interestingly, the same example fails to be embeddable in (3+1)-Minkowski space-time. In the associated paper [34], we have characterised a large class of operationally detectable causal loops within our framework. Whether any of these may be embeddable in (3+1)dimensions remains an interesting open question.

Fine-tuned causal explanations are often undesirable as fundamental explanations of physical phenomena [35], but can be crucial in practical information processing tasks. For instance, cryptographic protocols (such as the one time pad) rely on fine-tuned correlations, and the security of relativistic cryptographic protocols [36, 37] combines both relativistic notions of causality and information-theoretic concepts. Hence, understanding the extent to which compatibility with Minkowski spacetime restricts the possible operational causal models also has practical implications.

Causal models: correlations, interventions and finetuning—We begin by reviewing the essentials of the causal modelling framework (see [34] for details). A causal structure is modelled as a directed graph \mathcal{G} whose nodes correspond to observed or unobserved systems and directed edges \longrightarrow denote causation between these systems [38]. The set of observed nodes (denoted S_{obs}) comprises classical random variables (RVs) such as measurement settings or outcomes, while unobserved nodes may be classical, quantum or post-quantum systems (such as those from a generalised probabilistic theory).

Implicitly the nodes are associated with causal mechanisms that specify how information from their incoming edges is processed. For instance, in the classical case this processing corresponds to a function f_{N_i} from the parents of the node and possibly an additional, parentless [39] RV E_{N_i} (to allow for situations where N_i is not deterministically dependent on its parents) to the node variable N_i itself. "A is a direct cause of B" then corresponds to the function f_B having A as a (non-trivial) argument. [In the non-classical case, these functions would be replaced by valid maps between systems in the theory (e.g., CPTP maps in quantum theory)].

Often these causal mechanisms are not explicitly known and hence our treatment has to work at the level of observed probability distributions rather than the causal mechanisms. The causal structure imposes constraints on the possible distributions that may arise over the observed nodes. One set of such constraints can be expressed using the notion of *d-separation*. For two disjoint subsets X and Y of observed nodes S_{obs} of a causal structure \mathcal{G} , X and Y are said to be d-separated, denoted $(X \perp^d Y)_{\mathcal{G}}$, if there are no directed paths between variables in X and Y and if no variables in X and Y have common ancestors in \mathcal{G} . More generally, d-separation $(X \perp^d Y|Z)_{\mathcal{G}}$ is defined for three disjoint subsets X, Y and Z of the observed nodes (see Section A of the Supplemental Material for details).

We say that a distribution $P_{\mathcal{G}}(S_{obs})$ satisfies the dseparation property with respect to a causal structure \mathcal{G} if whenever we have d-separation between observed nodes in \mathcal{G} , then we have a corresponding conditional independence in the observed distribution $P_{\mathcal{G}}(S_{\text{obs}})$ i.e., $(X \perp^d Y|Z)_{\mathcal{G}}$ implies $P_{\mathcal{G}}(XY|Z) = P_{\mathcal{G}}(X|Z)P_{\mathcal{G}}(Y|Z).$ If the converse holds for all disjoint subsets of observed nodes, then the distribution is said to be *faithful* or, equivalently, not *fine-tuned*. In the present work, we adopt a minimal definition of a causal model which corresponds to a directed graph \mathcal{G} and an observed distribution $P_{\mathcal{G}}(S_{\text{obs}})$ that satisfies the d-separation property with respect to \mathcal{G} [40]. In this work we allow fine-tuned causal models in which conditional independences can occur in $P_{\mathcal{G}}(S_{obs})$ without the corresponding d-separation in \mathcal{G} (i.e., we allow causation without correlation).

In cases where we know the causal mechanisms, a causal model can be specified by a causal structure \mathcal{G}^C , causal mechanisms $\{f_{N_i}, P_{\mathcal{G}^C}(E_{N_i})\}_i$ and an associated observed distribution $P_{\mathcal{G}^C}(S_{obs})$ for $S_{obs} \subseteq \{N_i\}_i$. Here the observed distribution must be consistent with relationships specified by the causal mechanisms. Methods for modelling operational causal structures beyond the classical case are proposed in our associated paper [34]. In the present work, we restrict to the classical case which suffices to illustrate our main claims.

So far, we have only discussed the possible correlations that fit with a causal structure. Inferring causation requires more, and the concept of an intervention has been introduced to deal with this [1]. If intervening on Xchanges the distribution on Y, then we can deduce that X is a cause of Y.

In the case where X is parentless in a causal structure \mathcal{G} , correlation between X and another variable Y i.e., $P_{\mathcal{G}}(Y|X) \neq P_{\mathcal{G}}(Y)$ suffices to conclude that X is a cause of Y. More generally, an intervention on X corresponds to forcing X to take a certain value, x, irrespective of its parents. This results in a post-intervention causal structure $\mathcal{G}_{do(X)}$ obtained from \mathcal{G} by removing all the incoming edges to X, while maintaining the causal mechanisms for $\mathcal{G}_{do(X)}$ from \mathcal{G} , except the mechanism for the intervened node X, which is replaced by X = x. We say that X affects Y if there exist values x and y such that

$$P_{\mathcal{G}_{\mathrm{do}(X)}}(Y=y|X=x) \neq P_{\mathcal{G}}(Y=y). \tag{1}$$

In general $P_{\mathcal{G}_{d_0(X)}}(Y|X)$ is not the same as $P_{\mathcal{G}}(Y|X)$ (but they are equal when X is parentless in \mathcal{G}).

Acyclic and cyclic causal models embedded in Minkowski space-time—The affects relation defined in (1) naturally extends to joint interventions on sets of observed nodes. In the presence of fine-tuning, a set of RVs can jointly affect another set without any affects relations between individual pairs of elements in the sets. We will illustrate three such examples here and use these to establish our main claim. All three causal models have observed nodes $S_{obs} \coloneqq \{A, B, C\}$ which we will take to be classical and binary. The same observed correlations will be used in all three cases. The distinguishing feature at the observed level will be the affects relations which model what happens under intervention. The causal structures and compatible space-time embeddings for the three examples are given in Figure 1.

Example 1 (The one-time pad). Consider the causal structure of Figure 1a, \mathcal{G}^{OTP} with the causal mechanisms $A = E_A, C = E_C, B = A \oplus C$, with E_A and E_C being independent and uniformly distributed. Then, B is uniformly distributed and we have $P_{\mathcal{G}^{OTP}}(B|AC) \neq P_{\mathcal{G}^{OTP}}(B)$, $P_{\mathcal{G}^{OTP}}(B|A) = P_{\mathcal{G}^{OTP}}(B)$ and $P_{\mathcal{G}^{OTP}}(B|C) = P_{\mathcal{G}^{OTP}}(B)$. Since A and C are parentless, these statements are equivalent to $\{A, C\}$ affects B, A does not affect B, and C does not affect B [41]. A Minkowski space-time embedding of the RVs that does not enable any signalling outside the future is one in which B is assigned a space-time location in the joint future of the space-time locations of A and C, as illustrated in Figure 1a.

Example 2 (A simplified jamming scenario (cf. [42])). Consider the causal structure of Figure 1b, \mathcal{G}^{jam} with the observed nodes $\{A, B, C\}$ and a classical unobserved node Λ . Suppose we have the causal mechanisms $\Lambda = E_{\Lambda}$, $A = \Lambda, B = E_B, C = B \oplus \Lambda$ with E_Λ and E_B uniformly distributed. This gives the same observed correlations as the previous example, with $B = A \oplus C$ and with all observed variables uniformly distributed. Additionally, $P_{\mathcal{G}^{jam}}(AC|B) \neq P_{\mathcal{G}^{jam}}(AC), P_{\mathcal{G}^{jam}}(A|B) = P_{\mathcal{G}^{jam}}(A)$ and $P_{\mathcal{G}^{jam}}(C|B) = P_{\mathcal{G}^{jam}}(C)$. Since B is parentless, this is equivalent to B affects $\{A, C\}$, B does not affect A and B does not affect C. A compatible Minkowski space-time embedding in this case requires that the joint future of the space-time locations of A and C is contained entirely within the space-time future of B, as shown in Figure 1b. This is because B affects $\{A, C\}$ can only be verified when A and C are brought together, which is possible only in their joint future. Since we have no pairwise affects relations, there is no pairwise signalling between the RVs and no RV is required to be in the future of any other. Note that the causal influence $B \longrightarrow C$ is superluminal, even though there is no superluminal signalling.

Example 3 (A fine-tuned causal loop). Consider the causal structure of Figure 1c, $\mathcal{G}^{\text{loop}}$, with the same observed and unobserved nodes as the previous example, but with the causal mechanisms $\Lambda = E_{\Lambda}$, $A = \Lambda$, $C = B \oplus \Lambda$, $B = A \oplus C$ with E_{Λ} uniformly distributed. Note that these causal mechanisms do not admit a unique solution. Nevertheless, in Section B of the Supplemental Material we apply a method proposed in [34] to uniquely determine the observed distribution based on these mechanisms, and show that the same distribution as the previous two examples is obtained. The effect of interventions

in this case cannot be directly inferred from the observed correlations since none of the observed nodes are parentless. The post-intervention causal structure $\mathcal{G}_{do(AC)}^{loop}$ is identical to \mathcal{G}^{OTP} and the post-intervention causal structure $\mathcal{G}_{do(B)}^{loop}$ is identical to \mathcal{G}^{jam} . Applying (1), we find $\{A, C\}$ affects *B* and *B* affects $\{A, C\}$ and no pairwise affects relations between *A*, *B* and *C*. The space-time embedding must satisfy the compatibility conditions of both the previous examples, and Figure 1c illustrates an embedding with these properties i.e., this causal loop can be compatibly embedded in (1+1)-Minkowski space-time.

All three examples above lead to the same correlations $B = A \oplus C$ where B is correlated jointly with A and C but not individually, hence there is fine-tuning. Correlation between B and $\{A, C\}$ implies (by the d-separation property) that B and $\{A, C\}$ must be d-connected i.e., B is d-connected with A and/or C. However, the RVs A, B and C are pairwise uncorrelated so there must be a pair of variables that are d-connected and yet independent, which constitutes fine-tuning (these independences disappear for small changes in the distribution of one of the variables, e.g., if Λ is non-uniform in Examples 2 and 3).

The causal loop of Example 3 exhibits many curious features. It is an operationally detectable causal loop i.e., any causal model that gives rise to the affects relations of this example must necessarily be associated with a cyclic causal structure. This is proven in Section C of the Supplemental Material, but the intuition is relatively simple: Consider three parties in possession of the 3 observed RVs A, B and C and two types of experiment: E1) Alice and Charlie perform all possible interventions on A and C and Bob observes B without intervening; E2) Bob performs all possible interventions on B while Alice and Charlie observe A and C without intervening. After both experiments the parties can get together to verify whether $B = A \oplus C$ holds. Here, "all possible interventions" on a variable corresponds to running through all possible values of that variable, setting these independently of their parents - see (1), and collecting statistics for each choice. These statistics differ between the three causal models, as they have different sets of affects relations. These interventions do not enable the parties to signal outside the space-time future as the affects relations of all these causal models are compatible with the given space-time embedding. They nevertheless allow the parties to operationally verify the existence of a causal loop as we show in the Supplemental Material.

In Example 3, experiment E1 shows that A and C are causes of B while E2 shows that B is a cause of at least one of A and C. Given the space-time embedding from Figure 1c, these interventions would enable the agents to operationally detect retro-causation. By contrast, for the first model, E2 would correspond to a post-intervention scenario with no edges and therefore lead to no correlations between the RVs, while in the second model, E1 would also lead to no correlations. In other words, these two experiments enable the parties to operationally dis-



FIG. 1: Three examples of causal models and their compatible embeddings in (1+1)-Minkowski space-time. In each case, the operational causal structure associated with the model is given in black, circled variables are observed nodes, while uncircled ones are not and the black arrows denote causation. Space-time information is given in blue with time along the vertical and space along the horizontal axis. The solid lines represent light-like surfaces and the shaded region corresponds to the joint future of A and C in all cases.

tinguish between the causal models of Examples 1-3 in spite of them having the same observed correlations.

The mathematical possibility of an operationally detectable causal loop being embedded in Minkowski spacetime without signalling necessarily involves fine-tuning (see [34]) since in the absence of fine-tuning, signalling and causation are equivalent. Note that in Example 3 the locations of the random variables in Minkowski spacetime have to be carefully chosen to allow compatibility: B and C must be light-like separated, and arbitrarily small adjustments to the location of either B or C will remove compatibility. This requires B to be embedded exactly at the earliest location in the joint future of Aand C. In a sense this is another kind of fine-tuning, but at the level of the space-time embedding. Furthermore, such an embedding is not possible in (3+1)-Minkowski space-time because the intersection of two lightcones is not itself a cone and consequently there does not exist a frame-independent earliest location in the joint future of two given points, unlike in (1+1)-D [43].

Beyond these examples, a causal model may give rise to a complicated set of affects relations between various subsets of the observed nodes, making characterising compatibility with a space-time a more complicated task. For instance, if A affects B and $\{A, C\}$ affects B but C does not affect B, does B have to be embedded in the joint future of A and C, or only in the future of A? The novel causal modelling concept required to answer such questions in fine-tuned models is the notion of a higher-order affects relation, which we introduce in our framework [34]. Here we have illustrated one example, but there is a general class of physical theories involving cyclic causal structures that are compatible with no superluminal signalling in Minkowski space-time as discussed in our associated paper [34].

Discussion and outlook.—In relativistic physics, a causal influence from one space-time location to another implies

that the latter is in the future light cone of the former. This means that within relativistic physics, it is not possible to have causal loops or closed timelike curves in space-times whose light cone structures form a partial order. On the other hand, the approach adopted here, where causal structure and space-time structure are distinct notions, only requires observable signalling (and not causal influences) to stay within the space-time future for compatibility. Maintaining a clear separation between causation and space-time structure and characterising their interdependence is useful for considering formulations of physics without a fixed background space-time structure (e.g., in quantum gravity [44, 45]), as well as practical information processing tasks in space-time.

We note that a previous work [46] proposes a set of necessary and sufficient conditions for ruling out causal loops in Bell-type scenarios. Our framework [34] identifies implicit assumptions in this claim. In particular, one of the claims in [46] assumes that no superluminal signalling in Minkowski space-time rules out causal loops, which we have shown does not hold in general.

We have established that the principle of no superluminal signalling in Minkowski space-time alone is insufficient to rule out causal loops, through an explicit construction in (1+1) dimensions. We also found that the conditions for ruling out causal loops can depend on the space-time dimension. Previous works have established that logical consistency [47, 48] or familiar quantum properties such as linearity, no-cloning [49] etc. are not sufficient. Finding underlying principles that can do so remains an interesting open problem, a pertinent question being whether the principle of no superluminal signalling rules out causal loops in the case of (3+1)-Minkowski space-time.

Acknowledgements— VV acknowledges support from the Department of Mathematics, University of York and the ETH Postdoctoral Fellowship from ETH Zürich.

- [1] Pearl, J. Causality: Models, reasoning, and inference. Second edition, Cambridge University Press (2009). link.
- [2] Spirtes, P., Glymour, C. N. & Scheines, R. Causation, prediction, and search (The MIT Press, 2nd ed., 2001). link.
- [3] Kleinberg, S. & Hripcsak, G. A review of causal inference for biomedical informatics. *Journal of Biomedical Informatics* 44, 1102–1112 (2011). link.
- [4] Raita, Y., Camargo, C. A., Liang, L. & Hasegawa, K. Big data, data science, and causal inference: A primer for clinicians. *Frontiers in Medicine* 8, 998 (2021). link.
- [5] Spirtes, P. Graphical models, causal inference, and econometric models. *Journal of Economic Methodology* 12, 3–34 (2005). link.
- [6] Petersen, M. J., Maya L.; van der Laan. Causal models and learning from data. *Epidemiology* 25, 418–426 (2014). link.
- [7] Arti, S., Hidayah, I. & Kusumawardani, S. S. Research trend of causal machine learning method: A literature review. *IJID (International Journal on Informatics for Development)* 9, 111–118 (2020). link.
- [8] Liu, T., Ungar, L. & Kording, K. Quantifying causality in data science with quasi-experiments. *Nature Computational Science* 1, 24–32 (2021). link.
- [9] Bell, J. S. Speakable and unspeakable in quantum mechanics (Cambridge University Press, 1987).
- [10] Wood, C. J. & Spekkens, R. W. The lesson of causal discovery algorithms for quantum correlations: causal explanations of Bell-inequality violations require finetuning. *New Journal of Physics* 17, 33002 (2015). link.
- [11] Tucci, R. R. Quantum Bayesian nets. International Journal of Modern Physics B 09, 295–337 (1995). link.
- [12] Leifer, M. S. Quantum dynamics as an analog of conditional probability. *Physical Review A* 74 (2006). link.
- [13] Laskey, K. B. Quantum causal networks. arXiv:0710.1200 (2007). link.
- [14] Leifer, M. & Poulin, D. Quantum graphical models and belief propagation. Annals of Physics **323**, 1899–1946 (2008). link.
- [15] Leifer, M. S. & Spekkens, R. W. Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference. *Physical Review A* 88, 052130 (2013). link.
- [16] Henson, J., Lal, R. & Pusey, M. Theory-independent limits on correlations from generalized Bayesian networks. *New Journal of Physics* 16, 113043 (2014). link.
- [17] Pienaar, J. & Brukner, Č. A graph-separation theorem for quantum causal models. New Journal of Physics 17, 73020 (2015). link.
- [18] Ried, K. et al. A quantum advantage for inferring causal structure. Nature Physics 11, 414–420 (2015). link.
- [19] Costa, F. & Shrapnel, S. Quantum causal modelling. New Journal of Physics 18, 63032 (2016). link.
- [20] Fritz, T. Beyond Bell's theorem II: Scenarios with arbitrary causal structure. *Communications in Mathematical Physics* 341, 391–434 (2015). link.
- [21] Allen, J.-M. A., Barrett, J., Horsman, D. C., Lee, C. M. & Spekkens, R. W. Quantum common causes and quantum causal models. *Physical Review X* 7, 031021 (2017). link.
- [22] Weilenmann, M. & Colbeck, R. Analysing causal structures with entropy. Proceedings of the Royal Society A

473, 20170483 (2017). link.

- [23] Barrett, J., Lorenz, R. & Oreshkov, O. Quantum causal models. arXiv:1906.10726 (2020). link.
- [24] Pienaar, J. Quantum causal models via quantum Bayesianism. *Physical Review A* **101** (2020). link.
- [25] Schmid, D., Selby, J. H. & Spekkens, R. W. Unscrambling the omelette of causation and inference: The framework of causal-inferential theories. arXiv:2009.03297 (2021). link.
- [26] Barrett, J. Information processing in generalized probabilistic theories. *Phys. Rev. A* 75, 032304 (2007). link.
- [27] Weilenmann, M. & Colbeck, R. Analysing causal structures in generalised probabilistic theories. *Quantum* 4, 236 (2020). link.
- [28] Bohm, D. A suggested interpretation of the quantum theory in terms of "hidden" variables. I. *Physical Review* 85, 166–179 (1952). link.
- [29] Wharton, K. B. & Argaman, N. Colloquium: Bell's theorem and locally mediated reformulations of quantum mechanics. *Rev. Mod. Phys.* 92, 021002 (2020). link.
- [30] Baumeler, A. & Wolf, S. The space of logically consistent classical processes without causal order. New Journal of Physics 18, 013036 (2016). link.
- [31] Baumeler, A. & Wolf, S. Causality-complexityconsistency: Can space-time be based on logic and computation? *Tutorials, Schools, and Workshops in the Mathematical Sciences* 69–101 (2017). link.
- [32] Baumeler, Ä., Degorre, J. & Wolf, S. Bell correlations and the common future. *Quantum Foundations, Probability and Information* 255–268 (2018). link.
- [33] We reference Minkowski space-time here for concreteness, but our results also apply to any other partially ordered space-time.
- [34] Vilasini, V. & Colbeck, R. Possibility of causal loops without superluminal signalling – a general framework. arXiv:2109.12128 (2021). link.
- [35] Since one then has to explain how the fine-tuning comes about.
- [36] Kent, A. Secure classical bit commitment over finite channels. *Journal of Cryptology* 18, 313–335 (2005). link.
- [37] Colbeck, R. & Kent, A. Variable Bias Coin Tossing. *Physical Review A* 73, 032320 (2006). link.
- [38] Such causal structures are formally independent of spacetime, although, as we will later discuss, they may be embedded in a separate space-time structure.
- [39] The functional dependences are typically described in terms of an augmented causal structure that is formed from the original causal structure by including $E_{N_i} \longrightarrow N_i$ for each node N_i . Note that the E_{N_i} nodes have no other edges associated with them and are parentless in the augmented graph.
- [40] Here, we make no assumptions about the causal mechanisms or the physical theory describing them, but instead derive the consequences of imposing d-separation at the observed level. We note that one may consider other topdown methods for describing causal structures without imposing d-separation, but this is yet to be explored.
- [41] This is analogous to a one-time pad protocol where a message M (a random classical bit) is encrypted using a uniformly distributed key K to obtain a cipher text

 $C = M \oplus K$. Even though the message M is a cause of the cipher text C, we have P(C|M) = P(C) and C contains no information about M in the absence of the key. In causal modelling terms, C and M are independent even though they are d-connected; the model is fine-tuned.

- [42] Grunhaus, J., Popescu, S. & Rohrlich, D. Jamming nonlocal quantum correlations. *Phys. Rev. A* 53, 3781–3784 (1996). link.
- [43] Although we have focused here on Minkowski space-time, the framework of our associated paper [34] can be applied to consider compatible embeddings in any partially ordered space-time. Applying this framework, one can find other examples of embeddable causal loops in space-times (not necessarily Minkowski) where the property of having a frame-independent earliest location in the joint future of any two given points fails. We saw that this property holds in (1+1) but not in (2+1) or (3+1) Minkowski space-time.
- [44] Oreshkov, O., Costa, F. & Brukner, Č. Quantum correlations with no causal order. *Nature Communications* 3, 1092 (2012). link.
- [45] Zych, M., Costa, F., Pikovski, I. & Brukner, Č. Bell's theorem for temporal order. *Nature Communications* 10, 3772 (2019). link.
- [46] Horodecki, P. & Ramanathan, R. The relativistic causality versus no-signaling paradigm for multi-party correlations. *Nature Communications* 10, 1701 (2019). link.
- [47] Deutsch, D. Quantum mechanics near closed timelike lines. Phys. Rev. D 44, 3197–3217 (1991). link.
- [48] Lloyd, S. et al. Closed timelike curves via postselection: Theory and experimental test of consistency. *Phys. Rev. Lett.* **106**, 040403 (2011). link.
- [49] Araújo, M. et al. Witnessing causal nonseparability. New Journal of Physics 17, 102001 (2015). link.
- [50] Forré, P. & Mooij, J. M. Markov Properties for Graphical Models with Cycles and Latent Variables. arXiv:1710.08775. link.
- [51] Neal, R. M. On Deducing Conditional Independence from d-Separation in Causal Graphs with Feedback (Research Note). *Journal of Artificial Intelligence Research* 12, 87-91 (2000). link.
- [52] Barrett, J. & Lorenz, R. & Oreshkov, O. Cyclic Quantum Causal Models. *Nature Communications* **12**, 885 (2021). link.
- [53] Bongers, S. et al. Foundations of structural causal models with cycles and latent variables. The Annals of Statistics 49, 5 (2021). link.

I. FURTHER DETAILS OF THE CAUSAL MODELLING FRAMEWORK

The d-separation property [1, 2] was motivated in the main text and can be used as a convenient tool to read off conditional independences in the observed distribution $P_{\mathcal{G}}(S_{\text{obs}})$ associated with a causal structure \mathcal{G} , from the topology of \mathcal{G} . We provide the formal definition below.

Definition I.1 (Blocked paths). Let \mathcal{G} be a directed graph in which X and $Y \neq X$ are nodes and let Z be a set of nodes not containing X or Y. A path from X to Y (not necessarily directed) is said to be *blocked* by Z if it contains either $A \longrightarrow W \longrightarrow B$ with $W \in Z$, $A \longleftarrow W \longrightarrow B$ with $W \in Z$ or $A \longrightarrow W \longleftarrow B$ such that neither W nor any descendant of W belongs to Z, where A and B are arbitrary nodes in the said path between X and Y.

Definition I.2 (d-separation). Let \mathcal{G} be a directed graph in which X, Y and Z are disjoint sets of nodes. X and Y are *d*-separated by Z in \mathcal{G} , denoted as $(X \perp^d Y | Z)_{\mathcal{G}}$ if every path from a variable in X to a variable in Y is blocked by Z, otherwise, X is said to be *d*-connected with Y given Z.

For acyclic causal structures (both classical [1] and non-classical [16, 23]), it is known that any distribution arising from the causal mechanisms satisfies the dseparation property. This is also known to hold in several classical cyclic causal models [50] and our accompanying paper [34] provides an example of a quantum cyclic causal model where this holds as well. However, in contrast to the acyclic case, in some classical cyclic causal structures these mechanisms do not imply the d-separation property fails [51]. Such examples are ruled out in the definition of causal model we use which imposes the dseparation property. In the classical case, [50] introduces a more general graph separation criterion (σ -separation) that applies to a larger class of cyclic causal models and reduces to d-separation in the acyclic case, and one can in principle generalise our definition of a causal model in terms of σ -separation to account for this larger class of models. There nevertheless remain classical cyclic causal models for which no graph separation property is known, highlighting the counter-intuitive nature of cyclic causal structures. Characterising the set of classical or quantum cyclic models, defined from causal mechanisms, for which d- or σ -separation holds is an interesting question for future work, and Appendix C of our accompanying paper [34] outlines a possible method for doing so.

Another useful framework for modelling cyclic quantum causal structures was proposed in [52]. This follows a different, split-node approach to causal modelling where all nodes are modelled quantum mechanically. This is a bottom-up approach that describes the observed probabilities in terms of the underlying quantum causal mechanisms. Unfaithful causal models, post-quantum causes already suffices to solve the problem regarding space-time embeddings of unfaithful, cyclic and non-classical causal models and reproduces a number of results from the classical causal modelling literature such as Pearl's rules of do-calculus. Even though the causal models in our framework are defined through the d-separation property, the techniques and space-time compatibility condition we propose relies only on the set of affects relations generated by a causal model and can hence also be applied to cyclic causal models that do not satisfy the d-separation condition.

II. A METHOD TO DERIVE THE DISTRIBUTION OF EXAMPLE 3

In acyclic (classical and non-classical) causal models, the observed distribution is uniquely determined once all the causal mechanisms are specified [1, 16, 23]. However this is not the case in cyclic models. The simplest example illustrating this is the 2-node cycle with $X \longrightarrow Y$ and $Y \longrightarrow X$ where the causal mechanisms are simply $X = f_X(Y) = Y$ and $Y = f_Y(X) = X$. In this case any distribution P(XY) = P(X) = P(Y) agrees with the causal mechanisms, and using the standard methods available in the previous literature on cyclic causal models [50, 53], this distribution is not uniquely determined by the mechanisms. In the associated paper [34](Appendix C), we have proposed a possible method for uniquely determining the observed distribution in a class of classical and non-classical cyclic models. The method can be applied to models where each directed cycle includes at least one observed node, and involves splitting this observed node in each cycle into two nodes, which results in a directed acyclic graph. We can then use the known methods for acyclic models (which always vield a unique distribution), post-select on the values of the split nodes being equal, and finally renormalise the overall distribution to ensure that the final distribution for the cyclic model is a valid normalised distribution.

We now apply this method to the cyclic causal model of Example 3 from the main text, for which standard methods leave the observed distribution undetermined. The directed cycle in this case is between the nodes Band C and we can choose to split either one of these. Splitting B into two nodes B and B' such that B has all the same incoming arrows as B (and no outgoing arrows) while B' has all the same outgoing arrows as B (and no incoming arrows), we obtain the following acyclic causal model: $P(\Lambda)$ is uniformly distributed, $A = \Lambda$, $C = B' \oplus \Lambda$, $B = A \oplus C$. The corresponding causal structure \mathcal{G}^{acyc} is given in Figure 2a. Notice that the above causal mech-



Λ	A	В	B'	C	$P_{\mathcal{G}^{\mathrm{acyc}}}(ABC B')$	$P_{\mathcal{G}^{\mathrm{loop}}}(ABC)$	
0	0	0	0	0	$\frac{1}{2}$	$\frac{1}{4}$	
0	0	1	1	1	$\frac{1}{2}$	$\frac{1}{4}$	
1	1	0	0	1	$\frac{1}{2}$	$\frac{1}{4}$	
1	1	1	1	0	$\frac{1}{2}$	$\frac{1}{4}$	
(b)							

FIG. 2: Calculating the observed distribution of Example 3 (a) Acyclic causal structure \mathcal{G}^{acyc} obtained by splitting the node *B* of the cyclic causal structure \mathcal{G}^{loop} of Figure 1(b). (b) Table showing the conditional probability distribution $P_{\mathcal{G}^{acyc}}(ABC|B')$ of the induced causal model over (a). We then post-select on B = B' to recover the original cyclic model and then ignore B'. But this yields an unnormalised distribution, and we must therefore renormalise the distribution to obtain the required observed distribution $P_{\mathcal{G}^{loop}}(ABC)$ as given in the last column.

anisms for this acyclic graph which are induced by the original cyclic model of Example 3 are not a complete set of mechanisms for the acyclic case, since we are missing a specification of the distribution over the new exogenous node B'. Therefore, we can only calculate the conditional distribution $P_{\mathcal{G}^{acyc}}(ABC|B')$ associated with the observed nodes of the acyclic model. We can then post-select on B = B' and renormalise the distribution to obtain the observed distribution $P_{\mathcal{G}^{\text{loop}}}(ABC)$ of the original cyclic model. This is explicitly shown in Figure 2b, and one can see that this procedure yields exactly the same distribution stated in the main text, and in particular this is also the observed distribution obtained in Examples 1 and 2 of the main text. It is easy to verify that splitting the node C instead and applying this procedure also yields the same distribution. Generalising this method to arbitrary non-classical cyclic causal structures is a subject of future work.

III. OPERATIONALLY CERTIFYING THE CYCLICITY OF THE CAUSAL MODEL

We have seen that the cyclic causal model of Example 3 leads to the affects relations B affects $\{A, C\}$, and $\{A, C\}$ affects B, with no pairwise affects relations between A, B and C. We now show that it is impossible for any acyclic causal model to give rise to the same affects relations i.e., these affects relations operationally certify the cyclicity of the causal model and hence the causal loop giving rise to these relations is operationally detectable.

We show this by contradiction, suppose that there exists an acyclic causal model with an associated causal structure \mathcal{G} , which is a directed acyclic graph (DAG) that produces these affects relations. First consider the affects relation B affects $\{A, C\}$. We will show that this implies the existence of a directed path from B to A or from B to C in \mathcal{G} . Since every acyclic causal model satisfies the d-separation property [1, 16], we can apply it to \mathcal{G} . If there were no directed paths from B to $\{A, C\}$ in \mathcal{G} , then B would be d-separated from $\{A, C\}$ in $\mathcal{G}_{do(B)}$ i.e., $(B \perp^d AC)_{\mathcal{G}_{\mathrm{do}(B)}}$. It is shown in [34] that for any \mathcal{G} satisfying the d-separation property, $(X \perp^d Y)_{\mathcal{G}_{\mathrm{do}(X)}}$ for any two disjoint sets X and Y of observed nodes implies that X does not affect Y. Applying this result here gives Bdoes not affect $\{A, C\}$ which contradicts the given affects relations, so we must have a directed path from Bto $\{A, C\}$ in \mathcal{G} .

Next consider a causal model obeying the affects relation $\{A, C\}$ affects B along with A does not affect Band C does not affect B. We will now show that this implies the existence of a directed path from A to B and from C to B in \mathcal{G} . For this we will use the concept of a *higher-order affects relation* introduced in [34]. For three disjoint subsets X, Y and Z of the observed nodes, we say that X affects Y given do(Z) if there exist values xof X, y of Y and z of Z such that

$$P_{\mathcal{G}_{\mathrm{do}(XZ)}}(Y=y|X=x,Z=z) \neq P_{\mathcal{G}_{\mathrm{do}(Z)}}(Y=y|Z=z) \quad (2)$$

Note that $\{A, C\}$ affects B, and C does not affect B imply the higher-order affects relation A affects B given do(C). Explicitly, the first two conditions are

$$P_{\mathcal{G}_{\mathrm{do}(AC)}}(B|AC) \neq P_{\mathcal{G}}(B)$$

$$P_{\mathcal{G}_{\mathrm{do}(C)}}(B|C) = P_{\mathcal{G}}(B).$$
(3)

When combined, this yields $P_{\mathcal{G}_{do(AC)}}(B|AC) \neq P_{\mathcal{G}_{do(C)}}(B|C)$ which precisely expresses the higherorder affects relation A affects B given do(C). Similarly $\{A, C\}$ affects B along with A does not affect B imply the higher-order affects relation C affects B given do(A). It was shown in [34] that X affects Y given do(Z) implies that there is a directed path from Xto Y in any underlying causal structure \mathcal{G} producing the higher-order affects relations (and satisfying the d-separation property). Applying this to the two higher-order affects relations obtained above establishes the existence of directed paths from A to B and from C to B in \mathcal{G} (which we have assumed to be a DAG, and this assumption implies that the associated model must satisfy the d-separation property).

To summarise, we have established above that there

exists a directed path from A to B and a directed path from C to B in \mathcal{G} as well as a directed path from B to at least one of A and C. This implies the existence of a directed cycle in \mathcal{G} , contradicting the assumption that it is a directed acyclic graph, and certifying the cyclicity of the model.