

Causal structure in the presence of sectorial constraints, with application to the quantum switch

Nick Ormrod¹, Augustin Vanrietvelde^{1,2,3}, and Jonathan Barrett¹

¹Quantum Group, Department of Computer Science, University of Oxford

²Department of Physics, Imperial College London

³HKU-Oxford Joint Laboratory for Quantum Information and Computation

Existing work on quantum causal structure assumes that one can perform arbitrary operations on the systems of interest. But this condition is often not met. Here, we extend the framework for quantum causal modelling to situations where a system can suffer *sectorial constraints*, that is, restrictions on the orthogonal subspaces of its Hilbert space that may be mapped to one another. Our framework (a) proves that a number of different intuitions about causal relations turn out to be equivalent; (b) shows that quantum causal structures in the presence of sectorial constraints can be represented with a directed graph; and (c) defines a fine-graining of the causal structure in which the individual sectors of a system bear causal relations. As an example, we apply our framework to purported photonic implementations of the quantum switch to show that while their coarse-grained causal structure is cyclic, their fine-grained causal structure is acyclic. We therefore conclude that these experiments realize indefinite causal order only in a weak sense. Notably, this is the first argument to this effect that is not rooted in the assumption that the causal relata must be localized in spacetime.

1 Introduction

Causal structure in quantum theory has recently been the subject of intense study, both theoretically [1, 2, 3, 4, 3, 5, 6, 7, 8, 9, 10, 11, 12, 13] and in the laboratory [14, 15, 16, 17, 18, 19] (see [20]

Nick Ormrod: nicholas.ormrod@cs.ox.ac.uk

Augustin Vanrietvelde: augustin.vanrietvelde@inria.fr

Jonathan Barrett: jonathan.barrett@cs.ox.ac.uk

for a review of the experiments). An impetus for this was the discovery that a simple extension of quantum theory allows for *indefinite causal order* [1, 21, 22, 23, 2, 24, 25, 3, 26, 27, 28, 29], in which the causal structure of a scenario inherits some of the quantum indeterminacy associated with superpositions of states. Such structures should be important in any theory of quantum gravity in which one can have a superposition of different classical solutions of Einstein’s field equations [21].

However, even when the causal order is definite, quantum theory poses significant challenges to classical notions of cause and effect. For example, no common cause can explain the Bell correlations in accordance with Reichenbach’s principle¹ unless superluminal, retrocausal, or superdeterministic causal influences are invoked. Even allowing such exotic influences, any classical causal explanation of the Bell correlations would have to be fine-tuned [31].

One is left with two options: abandon the notions of cause and effect altogether, or update them for a quantum setting. The latter approach was taken in [32, 5, 8] by generalizing classical causal models [33] to create an intrinsically quantum theory of causal structure.² The resulting framework of ‘quantum causal models’ provides answers to the basic foundational questions about causality – what are the relata of causal relations? what does it mean for one thing to be a direct cause of another? are quantum causal relations

¹This states that conditioning on the common cause C should eliminate the correlations between the correlated variables, A and B [30]. Formally, $P(AB|C) = P(A|C)P(B|C)$.

²See [34, 35, 36, 37] for alternative approaches to quantum causal modelling, or [38] for a theory-independent approach to generalizing classical causal models.

time-symmetric? Beyond this, it lays out the implications of a postulated causal structure for the correlations we can actually observe. This approach has also borne fruit for the question of indefinite causal order by rigorously analysing [8] the causal structures of the quantum switch [2] and the Lugano process [26].

However, the existing framework makes a highly nontrivial assumption about the scenarios it serves to model. Namely, that any quantum operation on the Hilbert space of one of the systems (and possibly some local ancillas) is a possible intervention. Mathematically, this assumption is manifested in the use of a process matrix to represent the physical scenario, which can be combined with any set of local quantum operations for each system to yield a valid probability distribution.

The assumption is important since the framework defines causal relations in terms of the possibility of signalling given the underlying unitary process between agents that each have access to one of the relevant systems. But if the interventions that the agents can perform are restricted, then this might limit their ability to signal to one another. Thus restrictions on the possible interventions should change the causal structure. It follows that the existing framework for quantum causal models is inadequate for modelling scenarios in which there are restrictions on the possible interventions on a system.

Some particularly interesting restrictions are *sectorial constraints*. These are restrictions on which sectors (i.e., orthogonal subspaces) of the input space can be mapped to which sectors of the output space by some transformation. They apply to the common quantum optical technique of sending a photon into a superposition of trajectories (in which they arise from the practical impossibility of creating a photon from the vacuum) [39]; to various superselected systems; to the Aharonov-Bohm effect [40, 41]; the del Santo-Dakić protocol [42, 43, 44, 45]; recent alleged implementations [16, 14, 15, 17] of the quantum switch [23, 2]; and many more quantum protocols. More abstractly, sectorial constraints are at the heart of the various powerful theorems and techniques in quantum theory that make use of non-factor C^* algebras, whose structure is specified by the sectorial constraints suffered by their elements. Examples include the Schur-Weyl duality

and its generalisations [46, 47], or the existence of decoherence-free subspaces [48, 49, 50, 51, 52, 53].

What this shows is that the causal structure of many foundationally interesting and practically significant scenarios cannot be adequately analysed by the existing framework for quantum causal models, which fails to accommodate the sectorial constraints that help define them. The goal of this paper is to present an appropriate framework for assessing the causal structure in these situations.

Our framework is self-contained, scalable, and applicable to a wide range of quantum-theoretical scenarios. It shows how various intuitions of causal influence turn out to be equivalent in the unitary case. It does not only extend standard quantum causal modelling to situations featuring sectorial constraints; it also shows that in such situations, a more fine-grained account of the causal structure can be given by leveraging on the sectorial structure of the operations. The resulting *sectorized causal structure* possesses a natural conceptual interpretation, and is strictly more detailed than the standard, unsectorized causal structure.

The sectorized causal structure can shed light on enigmatic features of the more standard sort of causal structure. To illustrate this, we end this paper with a detailed analysis of recent purported photonic implementations [16, 14, 15, 17] of the quantum switch [23, 2]. Weighing in on a recent debate over whether these experiments realize indefinite causal order [54, 14, 6], we argue that they only do so at the coarse-grained, unsectorized level, but not at the more detailed, sectorized level. We conclude that they are weak realizations of indefinite causal order. Unlike previous arguments to a similar effect [6], ours is not at all rooted in the assumption the objects that bear causal relations should be localized in spacetime. In fact, our basic assumptions resemble those that are usually taken to motivate the view that the experiments do realize indefinite causal order.

The structure of the paper is as follows. First, we recap in Section 2 the framework of *routed maps and supermaps*, which was recently developed [55, 56, 57] as a way of hardcoding sectorial constraints into models of quantum processes, providing a basis for our analysis. In Section 3, we provide a review of the canonical notion of a causal relation between the inputs and outputs

of a standard unitary transformation, then show how it can be naturally generalised to the case of *routed unitary transformations*, satisfying sectorial constraints. We then define fine-grainings of causal relations through routed unitaries, and thus define the sectorized causal structure.

Building on this, in Section 4 we first summarise how the existing quantum causal models framework treats the causal structure of a more complicated sort of scenario, represented by a higher-order transformation known as a supermap³, then generalise this analysis to *routed supermaps*, in which operations satisfy sectorial constraints. We show that in this case again, a sectorized causal structure can be defined. In Section 5, we apply the resulting framework to the alleged photonic implementations of the switch. Finally, in Section 6, we fend off possible objections to our analysis of those experiments.

2 Sectorial constraints and routed quantum circuits

In this section, we review the basics for modelling operations that feature sectorial constraints.

2.1 Why routed circuits?

The routed circuits framework [55] was devised in response to situations [10, 58, 39, 59] where the relevant Hilbert space of a system is a proper subspace of the tensor product of the spaces of each individual subsystem. For example, in the superposition of paths scenario, a particle described by a d -dimensional Hilbert space P is sent down Alice's or Bob's transmission line depending on the logical value of a qubit C . Alice's Hilbert space $A = A_{\star}^{vac} \oplus A_{\star}^{par}$ is a direct sum of a one-dimensional vacuum sector A_{\star}^{vac} and a d -dimensional particle sector A_{\star}^{par} . Likewise, Bob's space can be written $B = B_{\star}^{vac} \oplus B_{\star}^{par}$. The transformation can be represented as follows, where $|\Omega\rangle$ is the vacuum state:

$$\begin{aligned} |0\rangle_C |\psi\rangle_P &\rightarrow |\psi\rangle_A |\Omega\rangle_B \\ |1\rangle_C |\psi\rangle_P &\rightarrow |\Omega\rangle_A |\psi\rangle_B \end{aligned} \quad (1)$$

³Note that supermaps can equivalently be seen as *process matrices*, used elsewhere in the literature on indefinite causal order. More precisely, process matrices are essentially a Choi-Jamiołkowski representation of supermaps.

It is easily checked that (1) defines a unitary transformation $U : C \otimes P \rightarrow (A_{\star}^{par} \otimes B_{\star}^{vac}) \oplus (A_{\star}^{vac} \otimes B_{\star}^{par})$. However, one cannot provide a standard unitary quantum circuit representing this transformation in which A and B are each associated with one output wire, unless one extends U to be defined on the physically irrelevant sectors $A_{\star}^{vac} \otimes B_{\star}^{vac}$ and $A_{\star}^{par} \otimes B_{\star}^{par}$, as well as a larger input space. One is left with a dilemma: either accept a non-unitary diagrammatic representation of a scenario which should be physically understood as unitary, or allow physically redundant degrees of freedom in the representation.

The dilemma is a direct consequence of a very fundamental fact about standard quantum circuit diagrams – that putting two wires next to each other means taking the tensor product of the associated Hilbert spaces. When there are constraints that prevent certain subspaces of this tensor product space from being populated, this means including some redundancy in our representation. More generally, the problem with standard circuits is that they lack the resources to describe constraints on how transformations are allowed to act on certain relevant subspaces: they cannot explicitly represent *sectorial constraints*. The very same deficiency leads to problems in completely different areas; for example, it makes it impossible for standard circuits to provide *causal decompositions* of certain unitary transformations [10].

The framework of routed quantum circuits solves this problem by supplementing standard quantum circuits with classical (more precisely, possibilistic) information representing sectorial constraints. We now describe that framework.

2.2 Routed maps

We start off by adding structure to the Hilbert spaces. A *sectorized Hilbert space*⁴ is a Hilbert space with a preferred decomposition into a set of orthogonal sectors (i.e. orthogonal subspaces), which we call a *sectorization*. Not too formally, given a Hilbert space A and a complete orthogonal family of projectors π_A^i , we can construct a sectorized version of the space by appealing to a

⁴In [55], this was referred to as a ‘partitioned Hilbert space’. Here we are aiming for more precise terminology.

preferred sectorization:

$$A^i = \bigoplus_{i \in \mathcal{Z}_A} \pi_A^i(A) \quad (2)$$

Note that A^i represents the whole sectorized space; its superscript does not stand for some specific sector, but simply serves to remind us that the space is sectorized. We will represent specific sectors in the sectorization of the space using a star subscript $A_\star^i := \pi_A^i(A)$.

More formally, a sectorized space can be defined as a tuple including the original space A , the projectors π_A^i , and the possible values \mathcal{Z}_A of i .

The tensor product of two sectorized Hilbert spaces A^i and B^j is rather natural, corresponding to the decomposition:

$$A^i \otimes B^j = \bigoplus_{i \in \mathcal{Z}_A, j \in \mathcal{Z}_B} A_\star^i \otimes B_\star^j \quad (3)$$

That takes care of the Hilbert spaces; now we must address the transformations. We supplement the transformations with *routes*, which encode sectorial constraints. A route is a *relation* λ , a mathematical object which can be thought of as a possibly multi- or empty-valued function; it can be represented as a Boolean matrix λ_i^j . Its Boolean element λ_i^j determines whether a transformation is permitted to map the i th input sector to the j th output sector.

To state this more formally, consider a linear map $f : A \rightarrow B$ between Hilbert spaces which are sectorized to form A^i and B^j as described above. We say that $f : A \rightarrow B$ *follows* a route λ just in case

$$f = \sum_{ij} \lambda_i^j \cdot \pi_B^j \circ f \circ \pi_A^i \quad (4)$$

As proven in [55], this is equivalent to the statement that for all i and j , if $\lambda_i^j = 0$ then $\pi_B^j \circ f \circ \pi_A^i = 0$. Hence if $\lambda_i^j = 0$ then f cannot map any state with null support outside the i th sector of A , $\psi_A \in A_\star^i$, to a state with support in the j th sector of B , B_\star^j . Another equivalent formulation is that, given a matrix representation of f , the corresponding block in the block-decomposition provided by the projectors is zero whenever $\lambda_i^j = 0$. We thus see how routes (and specifically, their zeroes) enforce sectorial constraints on linear maps.

Routes can be sequentially composed using matrix multiplication and parallel-composed using the cartesian product. Their compositions play

well with the compositions of linear maps, as well as their adjoints: if f follows λ and g follows ϵ , then $g \circ f$ and $f \otimes g$ follow $\epsilon \circ \lambda$ and $\epsilon \times \lambda$ respectively, and f^\dagger follows λ^\top .

We can now define a *routed linear map* from A^i to B^j as a pair (λ, f) for a route λ and a linear map f that follows it. Routed linear maps (hereon just ‘routed maps’) can be sequentially and parallelly composed by composing their elements pairwise, and one can take their Hermitian adjoint by adjoining their elements.

It will also be useful to define the notion of practical input or output spaces. The idea is that there can be elements i of λ ’s input space that are not connected by λ to anything in the output space, i.e. satisfying $\lambda_i^j = 0 \forall j$. This means that a linear map f that follows λ will necessarily be zero on the corresponding subspaces A_\star^i . As these A_\star^i ’s are in practice never used by the routed map (λ, f) , we define the *practical input space* of (λ, f) as corresponding to the rest of the subspaces.

Formally, given a routed map $(\lambda, f) : A^i \rightarrow B^j$, f ’s practical input space is given by

$$A^{\text{prac}} := \bigoplus_i \alpha_i \cdot A_\star^i \quad (5)$$

where α_i is given by

$$\alpha_i := \sum_j \lambda_i^j. \quad (6)$$

In (6), the sum is Boolean, i.e. given by $0+0=0$, $0+1=1$, $1+0=1$, $1+1=1$. A^{prac} is the orthogonal subspace of A in which f is not constrained by λ to have null support. One defines the practical output space in an exactly analogous way, this time by taking the Boolean sum $\sum_i \lambda_i^j$.

2.3 Routed circuits

It is proven in [55] that sectorized Hilbert spaces and routed maps form a dagger symmetric monoidal category, which in turn implies that routed maps have a sound representation in terms of circuit diagrams.

For example, we can give a routed circuit diagram for the superposition-of-paths transformation discussed above. To this end, we will say that Alice and Bob have the *sectorized* Hilbert spaces:

$$\begin{aligned} A^i &:= A_\star^0 \oplus A_\star^1 \\ B^j &:= B_\star^0 \oplus B_\star^1 \end{aligned} \quad (7)$$

where we have defined:

$$\begin{aligned} A_\star^0 &:= A_\star^{vac} \\ A_\star^1 &:= A_\star^{par} \\ B_\star^0 &:= B_\star^{par} \\ B_\star^1 &:= B_\star^{vac} \end{aligned} \quad (8)$$

The scenario can be represented as a routed transformation (δ, U) , where U is the transformation (1), and δ is the Kronecker delta function. Diagrammatically, this is represented by a box for U and a floating label for δ :

$$\begin{array}{ccc} A^i & | & B^j \\ & \boxed{U} & \\ C & | & P \end{array} \quad (9)$$

In the diagram, wires represent sectorized Hilbert spaces; putting two wires next to each other represents the tensor product of those spaces as previously defined (likewise for transformations). We omit superscripts on wires with trivial (i.e. one-sector) sectorizations below, as in the rest of the paper, but it should be remembered that the mathematical object represented is technically a sectorized space.

The role of the route δ is to enforce the constraint that U has no support on its output space outside the one-particle sector $(A \otimes B)^{prac} := (A_\star^0 \otimes B_\star^0) \oplus (A_\star^1 \otimes B_\star^1)$, which serves as its practical output space. Since such a ‘delta-route’ simply matches up the values of the indices, we can equivalently write (9) using the convenient shorthand of *index-matching*:

$$\begin{array}{ccc} A^i & | & B^i \\ & \boxed{U} & \\ C & | & P \end{array} \quad (10)$$

Now, in what sense does (10) offer an improved representation of superposition of paths scenario? By matching up the indices, the diagram indicates the Hilbert space $(A \otimes B)^{prac}$ upon restricting to which U is unitary.⁵ And the index-matching makes it immediately clear that sectors

⁵In more detail: the *accessible space* [55], which can always be defined formally using the routes, at the top of our diagram coincides with the practical output space of the unitary. This means that the diagram unambiguously and formally defines a unitary transformation with this output space.

of the larger Hilbert space $A \otimes B$ in which there is not exactly one particle are irrelevant. The dilemma with which we introduced this section has been overcome: we have achieved a representation that makes clear the physical unitarity without encouraging us to consider irrelevant degrees of freedom.

More generally, routed linear maps (λ, U) with the property that U defines a unitary upon restriction to the practical input and output spaces are called *routed unitaries*.⁶ These transformations are very important for us, since we will use them to define causal relations in the presence of sectorial constraints.

Finally, we introduce *mixed* routed quantum circuits, in which the boxes become routed quantum *channels*, and the wires become spaces of linear operators over Hilbert spaces. The routes then work in the same way and have the same meaning, forbidding input spaces that are in (i.e. have support only in) specific subspaces to be sent to other subspaces.⁷ To differentiate between unitary operators and unitary channels, we reserve calligraphic letters for the latter. For example,

$$\begin{array}{ccc} A^i & | & B^i \\ & \boxed{\mathcal{U}} & \\ C & | & P \end{array} \quad (11)$$

is a routed unitary channel.

Before moving on, a quick note on notation and terminology. Throughout this paper, we will stick to using roman letters like A to denote a Hilbert space. In a helpful abuse of notation, we will sometimes use the same letter to denote the system that that Hilbert space represents. Similarly, A^i is used both to represent a sectorized Hilbert space and the corresponding *sectorized system* – that is, a system that suffers from sectorial constraints.

⁶Or, in the terminology of [55], *practical unitaries*.

⁷In [55], routes in the mixed case were made more expressive by having them also encode so-called *coherence constraints*. This procedure, which involves a doubling of the spaces’ and routes’ indices, is unnecessary for our needs in this paper; we will therefore keep working with standard routes in the mixed theory as well, to avoid clutter, as was already done in [56]. To translate the notations here to those of [55], one just has to replace A^i ’s with $A^{ii'}$ ’s and λ_i^j ’s with $\lambda_i^j \lambda_{i'}^{j'}$ ’s.

3 Causal influence through transformations

What does it mean to say that some system A exerts a causal influence on another system D ? Let us assume, for the time being, that we can model A as an input and D as an output of some transformation. In this case, an obvious idea is:

A influences D just in case it is possible for information to flow from A to D through the transformation that connects them.

In the quantum case, this suggests we should define causal relations as signalling relations through quantum operations. However, in our view, not all quantum operations are suitable to define causal relations. This is because we assume that the existence of a causal relation should be independent of our state of knowledge, but general quantum transformations – and the opportunities for signalling that they afford – do depend on what we know.

For example [5], if A is the control input and D is the target output of a quantum CNOT channel, then A can signal to D through this transformation. Hence we can reasonably say that A is a cause of D . But then suppose that we have no idea about what happens at the target input, so we describe the situation with the channel obtained by preparing it in a maximally mixed state. Through *this* channel, A does *not* signal to D . If we took signalling relations in general quantum channels (or instruments) to define causal relations, then we could legitimately say that A is *not* a cause of D .

The upshot is that, if general quantum operations defined causal relations, then there could be disagreement about whether A is a cause of D , without either of the options being objectively wrong: they would just correspond to different states of knowledge. This relativity is incompatible with our preferred perspective that causal relations are ‘out there’ as objective features of the natural world, waiting to be discovered.⁸

To avoid this relativization of causal structure to states of knowledge, we assume that quantum causal relations correspond to signalling relations

⁸Or at least, as objective features relative to some division of the world into subsystems.

in a *unitary* channel. This approach is consistent with the view that unitary processes (or at some suitable properties thereof) are objective features of the world.

Such causal relations are the subject of this section. Firstly, we will focus on standard unitary channels whose input and output subsystems correspond to tensor factors of the overall Hilbert spaces, summarising the salient features of the framework presented in [5]. We will then generalize to the case of the routed unitaries described above.

3.1 No-influence through an unrouted unitary channel

We have argued that quantum causal relations should be understood in terms of the possibility of signalling through unitary transformations. But we have not spelled out precisely what this means, either by laying out exactly what resources one should be granted in attempting to send messages from A to D , or by translating the idea into mathematics. In this subsection, we will offer six distinct notions of the lack of causal relation, which we call *no-influence through an unrouted unitary channel*.

Some of the notions correspond to quite distinct ideas of what it means to signal, and some are not obviously formulations of no-signalling at all. Remarkably, however, they all turn out to be equivalent. This is good news for our programme of defining causal relations as signalling relations through unitaries, since it means that there is one notion of no-influence that stands out as particularly natural, rather than a number of competing alternatives.

After defining no-influence, we will describe some of the most important consequences of the definition. In doing so, we will provide a compact summary of those features of the framework for quantum causal models developed in [32, 5] that are most relevant to our needs, as a preparation for their subsequent extension to systems suffering sectorial constraints.

An obvious candidate formulation of the claim that A exerts no causal influence on D through a unitary $U : A \otimes B \rightarrow C \otimes D$ is given by the following diagram, in which Tr represents the trace

operation [32]:

(U1)

This says that the channel obtained by tracing out C is equivalent to a channel that traces out A . Intuitively, the notion of no-influence being presented here is that if we only care about the non-effect, we may as well forget about the non-cause.

A closely related notion is that the channel obtained by tracing out C should be independent of any local channel Φ applied to system A before U :

$$\forall \Phi, \quad \begin{array}{c} \text{---} \\ \text{---} \\ C \end{array} \begin{array}{|c|} \hline U \\ \hline \end{array} \begin{array}{c} A \\ B \end{array} \begin{array}{c} \Phi \\ \hline \end{array} \begin{array}{c} D \\ \text{---} \\ \text{---} \end{array} = \begin{array}{c} \text{---} \\ \text{---} \\ C \end{array} \begin{array}{|c|} \hline U \\ \hline \end{array} \begin{array}{c} A \\ B \end{array} \begin{array}{c} D \\ \text{---} \\ \text{---} \end{array} \quad (\text{U2})$$

This condition has a very clear operational meaning – it says that no local interventions on a non-cause can be used to send signals to the non-effect.



A third notion of no-influence from A to D is based on the intuition that if a (mixed) product state is fed into the channel, then the resulting reduced state at D is independent of the choice of state at A :

$$\forall \rho, \rho', \sigma, \quad \begin{array}{c} \text{---} \\ \text{---} \\ C \end{array} \begin{array}{c} \text{---} \\ D \end{array} \begin{array}{c} \boxed{\mathcal{U}} \\ \begin{array}{c} A \quad B \end{array} \begin{array}{c} \triangle \quad \triangle \\ \rho \quad \sigma \end{array} \end{array} = \begin{array}{c} \text{---} \\ \text{---} \\ C \end{array} \begin{array}{c} \text{---} \\ D \end{array} \begin{array}{c} \boxed{\mathcal{U}} \\ \begin{array}{c} A \quad B \end{array} \begin{array}{c} \triangle \quad \triangle \\ \rho' \quad \sigma \end{array} \quad (\text{U3})$$

Note that this is a conceptually distinct notion of signalling to (U2), since one might naively imagine a situation in which one can only signal from A to D via applying a local channel on A to a *nonseparable* state on $A \otimes B$.

While they have the advantage of being obviously related to signalling, all of the candidate no-influence relations so far make essential use of the

trace operation $\overline{\text{---}}$. This might give these conditions a somewhat anthropocentric feel, given that the trace is often taken to represent the ignorance of an agent. Given that we are trying to get away from the idea that causal relations have some essential connection to agents, one might prefer to give a no-influence relations at the level of pure unitary quantum theory. One way of doing this is to require the unitary to decompose in the following way.

$\exists V, W,$

 $=$

 $(U4)$

The intuition here is that A does not influence D just in case there is some decomposition of the unitary into a unitary circuit in which there is no path of wires connecting A and D . Here, a causal relation is understood as primarily a *compositional* property of a unitary channel, rather than a signalling relation.

Another way to define no-influence without the use of the trace is to consider how the unitary channel $\mathcal{U} := U(\cdot)U^\dagger$ acts on the local unitary operators on A . Specifically, we might require that for all unitaries $V \in \mathcal{L}(A)$, there exists a unitary $V' \in \mathcal{L}(C)$ such that:

(U5)

The intuition behind (U5) is explained as follows. A local unitary intervention applied to A before the implementation of U results in a transformation $U(V \otimes I)$. By the unitarity of U , this is equivalent to $\mathcal{U}(V \otimes I)U = (V' \otimes I)U$. So (U5) corresponds to the intuition that A is not a cause of D just in case the result of applying a local unitary intervention on A before the implementation of U is equivalent to that of doing nothing before the implementation of U , and instead just

unitarily operating on output systems other than D .

Our final notion of no-influence concerns the operator algebras $\mathcal{A} := \mathcal{L}(A)$ and $\mathcal{D} := \mathcal{L}(D)$. Specifically, we require that the following commutation relation holds:

$$[U(\mathcal{A} \otimes I_B)U^\dagger, I_C \otimes \mathcal{D}] = 0 \quad (\text{U6})$$

In identifying causal relations with commutation relations, this last notion of no-influence corresponds to well-established intuitions about causality in quantum field theory.

Remarkably, every one of the intuitions discussed here leads to the precisely the same no-influence relation:

Theorem 3.1. *Conditions (U1 - U6) are all equivalent.*

Theorem 3.1 is proven in Appendix A. This allows for a formal definition of no-influence through an unrouted unitary channel.

Definition 3.1. *No-influence through an unrouted unitary channel. Consider a unitary transformation $U : A \otimes B \rightarrow C \otimes D$. We say that A has no influence on D through U , written $A \not\rightarrow D$, just in case conditions (U1 - U6) hold.*

We say that A is a cause of D through U , written $A \xrightarrow{U} D$, just in case it is not true that $A \not\rightarrow D$.

We thus have a notion of no-influence that is motivated by a number of distinct intuitions. It turns out that this notion leads to two very interesting and useful properties, to which we now turn.

Proposition 3.1 (Time symmetry). *A is not a cause of D through U if and only if D is not a cause of A through U^\dagger . That is,*

$$A \not\rightarrow D \iff D \not\rightarrow A, \quad (12)$$

This property is immediate from (U4), as well as (U6). The property feels quite intuitive in a time symmetric universe – if the dynamics are reversible, and we define causal relations in terms of those dynamics, then it appears natural that those causal relations should be reversible too.

However, remarkably, the dynamical time symmetry of a theory does not always imply that its causal structure is time symmetric, meaning that

the latter is a distinctive feature of quantum theory. In particular, we note that classical causal relations, defined as dependencies in a reversible function, are not time-reversible. For example, in a classical CNOT gate, there is signalling from the control to the target. But in the time-reversed function (which is just the CNOT), there is no signalling from the target to the control. Proposition 3.1 thus provides a sense in which quantum theory actually feels much more natural than classical theory⁹.

The second useful property is the *atomicity* of causal structure in unitary transformations [5, 57].¹⁰

Proposition 3.2 (Atomicity). *No-influence relations in a unitary transformation can be composed and decomposed in the following ways:*

$$(A \not\rightarrow E) \wedge (A \not\rightarrow F) \iff (A \not\rightarrow E \otimes F) \quad (13)$$

$$(A \not\rightarrow F) \wedge (B \not\rightarrow F) \iff (A \otimes B \not\rightarrow F) \quad (14)$$

This proposition is proven in Appendix B.


Atomicity says that if a system A cannot influence either system E or system F individually, it cannot influence the composite system $E \otimes F$. This is a remarkable property since there are pure states on $E \otimes F$ for which much (if not all) of the information is only present in the nonlocal correlations between the subsystems, and does not show up in the individual reduced density matrices of the subsystems. Thus, naively, one could imagine that A could influence $E \otimes F$ without influencing its subsystems by only affecting the nonlocal correlations (indeed, this can happen for signalling relations in general quantum channels). Yet atomicity shows that this is impossible: the nonseparability of pure states does not lead to a similar sort of nonseparability at the level of causal structure.

Atomicity is also very useful since it enables us to represent the overall *causal structure* of a unitary transformation using a directed graph between its individual, ‘atomic’ input and output


⁹Note, however, that time-symmetry of causal relations holds for the theory of classical *symplectic transformations* [60].

¹⁰The term ‘atomicity’ is not used in [5], even though the notion is introduced there (see Remark 4.4). The name was coined later in [57].

subsystems. This is a very important theme for us, so it is worth dwelling on. Consider a unitary transformation of the following form:



We now know how to tell whether there is causal influence between any pair of input/output subsystems – we just check using whichever of (U1–U6) is most convenient. We can thus write down six causal influence or no-influence relations, which can conveniently be represented with a directed acyclic graph with arrows representing the influence relations. For example, if our only *no*-influence relations are $A \not\rightarrow D$, and $A \not\rightarrow E$, we can represent all six relations as follows:



If we didn't know about atomicity, then we might think that there are still more facts about the causal structure that aren't captured by this graph. In particular, we might imagine that A is a cause of the composite system $D \otimes E$. But if we bear atomicity in mind, this graph tells us that $A \not\rightarrow D \otimes E$, even though there is no node representing $D \otimes E$. Thus a directed acyclic graph between the five subsystems from the original decomposition of the unitary represents a much larger set of facts about the causal structure.

More generally, atomicity implies that the causal relations among arbitrary subsets of the input/output subsystems are fixed by the causal relations among the individual subsystems. This justifies us in simply *defining* the causal structure of a unitary transformation as this directed acyclic graph:

Definition 3.2. *The causal structure of a unitary transformation is directed acyclic graph representing the causal relations between its individual input/output subsystems.*

While this notion of causal structure is good as far as it goes, it does not apply to situations where


sectorial constraints are at play. In the next subsection, we generalize the notion to *routed* unitaries, in order to accommodate them.

3.2 No-influence through a routed unitary channel

As (U2) makes clear, no-influence can be understood as the independence of measurements on the non-effect on transformations on the non-cause. This suggests that if the possible transformations or measurements on subsystems are *restricted*, different causal relations will arise. But our current definition of no-influence does not accommodate such restrictions; indeed, (U2) quantifies over *all* quantum channels on A . The unfortunate implication is that the existing no-influence condition will attribute the wrong causal relations to scenarios featuring restrictions on transformations and measurements.

As a concrete example, suppose that A and D are the control input and the target output respectively of a quantum CNOT gate. According to our current definition, A is a cause of D . But now suppose that the only measurements one can perform on D are measurements of the orthonormal basis vectors $\frac{|0\rangle \pm |1\rangle}{\sqrt{2}}$ (perhaps because of a superselection rule). In this case, all the possible measurements on D are independent of any transformations on A . Then, presumably, we should say that A is *not* a cause of D . We need a more general notion of no-influence to accommodate this.

We can understand this restriction on measurements as the result of a sectorial constraint. Note that, due the Naimark dilation, any measurement on a system D can be performed by implementing a channel with an ancillary output D' , and then measuring D' . Diagrammatically, this channel has the form



Implicitly, (U2) assumed that any measurements of this form on D can be performed, since it requires that its reduced state should be independent of the channel Φ on A . Let us now restrict ourselves to the *sector-preserving* measurements on a sectorized version of the system, D^l . These

are the measurements that can be implemented using a channel of the form (15) that maps each of D^l 's sectors to itself. In other words, sector-preserving measurements are the ones that can be performed using a routed channel of the form:

$$\begin{array}{c} D^l \quad D' \\ | \quad | \\ \boxed{\Psi} \\ | \quad | \\ D^l \end{array} \quad (16)$$

Suppose D^l is a sectorized qubit with two sectors, spanned by $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$ and $\frac{|0\rangle-|1\rangle}{\sqrt{2}}$ respectively. In this case, the measurements on D^l that one can perform via measurements of D' with this sort of channel are precisely the measurements of (probabilistic mixtures of) these two states. We can therefore accommodate the constraint from the previous example by restricting to sector-preserving channels.

More generally, we also want to consider sectorial constraints on the putative cause systems. We thus restrict to the sector-preserving channels Φ on the sectorized system A^i . These are the ones for which

$$\begin{array}{c} A^i \\ | \\ \boxed{\Phi} \\ | \\ A^i \end{array} \quad (17)$$

is a valid routed channel.

We can therefore consider a causal relation between sectorized systems as the possibility of signalling using sector-preserving transformations on the cause and sector-preserving measurements on the effect. Since we are now dealing with sectorized systems, we can define the no-influence relations with respect to routed unitaries of the following form:

$$\begin{array}{c} C^k \quad D^l \\ | \quad | \\ \boxed{U} \\ | \quad | \\ A^i \quad B^j \end{array} \quad \lambda_{ij}^{kl} \quad (18)$$

We can formulate the resulting notion of *no-influence through a routed unitary* (λ, U) from A^i to D^l as follows:

$$\forall \Phi, \Psi, \lambda_{ij}^{kl} \quad \begin{array}{c} C^k \quad D^l \quad D' \\ | \quad | \quad | \\ \boxed{\Psi} \\ | \quad | \quad | \\ C^k \quad D^l \\ | \quad | \\ \boxed{U} \\ | \quad | \\ A^i \quad B^j \\ | \quad | \\ \boxed{\Phi} \\ | \quad | \\ A^i \quad B^j \end{array} = \lambda_{ij}^{kl} \begin{array}{c} C^k \quad D^l \quad D' \\ | \quad | \quad | \\ \boxed{\Psi} \\ | \quad | \quad | \\ C^k \quad D^l \\ | \quad | \\ \boxed{U} \\ | \quad | \\ A^i \quad B^j \end{array} \quad (R1)$$

Applying this definition of no-influence to the example above, we indeed find that there is no influence from A (which has a trivial sectorization) to the now-sectorized system D^l .

As in the case of unrouted unitary channels, this definition of no-influence turns out to be equivalent to a number of conceptually distinct formulations. We now lay these out, beginning with a generalization of (U5). This notion is once again motivated by the intuition that applying (now sector-preserving) local unitaries on the non-cause and then applying U should be equivalent to applying U and then some unitary that only acts trivially on the non-effect.

However, since the non-effect is now a sectorized Hilbert space, we do not require that that this last unitary is confined to the C wire. Instead, it can take the more general form of a family of unitaries on C that is coherently controlled on the sector of D . More precisely, we can have something of the form $V' = \sum_l V_C^l \otimes \pi_D^l$, where each V_C^l is a unitary on C , and each π_D^l projects onto the l th sector of D^l . We can give this second notion of no-influence through a routed unitary the following diagrammatic representation:¹¹

¹¹We note that the equality here only holds between the unitary circuits, and not necessarily between the routes. That is why we allow a D^l at the bottom right of the right side of the equation, which is index-matched with the D^l at the top right, and a D^n in the corresponding position on the left side, which may or may not be matched up with D^l via the route λ .

$$\begin{array}{c}
\forall V, \exists V', \\
\begin{array}{c}
C^k \quad D^l \\
\lambda_{ij}^{kl} \quad U \\
A^i \quad B^j \\
V \\
A^i \quad B^j \\
(\lambda^T)^{ij}_{mn} \quad U^\dagger \\
C^m \quad D^n
\end{array}
=
\begin{array}{c}
C^k \quad D^l \\
V' \quad \text{---} \quad \text{---} \quad W^\dagger \\
C^k \quad \dot{D}^l \quad D^l \\
C^m \quad \dot{D}^l \quad D^l \\
C^m \quad D^l \\
W
\end{array}
\end{array}
\quad (R2)$$

Above, V is any unitary operator on A that is block diagonal in its sectorization. λ^T denotes the transpose of λ . W is a unitary matrix that maps each sector D_\star^l to $\dot{D}_\star^l \otimes D_\star^l$, where \dot{D}_\star^l is one-dimensional. Accordingly, if D^l has n sectors, then \dot{D}^l is a sectorized Hilbert space with n one-dimensional sectors. The map between W and its adjoint on the RHS is a unitary operation V' on the left wire coherently controlled by the value of \dot{D}^l in the basis given by its sectorization.

A third notion of no-influence through a routed unitary is a generalization of the commuting-algebras formulation (U6). For this, we need to consider *sector-preserving* operator algebras, which are not generally factor algebras of the form $\mathcal{L}(H)$. For a sectorized Hilbert space A^i , the corresponding sector-preserving algebra has the form:

$$\mathcal{A} := \bigoplus_i \mathcal{L}(A_\star^i) \quad (19)$$

\mathcal{A} can be embedded into an algebra of operators \mathcal{A}^* on our routed unitary's entire input space. \mathcal{A}^* is defined from \mathcal{A} using an 'embedding map', defined in Appendix C. The third notion of no-influence can then be written:

$$[U\mathcal{A}^*U^\dagger, \mathcal{D}^*] = 0 \quad (R3)$$

where \mathcal{A}^* is the algebra for A^i embedded into U 's input space, and \mathcal{D}^* is the algebra for D^l embedded into U 's output space.

Finally, there is a last notion of no-influence which uses a special unitary transformation called the *exchange gate*. We leave the formal definition of this gate to the next subsection, in which it will play a more central role. Suffice it to say for now that the exchange gate for some sectorized system can be intuitively regarded as the unitary acting on that system together with and unsectorized ancillary that allows them to exchange the

maximum possible amount of information while respecting the sectorial constraints. Our final notion of no-influence through a routed unitary is that there should be no signalling from the ancillary input of the exchange gate on A^i to the ancillary output of the one on D^l :

$$\begin{array}{c}
\begin{array}{c}
\equiv \quad \equiv \quad \equiv \\
C^k \quad D^l \\
\mathcal{E}\mathcal{X}\mathcal{C}\mathcal{H} \\
C^k \quad D^l \\
\lambda_{ij}^{kl} \quad U \\
A^i \quad B^j \\
\mathcal{E}\mathcal{X}\mathcal{C}\mathcal{H} \\
A^O \quad A^i \quad B^j
\end{array}
=
\begin{array}{c}
\equiv \\
A^O \\
\Phi \\
A^i \quad B^j \quad D^l \\
\mu_{ij}
\end{array}
\end{array}
\quad (R4)$$

Above, the route $\mu_{ij} := \sum_{kl} \lambda_{ij}^{kl}$ simply restricts to U 's practical input space, and Φ is some quantum channel.

In the special case where the sectorization is trivial (i.e. involves just one sector), the exchange gate is just a swap. Hence, (R4) is a generalization of (U1).¹²

Although motivated by distinct intuitions, each of these notions turns out to be equivalent, just like in the unrouted case.

Theorem 3.2. *(R1 - R4) are all equivalent.*

Theorem 3.2 is proven in Appendices E and F. Thus there is a clear, natural choice for notion of no-influence when sectorial constraints are at play, which we capture as follows.

Definition 3.3. *Given a routed unitary transformation (18), we say that A^i has no influence on D^l in U , written $A^i \not\stackrel{U}{\rightarrow} D^l$, just in case the conditions (R1 - R4) hold.*

We say that A^i is a cause of D^l through U , written $A^i \stackrel{U}{\rightarrow} D^l$, just in case it is not true that $A^i \not\stackrel{U}{\rightarrow} D^l$.

The reader will notice that we have not provided generalizations of (U3) or (U4) for the

¹²Formally, the diagram obtained from (R4) by assuming all routed and sectorizations are trivial (and therefore substituting swaps for exchanges) and (U6) are slightly different, since the former includes ancillary systems such as A^O that do not feature in the latter. But it is easily shown that each one can be derived from the other.

routed case. There does not appear to be any reasonable generalization of (U3). One reason for this is that the practical input space of a routed unitary is not generally the tensor product of its subsystems, and, accordingly, starting with a product state $\rho_A \otimes \sigma_B$ that lies in the space and then changing ρ_A to some ρ'_A might take you out of that space. And anyway, this sort of switching should not be permitted since it cannot be brought about with a sector-preserving channel. We draw the moral that in the general case, causal relations are best thought of in terms of transformations rather than states. The question of whether (U4) can be generalized, and shown to be equivalent to (R1 – R4), is left open for future work.

The next obvious question is whether this more general notion of no-influence leads to the same nice properties as its unrouted counterpart. It turns out that this is indeed the case: the no-influence relation through a routed unitary is time-symmetric and atomic.

Proposition 3.3 (Time symmetry – routed case). *Given the routed unitary (18), A^i is not a cause of D^l in U if and only if D^l is not a cause of A^i in U^\dagger . That is,*

$$A^i \not\rightarrow U D^l \iff D^l \not\rightarrow U^\dagger A^i, \quad (20)$$

Time-symmetry follows immediately from (R3).

For the atomicity property, consider a unitary transformation of the following form (expressed in the pure theory):

$$\begin{array}{c} D^l \quad \quad E^m \quad F^n \\ \lambda_{ijk}^{lmn} \quad \boxed{U} \\ A^i \quad \quad B^j \quad C^k \end{array} \quad (21)$$

Recall that, just like unsectorized Hilbert spaces, we can take a tensor product of sectorized Hilbert spaces as in 3. We then have the following proposition, proven in Appendix G.

Proposition 3.4 (Atomicity – routed case). *No-influence relations in the routed unitary (21) can be composed and decomposed in the following ways:*

$$(A^i \not\rightarrow U E^m) \wedge (A^i \not\rightarrow U F^n) \iff (A^i \not\rightarrow U E^m \otimes F^n) \quad (22)$$

$$(A^i \not\rightarrow U F^n) \wedge (B^j \not\rightarrow U F^n) \iff (A^i \otimes B^j \not\rightarrow U F^n) \quad (23)$$

Just like the unrouted case, this property allows us to represent the causal structure of a routed unitary as a directed acyclic graph between its input and output subsystems. Since we will define a fine-graining of this structure in the next subsection, called the ‘sectorized’ causal structure, we call this the ‘unsectorized’ causal structure of a routed unitary.

Definition 3.4. *The unsectorized causal structure of a routed unitary transformation is a directed acyclic graph representing the causal relations between its individual input and output subsystems.*

3.3 Sectorized no-influence through a routed unitary transformation

Suppose we know that A^i influences D^l through a routed unitary. (R1) then implies that it is possible for associated agents to send and receive messages by performing local, sector-preserving interventions on A^i and D^l . However, one might also want to know whether it is possible to send messages via a restricted set of interventions on A^i that only act nontrivially on a specific sector, such as A_\star^0 . If so, one could further ask whether sending such messages is still possible if we additionally require that the local effect of the intervention on D^l is only to modify the relative phases between the different sectors, and to do nothing to states that have support within only one sector.

More generally, we can ask about not only whether one entire sectorized system influences another, but also whether there is influence between the systems’ specific sectors and/or the relative phases between them. To tackle such questions, this section introduces *sectorized* no-influence relations.

The essential tool here is the exchange gate, which we already mentioned in the previous subsection. Let us now describe its role and definition in detail, beginning with an intuitive account in terms of agents and their actions. Suppose that an agent Alice has unconstrained access to a system A . Alice wants to exploit, to the best of her

ability, the possibilities for sending and receiving messages afforded by A . A natural strategy is to swap A with an ancillary system, A' , of the same dimension. Then preparing A' in her choice of initial state allows her to send arbitrary messages through A , and measuring the final state of A' allows her to probe all the information that flowed into A . Alice's use of the swap thus allows her to extract as much information as possible from A , and likewise insert as much information as possible into it.

In the presence of sectorial constraints, Alice is forbidden to use this strategy: the swap with an ancillary system is not sector-preserving on A^i . Intuitively, the point of the exchange gate is to provide her with the *next best strategy*. In other words, the gate allows her to exchange the maximum amount of information possible between A^i and an ancilla without violating the sectorial constraints.

How can Alice do this? Taking A^i to have N sectors, one sensible approach is for her to keep in store not one, but N ancillary systems, each corresponding to one of the sectors; she can then swap A^i with the corresponding ancilla depending (coherently) on which sector is populated, e.g. swapping it with the first ancilla if it is in the first sector, etc. These ancillas will be called the *sectorial ancillas*.

But swapping the individual sectors is not enough; one also needs to exchange the information contained in the relative phases between them. To this end, Alice will also need an $(N+1)$ th ancilla of dimension N , on which a generalised CNOT gate is performed, whose control is the 'which sector of A^i is populated?' information. This $(N+1)$ th ancilla will be called the 'which-sector' ancilla A_w , as, when its input state was prepared in the computational basis, reading its output tells Alice about which sector of A^i was populated (and decoheres this information). Alternatively, by preparing the input states of A_w in the Fourier basis, Alice can introduce phase shifts between the sectors of A^i through phase back-action.

Using both the sectorial ancillas and the which-sector ancilla is the best strategy for Alice: this is what the exchange gate represents.

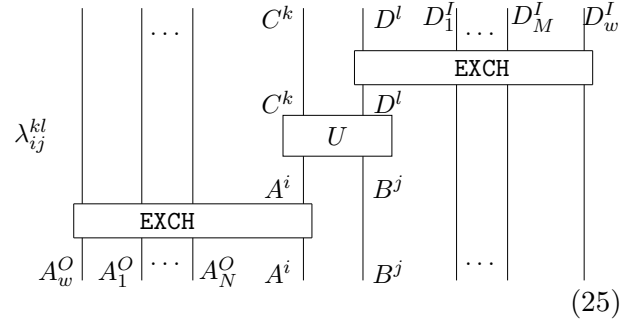
Let us now turn to its mathematical definition.¹³ The exchange gate acts on A^i and $A' =$

$A_1 \otimes \dots \otimes A_N \otimes A_w$, and decomposes as a direct sum of N terms, each acting on one sector. The i th term swaps the i th sector A_\star^i of A with A_i , increases the logical value of A_w by i , and does nothing to the remaining sectorial wires. In other words, we have:

$$|m\rangle_{A_w} |\phi\rangle_{\overline{A_i}} |\xi\rangle_{A_i} |\psi_i\rangle_A \xrightarrow{\text{EXCH}} |m+i\rangle_{A_w} |\phi\rangle_{\overline{A_i}} |\psi_i\rangle_{A_i} |\xi\rangle_A \quad (24)$$

for a state $|\psi_i\rangle_A \in A_\star^i$ in the i th sector, a state $|\phi\rangle_{\overline{A_i}}$ on $\overline{A_i} := \bigotimes_{j \neq i} A_j$, and the logical state $|m\rangle_{A_w}$. The exchange gate on a system with only one sector is just the swap. Hence the exchange gate is a generalization of the swap formally, as well as conceptually.

When writing the exchange gate as part of a circuit diagram, it is helpful to distinguish between A' before and after the implementation of the exchange gate using the notations A^O and A^I respectively, and similarly for its subsystems. We can now rewrite the channel from (R4) in terms of those subsystems:



This routed unitary is all we need to define the more fine-grained causal structure that we were looking for, as the exchange gate allows us to neatly differentiate between the causal influences that are mediated by each individual sector, and also by the 'which-sector' information. As we already know how to define no-influence relations among the input/output subsystems of a routed unitary transformation, we can simply define a sectorized no-influence relation as the no-influence relation between a subsystem of A' and D' in the routed unitary above. For example, if $A_3^O \not\rightarrow D_w^I$ through the routed unitary in (25), then we take this to mean that the third sector of A^i has no influence on which-sector information in D^I . If A^i and D^I have N and M sectors respectively, then we have a total of $(N+1)(M+1)$ sectorized no-influence relations.

proofs, is provided in Appendix D.

¹³An equivalent definition, that is useful for following

Since they were defined in terms of our existing notion of no-influence through a routed unitary, the sectorized no-influence relations immediately inherit time-symmetry and atomicity. This means that we can define the sectorized causal structure as a directed acyclic graph, where a system with N sectors is represented by $N + 1$ vertices.

Definition 3.5. *The sectorized causal structure of a routed unitary is a directed acyclic graph with a vertex for each sector of each input/output subsystem, and a vertex for the which-sector information in each system. The arrows represent the causal relations between them.*

We also note that there is an unsectorized causal influence between an input and output sectorized system if and only if there is at least one corresponding sectorized causal relation. This follows from (R4), atomicity, and the fact that the ancillary wires of the exchange gate are the tensor products of the $n + 1$ wires that define the sectorized causal structure. This is equivalent to the following proposition, which makes clear the sense in which the sectorized causal structure is a fine-graining of the unsectorized one.

Proposition 3.5. *The unsectorized causal structure of a routed unitary transformation is obtained from the sectorized causal structure by combining all the vertices that correspond to the same sectorized system, while preserving the edge structure.*

It is worth elaborating on what the sectorized causal relations mean, and why we need to consider the which-sector ancilla A_w . Causal influence on a sector means the possibility of affecting the information within that sector. Now, since the total state space of the sectorized system is just the direct sum of its sectors, one might naively think that if none of its sectors are influenced by some potential cause, then the sectorized system itself is not influenced. But a moment's reflection shows that this is false, since a cause might be able to affect *which* sector happens, and yet have no effect at all on the information contained within the sectors. And the question of which sector happens can indeed be settled with a sector-preserving measurement.

All of this is particularly obvious if all the sectors are one-dimensional, so that they contain no

information; an example is provided in Figure 1. This shows that accounting for the which-sector information via the which-sector ancilla is crucial for Proposition 3.5 to hold, and therefore for our sectorized causal structure to deserve the status of a fine-graining of the unsectorized one.

Similarly, the causal influence exerted by a sectorized system is not just the combination of that exerted by its sectors, since one can also change the relative phases between the sectors via a sector-preserving transformation. This is once again accounted for by the which-sector ancilla. An example is provided in Figure 2, which is just the time-reversed version of the previous example.

We briefly note that a number of alternative formulations of the sectorized no-influence relations are available. It is possible to associate each of the vertices in a sectorized causal structure with an algebra of operators – in fact, a subalgebra of the sector-preserving algebra for the corresponding system. This algebra defines the class of channels and unitaries one quantifies over in expressions like (R1) and (R2), which then turn out to be equivalent to the no-influence relation provided in this subsection, assuming that any sectors involved have more than one dimension. This is discussed in Appendix E.

4 Causal structure in supermaps

The no-influence relations introduced so far only allow us to consider a very limited range of causal structures. For it seems perfectly reasonable to consider causal structures, such as $A \rightarrow B \rightarrow C$, in which some systems act as both causes and effects. But in a unitary transformation, each subsystem can only act as a cause *or* as an effect, depending on whether it is an input or an output of the unitary.

For an example of a scenario which might exhibit this causal structure, suppose that Alice and Charlie are separated not just by unitary channels, but by an intervention that can be freely chosen by a third agent, Bob. The situation is depicted below, where the gap in the wire represents Bob's ability to choose an arbitrary operation with which to fill it.





Figure 1: A routed unitary illustrating the need to include the which-sector ancilla in the exchange gate. The routed unitary on the left of (a) is just an identity channel from a qubit A with only one sector (so that we can omit an index) to a qubit B^i sectorized into two one-dimensional sectors. The exchange gate on an unsectorized system is just the swap, and the exchange gate on a sectorized qubit is the CNOT controlled on that qubit. Hence putting in exchange gives us the right side of (a) where the control input and output to the CNOT is A^O and B^i respectively (the wires of the exchange gate corresponding to one-dimensional sectors are themselves one-dimensional, so they can be ignored). Evaluating the no-influence conditions on this second transformation gives us the more fine-grained sectorized causal structure on the left of (b), which yields the coarse-grained, unsectorized causal structure on the right when we combine the vertices associated with B^i .

Since they are both one-dimensional, there is obviously no influence from A to either of the sectors. Thus, if we omitted the which-sector information B_w , we would conclude that there is no influence at all from A to B^i . But this would be wrong, since Alice can send a message to Bob by preparing A in either $|0\rangle$ or $|1\rangle$, which Bob can then read by performing a (sector-preserving) computational basis measurement on B^i . Therefore, not all the information in B^i is encoded within its sectors: there is also information about which sector happened.





Figure 2: A routed unitary that illustrates the need to consider influence through relative phases, hence the which-sector ancilla. The routed unitary on the left of (a) is just the time reverse of that from Figure 1. From time-symmetry, it follows that the sectorized and unsectorized causal structures are obtained by reversing the arrows in Figure 1. Again, the which-sector information turns out to be essential for getting the right unsectorized structure; this time, because Bob can signal to Alice only by acting on the relative phase between the two sectors.



Our existing notion of no-influence might help us determine whether A influences B , and whether B influences C . But how do we tell whether A influences C ? Our existing notion of no-influence doesn't apply, since the systems are not the input and output of a single unitary transformation.

This section extends our framework to deal with this more general sort of scenario. The diagram (26) can be formally interpreted as a quantum *supermap* [61]¹⁴. Taking this as a hint, we start by characterizing the causal structure of supermaps, once again re-expressing compactly the central ideas of [5]. This involves reformulating these ideas in the language of supermaps, rather than in terms of the process matrices used in [5]. We then generalize to *routed supermaps* [56], in which sectorial constraints can be accounted for, and in which a more fine-grained, sectorized causal structure can also be displayed.

4.1 The causal structure of a standard supermap

A quantum supermap transforms a set of quantum operations as an input into a single quantum operation as an output: it is an 'operation on operations'. We will exclusively consider the deterministic supermaps defined precisely in [2]; the rough idea is as follows. A bipartite, deterministic quantum supermap on product channels, \mathcal{S} , is a linear map of the form

$$\mathcal{S} : \text{Herm}(A^I \otimes B^I \rightarrow A^O \otimes B^O) \rightarrow \text{Herm}(P \rightarrow F) \quad (27)$$

where each of the italicised letters are Hilbert spaces, and $\text{Herm}(X \rightarrow Y)$ is the space of

¹⁴An alternative, equivalent representation is that of *process matrices* [3]. See footnote 17 for a brief comment on the connection between the two pictures, or Appendix A of [29] for an extensive conceptual and mathematical treatment of this connection.


Hermitian-preserving maps from $\mathcal{L}(X)$ to $\mathcal{L}(Y)$. To state a further requirement we consider an arbitrary collection of ancillary Hilbert spaces $\gamma_A^I, \gamma_A^O, \gamma_B^I, \gamma_B^O$, and an arbitrary pair of quantum channels:

$$\begin{aligned} \mathcal{M}_A &: \mathcal{L}(\gamma_A^I \otimes A^I) \rightarrow \mathcal{L}(\gamma_A^O \otimes A^O) \\ \mathcal{M}_B &: \mathcal{L}(\gamma_B^I \otimes B^I) \rightarrow \mathcal{L}(\gamma_B^O \otimes B^O) \end{aligned} \quad (28)$$

A bipartite deterministic supermap on channels is any linear map of the form (27) for which

$$(\mathcal{I} \otimes \mathcal{S} \otimes \mathcal{I})[\mathcal{M}_A \otimes \mathcal{M}_B] \quad (29)$$

is always a quantum channel, where the left and right \mathcal{I} 's are the identities on $\text{Herm}(\gamma_A^I \rightarrow \gamma_A^O)$ and $\text{Herm}(\gamma_B^I \rightarrow \gamma_B^O)$ respectively. An n -partite, deterministic quantum supermap on product channels is then the obvious generalization of this idea. A general such supermap is depicted below, in which we consider an arbitrary factorization of its 'past' P and its 'future' system F .



Each node can be associated with an agent who chooses the intervention at that node. The commonly invoked intuition [3] is that each agent is in an isolated laboratory, the doors of which open only once to receive the input system associated with A_l^I , and then once more to eject the output system associated with A_l^O , leaving some time in between for the agent to perform an intervention.¹⁵ These laboratories are what is represented by the nodes. Outside the labs these systems are allowed to interact, in a way described by the internal structure of the supermap.

Although we refer to agents for pedagogical purposes, we do not want the notion of agency to be essential to our framework. We should therefore always allow for the possibility that an agent does nothing at all, in which case it should be

¹⁵If ancillas are used then these may be pictured as either residing in the lab the entire time, or as being received before the rest of the experiment and ejected afterwards.

possible to regard the outgoing system as essentially the same as the ingoing one. For this reason, we assume here that A_l^I and A_l^O are of the same dimension for all l .

As in the case of transformations, we reserve the notion of causal structure for *unitary* supermaps.¹⁶ Our reason for restricting to unitaries is the same as before – we assume that unitary supermaps are the ontic supermaps that allow us to keep causal relations objective [5, 32].

The objects that bear causal relations in a supermap are its nodes and its past/future subsystems. The relations can be defined in terms of the unitary channel obtained by pulling the outgoing wire of each node to the bottom of the page, and the ingoing wire to the top. This is called the *process channel*¹⁷:

(31)

This corresponds to each agent choosing a special intervention associated with each node: a swap onto an ancillary system of the same dimension as A_l^I .

At this point, we already know what it means for an input to a unitary channel to influence some output. Hence we can define a *direct cause* relation in terms of this influence relation. Specifically, the i th node is a direct cause of the j th node if and only if A_i^O is a cause of A_j^I in the process channel. We can also define direct cause relations involving the past/future subsystems: for example, the i th past system P_i is a direct cause of the j th node if and only if P_i is cause of A_j^I in the process channel.

¹⁶For supermaps this means either that it transforms n unitary channels to another unitary channel, or that the associated process matrix (defined below) is unitary – the two notions are equivalent [28].

¹⁷The operator representing this channel via the Choi-Jamiołkowski isomorphism is commonly known as the *process matrix* or *process operator*, introduced in [3] and commonly used in the literature on indefinite causal order.

There is a nice intuition behind this definition. When each agent implements a swap, they prevent any information from flowing from the input systems of the nodes to the outputs. This means that any possibilities for signalling between agents at the nodes via their ancillas are entirely due to the internal structure of the supermap. In this case, the ability of one agent to signal to another is not dependent on, or mediated by, the other agents allowing information to flow through the node.

Moreover, when we define direct cause relations in this way, they inherit the atomicity property from the relation of no-influence through a unitary transformation. The direct causal relations between arbitrary subsets of the nodes and past/future subsystems are therefore fixed by the direct causal relations between the individuals. This enables us to define the causal structure of a supermap as a directed graph, where the vertices are the individual supermap nodes (since the A_i^O and A_i^I vertices can be merged together) and the past/future subsystems. Indirect cause relations are then defined as directed paths in this graph.

The causal structure of the supermap is easily obtained from the causal structure of its process channel:

Definition 4.1. *The causal structure of a standard unitary supermap is a directed graph, obtained by combining the vertices corresponding to the same supermap node in the causal structure of the process channel, while preserving the edge structure.*

An example of this procedure is provided in Figure 3 for the causal structure of the quantum switch [2].

We note that, while the causal structure of the process channel will always be acyclic, the process of combining vertices might induce cycles in the causal structure of the supermap, as in Figure 3. This implies that the vertices cannot be ascribed a partial order such that effects are always higher up than their causes. Accordingly, if a supermap has a cyclic causal structure, we say that it exhibits *indefinite causal order* [8].

4.2 The causal structure of a routed unitary supermap

The above notion of the causal structure of a supermap is good as far as it goes, but it fails to ac-




Figure 3: Moving from the causal structure of the process channel to the causal structure of the supermap. On the left-hand side, we have a possible causal structure of the process channel for a bipartite unitary supermap, with nodes A and B (in particular, this is what we get for the quantum switch [2, 8]). To get the causal structure of the supermap, we just need to combine A^I and A^O into a single vertex representing the node A , and B^I and B^O into a single vertex representing B .

commodate sectorial constraints. Yet, as we aim to demonstrate with our analysis of purported implementations of the quantum switch, accounting for these constraints can be important for understanding the causal structure of a physical scenario. We need a more general framework.

Our approach will be to extend the framework sketched above so that it applies to supermaps acting on routed maps, which we call *routed supermaps*. Defined in [56], routed supermaps are basically the same as unrouted supermaps, except that the input channels to each node of a routed supermap are required to follow a route associated with that node.

Since we still do not want to rely on a notion of agency, we will assume that the ingoing and outgoing sectorized Hilbert spaces associated with the l th node, $A_{I,l}^{i_l}$ and $A_{O,l}^{i'_l}$, are copies of the same sectorized Hilbert space. By this we mean that they are sectorized into the same number of sectors, and that the dimensions of the sectors $A_{I,l}^{i_l}$ and $A_{O,l}^{i'_l}$ are the same for all $i_l = i'_l$. We further assume that each node only accepts sector-preserving channels; that is, the ones that follow the route $\delta_{i_l}^{i'_l}$.¹⁸

As before, we assume that the routed su-

¹⁸The causal structure of unitary supermaps with more general routes is an interesting problem for future work.

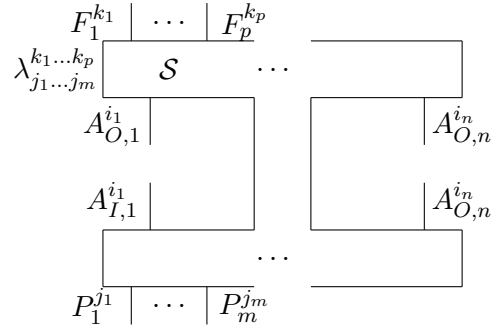



Figure 4: An n -partite routed unitary supermap, represented in the pure theory. The l th node takes as inputs routed unitary transformations that map each i_l th sector to itself. Then the whole supermap outputs a routed unitary of the form (λ, V) , where λ is some fixed route matrix, and V is a unitary transformation that depends on the choice of inputs.

permaps that have causal structure are the routed unitary supermaps; that is, the ones that return routed unitary channels when a unitary channel is inserted into each one of its nodes. An example of the sort of supermap we are considering is illustrated in Figure 4, in which the matched indices represent the routes at the nodes.

How can we define direct causes in such a supermap? In the case of standard supermaps, we did this by inserting swaps into each node. But this trick won't work here, since the swap is not a sector-preserving intervention. Hence there is

no guarantee that this will leave us with a routed unitary transformation – in fact, there is no guarantee it will be a channel at all. Accordingly, inserting swaps does not necessarily leave us with an object with a causal structure that we know how to analyse.

This is where the exchange gate introduced in Section 3.3 once again comes in handy. As previously discussed, the exchange gate is the closest thing to the swap on a sectorized system that preserves its sectors. This suggests a promising strategy. We can form a *routed process channel* by inserting exchange gates into each node:



This is a routed unitary transformation, whose causal structure we know how to analyse. It reduces to the usual process channel for the special case of standard supermap (which can be regarded as a routed supermap where all the sectorizations are trivial). If we define direct cause relations as signalling relations among the ancillary inputs and outputs of the exchange gates in this channel, we recover the intuition of a direct cause relation as the possibility of signalling even when agents extract and replace the maximum amount of information possible from their node.

This will be our approach. We say that the l th node is a direct cause of the r th node if and only if A_l^O is a cause of A_r^I in the routed process channel. We can similarly define direct causes involving past/future subsystems. For example, the r th past subsystem is a direct cause of the l th node if and only if $P_r^{j_r}$ is a cause of A_l^I in the routed process channel.

Recall that the atomicity of causal relations in a unitary channel meant that we could define the causal structure of an unrouted supermap as a directed graph, via the procedure sketched in Figure 3. A similar discussion applies in the routed case, leading to the following definition.

Definition 4.2. *The unsectorized causal structure of a routed unitary supermap is a directed graph. The graph is obtained by combining the vertices corresponding to the same node of the supermap in the sectorized causal structure of the routed process channel, while preserving the edge structure.*

As in the unrouted case, the process of combining pairs of vertices might lead to cycles, which we take to imply indefinite causal order.

Now recall that, for routed unitary transformations, the way the ancillary wires on the exchange gate factorized provided us with a way to define more fine-grained causal relations. We can provide a similar fine-graining of the causal structure of a routed unitary supermap.

The first step is to write the process channel as in Figure 5, where we have decomposed the ancillary systems of each node's exchange gate into its subsystems. We then take the sectorized causal structure (from Definition 3.5) for the decomposition of the routed unitary. (In this causal structure, for a past/future sectorized subsystem with N sectors, there are $N + 1$ vertices, whereas for a node with N sectors, there are $2(N + 1)$ vertices, since it is associated with the input and output wires on its exchange gate.) Then, we just merge the input and output vertex for each ancillary system.

Definition 4.3. *The sectorized causal structure of a routed unitary supermap is a directed graph. The graph is obtained from the sectorized causal structure of the routed process channel. It is obtained by combining each pair of vertices corresponding to the same sectorial ancilla, or to the which-sector ancilla, of the same node of the supermap, while preserving the edge structure.*

The atomicity of causal relations through a routed unitary transformation then implies the following proposition, which justifies our characterisation of the sectorized causal structure as a fine-graining of the unsectorized one.

Proposition 4.1. *The unsectorized causal structure of a routed unitary supermap is obtained from its sectorized structure by combining the vertices that correspond to the same node or input/output system of the supermap.*

The distinction between the unsectorized and sectorized causal structures of a unitary supermap will be crucial for our analysis of alleged




Figure 5: The routed process matrix with a fine-grained factorization of its input and output spaces. Direct causes between the sectorized relata of \mathcal{S} are defined in terms of no-influence relations between algebras over systems associated with the input and output wires in this circuit.

implementations of the switch. As we will soon argue, roughly speaking, the experiments only exhibit indefinite causal order at the unsectorized level.

First though, we mention the time-symmetric nature of the causal structure of routed unitary supermaps. We can define the time-reversed version \mathcal{S}^\dagger of a routed unitary supermap \mathcal{S} as the supermap satisfying

$$\mathcal{S}^\dagger(\mathcal{U}_1^\dagger, \dots, \mathcal{U}_n^\dagger) = \mathcal{S}(\mathcal{U}_1, \dots, \mathcal{U}_n)^\dagger \quad (33)$$

for any set of unitary channels (with ancillas) $\{\mathcal{U}_k\}_k$ at the nodes. The routed process channel of \mathcal{S}^\dagger is just the adjoint of the routed process channel for \mathcal{S} , up to local unitary transformations applied to its unsectorized input/output subsystems. As explained in Appendix H, this implies the following.

Proposition 4.2. *The causal structure of the time-reverse of a routed unitary supermap is obtained from that of the original supermap by reversing the direction of the arrows.*

5 An application: the causal structure of purported implementations of the switch

Lately, the quantum switch [23, 2] has received considerable theoretical [3, 6, 8, 2, 4, 62, 63] and experimental [16, 14, 15, 17, 19, 18] attention. In particular, there has been an interesting debate [54, 14, 6] about whether recent purported implementations of the quantum switch realize indefinite causal order.

In this section and the next one, we apply our framework to offer a fresh perspective on the debate. We provide a new model for the experiments called the *routed switch*, which we argue better captures their causal structure than the standard quantum switch does. Our analysis allows us to conclude that there is a distinction to be made between a ‘weak’ and a ‘strong’ realization of indefinite causal order, with the current experiments [14, 15, 17] only realizing indefinite causal order in the weak sense. Importantly, to reach this conclusion, we do not need the assumption that causal relations should be defined as obtaining between objects that are localized in space and time. In fact, we argue using assumptions that were originally taken to motivate the view that the experiments do realize indefinite causal order, and show that these assumptions only point to a weak realization. Thus, we push forward the debate by arguing that the conflicting *a priori* assumptions usually thought to lead to different conclusions actually converge on the same one: that the experiments are not strong realizations of indefinite causal order.

5.1 The experiments and the switch

Let us describe the relevant experiments. The ones [14, 15, 16, 17] that have excited the most controversy involve remarkably simple optical setups. The experimental setup common to [14, 15, 17] are depicted schematically in Figure 6. There are two sets of waveplates, each designed to implement a unitary transformation on the polarization degree of freedom of any photon passing through it. We call these transformations U_A and U_B , corresponding to the hypothetical agents Al-




Figure 6: The photonic experiments.

ice and Bob respectively.¹⁹

At the start of the experiment, a single photon is fired through a 50/50 beamsplitter, which sends it into a superposition of two paths. On the red path in the figure, the photon passes through Alice’s waveplates first, and then through Bob’s. But on the blue path, the reverse is true: the photon passes through Bob’s waveplates first, and then Alice’s. The result is that the order in which the unitaries are applied to the photon’s polarization is *coherently controlled* on its spatial degree of freedom. More concretely, the unitary transformation on the overall Hilbert space $\mathcal{H}^{\text{path}} \otimes \mathcal{H}^{\text{polarization}}$ is given by the following function of U_A and U_B :

$$\begin{aligned} \text{SWITCH}(U_A, U_B) := & |\text{red}\rangle \langle \text{red}| \otimes U_B U_A \\ & + |\text{blue}\rangle \langle \text{blue}| \otimes U_A U_B \end{aligned} \quad (34)$$

In other words, the overall unitary transformation applied to the photon is precisely what one gets by modelling the experiment using the quantum switch supermap [2]. The isometric transformation applied to its polarization is obtained by inserting the state $\frac{|\text{red}\rangle + |\text{blue}\rangle}{\sqrt{2}}$ to the left factor of $\text{SWITCH}(U_A, U_B)$.

At the end of the experiment, the two paths of the photons are recombined by another beamsplitter so that tests can be performed. In particular, causal witness tests [4] are performed to verify that any bipartite unrouted supermap that accurately models the experiment must have a cyclic causal structure.

5.2 The dynamical view vs. the spatiotemporal view

The question is whether any of this implies that the experiments are physical realizations of indefinite causal order. We will reconstruct two opposing perspectives from the literature, which we call the *dynamical view* and the *spatiotemporal view*.²⁰ At the heart of the disagreement is how

¹⁹In fact, the experiment in [17] involves two copies of this setup, where the ingoing photons are prepared in an entangled state. Accounting for this in the analysis does not lead to any substantially different conclusions.

²⁰Note that, in providing this reconstruction, our primary objective is to articulate what seem to us to be the most plausible detailed versions of the perspectives and

one should carve up the experiment into different *causal relata*; that is, into different objects that bear the causal relations.

The *dynamical view* is that the experiments realize indefinite causal order between two relata, one corresponding to Alice’s transformation, U_A , and the other to Bob’s transformation, U_B . The intuition is depicted on the left side of Figure 7. More specifically, the dynamicist assumes that a causal relatum is²¹ a single ‘use’ of a channel, where ‘use’ is somehow defined ‘operationally’ [12, 14, 54]. The challenge is then to say exactly just what this operational definition of the use of a channel could be. Unfortunately, this is far from straightforward in experiments that appear to realize indefinite causal order.

One approach defines the number of ‘uses’ as the logical value of a hypothetical, auxiliary ‘flag’ system affixed to the physical device responsible for implementing the channel [12]. The device is hypothetically modified such that every time the channel is used, it raises the flag’s logical value by one. Now, this definition might appear circular, since the definition of the modification seems to beg the question of when the channel is used. But the idea is that we already know how to define channel use in the simple case of experiments with definite causal order, in which the channel can be considered as acting at a definite time. If we define the modification so that the flag correctly tracks the channel uses in these cases, then we more or less fix its behaviour in the cases where there is indefinite causal order.²²

arguments, in order to make progress in the debate. Our primary goal is not, therefore, to achieve perfect fidelity to the actual positions defended by the authors that inspired them. Although we think we have not strayed much from those actual positions, we leave it up to the reader to decide whether our eventual criticisms apply to those actual positions, or just to our particular attempt to reconstruct them in detail.

²¹In fact, the dynamical view admits a number of possible variations, in which the causal relata are different. For example, they could be the sets of possible channels Alice and Bob could implement, or the (time-delocalized) systems that their transformations act on. Since these alternative relata are in one-to-one correspondence with the channel uses from the version of the view that we discuss, we have still have a total of two relata. And, by the same arguments we discuss below for the case of channel uses, these two relata can be held to give rise to the same causal relations and structures.

²²There is some residual ambiguity in the flag’s behaviour depending on (i) whether the raising of the flag

We then take its behaviour in these conceptually less transparent scenarios to define the notion of channel use there. Hence the flag definition of channel use provides a well-defined way of extrapolating our intuition from definite-order scenarios to indefinite-order scenarios.

This definition is ‘operational’ in the sense that the flag could be constructed, and its behaviour could be observed. However, one might hope that the definition of channel use could be made operational in another, perhaps stronger, sense. In an experiment that can be described by a standard, acyclic quantum circuit, the circuit model makes clear what the operational significance of each channel is. That is, it tells the modeller which measurements one could perform to learn which channel was performed at a given point of circuit, or, conversely, what implications a given channel has for various possible measurements. One might hope for a similar characterization of the operational significance of channels used in an indefinite causal order, even though they cannot be represented as elements of a standard, acyclic circuit.

Considerable progress has been made in this direction in the bipartite and tripartite cases [54, 64]. A key idea is that we can represent an experiment with a circuit that has a feedback loop; that is, a circuit with one of its outputs plugged into one of its inputs. (And we can relate this cyclic circuit to the standard, acyclic circuit model via a change in variables.) We can then take the individual mathematical transformations that build up the circuit to represent single uses of channels. This cyclic circuit can also be used to characterize the input and output subsystems of the transformations, and determine precisely which measurements tell us about these subsystems. We can thereby precisely characterize the operational significance of the channels themselves. We can then think of the use of channel as defined by the operational implications that we have carefully laid out.

Such operationally defined channels can act on subsystems that are delocalised in time, and hence the channels themselves can be considered delocalized – this is indeed the case for Alice’s

is coherently or incoherently controlled on the use of the channel, and (ii) what the relative phase is if it is coherently controlled. But the precise definition one chooses does not matter much for the following discussion.

and Bob’s channels in the photonic experiments.

Once one assumes that the relata of causal relations are Alice’s and Bob’s operationally defined, time-delocalized channels, it seems fairly easy to argue that the experiments realize indefinite causal order. For these channels are held to serve as inputs to the quantum switch, upon which they form an accurate model of the experiments. Applying the framework for analysing the causal structure of standard supermaps sketched in Section 4, one can show that the quantum switch has the cyclic causal structure on the right side of Figure 3 [8]. It seems to follow that the experiments themselves have a cyclic causal structure.

In the actual experiments, a full tomography of the supermap is not performed, so it is not explicitly verified that the switch is a completely accurate model for the experiments (even if there are good theoretical reasons for believing as much). However, causal witness tests [4] are performed, which verify that *if* any bipartite unrouted supermap accurately models the experiment, *then* that supermap must have a cyclic causal structure. So the dynamicist can argue that if the relata are two time-delocalized channels, the empirical data strongly suggests that there is indefinite causal order.

We can now precisely formulate the dynamical point of view in argument form.

1. The experiment consists of two operationally defined, time-delocalized channels.
2. These channels serve as the relata of causal relations.
3. Moreover, these channels should be thought of as inputs to some bipartite unrouted supermap that accurately models the experiment.
4. The causal structure of the experiment is given by the causal structure of this supermap.
5. Since the experiments certify a causal witness tests, it follows that they realize indefinite causal order.

(1) and (2) are the central motivating assumptions of the dynamical view. (3) and (4) are auxiliary assumptions needed to make its conclusion (5) follow.



Figure 7: Intuitive depiction of the causal structure associated with the two- and four-relata views of alleged implementations of the switch. On the left-hand side, there is one relatum per gate, and each term in the superposition is associated with one of the arrows. On the right-hand side, there are two relata per gate, one for each time of implementation, and each term in the superposition is associated with one of the two islands in the graph.

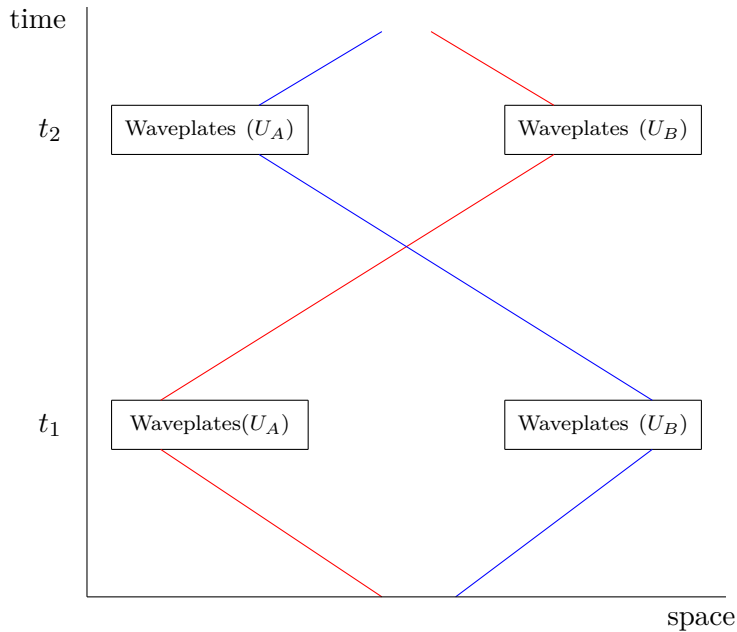


Figure 8: The spatiotemporal structure of the photonic experiments.

Let us now turn to a reconstruction of the rival perspective. The *spatiotemporal view* [6] is that the causal relata should be understood as spacetime events corresponding to the implementation of channels, so that they form a *definite* causal order.²³ In the experiments from [14, 15, 17], the photon enters Alice’s gate at a different spacetime location depending on which of the superposed paths it is on, and likewise for Bob’s, as illustrated in Figure 8. So, even if we imagine that only two channels are implemented, each of them is implemented at two different spacetime events, which we call *implementation events*. The four implementation events are taken to be the relata of causal relations. Assuming causal influences can only propagate inside the forward lightcone, it is not hard to see that this implies a definite causal order, as sketched intuitively on the right-hand side of Figure 7. Note that the spatiotemporalist can consistently agree with the dynamicist on assumptions (1) and (3), while nevertheless rejecting (2) and (4).

The dynamicist might argue that the different implementation events have no operational meaning [6]. One attempt to rebut this is as follows. In principle, one could always check with a photon detector which of the two spacetime locations Alice’s (or Bob’s) channel was implemented at [6]. Thus the two events are operationally meaningful, in the sense that they correspond to the two possible outcomes of this measurement.

However, the dynamicist can object [14] that performing this measurement invariably destroys the interference between the paths. But it is only this interference that makes the causal witness tests that certify indefinite causal order possible. If the phenomenon that we are interested in is defined by the possibility of making this certification, it follows that it is impossible to measure which location the channel was implemented at *without destroying the phenomenon of interest*. Conversely, assuming that the phenomenon of interest actually takes place, it is impossible to perform this measurement. By assum-

²³There exist variants of the spatiotemporal view in which the relata are not these spacetime events, but rather other objects in one-to-one correspondence with them. The approach in [65, 66], which takes signalling relations between localized systems as the most fundamental sort of causal relation, arguably falls into this category. The following discussion also applies to these variants (c.f. note 21).

ing that (a) this measurement is the only thing that could make the implementation events operationally meaningful, and (b) that measurements do not make something operationally meaningful if they destroy the phenomenon of interest, the dynamicist concludes that the implementation events are operationally meaningless.

But (a) is false: the operational content of the implementation events is not exhausted by measurements to find out which implementation event ‘happened’ [6]. In particular, [6] proposes an experiment involving erasure in which the interference is not ultimately destroyed, and whose outcome is different depending on whether there are four implementation events or two (as in a gravitational realization of the switch [63]). Also, as we will argue below, the multiplicity of implementation events for each channel ensures that a larger range of operations can be performed by Alice and Bob than are typically accounted for.

Thus it cannot be maintained that the implementation events have no operational meaning. But this does not prove the spatiotemporal view right, since it does not imply that the implementation events should be regarded as causal relata. After all, the dynamicist’s operationally-defined channels also have operational meaning, so we have not demonstrated any sense in which the spatiotemporal view is superior. While the two channels might each correspond to two operationally meaningful implementation events, the dynamicist can still legitimately maintain that the causal relata are those two channels. In that case, there are still only two relata, and they still exhibit an indefinite causal order.

We appear to have reached an impasse: unable to settle the dispute by appeal to operational meaningfulness, the dynamicist falls back on the *a priori* assumption that causal relata correspond to channels, while the spatiotemporalist falls back on the *a priori* assumption that causal relata are spacetime events. The dynamicist’s assumption seems to suggest there is indefinite causal order; the spatiotemporalist’s assumption implies there is no such thing.

Here, we aim to break this impasse. We will argue that the dynamicist’s central assumptions (1) and (2) – that there are two time-delocalized channels, and that these serve as causal relata – do not in fact suggest that the experiments [14, 15, 17] should be understood as realizations

of indefinite causal order in any strong sense. By the same token, we will show that the spatiotemporalist’s central assumption – that the causal relata should be localized in spacetime – needs not be made in order to argue that the experiments are not strong realizations of indefinite causal order. Therefore, whichever starting assumptions one prefers, one should at most view the experiments as weak realizations of indefinite causal order. We will make our case by motivating and analyzing a new model for the experiment called the *routed switch*, to which we now turn.

5.3 The routed switch and its causal structure

One of our most important claims is the following. The quantum switch is both

- a) an accurate model of the actual transformations experienced by the photon in the experiments; and
- b) a bad supermap for reasoning about the causal structure of those experiments.

How can this be? If the switch tells an accurate story about the transformations that take place, why should we not trust what it says about causal structure? The answer is to be found in the intuition for causal influence sketched in Section 3. In a nutshell: causal influence is about how information *can* flow, and not just about how it does flow. Therefore, one must take into account which transformations are *possible*, rather than just those transformations that actually took place. This idea was reinforced by the definitions and examples of influence between sector-preserving channels from Section 3, where it was shown that the causal relations one attributes to a scenario can change due to a restriction on the set of transformations one calls ‘possible’, even if the transformations that were actually performed are the same in both cases.

Moreover, to the extent that ‘indefinite causal order’ is a concept of the foundations of physics, the relevant notion of possibility here must be something like *physical possibility*. Roughly, we should be concerned with what transformations could take place in keeping with the relevant physical laws, rather than with what transformations could take place given our choices, or given the particular tools we are using, or given our current level of technological sophistication.

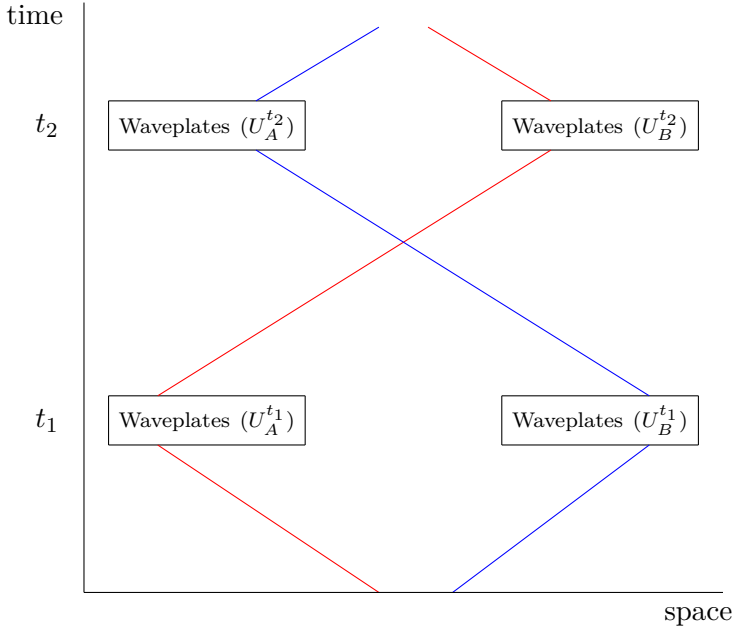


Figure 9: More general physically possible transformations, as implied by the spatiotemporal structure of the experiments.

And it turns out that the standard switch misses a wide range of physically possible transformations. As the spacetime diagram in Figure 8 makes clear, it is physically possible in the experiments from [14, 15, 17] for Alice to vary her transformation depending on the logical value of the control qubit. All she has to do is reach in after the first implementation event (when the photon enters her gate on the first path), and change her gate before the second implementation event (when the photon enters her gate on the second path). Bob could do the same thing, leading to the more general scenario depicted in Figure 9.

This possibility turns out to have important consequences for causal structure. Therefore, we generalize the switch to provide a model of the scenario described in Figure 9. We want to allow transformations of the following form, where the agents can coherently control their choice of unitary based on the path.

$$\begin{aligned} U_A &= |\text{red}\rangle \langle \text{red}| \otimes U_{t_1}^A + |\text{blue}\rangle \langle \text{blue}| \otimes U_{t_2}^A \\ U_B &= |\text{red}\rangle \langle \text{red}| \otimes U_{t_2}^B + |\text{blue}\rangle \langle \text{blue}| \otimes U_{t_1}^B \end{aligned} \quad (35)$$

(The subscripts refer to the time coordinates $t_1 < t_2$ of the different implementation events.) We do not, on the other hand, want to consider transformations that map the sector of $\mathcal{H}^{\text{path}} \otimes \mathcal{H}^{\text{polarization}}$ containing $|\text{red}\rangle$ to the one contain-

ing $|\text{blue}\rangle$ or vice versa, since it is not physically possible, or even comprehensible, for the agents to transform one path of the photon into another. So we only want to consider the transformations of the form (35); that is, the sector-preserving transformations in the sectorization induced by the path. We allow these to serve as the inputs to a routed supermap called the routed switch, which is defined as follows.²⁴

$$\begin{aligned} \text{RSWITCH}(U_A, U_B) &= |\text{red}\rangle \langle \text{red}| \otimes U_{t_2}^B U_{t_1}^A \\ &\quad + |\text{blue}\rangle \langle \text{blue}| \otimes U_{t_2}^A U_{t_1}^B \end{aligned} \quad (36)$$

A circuit decomposition of the routed switch is provided in Figure 12. Note that if we make the agents lose access to the control qubit, by requiring that $U_{t_1}^A = U_{t_2}^A$ and $U_{t_1}^B = U_{t_2}^B$, then we recover the standard switch, so the routed switch is strictly more general than the standard switch. Note also that the routed switch can in turn be recovered from even more detailed models underlying the standard switch, accounting for even more physical possibilities, such as its ‘causal box’ [67, 65, 66].²⁵

We now have a model of the experiment that accounts for a more appropriately broad range of physically possible transformations than the standard switch, which therefore tells a more trustworthy story about the causal structure of the experiments. So what is this story? Since the routed switch is a routed supermap, we can apply our framework to analyse causal structure at both

²⁴More generally, the routed switch acts on unitaries with ancillas as $\text{RSWITCH}(U_A, U_B) = |\text{red}\rangle \langle \text{red}| \otimes (I_X \otimes U_{t_2}^{BY})(U_{t_1}^{AX} \otimes I_Y) + |\text{blue}\rangle \langle \text{blue}| \otimes (U_{t_2}^{AX} \otimes I_Y)(I_X \otimes U_{t_1}^{BX})$. This expression also allows us to calculate the action of the routed switch on arbitrary channels via their stinespring dilation.

²⁵Should we then be using one of these more detailed models, since we have just emphasised the importance in accounting for physical possibilities? There are two reasons we stop at the routed switch. Firstly, the routed switch can be recovered from the more detailed models by assuming that agents cannot create or destroy particles, which is appropriate if one wants to think about the causal structure of the scenario as the possibilities for signalling afforded by the *particular* photon that is fired through the beamsplitter at the start of the experiment. Secondly, as we will discuss below, the fact that the routed switch still takes time-delocalized channels as inputs makes its use compatible with the dynamicist’s central assumptions, which in turn makes the conclusions of our analysis more significant.

the unsectorized and sectorized levels. The unsectorized causal structure, depicted in Figure 10, is identical to the causal structure of the standard switch uncovered in [8], which has an indefinite causal order, and is precisely the causal structure that the dynamicist is likely to attribute the scenario. However, the sectorized causal structure, depicted in Figure 11, does admit a definite causal order.

So we only have indefinite causal order at the coarse-grained, unsectorized level, that vanishes when we perform the sectorization. It is worth capturing this property with a general definition.

Definition 5.1. *A routed supermap exhibits **strong indefinite causal order** if both its sectorized and unsectorized causal structures are cyclic. It exhibits **weak indefinite causal order** if its unsectorized structure is cyclic, but its sectorized structure is acyclic.*

Recall that, by Proposition 4.1, the sectorized causal structure fixes the unsectorized causal structure. More precisely, the unsectorized directed graph is obtained by combining the sectorized vertices corresponding to the same unsectorized vertex. Therefore, if a routed supermap exhibits weak indefinite causal order, then its indefinite causal order can be ‘explained away’ as the result of coarse-graining a more detailed structure in which there is no indefinite causal order. On the other hand, the indefiniteness of a routed supermap with strong indefinite causal order admits no such explanation, and accordingly must be regarded as a brute property of its causal structure.

The routed switch exhibits weak indefinite causal order. One can easily check that Figure 10 is obtained from Figure 11 by combining the A nodes and combining all the B nodes, and thus one can understand the indefiniteness of the supermap as being entirely a consequence of this coarse-graining.

5.4 So, are the experiments physical realizations of indefinite causal order?

In the context of the experiments, the inputs to the routed switch are two time-delocalized channels, which admit an operational definition in the sense of [12] or [54]. Therefore, we claim that using the routed switch as a model for the experiments and their causal structure is consistent with the central assumptions (1) and (2) of the




Figure 10: The unsectorized causal structure of the routed switch. This is the same as the causal structure of the standard switch (Figure 3). The presence of a cycle indicates that A and B cannot be embedded into a definite causal order.




Figure 11: The sectorized causal structure of the routed switch. The directed path $P \rightarrow A_{\text{red}} \rightarrow B_{\text{red}} \rightarrow F$ corresponds to the branch where Alice’s intervention comes first; $P \rightarrow A_{\text{blue}} \rightarrow B_{\text{blue}} \rightarrow F$ corresponds to the branch where Bob’s comes first. As the lack of cycles indicates, once we break down the causal influences into those associated with each sector and the which-sector information, there is no need to ascribe an indefinite causal order to the routed switch.




Figure 12: The routed switch. This is a bipartite supermap, in which we have factorized each node's ingoing and outgoing spaces into a sectorized control qubit and target system, e.g. $A_I^i = (C_A)_I^i \otimes (T_A)_I$. Thus each party has access to a sectorized copy of the control qubit. The sectorization is into one-dimensional sectors spanned by logical states, so the index matching on $(C_A)_I^i$ and $(C_A)_O^i$ restricts Alice to transformations of the form 35. (δ, W) is the superposition-of-paths routed unitary from (10) (where we have relabelled $P \rightarrow T$). The wires bent into 'cup' or 'cap' shapes formally represent (the unnormalised) perfectly correlated entangled kets and bras respectively. When $i = 0$, the particle exits through the left side of each W . This leads to the particle passing through Alice's node first, and then Bob's. The opposite is true for $i = 1$. Inserting the transformations (35) recovers the transformation (36).

dynamical view; namely, that there are two such channels and that they serve as causal relata (we will defend this claim in Section 6.2).

And yet, the routed switch exhibits the formal property of weak indefinite causal order. Since we take the routed switch to provide a good model for the causal structure of the experiments, and, in particular, a better one than the standard switch, we take this to suggest that the experiments themselves are physical realizations of indefinite causal order only in the weak sense. If this is right, then it seems that the dynamicist's central assumptions (1) and (2) do not actually suggest the experiments are realizations of indefinite causal order in anything more than in a weak sense. By the same token, one needs not assume *a priori* that causal relations obtain between objects that are localized in spacetime in order to argue that indefinite causal order is not strongly realized. We conclude that there are no compelling *a priori* assumptions from which to argue that indefinite causal order is realized in more than a weak sense – or at least, if there are, they have not yet been clearly articulated in the literature.

Now, if all this is correct, then there must be something wrong with the dynamicist's argument. The issue is with propositions (3) and (4). The fact that some time-delocalized channels and a supermap acting on them accurately model the transformations that take place in an experiment does not imply that they provide a useful model for the causal structure of the experiment. This is because they may still not account for an appropriately broad range of physically possible transformations, or of possibilities for information flow. Hence the combination of (3) and (4) should be rejected, and (5) is unsupported.

Essentially, the issue with the dynamicist's argument is that just because an experiment *implements a supermap*, it does not automatically follow that it *realizes its causal structure*. And this is because the former just requires that the supermap models the actual transformations, while the latter requires that it also models the non-actualized but possible transformations. For future discussions of the causal structure of an experiment, we therefore recommend that the questions of whether a supermap is implemented and whether its causal structure is realized should always be carefully distinguished.

We note that our arguments do not establish

that all (potential) implementations of the switch other than [14, 15, 17] do not realize strong indefinite causal order. In particular, for an implementation in which it is physically impossible for the agents to access to the control qubit, there would be no reason to model it using the routed switch.

A gravitational implementation of the switch, such as the one proposed in [63], would probably fall into this category. Another example is the photonic experiment in [16], in which Alice and Bob cannot change their actions depending on the order in a straightforward way. This reason for this is that the photon there has a very long coherence length. Hence there are no longer two implementation events for each agent, but rather elongated implementation *regions* of spacetime, which are almost entirely overlapping. So Alice and Bob can no longer implement the routed switch by controlling their intervention on two distinct times of implementation. We recommend that any future debate over whether indefinite causal order has been realized yet should focus on this more challenging example.

6 Possible objections to our analysis

Here, we consider a number of possible objections to the analysis from the previous section. In particular, we discuss the criticisms that (i) our analysis only goes through relative to a particular notion of causal structure; (ii) we are wrong to claim that our approach is consistent with the central assumptions of the dynamical view; and (iii) there is something wrong with our formal definition of weak indefinite causal order. We will take each of these in turn.

6.1 ‘The analysis relied on a particular view of causal structure’

The first objection we consider is that our conclusion that the experiments are weak realizations of indefinite causal order is only true relative to our assumption that causal structure is to be understood as possibilities for information flow, as captured by some appropriate routed supermap and our framework for analysing its causal structure. There are other legitimate notions of causal structure, such as the possibilities for signalling between implementation events, or the partial or-

der associated with a Minkowski manifold, for which indefinite causal order is not even weakly realized by the experiment [6].

We are happy to accept this point. We chose our notion because it is information-theoretic; because it makes causal structure of a physical scenario intrinsic to its quantum description; because it enables us to describe causal structures precisely and in great detail; and because it makes it seem initially plausible that indefinite causal order could be realized on a Minkowski spacetime. Nevertheless, undoubtedly other notions of causal structure that are more spatiotemporal and less information-theoretic than ours are also interesting and important in their own right, and they do indeed lead to a more forceful rejection of indefinite causal order. So we are happy for our conclusion that the ‘causal structure’ of the experiments exhibits only weak indefinite causal order, rather than definite causal order, to be understood as true relative to our particular notion of causal structure.

However, our conclusion that the experiments are at most weak, rather than strong, realizations of indefinite causal order seems to be more robust. For our analysis is compatible with the *dynamicist’s* notion of causal structure. Hence the analysis also suggests that the experiments are not strong realizations of indefinite causal order according to *any* mainstream notion of causal structure.

6.2 ‘The analysis is not actually compatible with the dynamicist’s central assumptions’

Another possible criticism is that we were wrong to claim that our analysis was compatible with the central assumptions of the dynamical view. One way of defending this would be to point out that, although we do understand the experiment in terms of two time-delocalized channels (as per assumption (1)), we ended up considering a broader set of possible channels than those considered by the dynamicist. However, as we have argued, this extra generality is needed in order to account for the physical possibilities for information flow, and hence to properly capture the causal structure. So if the central assumptions of the dynamical view prohibit this extra generality, then we think they are wrong; while if they do not, then we think our analysis is consistent with them.

Another way to make this criticism would be to argue that our consideration of the individual sectors as relata in the sectorized causal structure can only be motivated by the spatiotemporalist's assumption that causal relata should be localized in spacetime. After all, we did appeal to the existence of two distinct implementation events in order to justify our use of the routed switch, in which the sectorial constraints apply.

However, this is a misunderstanding of the argument. We were led to consider the sector-preserving channels (35) on the basis that they account for a broader range of physically possible transformations than those accounted for by the switch. It is true that the physical possibility of these transformations is the *consequence* of spatiotemporal considerations, but spatiotemporal considerations did not motivate our desire to account for them. It is their physical possibility alone that motivated us to consider the full range of sector-preserving channels. Indeed, if their physical possibility were instead grounded in other, non-spatiotemporal considerations, then we would still be obliged to account for them.

This shows that our consideration of the sector-preserving channels and the routed switch is not motivated by any assumption of localization; it remains to show that our choice of causal relata does not rely on this either. To see this, note that, as a routed supermap, the routed switch comes with a distinct sectorized causal structure which, by Proposition 4.1, is strictly more detailed than its unsectorized counterpart. We consider this structure, and its relations between individual sectors and which-sector informations, because they help us understand the causal structure at a deeper level – in particular, how features of the unsectorized causal structure, such as cyclicity, come about in terms of the more detailed description. But the fact that the sectorized relata provide this understanding does not depend on them being localized, and indeed the which-sector informations are actually *delocalized*. So even though the individual sectors are localized, our use of them as causal relata does not depend on any assumption that causal relata must be localized.

Another spin on this criticism would be to say that we should have characterized the dynamicist's assumption (2) as requiring that the time-delocalized channels are the *only* causal relata, in

which case our analysis would not be consistent with it. If the assumption were so modified, then we would indeed reject it, since, as we have just argued, it is important to account for physically possible transformations, and doing so motivates the sectorized relata.

But perhaps there is still scope for the dynamicist to argue that there is something wrong or, at least misleading, about understanding the causal structure in terms of our sectorized relations and relata. This is the topic of the next subsection.

6.3 'An appropriate physical notion of weak indefinite causal order has not been successfully formalized'

We have relied on a formal notion that we chose to call 'weak indefinite causal order'. But perhaps the dynamicist could argue that our technical definition of this term fails to formalize any pre-formal, intuitive idea that its name might suggest. In that case, the routed switch exhibiting weak indefinite causal order in our formal sense would not obviously imply that the experiments are weak realizations of indefinite causal order in any interesting or relevant sense.

For example, the dynamicist could argue that our sectorized causal relations and relata are somehow illegitimate, or at least should be considered as conceptually secondary to the unsectorized relations and relata. Then the lack of cycles in the sectorized picture would arguably not much weaken the sense in which indefinite causal order is obtained between the time-delocalized channels in the experiments.

One way to argue that the sectors are not legitimate relata is to argue that they have no operational meaning. But our response to this charge can be similar to the spatiotemporalist's, discussed in the previous section. Alternatively, the dynamicist might deny that we can attribute our sectorized causal relations to the phenomenon of interest. For example, the arrow from A_{red} to B_{red} in Figure 11 means that the wire for the red sector on Alice's exchange gate signals to the wire for the red sector on Bob's exchange gate in the routed process channel for the routed switch. But if this is done experimentally, the Bob receiving information from Alice destroys the interference that makes the certification of indefinite causal order via a causal witness test possible. (This is because we end up effectively post-selecting on

the sector in which the control qubit's state is $|0\rangle$.) Thus this alleged causal relation might be held to have no operational meaning in the context of the phenomenon of interest.

However, this argument backfires, since it would also prevent us from attributing the causal relations between the channels that lead the dynamicist to see the experiments as exhibiting indefinite causal order. For example, the arrow from A to B in the causal structure of the switch means that Alice can signal to Bob in the process channel. But operationally witnessing this signalling would again destroy the interference. So the above argument would establish that this causal relation has no operational meaning, in which case we do not have a cyclic causal structure, and therefore we have a definite causal order. As long as one accepts our idea of causal relations as signalling relations through unitaries, attributing an indefinite causal order to this experiment necessarily involves considering possibilities for signalling which are not realized when the phenomenon of interest takes place.

Of course, one could reject this idea of causal relations, and instead simply define 'indefinite causal order' as the ability to pass a causal witness test. But since the routed switch can pass such a test, along with processes with a definite causal order and more parties, this does not seem well-motivated. There is only a justification for interpreting a causal witness test as an indicator of indefinite causal order when one also assumes the correct model for the experiment is an un-routed supermap with a certain number of parties.

Another possible response from the dynamicist is that our sectorized causal structures are somehow problematic because an individual sector isn't the sort of thing that can bear causal relations. The argument might run as follows. When we partition a Hilbert space into different sectors, we are not thereby dividing a system into constituent parts. Instead, we are considering different, mutually exclusive ways that system could be. But a way that a system might or might not be should not qualify as a relatum of causal relations. Rather, the causal relata in experiment should be the constituent parts of the experiment that are 'present' no matter what.

But this argument should be rejected, since it rules out a causal analysis of a classical version of

the experiments which seems obviously reasonable. Imagine a classical version of the routed switch that incoherently controls the functions $f_{t_2}^B \circ f_{t_1}^A$ and $f_{t_2}^A \circ f_{t_1}^B$ on a classical bit. We could perform an analysis similar to the one we have given for the quantum case, by using a classical version of the exchange gate. This analysis would also conclude that there is an indefinite causal order between A and B , but that this disappears when we define more fine-grained causal relata via a sectorization into subsets of the classical variables.

This analysis seems to be a fair reflection of the physical situation. True, it is possible that A can affect B on any particular run of the experiment, and it is possible that B can affect A on any particular run of the experiment. But, crucially, these possibilities turn out to be mutually exclusive: on any particular run of the experiment, causal influence between A and B propagates in only one direction. By letting the individual sectors of A and B define causal relata, we make this obvious but important physical fact clear at the formal level, by showing that their causal structure is acyclic.

Far from being an unreasonable thing to do, here breaking down the causal relata into different sectors helps us to formally articulate the key conceptual points about the causal structure of the experiment: that there is indefinite causal order between A and B , but only in a weak sense since on any particular run of the experiment causal influence passes only in one direction between them. But if it were true that the different ways a system could be should never define causal relata, then we would not be able to perform this useful analysis. From this, we conclude that if there is a reason for forbidding the quantum sectors from defining causal relata, it is not because they correspond to mutually exclusive possibilities.

A better argument against performing the sectorization is based on the assumption that the quantum sectors, unlike the classical ones, are *not* just mutually exclusive ways a system can be. This is evidenced by the fact that the relative phase between the sectors can be detected with a single measurement (and in fact is when indefinite causal order is verified with a causal witness test). But this relative phase is not associated within any particular sector – it is a *global*

property of the sectors. From this, it is tempting to conclude that all sectors – and the information they contain – are somehow ‘present’ on each run of the experiment. Applying this assumption to the photonic experiments, it seems we really do have causal influence in both directions between the gates on a single run, unlike in the classical version of the experiment.

Now, this might seem to bring us towards thinking about sectors as constituent parts of a system, like we think of tensor factors. And there is surely no problem with splitting up causal relata with the parts (commonly called ‘subsystems’) associated with tensor factors. This begs the question of why we can’t split up the causal relata into different parts associated with sectors, as we have in our analysis. But the dynamicist might want to argue that they cannot be understood as *independent* parts, so that we can’t attribute to them their own causal properties (or at least, if we can then this picture should be regarded as secondary to the unsectorized one).

To support this, the dynamicist might appeal to the idea that the information in a system is not just the combination of the information within each sector. For this does not include the information in the coherence between sectors, which, by the global phase symmetry, cannot be reconstructed from the information present on each sector. Mathematically, the point is that the equation

$$\pi_A^0 |\psi\rangle \langle \psi|_A \pi_A^0 + \pi_A^1 |\psi\rangle \langle \psi|_A \pi_A^1 = |\psi\rangle \langle \psi|_A \quad (37)$$

does not hold for all $|\psi\rangle_A$, where π_A^0 projects onto a sector of A and $\pi_A^1 := I - \pi_A^0$ projects onto the complement sector. This is reminiscent of how the pure state of a bipartite system does not always separate onto a state for each factor; i.e. that

$$\text{Tr}_X(|\psi\rangle \langle \psi|_{AX}) \otimes \text{Tr}_A(|\psi\rangle \langle \psi|_{AX}) = |\psi\rangle \langle \psi|_{AX} \quad (38)$$

does not hold for all $|\psi\rangle_{AX}$.

But the *sectorial nonseparability* captured by (37) is stronger than the usual notion of nonseparability across factors in the following sense. There are many pure states, called product states, for which (38) does hold. But the only pure states for which (37) holds are those for which one of the $\pi_A^i |\psi\rangle \langle \psi|_A \pi_A^i = 0$. In other words, a pure state can only be understood as the

combination of states on different sectors in the trivial case when there is only information associated with one of them. Intuitively, whenever both sectors are ‘present’, there are relational properties between them (i.e. the coherence) that prevent us from understanding the pure state as a simple combination of them. So one could argue that, if the sectors do correspond to parts, they do not correspond to independent parts, that can be properly separated out either mathematically or conceptually. Rather, to the extent that they are all ‘present’, they are always inextricably linked, or ‘entangled’. Given that the sectors are inextricably linked, it seems at least questionable that we can attribute them their own causal properties. For this reason, the dynamicist might either reject our sectorized perspective, or at least consider it as secondary to the unsectorized perspective from which there is indefinite causal order.

However, (37) was given for pure states, and it does not generalize to mixed states. And according to our model, Alice’s and Bob’s time-delocalized systems are always each in an incoherent mixture of the two different sectors. This is because their systems are entangled, in the sense that if Alice’s system is in the t_1 sector, then Bob’s is in the t_2 sector, and vice versa. But the information associated with an incoherent mixture of two different sectors, unlike a superposition of them, can indeed be understood as the combination of the information associated with each sector.

Still, in an *improper mixture* [68, 69] (that is, an incoherent mixture arising from tracing out one half of an entangled state), even if there is no coherence to be measured between two sectors, one can still induce local phase shifts between them that end up globally changing the entangled state; so sectorial nonseparability is still relevant in this sense. But Alice *can* still do this in our model, by using the which-sector ancilla of the exchange gate. So where sectorial nonseparability is relevant, we account for it, but where we do not account for it, it is not relevant.

Moreover, we note that the argument from sectorial nonseparability requires an intuitive leap from the structure of the state space to causal structure. But the atomicity property has already shown us that such leaps are unwarranted, since state entanglement does not lead to an analogous entanglement of causal relations. The gen-

eral moral that we take from this is that causal structure in quantum theory does not always behave in the way one would naively expect from studying its state space. So even if it is surprising that there exists a natural and rigorous sense in which sectors can have their own causal relations, one's astonishment does not constitute an objection to them.

7 Discussion

7.1 Summary

In this paper, we built a framework that attributes a causal structure to scenarios where sectorial constraints are at play. We started off with the assumption that the most basic sort of causal relation is a possibility for signalling afforded by a (routed or unrouted) unitary transformation. We then proceeded to show that a mathematical formulation of this idea is equivalent to a number of other conditions, expressing distinct intuitions about what causal relations are. These conditions imply that the causal structure of a (routed or unrouted) unitary transformation can be represented as a directed acyclic graph, due to the atomicity property [5]. All of this indicates that there is a stand-out, natural notion of a causal relation in quantum theory, even when sectorial constraints are at play. This notion is at the heart of our framework for causal structure.

When systems suffer sectorial constraints, we can break down their causal relations further. The resulting sectorized causal relations are between sectors and which-sector informations in preferred partitionings of each system's Hilbert space. These give rise to the sectorized causal structure of a transformation, which can also be represented as a directed graph. The sectorized causal structure is strictly more detailed than its unsectorized counterpart, and so can be regarded as a fine-graining.

We were able to use these notions of the causal structure of a transformation to define the sectorized and unsectorized causal structures of a (routed) supermap in a fairly straightforward way. The sectorized causal structure can be of considerable use in understanding scenarios which appear to display enigmatic causal properties, such as indefinite causal order. To illustrate this, we performed a sectorized and unsectorized analysis of recent photonic experiments

that purport to implement the switch, using our more general routed switch as a model. The analysis made clear that the cycles in the unsectorized causal structure of the experiment can be understood as the result of coarse-graining the more detailed, acyclic, sectorized causal structure. This provided an argument that the experiments realize indefinite causal order only in a weak sense. Significantly, this is the first argument to this effect that did not start off by assuming that the objects that bear causal relations should correspond to local regions of spacetime. Instead, the argument is based on acknowledging the impact that the sectorial constraints have on the causal structure.

7.2 Possibility and causality

Above all, this work demonstrates the importance of a simple dictum: **possibility is essential for causality**. More concretely, the point is that the question of which transformations are possible is of central importance for any information-theoretic account of causal structure. And the set of possible transformations might be nontrivial, in that it might not include all quantum operations on each of the systems of interest.

The reason the dictum matters so much is that failing to recognise it can lead one to attribute the wrong causal structure to a scenario. In the example sketched at the start of Section 3.2, ignoring the sectorial constraints on the system D leads one to incorrectly state that it is influenced by A . A more pertinent example is provided by photonic implementations of the switch. There, the attribution of indefinite causal order can result from mistakenly thinking that just because the switch provides an accurate model of the actual transformations, its causal structure must reflect the causal structure of the experiments.

7.3 Outlook

There are plenty of avenues for further research. One pressing question is what the causal structure of a routed supermap can tell us about its compositional structure, and about the observable correlations it can generate. Future work will seek to lay out this out in detail, most likely by attempting to extend Theorems 4.10 and 8.3 from [5] to the routed case. A satisfactory answer to such questions, in addition to the contributions

of this paper, would bring us most of the way to a fully-fledged theory of quantum causal models for sectorized systems.

Given that a part of the motivation for studying causal structure is to uncover the conceptual relationship between quantum theory and general relativity, it would be a useful step to investigate how frameworks for causal structure might generalize to the field-theoretic and relativistic cases. Most importantly, one could check how or whether the most striking features of causal structure in quantum theory are manifested in these more fundamental regimes, and determine whether any conceptually novel properties emerge.

A byproduct of our study of implementations of the quantum switch was the definition of a routed supermap of particular interest, the *routed quantum switch*, which essentially models the situation where agents in a switch can coherently control their actions on the order they are in. Even though we restricted our use of it here to the discussion of a foundational question, it could also have more practical applications. In particular, as an extension of the switch it might hold similar or additional advantages for communication, computation, or thermodynamical procedures.

Another avenue is to find further physical applications of our framework. Any scenario featuring sectorial constraints and interesting causal features is likely to provide a worthwhile application; examples include the Aharonov-Bohm effect [40] (in which charge superselection appears to play a key role [41]), and the recently proposed del Santo-Dakić protocol [42, 43, 44, 45]. It will also be interesting to see whether the understanding of causal structure in the presence of sectorial constraints can help one uncover new scenarios with exotic causal properties.

It is also likely possible to articulate a more systematic connection between routed maps and superselection rules, and hence also quantum reference frames [70]. Therefore, it might be possible to make some interesting general statements about the interplay between reference frames, superselection rules, and causal structure. It may be particularly interesting to explore the precise sense in which the causal structure one observes is dependent on one's quantum reference frame.

It would be interesting to provide an adaptation of our framework to the classical case. While

a similar approach can be taken, there are some interesting differences, such as the possibility of determining all the information associated with the target input to the classical version of the exchange gate by measuring its ancillary output, and the fact that the 'which-sector?' relation in the classical case never serve as causes, as a consequence of the absence of relative phases in classical theory.

Finally, it would be of interest to generalize the arguments we gave against the strong realization of indefinite causal order in [14, 15, 17] to the photonic experiment in [16]. The obvious way of doing this involves the use of infinite-dimensional Hilbert spaces for Alice and Bob, formed by directly summing together a qubit for each time at which a branch of the photon passes through the waveplates. This goes beyond the scope of our framework as it stands, but one could provide an analysis within our framework of a version of the experiment in which the photon only passes through the waveplates at a very large, but still finite, number of times. This might provide an argument that [16] does not realize strong indefinite causal order, which may be generalized even further to a wide range of experiments involving the superposition of various paths through space-time.

Acknowledgements

It is a pleasure to thank Giulio Chiribella, Hlér Kristjánsson, Ognyan Oreshkov, Nicola Pinzani, Tein van der Lugt, Matt Wilson, and two anonymous reviewers for helpful discussions and comments. AV is indebted to the organisers and participants of the 2021 Sejny Summer Institute – in particular Timothée Hoffreumon, Natália Móller and Eleftherios Tselentis – for discussions and ideas about implementations of the switch. We also thank Alexandra Elbakyan for her help to access the scientific literature.

NO acknowledges funding from the UK Engineering and Physical Sciences Research Council (EPSRC). AV is supported by the EPSRC Centre for Doctoral Training in Controlled Quantum Dynamics. This publication was made possible through the support of the grant 61466 'The Quantum Information Structure of Space-time (QISS)' (qiss.fr) from the John Templeton Foundation. The opinions expressed in this pub-

lication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

Note Added

Independent work appeared during the completion of the first version of this manuscript which provides a related fine-grained causal analysis of implementations of the switch [71].

References

- [1] L. Hardy, “Towards quantum gravity: a framework for probabilistic theories with non-fixed causal structure,” [*Journal of Physics A: Mathematical and Theoretical* **40** no. 12, \(2007\) 3081](#), [arXiv:gr-qc/0608043](#).
- [2] G. Chiribella, G. M. D’Ariano, P. Perinotti, and B. Valiron, “Quantum computations without definite causal structure,” [*Physical Review A* **88** no. 2, \(Aug, 2013\)](#), [arXiv:0912.0195 \[quant-ph\]](#).
- [3] O. Oreshkov, F. Costa, and Č. Brukner, “Quantum correlations with no causal order,” [*Nature communications* **3** no. 1, \(2012\) 1–8](#), [arXiv:1105.4464 \[quant-ph\]](#).
- [4] M. Araújo, C. Branciard, F. Costa, A. Feix, C. Giarmatzi, and Č. Brukner, “Witnessing causal nonseparability,” [*New Journal of Physics* **17** no. 10, \(2015\) 102001](#), [arXiv:1506.03776 \[quant-ph\]](#).
- [5] J. Barrett, R. Lorenz, and O. Oreshkov, “Quantum causal models,” (2020), [arXiv:1906.10726 \[quant-ph\]](#).
- [6] N. Paunković and M. Vojinović, “Causal orders, quantum circuits and spacetime: distinguishing between definite and superposed causal orders,” [*Quantum* **4** \(2020\) 275](#), [arXiv:1905.09682 \[quant-ph\]](#).
- [7] D. Felce and V. Vedral, “Quantum refrigeration with indefinite causal order,” [*Physical Review Letters* **125** \(Aug, 2020\) 070603](#), [arXiv:2003.00794 \[quant-ph\]](#).
- [8] J. Barrett, R. Lorenz, and O. Oreshkov, “Cyclic quantum causal models,” [*Nature Communications* **12** no. 1, \(2021\) 1–15](#), [arXiv:2002.12157 \[quant-ph\]](#).
- [9] A. Kissinger and S. Uijlen, “A categorical semantics for causal structure,” [*Logical Methods in Computer Science* **Volume 15, Issue 3** \(2019\)](#), [arXiv:1701.04732 \[quant-ph\]](#).
- [10] R. Lorenz and J. Barrett, “Causal and compositional structure of unitary transformations,” [*Quantum* **5** \(2021\) 511](#), [arXiv:2001.07774 \[quant-ph\]](#).
- [11] C. Branciard, M. Araújo, A. Feix, F. Costa, and Č. Brukner, “The simplest causal inequalities and their violation,” [*New Journal of Physics* **18** no. 1, \(2015\) 013008](#), [arXiv:1508.01704 \[quant-ph\]](#).
- [12] M. Araújo, F. Costa, and i. c. v. Brukner, “Computational advantage from quantum-controlled ordering of gates,” [*Physical Review Letters* **113** \(Dec, 2014\) 250402](#), [arXiv:1401.8127 \[quant-ph\]](#).
- [13] D. Felce, N. T. Vidal, V. Vedral, and E. O. Dias, “Indefinite causal orders from superpositions in time,” [*Physical Review A* **105** no. 6, \(2022\) 062216](#), [arXiv:2107.08076 \[quant-ph\]](#).
- [14] L. M. Procopio, A. Moqanaki, M. Araújo, F. Costa, I. A. Calafell, E. G. Dowd, D. R. Hamel, L. A. Rozema, Č. Brukner, and P. Walther, “Experimental superposition of orders of quantum gates,” [*Nature communications* **6** no. 1, \(2015\) 1–6](#), [arXiv:1412.4006 \[quant-ph\]](#).
- [15] G. Rubino, L. A. Rozema, A. Feix, M. Araújo, J. M. Zeuner, L. M. Procopio, Č. Brukner, and P. Walther, “Experimental verification of an indefinite causal order,” [*Science advances* **3** no. 3, \(2017\) e1602589](#), [arXiv:1608.01683 \[quant-ph\]](#).
- [16] K. Goswami, C. Giarmatzi, M. Kewming, F. Costa, C. Branciard, J. Romero, and A. G. White, “Indefinite causal order in a quantum switch,” [*Physical review letters* **121** no. 9, \(2018\) 090503](#), [arXiv:1803.04302 \[quant-ph\]](#).
- [17] G. Rubino, L. A. Rozema, F. Massa, M. Araújo, M. Zych, v. Brukner, and P. Walther, “Experimental entanglement of

- temporal order,” [Quantum](#) **6** (2022) 621, [arXiv:1712.06884 \[quant-ph\]](#).
- [18] X. Nie, X. Zhu, C. Xi, X. Long, Z. Lin, Y. Tian, C. Qiu, X. Yang, Y. Dong, J. Li, T. Xin, and D. Lu, “Experimental realization of a quantum refrigerator driven by indefinite causal orders,” [Physical Review Letters](#) **129** no. 10, (2022) 100603, [arXiv:2011.12580 \[quant-ph\]](#).
- [19] H. Cao, N.-n. Wang, Z.-A. Jia, C. Zhang, Y. Guo, B.-H. Liu, Y.-F. Huang, C.-F. Li, and G.-C. Guo, “Experimental demonstration of indefinite causal order induced quantum heat extraction,” (2021) , [arXiv:2101.07979 \[quant-ph\]](#).
- [20] K. Goswami and J. Romero, “Experiments on quantum causality,” [AVS Quantum Science](#) **2** no. 3, (Oct, 2020) 037101, [arXiv:2009.00515 \[quant-ph\]](#).
- [21] L. Hardy, “Quantum gravity computers: On the theory of computation with indefinite causal structure,” [Quantum Reality, Relativistic Causality, and Closing the Epistemic Circle](#) (2009) 379–401, [arXiv:quant-ph/0701019](#).
- [22] G. Chiribella, G. M. D’Ariano, and P. Perinotti, “Theoretical framework for quantum networks,” [Physical Review A](#) **80** no. 2, (Aug, 2009) , [arXiv:0904.4483 \[quant-ph\]](#).
- [23] G. Chiribella, G. D’Ariano, P. Perinotti, and B. Valiron, “Beyond quantum computers,” (2009) , [arXiv:0912.0195v1 \[quant-ph\]](#).
- [24] G. Chiribella, “Perfect discrimination of no-signalling channels via quantum superposition of causal structures,” [Physical Review A](#) **86** no. 4, (Oct, 2012) , [arXiv:1109.5154 \[quant-ph\]](#).
- [25] T. Colnaghi, G. M. D’Ariano, S. Facchini, and P. Perinotti, “Quantum computation with programmable connections between gates,” [Physics Letters A](#) **376** no. 45, (Oct, 2012) 2940–2943, [arXiv:1109.5987 \[quant-ph\]](#).
- [26] Å. Baumeler and S. Wolf, “The space of logically consistent classical processes without causal order,” [New Journal of Physics](#) **18** no. 1, (2016) 013036, [arXiv:1507.01714 \[quant-ph\]](#).
- [27] Å. Baumeler, A. Feix, and S. Wolf, “Maximal incompatibility of locally classical behavior and global causal order in multiparty scenarios,” [Physical Review A](#) **90** no. 4, (2014) 042106, [arXiv:1403.7333 \[quant-ph\]](#).
- [28] M. Araújo, A. Feix, M. Navascués, and Č. Brukner, “A purification postulate for quantum mechanics with indefinite causal order,” [Quantum](#) **1** (Apr, 2017) 10, [arXiv:1611.08535 \[quant-ph\]](#).
- [29] A. Vanrietvelde, N. Ormrod, H. Kristjánsson, and J. Barrett, “Consistent circuits for indefinite causal order,” (2022) , [arXiv:2206.10042 \[quant-ph\]](#).
- [30] H. Reichenbach, [The direction of time](#), vol. 65. Univ of California Press, 1956.
- [31] C. J. Wood and R. W. Spekkens, “The lesson of causal discovery algorithms for quantum correlations: causal explanations of bell-inequality violations require fine-tuning,” [New Journal of Physics](#) **17** no. 3, (Mar, 2015) 033002, [arXiv:1208.4119 \[quant-ph\]](#).
- [32] J.-M. A. Allen, J. Barrett, D. C. Horsman, C. M. Lee, and R. W. Spekkens, “Quantum common causes and quantum causal models,” [Physical Review X](#) **7** no. 3, (Jul, 2017) , [arXiv:1609.09487 \[quant-ph\]](#).
- [33] J. Pearl, [Causality](#). Cambridge university press, 2009.
- [34] J. Pienaar and Č. Brukner, “A graph-separation theorem for quantum causal models,” [New Journal of Physics](#) **17** no. 7, (2015) 073020, [arXiv:1406.0430v3 \[quant-ph\]](#).
- [35] F. Costa and S. Shrapnel, “Quantum causal modelling,” [New Journal of Physics](#) **18** no. 6, (June, 2016) 063032, [arXiv:1512.07106 \[quant-ph\]](#).
- [36] J. Pienaar, “A time-reversible quantum causal model,” (2019) , [arXiv:1902.00129 \[quant-ph\]](#).
- [37] J. Pienaar, “Quantum causal models via quantum bayesianism,” [Physical Review A](#)

- 101 no. 1, (2020) 012104, [arXiv:1806.00895 \[quant-ph\]](#).
- [38] S. Gogioso and N. Pinzani, “The topology and geometry of causality,” (2022) . <https://arxiv.org/abs/2206.08911>.
- [39] G. Chiribella and H. Kristjánsson, “Quantum shannon theory with superpositions of trajectories,” [Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences](#) **475** no. 2225, (May, 2019) 20180903, [arXiv:1812.05292 \[quant-ph\]](#).
- [40] Y. Aharonov and D. Bohm, “Significance of electromagnetic potentials in the quantum theory,” [Physical Review](#) **115** (Aug, 1959) 485–491.
- [41] N. Erez, “AB effect and aharonov–susskind charge non-superselction,” [Journal of Physics A: Mathematical and Theoretical](#) **43** no. 35, (Aug, 2010) 354030, [arXiv:1003.1044 \[quant-ph\]](#).
- [42] F. D. Santo and B. Dakić, “Two-way communication with a single quantum particle,” [Physical Review Letters](#) **120** no. 6, (Feb, 2018) , [arXiv:1706.08144 \[quant-ph\]](#).
- [43] L.-Y. Hsu, C.-Y. Lai, Y.-C. Chang, C.-M. Wu, and R.-K. Lee, “Carrying an arbitrarily large amount of information using a single quantum particle,” [Physical Review A](#) **102** (Aug, 2020) 022620, [arXiv:2002.10374 \[quant-ph\]](#).
- [44] F. Massa, A. Moqanaki, Ämin Baumeler, F. D. Santo, J. A. Kettlewell, B. Dakić , and P. Walther, “Experimental two-way communication with one photon,” [Advanced Quantum Technologies](#) **2** no. 11, (Sep, 2019) 1900050, [arXiv:1802.05102 \[quant-ph\]](#).
- [45] R. Faleiro, N. Paunkovic, and M. Vojinovic, “Operational interpretation of the vacuum and process matrices for identical particles,” [Quantum](#) **7** (2023) 986, [arXiv:2010.16042 \[quant-ph\]](#).
- [46] I. Marvian and R. W. Spekkens, “A generalization of Schur–Weyl duality with applications in quantum estimation,” [Communications in Mathematical Physics](#) **331** no. 2, (2014) 431–475, [arXiv:1112.0638 \[quant-ph\]](#).
- [47] A. W. Harrow, [Applications of coherent classical communication and the Schur transform to quantum information theory](#). PhD thesis, Massachusetts Institute of Technology, 2005. [arXiv:quant-ph/0512255](#).
- [48] G. M. Palma, K.-A. Suominen, and A. K. Ekert, “Quantum computers and dissipation,” [Proceedings of the Royal Society A](#) **452** (1996) 567–584, [arXiv:quant-ph/9702001](#).
- [49] L.-M. Duan and G.-C. Guo, “Preserving coherence in quantum computation by pairing the quantum bits,” [Physical Review Letters](#) **79** (1997) 1953–1956, [arXiv:quant-ph/9703040](#).
- [50] P. Zanardi and M. Rasetti, “Noiseless quantum codes,” [Physical Review Letters](#) **79** no. 17, (1997) 3306, [arXiv:quant-ph/9705044](#).
- [51] D. A. Lidar, I. L. Chuang, and K. B. Whaley, “Decoherence-free subspaces for quantum computation,” [Physical Review Letters](#) **81** no. 12, (1998) 2594, [arXiv:quant-ph/9807004](#).
- [52] A. Beige, D. Braun, B. Tregenna, and P. L. Knight, “Quantum computing using dissipation to remain in a decoherence-free subspace,” [Physical Review Letters](#) **85** no. 8, (2000) 1762.
- [53] P. G. Kwiat, A. J. Berglund, J. B. Altepeter, and A. G. White, “Experimental verification of decoherence-free subspaces,” [Science](#) **290** no. 5491, (2000) 498–501.
- [54] O. Oreshkov, “Time-delocalized quantum subsystems and operations: on the existence of processes with indefinite causal structure in quantum mechanics,” [Quantum](#) **3** (2019) 206, [arXiv:1801.07594 \[quant-ph\]](#).
- [55] A. Vanrietvelde, H. Kristjánsson, and J. Barrett, “Routed quantum circuits,” [Quantum](#) **5** (Jul, 2021) 503, [arXiv:2011.08120 \[quant-ph\]](#).
- [56] A. Vanrietvelde and G. Chiribella, “Universal control of quantum processes

- using sector-preserving channels,” [Quantum Information and Computation](#) **21** no. 15-16, (Dec, 2021) 1320–1352, [arXiv:2106.12463 \[quant-ph\]](#).
- [57] M. Wilson and A. Vanrietvelde, “Composable constraints,” (2021) , [arXiv:2112.06818 \[math.CT\]](#).
- [58] A. A. Abbott, J. Wechs, D. Horsman, M. Mhalla, and C. Branciard, “Communication through coherent control of quantum channels,” [Quantum](#) **4** (Sep, 2020) 333, [arXiv:1810.09826 \[quant-ph\]](#).
- [59] H. Kristjánsson, G. Chiribella, S. Salek, D. Ebler, and M. Wilson, “Resource theories of communication,” [New Journal of Physics](#) **22** no. 7, (Jul, 2020) 073014, [arXiv:1910.08197 \[quant-ph\]](#).
- [60] I. Friend, “Private communication,” (2022) .
- [61] G. Chiribella, G. M. D’Ariano, and P. Perinotti, “Transforming quantum operations: Quantum supermaps,” [EPL \(Europhysics Letters\)](#) **83** no. 3, (Jul, 2008) 30004, [arXiv:0804.0180 \[quant-ph\]](#).
- [62] M. Zych, F. Costa, I. Pikovski, and Č. Brukner, “Bell’s theorem for temporal order,” [Nature communications](#) **10** no. 1, (2019) 1–10, [arXiv:1708.00248 \[quant-ph\]](#).
- [63] N. S. Möller, B. Sahdo, and N. Yokomizo, “Quantum switch in the gravity of Earth,” [Physical Review A](#) **104** no. 4, (2021) 042414, [arXiv:2012.03989 \[quant-ph\]](#).
- [64] J. Wechs, C. Branciard, and O. Oreshkov, “Existence of processes violating causal inequalities on time-delocalised subsystems,” [Nature Communications](#) **14** no. 1, (2023) 1471, [arXiv:2201.11832 \[quant-ph\]](#).
- [65] V. Vilasini, “An introduction to causality in quantum theory (and beyond) (master’s thesis),” (2017) . https://foundations.ethz.ch/wp-content/uploads/2019/07/vilasini_master_thesis-v2.pdf.
- [66] V. Vilasini, “Causality in definite and indefinite space-times (extended abstract for qpl 2020),” (2020) . <https://wdi.centralesupelec.fr/users/valiron/qplmfps/papers/qs01t3.pdf>.
- [67] C. Portmann, C. Matt, U. Maurer, R. Renner, and B. Tackmann, “Causal boxes: quantum information-processing systems closed under composition,” [IEEE Transactions on Information Theory](#) **63** no. 5, (2017) 3277–3305. <https://doi.org/10.1109/TIT.2017.2676805>.
- [68] B. d’Espagnat, “An elementary note about ‘mixtures’,” [Preludes in Theoretical Physics in honor of VF Weisskopf](#) (1966) 185.
- [69] B. d’Espagnat, [Conceptual foundations of quantum mechanics](#). CRC Press, 2018.
- [70] S. D. Bartlett, T. Rudolph, and R. W. Spekkens, “Reference frames, superselection rules, and quantum information,” [Review of Modern Physics](#) **79** (Apr, 2007) 555–609, [arXiv:quant-ph/0610030](#).
- [71] V. Vilasini and R. Renner, “Embedding cyclic causal structures in acyclic spacetimes: no-go results for process matrices,” (2022) , [arXiv:2203.11245 \[quant-ph\]](#).
- [72] B. Schumacher and M. D. Westmoreland, “Locality and information transfer in quantum operations,” [Quantum Information Processing](#) **4** no. 1, (2005) 13–34, [arXiv:quant-ph/0406223](#).

A Proof of Theorem 3.1

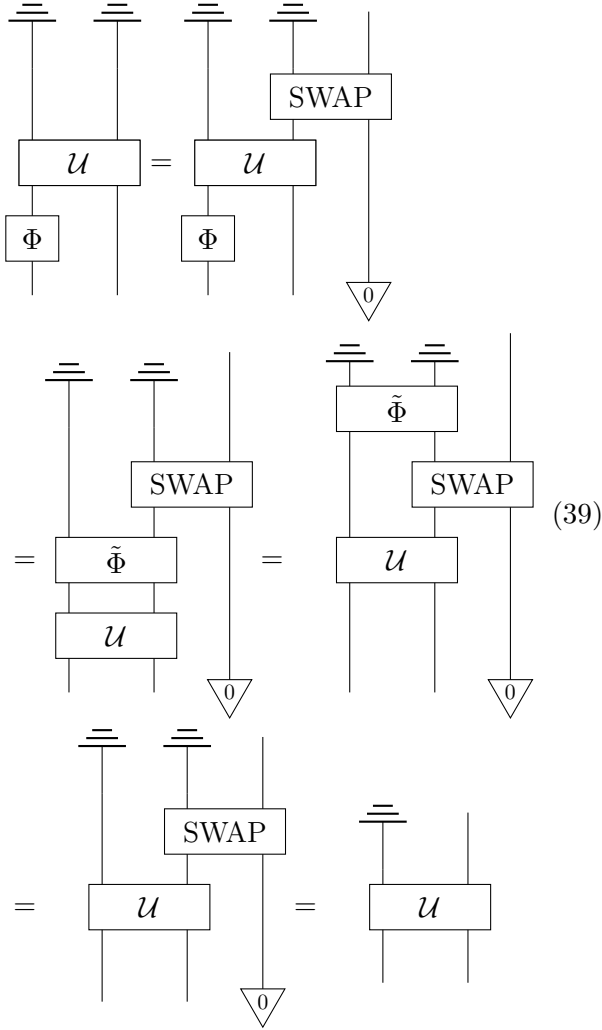
The implication from (U1) to (U2) is obvious. That (U2) implies (U3) follows from noting that, for any density matrices ρ and ρ' , one can always find a channel Φ such that $\rho = \Phi(\rho')$. Inserting $\rho' \otimes \sigma$ onto both sides of (U2) then gives (U3). The equivalence of (U1) and (U3) is proven in [72], and the equivalence of (U1) and (U4) is proven in [10].

At this point, we have proven that conditions (U1–U4) are all equivalent. It remains to show that all of these are equivalent to (U5) and (U6).

It is obvious that (U4) implies (U5). That (U5) implies (U6) follows from the facts that $\mathcal{A} := \mathcal{L}(A)$ is spanned by its unitary operators, and that the right side of (U5) commutes with $I_C \otimes \mathcal{D}$.

Finally, to show that (U6) implies (U2), we define the channel $\tilde{\Phi} := \mathcal{U} \circ (\Phi \otimes \mathcal{I}) \circ \mathcal{U}^\dagger$ on $C \otimes D$. We can then perform the following manipulations, in

which the triangle labelled '0' represents the state $|0\rangle\langle 0|$.



To see why the third equality holds, note that since Φ is a channel on A , all of its Kraus operators are elements of \mathcal{A} . It follows that the Kraus operators of $\tilde{\Phi}$ are elements of $U(\mathcal{A} \otimes I_B)U^\dagger$. By (U6), we know these operators commute with all the operators in $\mathcal{D} := \mathcal{L}(D)$ that only act on D . In the manipulations above, the swap channel only acts on D and its own ancilla, meaning that it will commute with $\tilde{\Phi}$.

B Proof of Proposition 3.2

It follows quite simply from (U6) that if $A \xrightarrow{U} E$, and $A \xrightarrow{U} F$, then $A \xrightarrow{U} E \otimes F$. Of course, if A can signal to a subsystem then it can always signal to the composite system, so the converse also holds. We therefore have the first part of atomicity (13), i.e. $(A \xrightarrow{U} E) \wedge (A \xrightarrow{U} F) \iff (A \xrightarrow{U} E \otimes F)$.

Taking into account the time symmetry property (12) leads to the second part of atomicity

(14), i.e. $(A \xrightarrow{U} F) \wedge (B \xrightarrow{U} F) \iff (A \otimes B \xrightarrow{U} F)$. Explicitly, we have:

$$\begin{aligned} (A \xrightarrow{U} F) \wedge (B \xrightarrow{U} F) &\iff (F \xrightarrow{U^\dagger} A) \wedge (F \xrightarrow{U^\dagger} B) \\ &\iff (F \xrightarrow{U^\dagger} A \otimes B) \\ &\iff (A \otimes B \xrightarrow{U} F) \end{aligned} \quad (40)$$

where in the first and third equivalences we have used time symmetry (12), and in the second equivalence we used the first part of atomicity, (13).

C The embedding map

Consider two sectorized Hilbert spaces A^i and B^j and some subspace $H = \bigoplus_{i,j} \lambda^{ij} A_\star^i \otimes B_\star^j$, where λ^{ij} is some route matrix. (H might be the practical input or output space of some routed unitary.) A *sector-preserving algebra* over A^i is one that can be written as:

$$\mathcal{X} = \bigoplus_i \mathcal{X}_i \quad (41)$$

where each \mathcal{X}_i is an algebra over A_\star^i . The point of the embedding map is to ‘embed’ this algebra into an algebra over the larger space H .

The embedding map acts on operators from \mathcal{X} . An operator from this algebra has the form $M = \bigoplus_i M_i$. The embedding map is defined as follows:

$$M \xrightarrow{*} M^* := \bigoplus_{i,j} \lambda^{ij} M_i \otimes I_{B_\star^j} \quad (42)$$

where $I_{B_\star^j}$ is the identity operator on B_\star^j . The codomain of $\xrightarrow{*}$ is given by:

$$\mathcal{X} \xrightarrow{*} \mathcal{X}^* := \bigoplus_{i,j} \lambda^{ij} \mathcal{X}_i \otimes I_{B_\star^j} \quad (43)$$

where the \mathcal{X}_i are algebras over the A_\star^i from (41).

We now prove that the embedding map has some properties that help with the proof of Theorem 3.2. Firstly, embedding map is invertible and linear. To see that it is invertible, note that any operator from \mathcal{X}^* can be uniquely written in the form $M^* = \bigoplus_{i,j} \lambda^{ij} M_i \otimes I_{B_\star^j}$, so we can easily define the inverse embedding map:

$$M^* \xrightarrow{*^{-1}} M := \bigoplus_i M_i \quad (44)$$

Now we prove that the embedding map is linear. Given the operators $M = \bigoplus_i M_i \in \mathcal{X}$ and $N = \bigoplus_i N_i \in \mathcal{X}$, consider some linear combination $R := aM + bN$. R can be written $R = \bigoplus_i R_i$ where each $R_i = aM_i + bN_i$. Applying the embedding map to a linear combination therefore has the following effect:

$$\begin{aligned} (aM + bN)^* &= \bigoplus_{i,j} \lambda^{ij} (aM + bN)_i \otimes I_{B_\star^j} \\ &= \bigoplus_{i,j} \lambda^{ij} (aM_i + bN_i) \otimes I_{B_\star^j} \\ &= \bigoplus_{i,j} \lambda^{ij} (aM_i \otimes I_{B_\star^j}) + \bigoplus_{i,j} (bN_i \otimes I_{B_\star^j}) \\ &= aM^* + bN^* \end{aligned} \quad (45)$$

which proves linearity.

The combination of the linearity and the invertibility of the embedding map has the useful consequence that if a set of operators $\{M_k\}_k$ spans \mathcal{X} , then $\{M_k^*\}_k$ spans \mathcal{X}^* , which we now prove. Denote by N^* an arbitrary member of \mathcal{X}^* . N^* can be mapped via $\xrightarrow{*^{-1}}$ to $N \in \mathcal{X}$, which by assumption can be written $N = \sum_k \alpha_k M_k$. We can apply $\xrightarrow{*}$ to both sides of this equation and use linearity to deduce $N^* = \sum_k \alpha_k M_k^*$.

D The exchange gate

Given a sectorized Hilbert space A^i , the exchange gate is a unitary operator

$$\text{EXCH} : A \otimes A^O \rightarrow A \otimes A^I \quad (46)$$

where A^O and A^I are ancillary systems of the same dimension with the factorizations:

$$\begin{aligned} A^O &= A_1^O \otimes \dots \otimes A_N^O \otimes A_w^O \\ A^I &= A_1^I \otimes \dots \otimes A_N^I \otimes A_w^I \end{aligned} \quad (47)$$

where N is the number of sectors. Each A_i^O and A_i^I has the dimension of the corresponding sector, and A_w^O and A_w^I are each of dimension N .

The exchange gate can be understood in terms of two more basic unitary operators. The first is called the ‘sectorial CNOT’, or **SCNOT**, gate, and it acts only on A and the which-sector ancilla. To define the **SCNOT** gate, we first note for every sector A_\star^i of A^i we can write down a basis $\{|i, j_i\rangle\}_{j_i}$. We can then combine these to get

a basis for A : $\{|i, j_i\rangle\}_{i, j_i}$. Then the gate **SCNOT**: $A \otimes A_w^O \rightarrow A \otimes A_w^I$ is defined as follows:

$$\text{SCNOT} |i, j_i\rangle_A |k\rangle_{A_w^O} = |i, j_i\rangle_A |k + i\rangle_{A_w^I} \quad \forall |i, j_i\rangle_A \quad \forall |k\rangle_{A_w^O} \quad (48)$$

where $\{|k\rangle_{A_w^O}\}$ and $\{|k\rangle_{A_w^I}\}$ are logical bases for A_w^O and A_w^I respectively, and the addition is modulo N . In other words, the **SCNOT** is a direct sum over the sectors of A^i of operators which leave A_\star^i alone send the ancillary basis $\{|k\rangle_{A_w^O}\}$ to $\{|k + i\rangle_{A_w^I}\}$.

The second gate is called the ‘routed swap’ gate, or **RSWAP**. **RSWAP**: $A \otimes A_1^O \otimes \dots \otimes A_N^O \rightarrow A \otimes A_1^I \otimes \dots \otimes A_N^I$ is a direct sum over the sectors of A^i of unitary operators that swap A_\star^i with the i th ancillary system. Let us make this more precise. We define $\text{SWAP}_i := I_{A_\star^i \rightarrow A_i^I} \otimes I_{A_i^O \rightarrow A_\star^i}$, where I represents the identity transformations between the Hilbert spaces denoted in the subscript. We denote by $I_{A_i^O \rightarrow A_i^I}$ the identity transformation from the tensor product of all except the i th ancillary input to the tensor product of all except the i th ancillary output. Then **RSWAP** is given by:

$$\text{RSWAP} = \bigoplus_i \text{SWAP}_i \otimes I_{A_i^O \rightarrow A_i^I} \quad (49)$$

Then the exchange gate is defined as

$$\text{EXCH} := (\text{RSWAP} \otimes I_{A_w^I}) \circ (\text{SCNOT} \otimes I_{A_1^O \dots A_N^O}), \quad (50)$$

or, equivalently, as $\text{EXCH} := (\text{SCNOT} \otimes I_{A_1^I \dots A_N^I}) \circ (\text{RSWAP} \otimes I_{A_w^O})$.

E A lemma for Theorem 3.2

In order to prove Theorem 3.2, we will first state and prove an important lemma. Since this will take a long time, we do it here, in a separate appendix to the rest of the proof. The lemma will show that our sectorized no-influence relation through a unitary channel – which is always equivalent to (FG4) below – is equivalent to three other conditions, except in the case where one of the relata is a one-dimensional sector.

E.1 Stating the Lemma

In order to state these conditions, we first associate the sectors and the which-sector information

with algebras of operators. We associate the i th sector of A^i with the algebra:

$$\mathcal{A}^i := \mathcal{L}(A_{\star}^i) \oplus \text{span}(I_{A_{\star}^i}) \quad (51)$$

where $\mathcal{L}(A_{\star}^i)$ is the algebra of linear operators on A_{\star}^i , $I_{A_{\star}^i}$ is the identity operator on $A_{\star}^i := \bigoplus_{i' \neq i} A_{\star}^{i'}$, and the span is over complex numbers (as it always will be subsequently). In other words, \mathcal{A}^i is the algebra of operators of the form $M_{A_{\star}^i} \oplus \alpha I_{A_{\star}^i}$, where M is an operator on A_{\star}^i , and α is a complex number.

We also define an algebra corresponding to the which-sector information:

$$\mathcal{A}^w := \text{span}(\{\pi_i^A\}_i) \quad (52)$$

where the π_i^A are the projectors used to define A^i 's sectorization: $\pi_i^A(A) = A_{\star}^i \forall i$.

The algebras associated with D^l 's sectors and which-sector information are defined in a precisely analogous way. We now proceed to give conditions for sectorized no-influence relations (in the absence of one-dimensional sectors) which largely resemble the unsectorized no-influence relations but for the fact that they are defined in terms of these smaller algebras, rather than the full sector-preserving algebras of the form (19)

Given an algebra \mathcal{X} corresponding to one of A^i 's sectors or which-sector information, and given an algebra \mathcal{Y} corresponding to one of D^l 's sectors or which-sector information, we are now ready to state the four conditions for no-influence between these two relata. In (FG1) below, Φ is a channel with Kraus operators from \mathcal{X} , and Ψ is a channel with Kraus operators from $\mathcal{Y} \otimes \mathcal{L}(* \rightarrow D^I)^{26}$ for some ancillary space D^I .

$$\lambda_{ij}^{kl} \left(\begin{array}{c} \text{Circuit with } \Psi \end{array} \right) = \lambda_{ij}^{kl} \left(\begin{array}{c} \text{Circuit with } U \end{array} \right) \quad \forall \Phi \forall \Psi \quad (\text{FG1})$$

²⁶These are the operators of the form $\sum_k M_k \otimes |\psi_k\rangle$ for $M_k \in \mathcal{Y}$ and $|\psi_k\rangle \in D^I$.

The second condition is given in the case that \mathcal{Y} is a which-sector algebra (FG2a), and separately in the case that it corresponds to a sector (FG2b). In both cases, V is any unitary operator from \mathcal{X} . In (FG2a) directly below, W , \dot{D} , and the routes all have the same meaning as in (R2). V' is a practical unitary controlled on \dot{D} of the form $\sum_l V'_l \otimes |l\rangle \langle l|_{\dot{D}}$, where $\{|l\rangle\}_l$ is the basis associated with \dot{D} 's sectorization.

$$\lambda_{ij}^{kl} \left(\begin{array}{c} \text{Circuit with } W \end{array} \right) = \lambda_{ij}^{kl} \left(\begin{array}{c} \text{Circuit with } V' \end{array} \right) \quad \forall V \exists V' \quad (\text{FG2a})$$

Note that this equality is only required to hold for the transformations, and not necessarily the routes.

On the other hand, in (FG2b), \dot{E}^u is a sectorized Hilbert space with two one-dimensional sectors, and (L, ϵ) is a routed unitary transformation that maps the specific sector that is associated with D_{\star}^q associated with \mathcal{Y} to $D_{\star}^q \otimes E_{\star}^0$, and maps the rest of the space D_{\star}^q to $D_{\star}^q \otimes E_{\star}^1$, where D_{\star}^q is the sector corresponding to \mathcal{Y} . The route matrix ϵ_n^{pu} is equal to 1 if $n = p = q$ and $u = 0$, or if $n = p \neq q$ and $u = 1$; otherwise, it is zero. V is again any unitary on \mathcal{X} , and the controlled practical unitary labelled V'' is of the form $V'' = V_0'' \otimes \pi_q^D \otimes |0\rangle \langle 0|_E + V_1'' \otimes |1\rangle \langle 1|_E$, where V_0'' acts on C , V_1'' acts on $C \otimes D$, π_q^D is the projector in D onto D_{\star}^q , and each $|u\rangle \langle u|_E$ spans the sector E_{\star}^u .

$$\forall V \exists V'' \quad (\text{FG2b})$$

Again, the equation is not required to hold for the routes.

In (FG3), \mathcal{X}^* is the algebra obtained by embedding \mathcal{X} into U 's practical input space. \mathcal{Y}^* is the algebra obtained by embedding \mathcal{Y} in U 's practical output space.

$$[U\mathcal{X}^*U^\dagger, \mathcal{Y}^*] = 0 \quad (\text{FG3})$$

(FG4) is a statement of a no-influence relation through an *unrouted* unitary transformation, conforming to Definition 3.1. It is given in terms of the following unitary, which we denote by S_U :

$$(53)$$

Here, $T_{\text{prac, in}}$ and $T_{\text{prac, out}}$ are practical unitaries that embed U 's practical input and output space respectively into an unsectorized system. The routes in this diagram are given by $\mu_{ij} := \sum_{kl} \lambda_{ij}^{kl}$ and $\gamma_{kl} := \sum_{ij} \lambda_{ij}^{kl}$. Also, in (FG4), X^O stands for the input subsystem related to \mathcal{X} ; either A_i^O

if $\mathcal{X} = \mathcal{A}^i$ or A_w^O if $\mathcal{X} = \mathcal{A}^w$. Similarly, $Y^I = D_l^I$ if $\mathcal{Y} = \mathcal{D}^l$ or $Y^I = D_w^I$ if $\mathcal{Y} = \mathcal{D}^w$.

$$X^O \not\stackrel{S_U}{\rightarrow} Y^I \quad (\text{FG4})$$

We define trivial sectorial relata as those algebras corresponding to sectors that are one-dimensional. We can now state the lemma:

Lemma E.1. *Assuming neither \mathcal{X} nor \mathcal{Y} correspond to trivial sectorial relata, conditions (FG1) – (FG4) are equivalent.*

We will now prove the lemma.

E.2 (FG2) is equivalent to (FG3)

We start by proving that (FG2) \Leftrightarrow (FG3). For the rightward implication, it is immediate that the transformations on the right sides of (FG2a) and (FG2b) commute with \mathcal{D}^{w*} and \mathcal{D}^{l*} respectively. The implication then follows from the fact that \mathcal{A}^{w*} and each \mathcal{A}^{i*} are spanned by their unitaries. This in turn follows from the fact that \mathcal{A}^w and each \mathcal{A}^i are spanned by their unitaries, and from the linearity and invertibility of the embedding map.

For the leftward direction, suppose first that $\mathcal{Y} = \mathcal{D}^{w*}$. Then (FG3) implies that, for any unitary $V \in \mathcal{X}$, the unitary transformation in:

$$(54)$$

commutes with each projector $Q_l \otimes \pi_l^D$, where π_l^D is the projector onto D_\star^l , and Q_l is the projector onto the largest orthogonal subspace of C , call it c_l , such that $c_l \otimes D_\star^l$ is part of U 's practical output space. This means that the transformation in (54) can be written in the form $V' = \bigoplus_l V'_l$, where each V'_l acts on $c_l \otimes D_\star^l$. It follows that it can be written as in the right-hand side of (FG2a).

Now suppose $\mathcal{Y} = \mathcal{D}^{l*}$. Set $H_l := c_l \otimes D^l$, and define $H_{\bar{l}}$ as the compliment of H_l in U 's practical output space, H_{acc} :

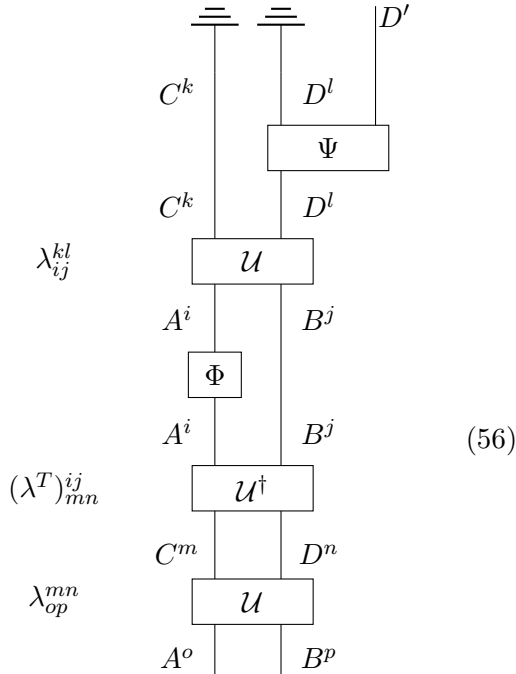
$$H_{\text{acc}} = H_l \oplus H_{\bar{l}} \quad (55)$$

(FG3) implies that any unitary $V \in \mathcal{U}(\mathcal{X}^*)$ commutes with the projector $Q_l \otimes \pi_l^D$ in H_{acc} onto H_l , and is therefore block diagonal in this two-sector partitioning: $V = V_l \oplus V_{\bar{l}}$. (FG3) implies that V must commute with $(Q_l \otimes \mathcal{L}(D_\star^l)) \oplus \mathbf{0}_{\bar{l}}$, meaning that V has the form $V = (V_{cl}' \otimes I_{D_\star^l}) \oplus V_{\bar{l}}$. Any unitary operator of this form can be written as in the right-hand side of (FG2b).

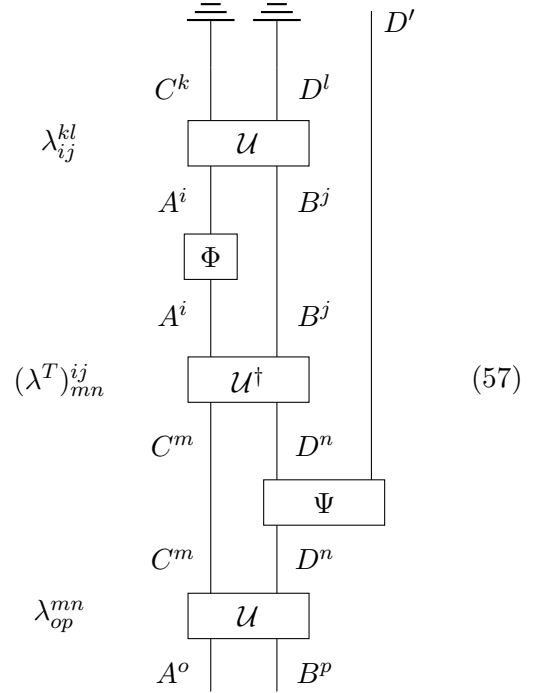
E.3 (FG3) implies (FG1)

We now prove (FG3) \implies (FG1). The reasoning is very similar to that in the proof that (U6) implies (U2), but it is carried out again explicitly here for the sake of completeness.


First we note that the transformation (although not necessarily the route) in the left-hand side of (FG1) is equal to the one in



equal to the one in



since it entails that all Kraus operators of Φ , once transformed by \mathcal{U} into operators acting on U 's output space, commute with all the Kraus operators of Ψ . Then we can trace out the channels after Ψ to show that the transformation we have been discussing all along is equal to the one in:



This implies (FG1).


E.4 Sublemmas

We state and prove two sublemmas required for the rest of the proof.

(FG3) implies the transformation in (56) is


Sublemma E.1. A_w^O is a not cause of D_w^I in

(53) if and only if A_w^O is a not cause of D_w^I in:



(59)


The ‘if’ part is easily shown by partially tracing all the outputs except D_w^I in (53) and noting from the definition of the EXCH gate that the SCNOT and RSWAP gates commute. The ‘only if’ part follows by considering the following channel:



(60)


which can be simplified by noting that the RSWAP on D will be traced out and the one on A will give the identity once combined with its adjoint. This leads to an expression from which one can trivially derive that A_w^O is not a cause of D^I in (59).

Sublemma E.2. For any density matrix ρ



(61)

and



(62)

have Kraus operators drawn from \mathcal{A}^w and $\mathcal{D}^w \otimes \mathcal{L}(* \rightarrow D_w^I)$ respectively.

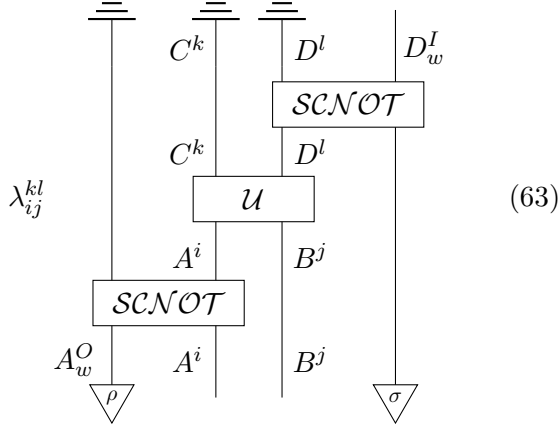
We denote the channels in (61) and (62) by \mathcal{M} and \mathcal{N} respectively. For any pair of states $|\psi_i\rangle$ and $|\phi_i\rangle$ both contained in A_\star^i , we have $\mathcal{M}(|\psi_i\rangle\langle\phi_i|) = |\psi_i\rangle\langle\phi_i|$. This means that not only does \mathcal{M} follow the delta-route depicted in (61) associated with the sectorization $H_{\text{acc}}^{\text{in}} = \bigoplus_i A_\star^i$, it also follows any similar delta-route obtained by also further sectorizing each sector A_\star^i into a set of one-dimensional sectors. By Theorem 6 of [55], this implies that the Kraus operators of \mathcal{M} must be block diagonal in all of the sectorizations, implying that they are each of the form $\sum_i \alpha_i \pi_i^A$ for complex numbers α_i , and hence a member of \mathcal{A}^w .

One can give a very similar argument that the Kraus operators \mathcal{N} must be block diagonal in the sense of belonging to $\bigoplus_l (\bigoplus_{m_l} \mathcal{L}((D_\star^l)^{m_l}) \otimes \mathcal{L}(* \rightarrow D_w^I))$ for any set of sectorizations $D_\star^l = \bigoplus_{m_l} (D_\star^l)^{m_l}$ of each D_\star^l into one-dimensional sectors. This implies that each Kraus operator must be of the form $\sum_l \alpha_l \pi_l^D \otimes |\xi_l\rangle_{D_w^I}$, hence a member of $\mathcal{D}^w \otimes \mathcal{L}(* \rightarrow D_w^I)$.

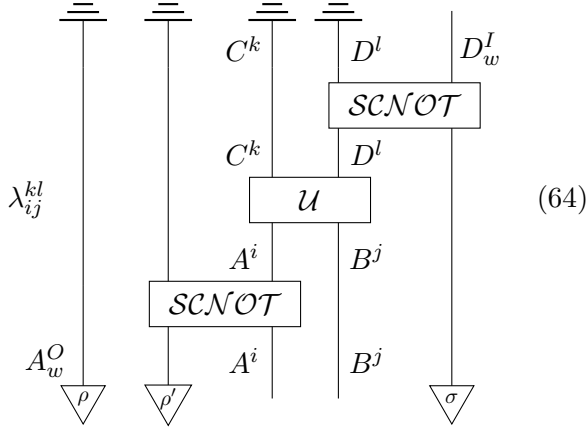
E.5 (FG1) implies (FG4)

Now we prove that (FG1) implies (FG4). We start by assuming $\mathcal{X} = \mathcal{A}^{w*}$ and $\mathcal{Y} = \mathcal{D}^{w*}$. Given Sublemma E.2, (FG1) implies that, for any nor-


malized density matrix σ ,




is independent of the density matrix ρ . Since ρ is normalized, it follows that, for an arbitrary state ρ' on A_w^O , the above is equal to:



for any ρ and all σ . We can use state tomography where we vary over ρ and σ but hold ρ' fixed to deduce that:



where μ is the matrix obtained by composing the route matrices on the left, and Φ is the following quantum channel:



Given Sublemma E.1, (65) implies (FG4).

The argument in the case that at least one of \mathcal{X} and \mathcal{Y} corresponds to a sector is closely related to the first one. One just has to note that the RSWAP gate on D^l (for example) can be written as a sequentially combination of commuting operators, on associated with each of D^l 's M sectors:

$$\text{RSWAP} = (I_{D_M} \otimes R_M) \circ \dots \circ (I_{D_1} \otimes R_1) \quad (67)$$

where I_{D_l} is the identity from $\bigoplus_{l' \neq l} D_{l'}^O$ to $\bigoplus_{l' \neq l} D_{l'}^I$, and, for all l ,

$$R_l := \text{SWAP}_l \oplus (I_{D_l} \otimes I_{D_l^*}) \quad (68)$$

where I_{D_l} is the identity operator from D_l^O to D_l^I , $I_{D_l^*}$ is the identity on $D_{l^*}^I = \bigoplus_{l' \neq l} D_{l'}^I$, and SWAP_l swaps the sector D_l^I with the ancillary wire D_l^O/D_l^I . Then the foregoing argument can be generalized to the case where at least one of \mathcal{X} and \mathcal{Y} is a sectorial relatum by proving a sublemma similar to E.1 where the SCNOT gates on any wires corresponding to the sectorial relata \mathcal{A}^i or \mathcal{D}^l are replaced with the corresponding R_i or R_l – these sublemmas are proved using similar arguments to the ones above. Then, maintaining the relevant substitutions, one repeats the argument encapsulated by equations (63 – 65), which, importantly, does not depend on any special properties of SCNOT.

E.6 (FG4) implies (FG3)

Finally, we prove that (FG4) implies (FG3), assuming the relata are nontrivial. The argument varies depending on whether the non-cause and

non-effect are sectors or which-sector information. However, it will suffice to prove the claim in the two cases where both relata are of the same kind, before indicating how the arguments are to be adapted in the other two cases.

We start by assuming that the no-influence relation is between the which-sector informations, so our algebras are $\mathcal{X} = \mathcal{A}^{w*}$ and $\mathcal{Y} = \mathcal{D}^{w*}$. Let F denote the quantum Fourier transform operator on the ancilla input A_w^O to the **SCNOT** gate on A . Consider the unitary operator on A obtained by inserting the state $F|-n\rangle = \frac{1}{\sqrt{N}} \sum_{m=0}^{N-1} \exp(\frac{-2\pi imn}{N}) |m\rangle$ into the ancilla input of the **SCNOT** gate, and ignoring the output state on A_w^I . This can be written as:

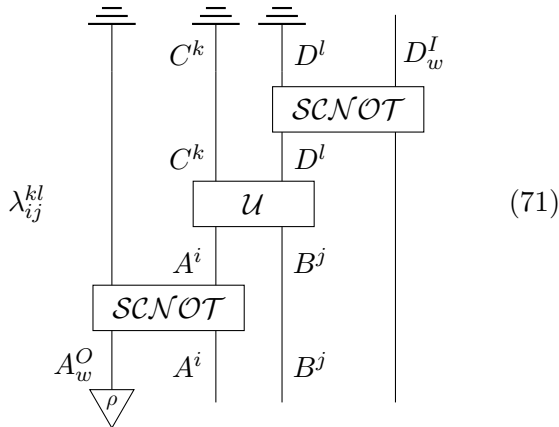
$$V_n = \sum_m \exp(\frac{2\pi imn}{N}) \pi_m^A \quad (69)$$

where each π_m^A is the projector onto A_m^* .²⁷ Denoting by P_m the projector onto the largest orthogonal subspace b_m of B such that $A_m^* \otimes b_m$ is inside U 's practical input space, V_n is embedded into the practical input space as follows (see Appendix C):

$$V_n^* := \sum_{m=0}^{N-1} \exp(\frac{2\pi imn}{N}) \pi_m^A \otimes P_m \quad (70)$$


Since, by the unitary of F , $\{\frac{1}{\sqrt{N}} \sum_{m=0}^{N-1} \exp(\frac{2\pi imn}{N}) |m\rangle\}_{n=0}^{N-1}$ spans any computational basis state $|n\rangle$ for A , it is clear that $\{V_n^*\}_{n=0}^{N-1}$ spans any projector $\pi_n^A \otimes P_n$, and thus that the same set spans \mathcal{A}^{w*} . It will therefore suffice to show that $\{\mathcal{U}(V_n^*)\}_{n=0}^{N-1}$ commutes with \mathcal{D}^{w*} , for any V_n^* .

By Sublemma E.1, (FG4) implies that



²⁷More precisely, the point is that $\text{SCNOT}(I_A \otimes F|-n\rangle_{A_w^O}) = V_n \otimes F|-n\rangle_{A_w^I}$.

is independent of ρ . Define $W_n := \mathcal{U}(V_n^*)$ and \mathcal{W}_n as the associated unitary channel, and set $\mathcal{P}_{\text{acc}} := P_{\text{acc}}(\cdot)P_{\text{acc}}$ for the projector P_{acc} onto U 's practical output space. Then, by equating the cases where $\rho = F|n\rangle\langle n|F^\dagger$ and $\rho = F|0\rangle\langle 0|F^\dagger$, one can easily show that:

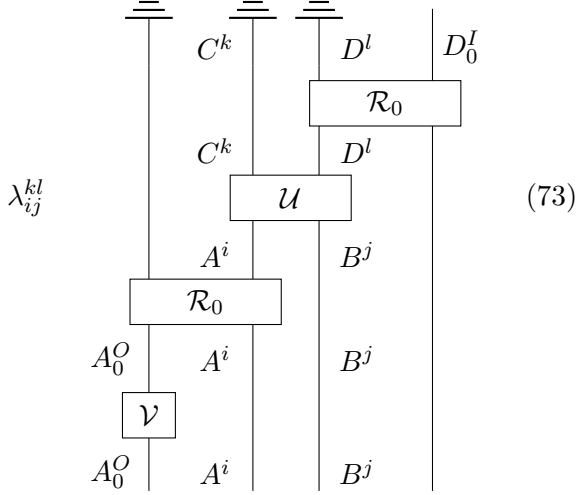


where we choose τ to be $\tau = |0\rangle\langle 0|$, and for all density matrices σ on U 's practical output space.


Consider the sectorization of the practical output space of U given by $H_{\text{acc}} = \bigoplus_l (c_l \otimes D_\star^l)$, where each c_l is an orthogonal subspace of C . In the following, we consider states in this space, where a subscript l on a state vector indicates that it is wholly contained in the corresponding sector $c_l \otimes D_\star^l$. Suppose, for contradiction, that $\exists |\psi_0\rangle : W_n |\psi_0\rangle = \sum_{l=0}^{M-1} a_l |\phi_l\rangle$, where at least one of the complex numbers a_m for $m \neq 0$ is non-zero. If we substitute $\sigma = |\psi_0\rangle\langle\psi_0|$ into (72), it is easy to show that the resulting state on the left-hand side is $\sum_{l=0}^{M-1} |a_l|^2 |l\rangle\langle l|$, whereas the right-hand one is $|0\rangle\langle 0|$, violating the equality. Hence $\mathcal{U}(V_n^*)$ maps $c_0 \otimes D_\star^0$ to itself. Of course, the choice of the 0th sector was arbitrary, meaning that $\mathcal{U}(V_n^*)$ maps each of the M subspaces to itself, and is thus block diagonal in the sectorization $H_{\text{acc}} = \bigoplus_l c_l \otimes D_\star^l$. Hence it commutes with \mathcal{D}^{w*} . Since this is true for all n , and $\{W_n\}_{n=0}^{D-1}$ spans \mathcal{A}^{w*} , it follows that $[U\mathcal{A}^{w*}U^\dagger, \mathcal{D}^{w*}] = 0$.

Now, suppose instead that both relata correspond to the zeroth sector, so our algebras are $\mathcal{X} = \mathcal{A}^{0*}$ and $\mathcal{Y} = \mathcal{D}^{0*}$. As discussed in the previous subsection, one can derive a version of

Sublemma E.1 for these relata in which the **SCNOT** channels in the sublemma are replaced with channels \mathcal{R}_0 corresponding to R_0 . This, together with (FG4), implies that the following is independent of the choice of unitary channel $\mathcal{V} = V(\cdot)V^\dagger$:



from which it follows that:



is also independent of \mathcal{V} .


Suppose that $|\psi\rangle$ is an eigenstate of the unitary operator V on A_0^O . We define a unitary operator V' on A :

$$V' = V_0 \oplus \lambda_\psi^V I_{\bar{0}} \quad (75)$$

where λ_ψ is the eigenvalue associated with $|\psi\rangle$, V_0 is a unitary operator of the form V on the sector A_\star^0 , and $I_{\bar{0}}$ is the identity operator on $A_{\bar{\star}}^0$.

We further define $W_{V'} := \mathcal{U}(V'^*)$ and denote by $\mathcal{W}_{V'}$ the corresponding unitary channel.

By inserting $|\psi\rangle\langle\psi|$ into the input A_0^O in (74) and equating the general case to the special case where V is the identity, we deduce that:




for arbitrary states σ and τ , where the route $\gamma_{kl}^{mn} := \sum_{ij} \lambda_{ij}^{kl} (\lambda^T)_{mn}^{ij}$.

We will proceed to show that $W_{V'}$ commutes with D^\star , starting by showing that it is block diagonal in the sectorization (55) of U 's practical output space for $l = 0$. Consider some states $|\psi_0\rangle$ and $|\phi_0\rangle$ in $H_0 := c_0 \otimes D_\star^0$, and $|\phi_{\bar{0}}\rangle$ in the complement space $H_{\bar{0}} := \bigoplus_{l \neq 0} c_l \otimes D_\star^l$. Suppose, for contradiction, that $W_{V'} |\psi_0\rangle = a |\phi_0\rangle + b |\phi_{\bar{0}}\rangle$ for $b \neq 0$. Then, substituting $\sigma = |\psi_0\rangle\langle\psi_0|$ into both sides of (76) gives that $|a|^2 \text{Tr}_C |\phi_0\rangle\langle\phi_0|_{CD_0^I} + |b|^2 \tau_{D_0^I} = \text{Tr}_C |\psi_0\rangle\langle\psi_0|_{CD_0^I}$. It is always possible to choose a $\tau_{D_0^I}$ so that this equality is violated. Hence $W_{V'}$ maps H_0 to itself. Then unitarity implies that $W_{V'}$ is block diagonal in the sectorization $H_{\text{acc}} = H_0 \oplus H_{\bar{0}}$.

We can therefore write

$$W_{V'} = W_{V'}^0 + W_{V'}^{\bar{0}} \quad (77)$$

where the operators on the right-hand side act on H_0 and $H_{\bar{0}}$ respectively. Now, by considering arbitrary choices of σ contained in $\mathcal{L}(c_0 \otimes D_\star^0)$, we can use tomography to deduce from (76) that



to which we can apply the essential uniqueness of purification to deduce that $W_{V'}^0 = T_{c_0} \otimes I_{D_\star^0}$ where $I_{D_\star^0}$ is the identity on D_\star^0 and T_{c_0} is some unitary on c_0 . Substituting this into (77), we can see that $W_{V'}$ commutes with \mathcal{D}^{0*} .

We now want to show that the unitary operators of the form V' from (75) span \mathcal{A}^0 . We will do this by explicitly showing how one can construct any operator from \mathcal{A}^0 by taking linear combinations of operators of the form V' .

In the case where V is the identity operator (up to global phase) on A_0^O , V' is the identity (up to global phase) on A . When V is not the identity (up to global phase), it has at least two distinct eigenvalues. We can therefore obtain the following two operators from V , assuming that A_0^O has a dimension greater than one (i.e. that the relatum A_\star^0 is nontrivial).

$$\begin{aligned} V_0 \oplus e^{i\alpha} I_{\bar{0}} \\ V_0 \oplus e^{i\beta} I_{\bar{0}} \end{aligned} \quad (79)$$

for $e^{i\alpha} \neq e^{i\beta}$. Subtracting the bottom operator from the top then gives us the operator $\mathbf{0}_0 \oplus (e^{i\alpha} - e^{i\beta}) I_{\bar{0}}$, which of course spans all operators of the form

$$\mathbf{0}_0 \oplus a I_{\bar{0}} \quad (80)$$

for arbitrary complex numbers a . These operators together with those in (79) span the operators of the form:

$$V_0 \oplus a I_{\bar{0}} \quad (81)$$

for unitary V_0 .

Note that since we can start with any unitary V , we can end up with any unitary V_0 in this expression. The operators of this form for unitary V_0 span the linear operators of the form:

$$M \oplus a I_{\bar{0}} \quad (82)$$

(as is easily seen from the fact that the unitary operators on a Hilbert space span the linear operators on that space). But these are precisely the operators in \mathcal{A}^0 . This shows that the operators of the form V' span \mathcal{A}^0 . As discussed in Appendix C, this implies that operators of the form V'^* span \mathcal{A}^{0*} . (FG3) follows.

The cases where the two relata are different in kind are proven by simply mixing arguments for the two cases in which they are the same. For example, if $\mathcal{X} = \mathcal{A}^{w*}$ and $\mathcal{Y} = \mathcal{D}^{l*}$, then one repeats the argument from the first case that one

can induce a set of unitary transformations of the form V_n from (69) on A by putting in states to the ancilla input of the SCNOT gate of the form $F|n\rangle$, and that the associated embedded unitaries V_n^* span \mathcal{A}^{w*} . Then one applies the argument above to show that (FG4) implies that these operators, once transformed by \mathcal{U} to the practical output space, must commute with \mathcal{D}^{l*} , implying that $[\mathcal{U}(\mathcal{A}^{w*}), \mathcal{D}^{l*}]$. Upon making a similar adaptation for $\mathcal{X} = \mathcal{A}^{i*}$ and $\mathcal{Y} = \mathcal{D}^{w*}$, the theorem is proved.

F Proof of Theorem 3.2

We will prove Theorem 3.2 by assuming Lemma E.1, which is proven independently in Appendix E.

Given the routed unitary (18), denote by (fg) the proposition that a relation of no-influence holds between each sectorized relatum associated with A^i and each sectorized relatum associated with D^l . We will show that (R2) is equivalent to (R3), that (R3) and (R4) are each equivalent to (fg), that (R3) implies (R1) and that (R1) implies (R4).

We start with (R2) \Leftrightarrow (R3). The rightward direction follows from the fact that the right side of (R2) commutes with \mathcal{D}^* , and the fact that $\mathcal{U}(\mathcal{A}^*)$ is spanned by its unitaries. For the leftward direction, consider an arbitrary unitary $V \in \mathcal{A}$, and the associated unitary V^* embedded in U 's practical input space (see Appendix C for the definition of the embedding map). Consider also the decomposition of U 's practical output space of the form $H_{\text{acc}} = \bigoplus_l c_l \otimes D_\star^l$, where each c_l is a subspace of C , and consider the projectors $Q_l \otimes \pi_l^D$ onto each term. (R3) implies that $\mathcal{U}(V^*)$ commutes with all of these projectors. Hence $\mathcal{U}(V^*)$ is block diagonal in the sectorization $H_{\text{acc}} = \bigoplus_l c_l \otimes D_\star^l$, and may be written $\mathcal{U}(V^*) = \bigoplus_l W_l$. (R3) also implies that each W_l commutes with an arbitrary operator of the form $I_{c_l} \otimes M_{D_\star^l}$, which implies that each W_l can be written $W_l = W'_{c_l} \otimes I_{D_\star^l}$. Hence $\mathcal{U}(V^*) = \bigoplus_l W'_{c_l} \otimes I_{D_\star^l}$, from which (R2) follows.

The fact that (R3) \Rightarrow (fg) follows simply from Lemma E.1, and the fact that the algebras associated with the sectors and which-sector information (see Appendix E) are subalgebras of the algebra associated with the sectorized system. For the converse direction, note that the union of

all those subalgebras spans the algebra associated with the sectorized system. In other words, denoting the union of the algebras associated to the multi-dimensional sectors and which-sector information of A^i as S_A , S_A spans \mathcal{A} . Similarly, S_D spans \mathcal{D} . It follows from S_A spanning \mathcal{A} that S_A^* spans \mathcal{A}^* (see Appendix C). By (fg) and Lemma E.1, we have $[\mathcal{U}(S_A^*), S_D^*] = 0$. (R3) follows. This proves that (R3) \Leftrightarrow (fg).

For (fg) \Leftrightarrow (R4), consider the transformation in (25). By considering U 's practical output space $\bigoplus_{kl} \sum_{ij} \lambda_{ij}^{kl} C_{\star}^k \otimes D_{\star}^l$ as a single Hilbert space, and similarly with its practical input space, one can think about (25) as an *unrouted* unitary transformation. The fact that (fg) \Leftrightarrow (R4) is then straightforward consequence of Proposition 3.2.

The fact that (R3) \Rightarrow (R1) is given by exactly the same argument as the proof that (FG3) \Rightarrow (FG1), except that now Φ and Ψ can now be arbitrary sector-preserving quantum channels as in (R1). One just needs to note that since Φ is sector-preserving, its Kraus operators are drawn from \mathcal{A} , and since Ψ is sector-preserving, its Kraus operators are drawn from $\mathcal{D} \otimes \mathcal{L}(* \rightarrow D')$. Both of these facts follow from Theorem 6 of [55].

The argument for (R1) \Rightarrow (R4) is very similar to the one encapsulated by equations (63 – 65) upon replacing $\mathcal{SCN}OT$ with exchange channel \mathcal{EXCH} ; in fact it is slightly simpler since there is no need for Sublemma E.1.

G Proof of Proposition 3.4

To prove (22), note that the practical output space of the routed unitary (21) can always be written in the form

$$H_{\text{acc}} = \bigoplus_{lmn} \omega^{lmn} D_{\star}^l \otimes E_{\star}^m \otimes F_{\star}^n \quad (83)$$

where ω is a route matrix. Consider the sector-preserving algebras $\mathcal{E} := \bigoplus_m \mathcal{L}(E_{\star}^m)$ and $\mathcal{F} := \bigoplus_n \mathcal{L}(F_{\star}^n)$ over E^m and F^n respectively. Then the sector-preserving algebra for the composite system $E^m \otimes F^n$ is $\mathcal{X} := \bigoplus_{mn} \mathcal{L}(E^m \otimes F^n) = \mathcal{E} \otimes \mathcal{F}$.

In Appendix C, we defined a map $\overset{*}{\rightarrow}$ that embeds sector-preserving algebras into a practical output space. The embedded algebras \mathcal{E}^* , \mathcal{F}^*

and \mathcal{X}^* have the form:

$$\begin{aligned} \mathcal{E}^* &= \bigoplus_{lmn} \omega^{lmn} \mathcal{E}_{lmn}^* \\ \mathcal{F}^* &= \bigoplus_{lmn} \omega^{lmn} \mathcal{F}_{lmn}^* \\ \mathcal{X}^* &= \bigoplus_{lmn} \omega^{lmn} \mathcal{X}_{lmn}^* \end{aligned} \quad (84)$$

where we have defined:

$$\begin{aligned} \mathcal{E}_{lmn}^* &:= I_{D_{\star}^l} \otimes \mathcal{L}(E_{\star}^m) \otimes I_{F_{\star}^n} \\ \mathcal{F}_{lmn}^* &:= I_{D_{\star}^l} \otimes I_{E_{\star}^m} \otimes \mathcal{L}(F_{\star}^n) \\ \mathcal{X}_{lmn}^* &:= I_{D_{\star}^l} \otimes \mathcal{L}(E_{\star}^m \otimes F_{\star}^n) \end{aligned} \quad (85)$$

where $I_{D_{\star}^l}$ is the identity operator on D_{\star}^l , etc.

Now, consider an arbitrary operator $M_{lmn} \in \mathcal{X}_{lmn}^*$. Since the product operators on the space $E_{\star}^m \otimes F_{\star}^n$ span $\mathcal{L}(E_{\star}^m \otimes F_{\star}^n)$, it is easily deduced from the algebra forms (85) that M_{lmn} can be written

$$M_{lmn} = \sum_i \alpha_i (N_i^{lmn} \circ R_i^{lmn}) \quad (86)$$

where $N_i^{lmn} \in \mathcal{E}_{lmn}^*$ and $R_i^{lmn} \in \mathcal{F}_{lmn}^*$ for each i .

Since any operator in \mathcal{X}^* is just a direct sum of operators each from some \mathcal{X}_{lmn}^* , it follows that any $M' \in \mathcal{X}^*$ can be written in the form:

$$\begin{aligned} M' &= \bigoplus_{lmn} M_{lmn} \\ &= \bigoplus_{lmn} \left(\sum_i \alpha_i (N_i^{lmn} \circ R_i^{lmn}) \right) \end{aligned} \quad (87)$$

where all $N_i^{lmn} \in \mathcal{E}_{lmn}^*$ and all $R_i^{lmn} \in \mathcal{F}_{lmn}^*$. If we define a new set of operators of the form $\tilde{N}_i^{lmn} := N_i^{lmn} \oplus (\bigoplus_{l'm'n' \neq lmn} \mathbf{0}_{l'm'n'})$, where $\mathbf{0}_{l'm'n'}$ is the zero operator in $\mathcal{E}_{l'm'n'}^*$, and similarly define \tilde{R}_i^{lmn} , then the direct sum above turns into a regular sum:

$$M' = \sum_{lmni} \alpha_i (\tilde{N}_i^{lmn} \circ \tilde{R}_i^{lmn}) \quad (88)$$

where each $\tilde{N}_i^{lmn} \in \mathcal{E}^*$ and each $\tilde{R}_i^{lmn} \in \mathcal{F}^*$.

Now, suppose that $(A^i \xrightarrow{Y} E^m) \wedge (A^i \xrightarrow{Y} F^n)$. By (R3), this means that

$$[\mathcal{U}(\mathcal{A}^*), \mathcal{E}^*] = [\mathcal{U}(\mathcal{A}^*), \mathcal{F}^*] = 0 \quad (89)$$

But (88) shows that each operator from \mathcal{X}^* can be written as a sum of sequential compositions of operators from \mathcal{E}^* and \mathcal{F}^* . It follows that $[\mathcal{U}(\mathcal{A}^*), \mathcal{X}^*] = 0$, implying that $A^i \xrightarrow{Y} E^m \otimes F^n$.

Conversely, suppose $A^i \not\stackrel{U}{\rightarrow} E^m \otimes F^n$, and thus $[\mathcal{U}(\mathcal{A}^*), \mathcal{X}^*] = 0$. It is easily deduced from the expressions (84) and (85) that \mathcal{E}^* and \mathcal{F}^* are subalgebras of $(\mathcal{E} \otimes \mathcal{F})^* = \mathcal{X}^*$. This means that $[\mathcal{U}(\mathcal{A}^*), \mathcal{E}^*] = 0$ and $[\mathcal{U}(\mathcal{A}^*), \mathcal{F}^*] = 0$, so we have $(A^i \stackrel{U}{\rightarrow} E^m) \wedge (A^i \not\stackrel{U}{\rightarrow} F^n)$. This proves (22).

The other part of atomicity then follows from time symmetry. More precisely, (23) then follows from (20) and (22) via essentially the manipulations in equation (40) from Appendix B we replace the Hilbert spaces with sectorized Hilbert spaces.

H Proof of Proposition 4.2

To show that the causal structure of a general routed unitary supermap is time symmetric, let us start by explaining why it is so for the special case of an unrouted unitary supermap. The process channel for an unrouted supermap is

$$\mathcal{S}^\dagger(\text{SWAP}, \dots, \text{SWAP}) = \mathcal{S}(\text{SWAP}, \dots, \text{SWAP})^\dagger \quad (90)$$

since the swap gate is self-adjoint. In other words, the process channel for the adjoint of the supermap is just the adjoint of the process channel for the original supermap. But we already know, from Proposition 12, that the causal structure of a unitary *channel* is time symmetric. It then follows from the definition of the causal structure of a unitary supermap in terms of its process channel that it is also time symmetric.

In summary, the fact that the swap gate is self-adjoint implies that the causal structure of an unrouted unitary supermap is time symmetric. If the exchange gate were self-adjoint, then we would be able to make an essentially identical argument. But it isn't: although **RSWAP** is self-adjoint, the other gate needed to define the exchange gate, **SCNOT**, is not (see Appendix D for the definition of **EXCH** in terms of these two unitaries). This apparently raises the possibility of time asymmetry in the routed case.

But in fact, **SCNOT** has a weaker property than self-adjointness, which is enough to secure time symmetry. Specifically, it is equal to its own adjoint up to a local unitary applied to both its ancillary input and output. Defining the negation unitary $U = \sum_j |-j\rangle\langle j|$, we have

$$\text{SCNOT} = (I_A \otimes U_{A'}) \text{SCNOT}^\dagger (I_A \otimes U_{A''}) \quad (91)$$

To prove (91), we note that $\text{SCNOT}^\dagger = \sum_{ij} \pi_A^i \otimes |j\rangle\langle j+i| = \sum_{ij} \pi_A^i \otimes |j-i\rangle\langle j|$, and calculate the action of the right-hand side to find

$$|i\rangle|j\rangle \rightarrow |i\rangle|-j\rangle \rightarrow |i\rangle|-j-i\rangle \rightarrow |i\rangle|j+i\rangle \quad (92)$$

Now, as (R1) makes clear, changing a unitary channel by composing it with another unitary that is both

- a) local (meaning that it acts only on one of the original unitary's input or output subsystems); and
- b) sector-preserving

does not alter its causal relations. The negation unitary U above is trivially sector-preserving since it acts on systems that only have one sector. Therefore, the process channel

$$\mathcal{S}(\text{EXCH}, \dots, \text{EXCH}) \quad (93)$$

for a routed unitary supermap \mathcal{S} has the same (sectorized and unsectorized) causal structure as

$$\mathcal{S}(\text{EXCH}^\dagger, \dots, \text{EXCH}^\dagger). \quad (94)$$

But by the time symmetry of the causal structure of routed unitary channels, $\mathcal{S}(\text{EXCH}^\dagger, \dots, \text{EXCH}^\dagger)$ has the causal structure obtained by reversing the arrows in the causal structure of $\mathcal{S}(\text{EXCH}^\dagger, \dots, \text{EXCH}^\dagger)^\dagger$ – the process operator for the time-reversed supermap \mathcal{S}^\dagger . It follows from the definition of the causal structure of routed unitary supermaps in terms of their process channels that the causal structure of \mathcal{S}^\dagger is obtained from the causal structure of \mathcal{S} by reversing the direction of the arrows.