

Giving Operational Meaning to the Superposition of Causal Orders

Nicola Pinzani

University of Oxford

nicola.pinzani@cs.ox.ac.uk

Stefano Gogioso

University of Oxford

stefano.gogioso@cs.ox.ac.uk

In this work, we give rigorous operational meaning to superposition of causal orders. This fits within a recent effort to understand how the standard operational perspective on quantum theory could be extended to include indefinite causality. The mainstream view, that of “process matrices”, takes a top-down approach to the problem, considering all causal correlations that are compatible with local quantum experiments. Conversely, we pursue a bottom-up approach, investigating how the concept of indefiniteness emerges from specific characteristics of generic operational theories. Specifically, we pin down the operational phenomenology of the notion of non-classical (e.g. “coherent”) control, which we then use to formalise a theory-independent notion of control (e.g. “superposition”) of causal orders. To validate our framework, we show how salient examples from the literature can be captured in our framework.

1 Introduction

When modelling the interaction of localised processes (aka “operations” or “experiments”) in spacetime, the causal structure between the corresponding events can be captured by circuit-like diagrams. Such interactions arise as combinations of parallel and sequential compositions of processes, relating to space-like and time-like separation of events respectively: if the processes themselves respect a notion of *causality*—defined as the impossibility of transmitting information from the future to the past—then it can be shown that causal relations between the events at which the processes take place can unambiguously be described by a directed acyclic graph, each edge capturing the forward-in-time information flow from the output of one process to the input of another [9, 12, 23]. For example, taking quantum instruments as processes/operations respects such a notion of causality, but adding post-selection breaks it, resulting in the possibility of signalling from the future.

A 1977 result by Malament [16] shows that—in the continuous limit—little more than the above would be needed to reconstruct the structure of spacetime itself: knowledge of the causal relationship between its events is enough to reconstruct a spacetime up to conformal equivalence, as long as that spacetime satisfies some mild requirement.¹ An extension of this identification of causal structure with directed graphs was recently used by the authors to study operational models of quantum information in the presence of causal anomalies, such as closed time-like curves [23].

In this work, we examine another scenario in which the operational approach to quantum theory needs to be extended in order to cope with exotic causal structure: that where the background spacetime is itself allowed to be in a superposition. In doing so, we also dispel some of the fog surrounding the notion of *control of causal order*: this is the situation in which the outcome of one process may influence the

¹Specifically, the requirement is for the spacetime to be both past-distinguishing and future-distinguishing: two events are the same if they have the same causal future or the same causal past. This excludes a number of causally problematic scenarios, such as the existence of closed time-like curves.

causal order of subsequent processes, both classically (as a probabilistic mixture) and coherently (as a superposition).

Reasoning about operational theories in the context where spacetime itself becomes a dynamical variable carries a number of additional complications. For example, thinking about the set-up and outcomes of an experiment presupposes the existence of causally stable surroundings, where the notions of cause and consequence take their familiar form independently of the specific processes being performed. Failing these assumptions, how can we make sure that a mathematical model of quantum theory in the presence of dynamical spacetimes is even empirically testable? This is an important question, upon which many others stumbled before us. For example, the following reflection can be found in a prominent piece of literature on the application of sheaves and topoi to quantum theory [11, 15]:

“[A]round fifteen years ago, I came to the conclusion that the use of standard quantum theory was fundamentally inconsistent, and I stopped working in quantum gravity proper [...] [W]hat could it mean to ‘measure’ properties of space or time if the very act of measurement requires a spatiotemporal background within which it is made?” (Chris J. Isham [15])

In this work, we set out to endow superposition of causal orders with rigorous operational meaning, through the development of suitable categorical semantics.

We start by considering a probabilistic theory of processes, modelling the operational aspects of some physical theory (such as quantum theory). Given such a theory, we want to provide a sound way of constructing and characterising localised processes which don’t take place against a fixed causal background, but rather against a “superposition” of causal backgrounds, determined by some “wave-function” over the set of all fixed causal backgrounds compatible with the processes in question. We do so by first defining a theory-independent notion of “control” of processes, accommodating both the classical case—where the choice of process to execute can be captured by some hidden variable—and the coherent case—where it is not possible to establish which one process was executed without reference to a specific measurement context. Armed with such a notion, we show how to construct superpositions of diagrams, giving categorical semantics to the execution of processes and operations against a dynamic causal background.

A significant body of literature exists which investigates the informational advantage of superposition of channels and causal orders [1, 2, 5, 19, 21, 24] and the possibility of experimentally detecting such superpositions [6, 17, 18, 26]. In particular, there has been some recent speculation about the possibility of realising a genuine “quantum switch” by coherently controlling the ‘spatial degree of freedom’ [22]. Unfortunately, however, the issue of delimiting a tight operational setting in which to interpret the results is not ordinarily viewed as a necessity, sometimes leading to misinterpretation of their physical significance. The goal of this paper is then to provide a rigorous standpoint from which to discuss the operational phenomenology of causal superposition and control. If quantum gravity is ever to be observed, we must first be in possession of the mathematical tools to rigorously understand, debate and communicate those observations.

The categorical semantics in this work are given within the process-theoretic framework for probabilistic theories introduced by [14], which is briefly summarised in Appendix A. The same semantics can easily be adapted to other popular frameworks, such as that of Operational Probabilistic Theories (OPTs) [3, 10]. The main differences between the framework presented here and that of OPTs are as follows:

- (i) normalisation is presented constructively, by choice of discarding map and an equation defining the normalised (aka “deterministic”) processes;
- (ii) the whole convex cone of processes is considered, including super-normalised components which can be used to decompose processes into simpler building pieces (e.g. as in the ZX calculus) [7];

- (iii) classical systems are explicitly considered as systems within the framework, e.g. allowing for diagrammatic treatment of the quantum-classical interface;
- (iv) the framework straightforwardly extends to semirings other than the probabilistic semiring \mathbb{R}^+ , allowing for discussion of quasi-probabilistic theories (semiring \mathbb{R}), possibilistic theories (boolean semiring) as well as more exotic examples including modal, hyperbolic or p -adic quantum theory [13].

Despite the presentational differences and the added flexibility, the framework used here contains all causal OPTs as special cases, allowing for all results to be transferred straightforwardly (as long as only normalised and sub-normalised maps are involved).

2 Controlled processes

As the first step towards our formulation of semantics for superposition of causal orders, we define a general notion of control of morphisms valid in arbitrary probabilistic theories. Our definition captures the idea of an agent being able to control the choice of morphisms by suitably encoding and/or decoding classical information about their choice into and/or from a physical system.

Definition 1. Let C be a probabilistic theory, let $A, B \in \text{obj}(C)$ be any two systems in the theory. Let $(F_x)_{x \in X}$ be a family of processes $F_x : A \rightarrow B$, not necessarily normalised or sub-normalised. A **controlled process** for the family is a triple (G, p, m) consists of a sharp preparation-observation (SPO) pair $(p : X \rightarrow H, m : H \rightarrow X)$ —where X is a classical system and H is a generic system—together with a process $G : H \otimes A \rightarrow H \otimes B$ satisfying the following equations:

$$\begin{array}{c}
 \begin{array}{c} H \quad B \\ | \quad | \\ \boxed{G} \\ | \quad | \\ \boxed{p} \\ \vdots \\ X \quad A \end{array} = \sum_{x \in X} \begin{array}{c} H \quad B \\ | \quad | \\ \boxed{p} \\ | \quad | \\ \triangleleft x \\ | \quad | \\ \triangleleft x \\ \vdots \\ X \quad A \end{array} \quad \boxed{F_x} \\
 \end{array}
 \qquad
 \begin{array}{c}
 \begin{array}{c} X \quad B \\ | \quad | \\ \boxed{m} \\ | \quad | \\ \boxed{G} \\ | \quad | \\ H \quad A \end{array} = \sum_{x \in X} \begin{array}{c} X \quad B \\ | \quad | \\ \triangleleft x \\ | \quad | \\ \triangleleft x \\ | \quad | \\ \boxed{m} \\ \vdots \\ H \quad A \end{array} \quad \boxed{F_x} \quad (1)
 \end{array}$$

Con conversationally, we will also say that (G, p, m) is a “control of” the family $(F_x)_{x \in X}$.

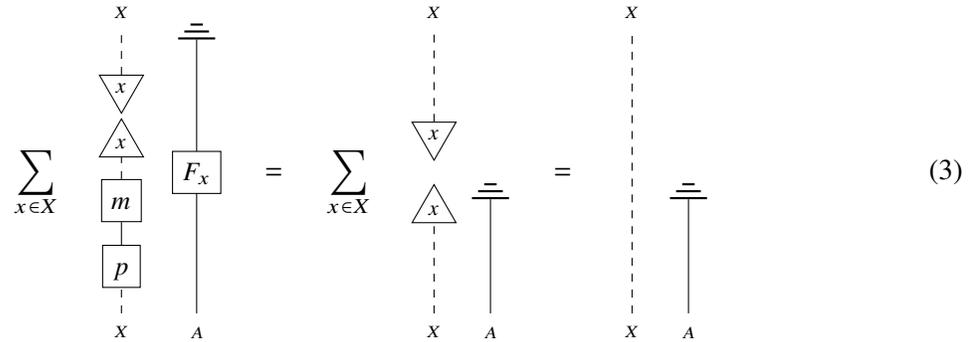
The system H acts as a physical control system, while the classical system X contains the logical information about the process choice. The SPO pair is used to encode the logical information into the physical system and/or to decode it from the physical system. From the point of view of an actor who is using the SPO pair to encode/decode the logical information, the controlled process is no different than classical control.

Example 2. The following **classically controlled process** always exists in every probabilistic theory:

$$\sum_{x \in X} \begin{array}{c} X \quad B \\ | \quad | \\ \triangleleft x \\ | \quad | \\ \triangleleft x \\ \vdots \\ X \quad A \end{array} \quad \boxed{F_x} \quad (2)$$

Con conversationally, we will also refer to the above as *the* “classical control of” the family $(F_x)_{x \in X}$. Note that if (G, p, m) is any control of $(F_x)_{x \in X}$ then the triple $((m \otimes \text{id}_B) \circ G \circ (p \otimes \text{id}_A), \text{id}_X, \text{id}_X)$ is always the classical control of the same $(F_x)_{x \in X}$.

The definition allows for much more general notions of control, as we shall shortly see, but it also limits the amount of leakage between the input/output systems A, B and the physical control system H . In particular, an agent without access to the output system B cannot, through the SPO pair alone, extract any information about the input state on system A if the maps F_x are normalised:



$$\sum_{x \in X} \begin{array}{c} x \\ \vdots \\ \triangleleft x \\ \triangleleft x \\ \square m \\ \square p \\ \vdots \\ X \end{array} \begin{array}{c} \overline{\overline{}} \\ \vdots \\ \square F_x \\ \vdots \\ A \end{array} = \sum_{x \in X} \begin{array}{c} x \\ \vdots \\ \triangleleft x \\ \triangleleft x \\ \vdots \\ X \end{array} \begin{array}{c} \overline{\overline{}} \\ \vdots \\ \square F_x \\ \vdots \\ A \end{array} = \begin{array}{c} x \\ \vdots \\ \vdots \\ X \end{array} \begin{array}{c} \overline{\overline{}} \\ \vdots \\ \square F_x \\ \vdots \\ A \end{array} \quad (3)$$

3 Coherent Control in Quantum Theory

In quantum information and computing, the idea of coherently controlling a family of unitary processes (and more generally pure CP maps) is certainly not a new one [1, 2, 19, 25].

Example 3. If $(F_x)_{x \in X}$ is a family of pure CP maps in quantum theory (e.g. isometries), a generic **coherently controlled process** for the family uses \mathbb{C}^X as a control system and takes the following form:



$$\begin{array}{c} | \\ \circ \alpha \\ | \end{array} \begin{array}{c} | \\ \square F \\ | \end{array} \quad (4)$$

where α denotes an arbitrary phase in the canonical basis $(|x\rangle)_{x \in X}$ for the control system. Conversationally, we will also say that the above is a “coherent control of” the family $(F_x)_{x \in X}$.

Proposition 4. Let $(F_x)_{x \in X}$ be a family of pure CP maps in quantum theory, not necessarily normalised (i.e. not necessarily trace-preserving). Assume that (G, p, m) is a controlled process for the family with control system \mathbb{C}^X , where (p, m) is the SPO for the canonical basis $(|x\rangle)_{x \in X}$ and where G is itself a pure CP map. Then G takes form (4) for some phase α .

Moving away from pure maps, the question arises what the “coherent” control for arbitrary CP maps should be. We adopt the following definition within our framework.

Definition 5. Let $(F_x)_{x \in X}$ be a family of CP maps in quantum theory, not necessarily normalised (i.e. not necessarily trace-preserving). A **coherently controlled process** for the family is a controlled process (G, p, m) for the family such that G is obtained as $\text{Tr}_E(G')$ for some pure CP map $G' : H \otimes A \rightarrow H \otimes B \otimes E$. Conversationally, we will also say that the above is a “coherent control of” the family $(F_x)_{x \in X}$.

If each map F_x in a family $(F_x)_{x \in X}$ of CP maps comes with a chosen purification $F_x = \text{Tr}_E(\hat{F}_x)$ —without loss of generality, using the same environment system E for all purifications—then it is easy to construct a

coherent control of the family, by taking (G', p, m) to be the coherent control of the family $(\hat{F}_x)_{x \in X}$ of purifications and then discarding the environment E of G' . However, it would be desirable for such a construction to be a function of the family $(F_x)_{x \in X}$ alone, without dependence on additional information. As shown by Proposition 4, this is indeed possible when all CP maps in the family are pure. However, the following no-go result shows this to no longer be the case when generic CP maps are considered.

Proposition 6. *Let $(F_x)_{x \in X}$ be a family of CP maps in quantum theory. It is not generally possible to construct a coherent control of the family which is a function of the family $(F_x)_{x \in X}$ alone, i.e. one which is independent of a choice of purification for the CP maps in the family. This is the same as the statement that it is not generally possible to construct a coherent control $(\text{Tr}_E(G'), p, m)$ of the family $(F_x)_{x \in X}$ in such a way that the following equation holds for all choices of unitaries $U_x : E \rightarrow E$:*

$$(5)$$

Note that the unitaries $(U_x)_{x \in X}$ correspond to all possible choices of purification for the CP maps $(F_x)_{x \in X}$.

The classical control of arbitrary families of CP maps is trivially possible, as shown in Example 2. The formulation of Proposition 6 shows that, on the other hand, the question of coherently controlling families of CP maps is much more sophisticated, leading to some confusion in the literature about its feasibility.

In [19], for example, the author interprets the failure to construct such a coherent control independently of the choice of purification (aka choice of Kraus operators) as a sign that an interferometric realisation of such coherent control would extract information about the underlying physical implementation of the CP maps themselves. We believe that this statement can be easily misinterpreted: the CP maps involved in the experiment are already the “physical” ones—defined on the direct sum of the vacuum sector and the 1-particle sector—and the results of the experiment are independent of the choice of purification for them. This is obvious, since the experiment itself can be easily written as a circuit.

What the results of the experiment actually depend on is the choice of purification for the “logical” CP maps involved, those restricted to the 1-particle sector. This is due to the specific design of the experiment: the implementation of the “physical” CP maps is such that they react to the vacuum state on their input by emitting a non-vacuum state $|e\rangle$ on the environment E . This is not physically unreasonable, e.g. if the environment system comprises some static massive particle which is made to interact with the photons passing in the interferometric setup. However, this dependence on the choice of purification for the “logical” CP maps goes away as soon as we allow the environment “rest” state $|e\rangle$ to be transformed covariantly with the choice of purification, i.e. if we set $|e\rangle \mapsto U|e\rangle$ whenever we change the purification by applying a unitary $U : E \rightarrow E$ to the environment.

It is this last observation which helps us frame the discussion by [19] within the context of Proposition 6: given the two CP maps, acting on the vacuum and non-vacuum sectors respectively, it is very much possible to find a coherent control which is invariant under application of the same unitary U to the environment of both purifications. This is always the case: if all $U_x : E \rightarrow E$ are chosen to be equal to some fixed U , then Equation (5) always holds (because U is trace-preserving). What is found to be impossible in the discussion by [19] is to choose such coherent control in a way which is invariant under application of U

to the environment of the non-vacuum sector map and of the identity to the environment of the vacuum sector map. This issue indeed generalises and formed the inspiration for our proof of Proposition 6.

4 Definite and Indefinite Causal Scenarios

When operational scenarios with definite causal order are depicted diagrammatically in the context of probabilistic theories, it is easy to conflate the boxes in the diagrams with processes happening locally at events (i.e. points in spacetime), and the wires in the diagrams with the information flow establishing the causal relationships between said events. It has been previously argued [23] that this practice—though natural and notationally pleasant—is not mathematically well-founded, as there need not be a canonical way to decide how a process should be decomposed into a diagram compatibly with a given definite causal structure. As a consequence, two ingredients are needed when talking about such operational scenarios:

- (i) a causal graph, representing the events in the scenario and their definite causal order;
- (ii) a map assigning each event in the scenario to the process happening at that event.

The mathematical structure introduced in [23] as the substrate for such operational scenarios is that of *framed causal graphs*. For reasons which will become clear later on, we generalise the original definition to include the possibility of multiple edges between the same pairs of event. Furthermore, we include some additional information about the classical interface of the local processes/experiments, in the form of finite sets of input values that can be used to control them and output values for their outcomes.

Definition 7. A *framed multigraph* is an directed multigraph Γ ² equipped with the following data:

- a sub-set $\text{in}(\Gamma) \subseteq \text{nodes}(\Gamma)$ of the nodes of Γ —the **input nodes**—such that each $x \in \text{in}(\Gamma)$ has zero incoming edges and a single outgoing edge;
- a sub-set $\text{out}(\Gamma) \subseteq \text{nodes}(\Gamma)$ of the nodes of Γ —the **output nodes**—such that each $x \in \text{out}(\Gamma)$ has zero outgoing edges and a single incoming edge;
- a **framing** for Γ , which consists of the following:
 - a total order on $\text{in}(\Gamma)$;
 - a total order on $\text{out}(\Gamma)$;
 - for each node $x \in \text{nodes}(\Gamma)$, a total order on the edges outgoing from x ;
 - for each node $x \in \text{nodes}(\Gamma)$, a total order on the edges incoming to x ;

We refer to nodes in $\text{in}(\Gamma)$ or in $\text{out}(\Gamma)$ as **boundary nodes** and to all other nodes in Γ as **internal nodes**. An **acyclic framed multigraph** is a framed multigraph which is acyclic (and in particular has no loops).

Remark 8. The input and output nodes of a framed multigraph are designed to behave as “half-edges”: when two framed multigraphs Γ and Γ' are composed sequentially, the outputs of Γ and the inputs of Γ' are joined and disappear, each pair of corresponding output/input resulting in a single edge of the composite framed multigraph $\Gamma' \circ \Gamma$. (We do not use such composition here.)

Definition 9. A *definite causal scenario* is a triple $\Theta = (\Gamma, \underline{I}, \underline{O})$ of an acyclic framed multigraph Γ with:

- a finite set I_ω of **classical inputs** for each $\omega \in \text{ev}(\Theta)$, i.e. the values available locally to control the process at the event;

²A directed multigraph Γ consists of a set $\text{nodes}(\Gamma)$, a set of edges (Γ) and a pair of functions $\text{tail} : \text{edges}(\Gamma) \rightarrow \text{nodes}(\Gamma)$ and $\text{head} : \text{edges}(\Gamma) \rightarrow \text{nodes}(\Gamma)$ specifying the tail and head of each edge respectively.

- a finite set O_ω of **classical outputs** for each $\omega \in \text{ev}(\Theta)$, i.e. the values that the process at the event can return locally as its outcome.

In the above, we have defined the **events** in the scenario as the set $\text{ev}(\Theta) := \text{nodes}(\Gamma) \setminus (\text{in}(\Gamma) \sqcup \text{out}(\Gamma))$ of internal nodes for Γ .

Remark 10. Compared to the original [23], we have restricted our attention to *chronology respecting* scenarios, i.e. those corresponding to acyclic framed multigraphs. However, one could easily extend Definition 9 to one for *chronology violating* scenarios, by allowing the framed multigraph to be cyclic and/or to have loops.

We now define exactly what it means to “draw a diagram over” one such definite causal scenario, with semantics valid in any probabilistic theory.

Definition 11. Let $\Theta = (\Gamma, \underline{I}, \underline{O})$ be a definite causal scenario and let C be a probabilistic theory. A **diagram over Θ in C** is a pair of functions $\text{sys} : \text{edges}(\Gamma) \rightarrow \text{obj}(C)$ and $\text{proc} : \text{ev}(\Gamma) \rightarrow \text{mor}(C)$, associating each $e \in \text{edges}(\Gamma)$ to a system $\text{sys}(e)$ in C and each $\omega \in \text{ev}(\Gamma)$ to a process $\text{proc}(\omega)$ in C with the following type:

$$\text{proc}(\omega) : I_\omega \otimes \bigotimes_{e \in \text{in}(\omega)} \text{sys}(e) \longrightarrow O_\omega \otimes \bigotimes_{e' \in \text{out}(\omega)} \text{sys}(e') \quad (6)$$

Above we denoted by $\text{in}(\omega) := \{e \in \text{edges}(\Gamma) \mid \text{head}(e) = \omega\}$ the edges of Γ coming into ω and we similarly denoted by $\text{out}(\omega) := \{e \in \text{edges}(\Gamma) \mid \text{tail}(e) = \omega\}$ the edges of Γ going out of ω .

Even though it specifies concrete processes in a probabilistic theory, the definition of diagram Δ above is still partly abstract, as it does not explicitly state how the processes fit together. What gives it fully concrete semantics is the following definition of the overall process $\llbracket \Delta \rrbracket$ associated to the diagram Δ . See Figure 1 for an exemplification.

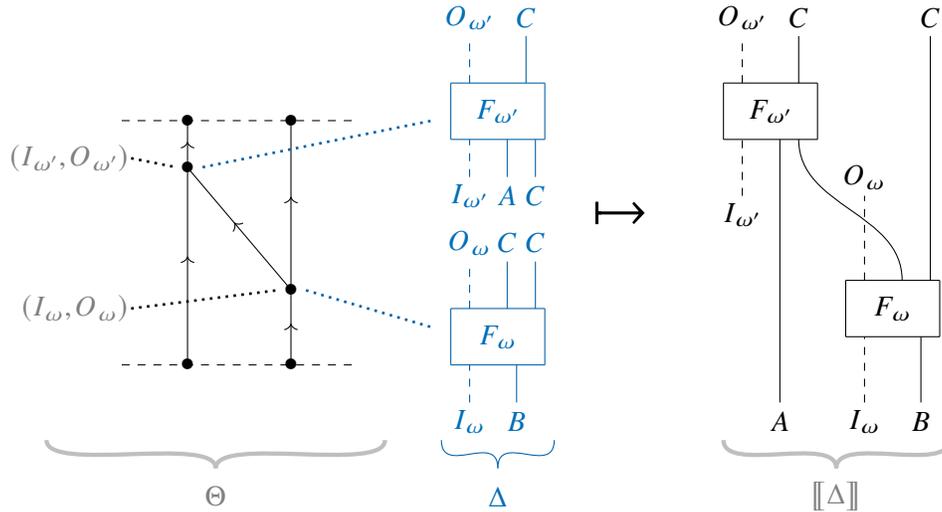


Figure 1: Graphical exemplification of the association of the process $\llbracket \Delta \rrbracket$ to a diagram Δ over a definite causal scenario Θ .

Definition 12. Let $\Theta = (\Gamma, \underline{I}, \underline{O})$ be a definite causal scenario, C be a probabilistic theory and $\Delta = (\text{sys}, \text{proc})$ be a diagram over Γ in C . The **process associated to** Δ is the unique process $\llbracket \Delta \rrbracket$ in C obtained by joining the outputs and inputs of the processes $\text{proc}(\omega)$ in the diagram Δ according to the directed multigraph Γ , resulting in a process with the following overall type:

$$\llbracket \Delta \rrbracket : \left(\left(\bigotimes_{\omega \in \text{ev}(\Gamma)} I_{\omega} \right) \otimes \left(\bigotimes_{e \in \text{in}(\Gamma)} \text{sys}(e) \right) \right) \rightarrow \left(\left(\bigotimes_{\omega \in \text{ev}(\Gamma)} O_{\omega} \right) \otimes \left(\bigotimes_{e' \in \text{out}(\Gamma)} \text{sys}(e') \right) \right) \quad (7)$$

Each classical input system I_{ω} of $\llbracket \Delta \rrbracket$ is wired to the classical input system I_{ω} of the process $\text{proc}(\omega)$. Similarly, each classical output system O_{ω} of $\llbracket \Delta \rrbracket$ is wired to the classical output system O_{ω} of $\text{proc}(\omega)$.

Definite causal scenarios and diagrams over them are perfectly adequate when it comes to discussion of operational scenarios over definite causal orders, but they don't have the necessary flexibility to accommodate operational scenarios where causality is indefinite. The main contribution of this work will now be to define diagrams over *indefinite causal scenarios*, giving them semantics in probabilistic theories using controlled processes. As a special case, we will be able to describe the idea of superposition of causal orders in quantum theory, i.e. we give solid mathematical foundations over which to discuss the possibility of operational scenarios where local quantum experiments are connected by a superposition of background spacetimes.

Firstly, we define the purely operational canvas against which the indefinite causal scenario takes places. This includes all the black-box information locally available to the actors in our scenario, but does not include any information about causal order nor any information about the specific implementation of the local processes.

Definition 13. An *indefinite causal scenario* Φ is specified by the following data:

- the set Ω of **events** at which the processes (operations, experiments, etc.) take place;
- a set Ξ of **system labels**, used to abstractly indicate which physical systems are guaranteed to be the same across different experiments;
- a finite set I_{ω} of **classical inputs** for each $\omega \in \Omega$, i.e. the values available locally to control the process at the event;
- a finite set O_{ω} of **classical outputs** for each $\omega \in \Omega$, i.e. the values that the process at the event can return locally as its outcome;
- a finite sequence $\underline{\Sigma}_{\omega}^{\text{in}}$ of elements of Ξ for each $\omega \in \Omega$, the system labels for the physical systems coming into the process from outside the event;
- a finite sequence $\underline{\Sigma}_{\omega}^{\text{out}}$ of elements of Ξ for each $\omega \in \Omega$, the system labels for the physical systems coming out of the process and leaving the event;
- a finite sequence $\underline{\Pi}^{\text{in}}$ of elements of Ξ , the system labels for the physical systems coming in from outside the region where the scenario is taking place;
- a finite sequence $\underline{\Pi}^{\text{out}}$ of elements of Ξ , the system labels for the physical systems going out from the region where the scenario is taking place.

Formally, the scenario is the tuple $\Phi = (\Omega, \Xi, \underline{I}, \underline{O}, \underline{\Sigma}^{\text{in}}, \underline{\Sigma}^{\text{out}}, \underline{\Pi}^{\text{in}}, \underline{\Pi}^{\text{out}})$, where the underlined letters indicate Ω -indexed families (e.g. $\underline{I} := \omega \mapsto I_{\omega}$).

Remark 14. Even though the definitions involved are formally different, the definite causal scenarios of Definition 9 arise naturally as a special case of the indefinite causal scenarios from Definition 13. Indeed, consider an indefinite causal scenario $\Phi = (\Omega, \Xi, \underline{I}, \underline{O}, \underline{\Sigma}^{in}, \underline{\Sigma}^{out}, \Pi^{in}, \Pi^{out})$ and assume that each symbol $e \in \Xi$ appears exactly twice as follows:

- it appears once either in an output set $\Sigma_{tail(e)}^{out}$ for some $tail(e) \in \Omega$ or otherwise at some place $tail(e) \in \{1, \dots, |\Pi^{in}|\}$ in the sequence Π^{in} ;
- it appears once either in an input set $\Sigma_{head(e)}^{in}$ for some $head(e) \in \Omega$ or otherwise at some place $head(e) \in \{1, \dots, |\Pi^{out}|\}$ in the sequence Π^{out} .

This defines a framed multigraph Γ with nodes $(\Gamma) := \Omega \sqcup \{1, \dots, |\Pi^{in}|\} \sqcup \{1, \dots, |\Pi^{out}|\}$ and edges $(\Gamma) := \Xi$. The internal nodes of Γ are the events in Ω , yielding a definite causal scenario $(\Gamma, \underline{I}, \underline{O})$. Conversely, each definite causal scenario $(\Gamma, \underline{I}, \underline{O})$ can be turned into an indefinite causal scenario by taking $\Xi := \text{edges}(\Gamma)$ and defining the sequences $\Sigma_{\omega}^{in}, \Sigma_{\omega}^{out}, \Pi^{in}$ and Π^{out} from $\text{in}(\omega), \text{out}(\omega), \text{in}(\Gamma)$ and $\text{out}(\Gamma)$ respectively.

Secondly, we define the set of definite causal scenarios which are *compatible* with a given indefinite causal scenario: they correspond exactly to all possible ways of joining the output and input physical systems into a multigraph in such a way as to respect the system labels for the indefinite causal scenario. Each indefinite causal scenario gives rise to a set of compatible definite causal scenarios, each definite causal scenario equipped with a labelling associating each edge of the multigraph to the corresponding system label from the indefinite causal scenario. See Figure 2 for an exemplification.

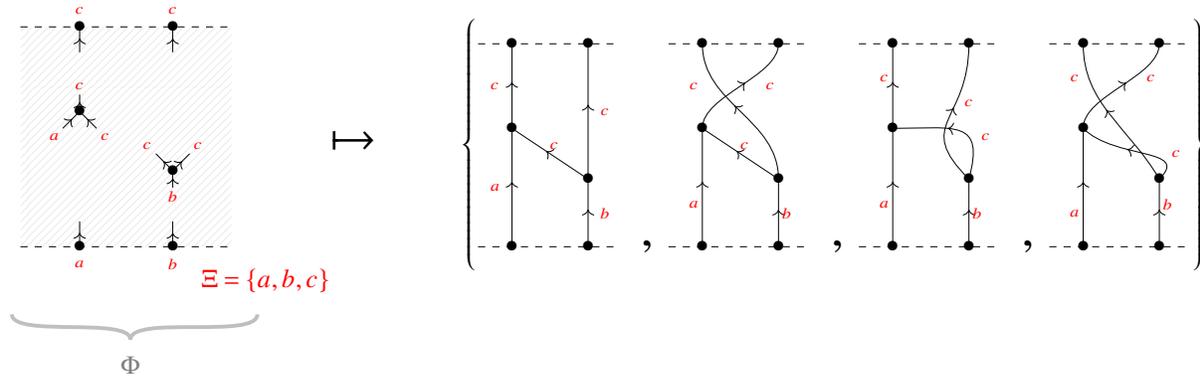


Figure 2: Graphical exemplification of the association of definite causal scenarios to an indefinite causal scenario Φ .

Definition 15. Let $\Phi = (\Omega, \Xi, \underline{I}, \underline{O}, \underline{\Sigma}^{in}, \underline{\Sigma}^{out}, \Pi^{in}, \Pi^{out})$ be an indefinite causal scenario. A definite causal scenario $\Theta = (\Gamma, \underline{I}', \underline{O}')$ is **compatible** with Φ if the following conditions hold:

- the events of Θ (internal nodes of Γ) are the events of the scenario, i.e. we have $\text{ev}(\Theta) = \Omega$;
- we have $\underline{I}' = \underline{I}$ and $\underline{O}' = \underline{O}$;
- for each $\omega \in \Omega$, $\text{in}(\omega)$ and Σ_{ω}^{in} have the same number of elements; write $\Sigma_{\omega}^{in}(e)$ for the system label in the totally ordered set Σ_{ω}^{in} at the same position as edge e in the totally ordered set $\text{in}(\omega)$;
- for each $\omega \in \Omega$, $\text{out}(\omega)$ and Σ_{ω}^{out} have the same number of elements; write $\Sigma_{\omega}^{out}(e)$ for the system label in the totally ordered set Σ_{ω}^{out} at the same position as edge e in the totally ordered set $\text{out}(\omega)$;
- the input nodes in (Γ) and Π^{in} have the same number of elements; write $\Pi^{in}(e)$ for the system label in the totally ordered set Π^{in} at the same position as $\text{tail}(e)$ in the totally ordered set $\text{in}(\Gamma)$;

- (vi) the output nodes $\text{out}(\Gamma)$ and Π^{out} have the same number of elements; write $\Pi^{\text{out}}(e)$ for the system label in the totally ordered set Π^{out} at the same position as $\text{head}(e)$ in the totally ordered set $\text{out}(\Gamma)$;
- (vii) for each edge $e \in \text{edges}(\Gamma)$, the system label at its tail and at its head coincide, i.e. we have $\Sigma_{\text{tail}(e)}^{\text{out}}(e) = \Sigma_{\text{head}(e)}^{\text{in}}(e)$; by convention, we set $\Sigma_{\text{tail}(e)}^{\text{out}}(e) := \Pi^{\text{in}}(e)$ when $\text{tail}(e)$ is an input node and $\Sigma_{\text{head}(e)}^{\text{in}}(e) := \Pi^{\text{out}}(e)$ when $\text{head}(e)$ is an output node (both can be true at the same time).

A definite causal scenario Θ which is compatible with the indefinite causal scenario Φ comes equipped with an edge labelling $\text{syslabel}_{\Phi, \Theta} : \text{edges}(\Gamma) \rightarrow \Xi$, sending each edge $e \in \text{edges}(\Gamma)$ to the system label $\text{syslabel}_{\Phi, \Theta}(e) \in \Xi$ which the indefinite causal scenario associates to the endpoints of the edge. We write $\text{DCaus}(\Phi)$ for the set of definite causal scenarios compatible with Φ .

Finally, we are in a position to define what it means to draw a diagram over an indefinite causal scenario Φ . This is a generalisation of the abstract notion of drawing a diagram over a definite causal scenario from Definition 11: processes are associated to the events, but now taking additional care that the induced diagrams over all definite scenarios compatible with Φ are well-defined. See Figure 3 for an exemplification.

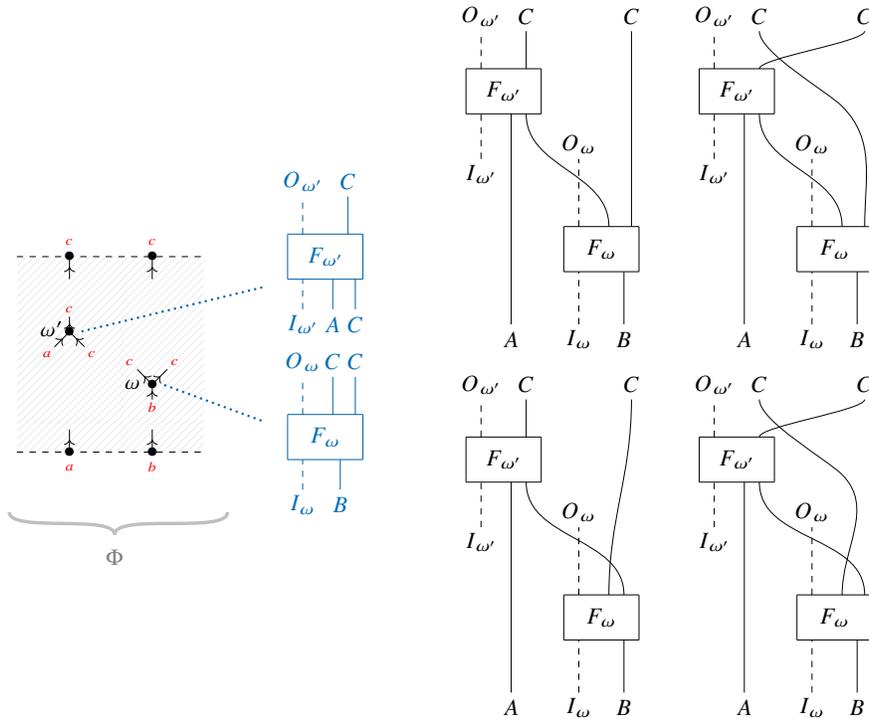


Figure 3: Graphical exemplification of how a diagram over an indefinite causal scenario Φ gives rise to diagrams over the compatible definite causal scenarios (cf. Figure 2).

Definition 16. Let $\Phi = (\Omega, \Xi, \underline{I}, \underline{O}, \underline{\Sigma}^{\text{in}}, \underline{\Sigma}^{\text{out}}, \Pi^{\text{in}}, \Pi^{\text{out}})$ be an indefinite causal scenario and let C be a probabilistic theory. A **diagram over Φ in C** is a pair of functions $\text{sys} : \Xi \rightarrow \text{obj}(C)$ and $\text{proc} : \Omega \rightarrow \text{mor}(C)$, associating each system label ξ to a system $\text{sys}(\xi)$ in C and each event $\omega \in \Omega$ to a process $\text{proc}(\omega)$ in C

with the following type:

$$\text{proc}(\omega) : I_\omega \otimes \bigotimes_{\xi \in \Sigma_\omega^{\text{in}}} \text{sys}(\xi) \longrightarrow O_\omega \otimes \bigotimes_{\xi' \in \Sigma_\omega^{\text{out}}} \text{sys}(\xi') \quad (8)$$

If Δ is such a diagram and $\Theta = (\Gamma, I, O)$ is a definite causal scenario compatible with Φ , then the **induced diagram** $\Delta|_\Theta = (\text{sys}|_\Theta, \text{proc}|_\Theta)$ over Θ is defined by setting $\text{sys}|_\Theta : \text{edges}(\Gamma) \rightarrow \text{obj}(C)$ to be $\text{sys}|_\Theta(e) := \text{sys}(\text{syslabel}_{\Phi, \Theta}(e))$ and setting $\text{proc}|_\Theta : \text{ev}(\Theta) \rightarrow \text{mor}(C)$ to be $\text{proc}|_\Theta(\omega) := \text{proc}(\omega)$.

The semantics for a diagram over an indefinite causal scenario Φ will no longer be given by a single process with a definite causal order—as was the case for the semantics of diagrams over definite scenarios—but rather a controlled process for the family of all diagrams over all definite scenarios compatible with Φ .

Definition 17. Let Φ be an indefinite causal scenario, C be a probabilistic theory and $\Delta = (\text{sys}, \text{proc})$ be a diagram over Φ in C . A **controlled process associated to Δ** is a controlled process (G, p, m) in C associated to the family $(\llbracket \Delta|_\Theta \rrbracket)_{\Theta \in \text{DCaus}(\Phi)}$ of induced diagrams over all definite causal scenarios Θ compatible with Φ .

It is worth noting that the semantics for diagrams over indefinite causal scenarios result in a *controlled* process (G, p, m) : the control system is left open on the side, allowing preparations and observations to be used to control the causal order in all possible ways. Regardless of the specific controlled process and regardless of the specific probabilistic theory, the following will always work.

If we pre-compose the controlled process G with p on the control system we are able to classically control the causal order. Specifically, feeding a specific value $\Theta \in \text{DCaus}(\Phi)$ as input to p results in the diagram $\Delta|_\Theta$ for the definite causal scenario Θ . More generally, feeding a probability distribution $\sum_{\Theta \in \text{DCaus}(\Phi)} \mathbb{P}(\Theta) \delta_\Theta$ as input to p results in a convex mixture of the diagrams associated to the definite scenarios, each diagram $\Delta|_\Theta$ happening with probability $\mathbb{P}(\Theta)$.

If we pre-compose the controlled process G with a normalised state ρ on the control system and we post-compose it with the observation m on the control system, we obtain again a convex mixture of the diagrams associated with the definite scenarios, with probability distribution given by the classical state $m \circ \rho$. If instead of the observation m we use the discarding map on the control system, we obtain the same convex mixture, but now without being able of extracting information about the causal order from the classical outcome of the observation m .³

The last observation, that discarding the control system always yields a convex mixture of causal orders, will play an important role in the next Section, when we construct superpositions of causal orders in quantum theory. Indeed, the observation implies that we cannot obtain superposition of causal orders by preparing the control system in a superposition and then discarding it after we are done. This is because the act of discarding is an *epistemic* one: it simply means that information about the system is not locally available, not that it can never be recovered. In order to obtain a *true* superposition, we will have to permanently *erase* the information about definite causal orders, by measuring in an unbiased basis (e.g. the Fourier one).

5 Superposition of Causal Orders

In this Section, we show how quantum superposition of causal orders can be modelled within our framework. We start by looking at the *quantum switch* and then proceed to generalise our construction to

³This is because applying the discarding map on the control system is the same as first applying the observation m and then discarding its classical outcome, resulting in a mixture.

arbitrary indefinite causal scenarios. We conclude by showing that superpositions of causal orders can be constructed purely as a function of the quantum instruments operated at the events, with no dependence on a choice of purification for the CP maps. This is somewhat surprising—in light of Proposition 6—and it is a consequence of the fact that each discarded environment refers to the same local CP map in all branches of the superposition (something which is not true in general coherent control of CP maps, e.g. in the circumstances considered by [19]).

The *quantum switch* is an indefinite causal scenario widely studied in recent literature [4, 20, 24]. In its simplest form it involves two parties—call them Alice and Bob—each operating some quantum instrument locally to their laboratory. Some physical system comes into the scenario from outside, enters the first instrument (Alice’s or Bob’s), comes out (possibly altered), enters the second instrument (Bob’s or Alice’s, respectively), comes out (possibly altered again) and finally leaves the scenario. The scenario involves the superposition of two definite causal orders: Alice-before-Bob and Bob-before-Alice. In the n -partite generalisation of the quantum switch, n parties operate their quantum instruments sequentially on the same physical system, resulting in a superposition of $n!$ causal orders (corresponding to all possible permutations of the parties). It is now straightforward to model such a scenario within our framework.

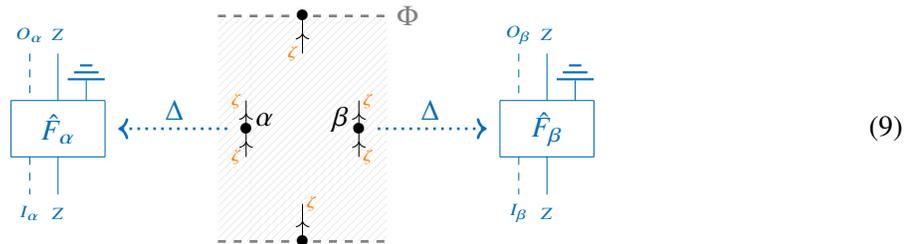
Definition 18. An n -partite switch is an indefinite causal scenario satisfying the following conditions:

- there are n events, corresponding to the n parties;
- the set Ξ has a single element, as the same physical system is operated upon by all parties;
- the sequences Σ_ω^{in} , Σ_ω^{out} , Π^{in} and Π^{out} each have a single element, forcing the parties to operate on the same physical system one after the other.

The classical inputs I_ω and classical outputs O_ω are free to choose.

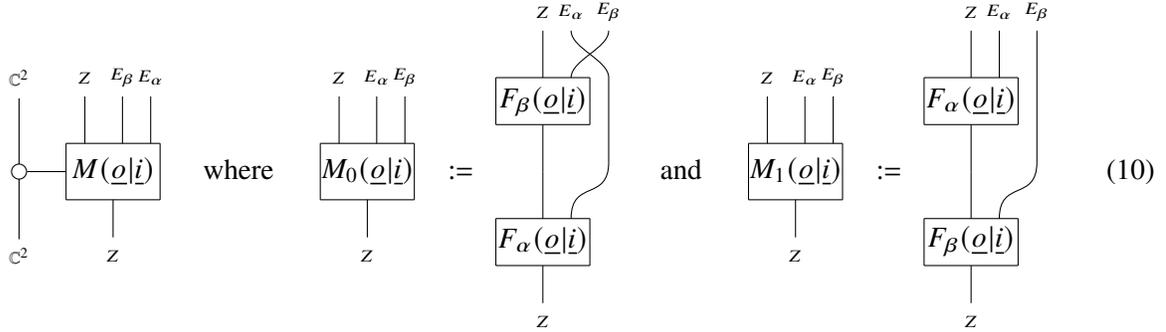
If Φ is an n -partite switch, the semantics of a diagram Δ over Φ in a probabilistic theory is given by a controlled process where each choice of permutation $\sigma \in S_n$ for the parties $\{1, \dots, n\}$ results in party $\sigma(1)$ acting first, followed by party $\sigma(2)$, followed by all other parties in order until party $\sigma(n)$.

For the sake of simplicity, we will now restrict our attention to the bipartite ($n = 2$) case. The two parties are our beloved Alice and Bob, the corresponding events are called α and β , the physical system will be some quantum system Z and the quantum instruments operated by Alice and Bob will be $F_\alpha : I_\alpha \otimes Z \rightarrow O_\alpha \otimes Z$ and $F_\beta : I_\beta \otimes Z \rightarrow O_\beta \otimes Z$ respectively. Each quantum instrument $F_\omega : I_\omega \otimes Z \rightarrow O_\omega \otimes Z$ (for $\omega \in \{\alpha, \beta\}$) is defined by a family of CP maps $F_\omega(o_\omega|i_\omega) : Z \rightarrow Z$ indexed by each possible classical input value $i_\omega \in I_\omega$ and classical output value $o_\omega \in O_\omega$, subject to the normalisation requirement that $\sum_{o_\omega \in O_\omega} F_\omega(o_\omega|i_\omega)$ be a CPTP map for each choice of classical input $i_\omega \in I_\omega$. Let $\hat{F}_\omega(o_\omega|i_\omega) : Z \rightarrow Z \otimes E_\omega$ be a family of purifications for the CP maps, chosen (without loss of generality) to all have the same environment E_ω . This results in the following scenario Φ and diagram Δ :



We can construct a controlled process associated to Δ by using our definition of coherently controlled processes from Section 3. Specifically, for each fixed $\underline{i} = (i_\alpha, i_\beta) \in \underline{I}$ and $\underline{o} = (o_\alpha, o_\beta) \in \underline{O}$ we can define

the following pure controlled processes from the purifications:



If we discard the two environment E_α and E_β and we reintroduce the classical inputs and outputs, we obtain the controlled process for Δ :



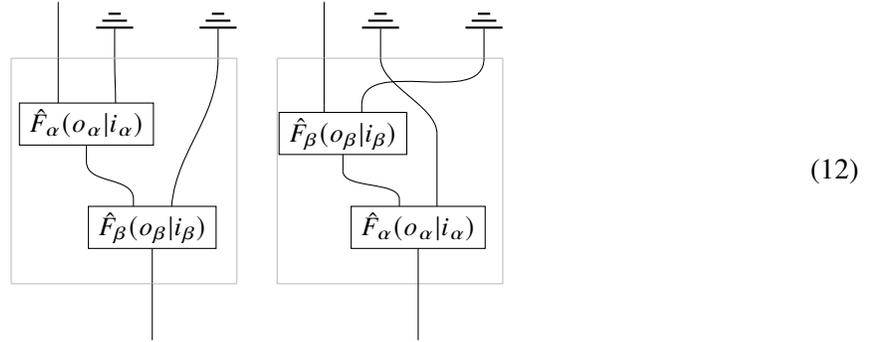
As mentioned in the previous Section, the controlled process above is very general: amongst other things, we can plug any qubit state into the control systems and perform any measurement on it afterwards. In order to obtain a true superposition of the two causal orders, we use some phase state $|\varphi\rangle := \frac{1}{\sqrt{2}}(|0\rangle + e^{i\varphi}|1\rangle)$ and measure in the Pauli X basis. The two measurement outcomes $\langle \pm |$ then correspond to families of processes—indexed by the classical inputs $\underline{i} \in \underline{I}$ and classical outputs $\underline{o} \in \underline{O}$ —which see a superposition with phase of the two causal orders. In traditional notation, the processes can be written as follows:

$$\frac{1}{4} \text{Tr}_{E_\alpha \otimes E_\beta} \left[\text{dbl} \left[\left(\hat{F}_\alpha(o_\alpha | i_\alpha) \otimes \text{id}_{E_\beta} \right) \circ \hat{F}_\beta(o_\beta | i_\beta) \right] \pm e^{i\varphi} \text{dbl} \left[\left(\text{id}_Z \otimes \sigma_{E_\beta, E_\alpha} \right) \circ \left(\hat{F}_\beta(o_\beta | i_\beta) \otimes \text{id}_{E_\alpha} \right) \circ \hat{F}_\alpha(o_\alpha | i_\alpha) \right] \right]$$

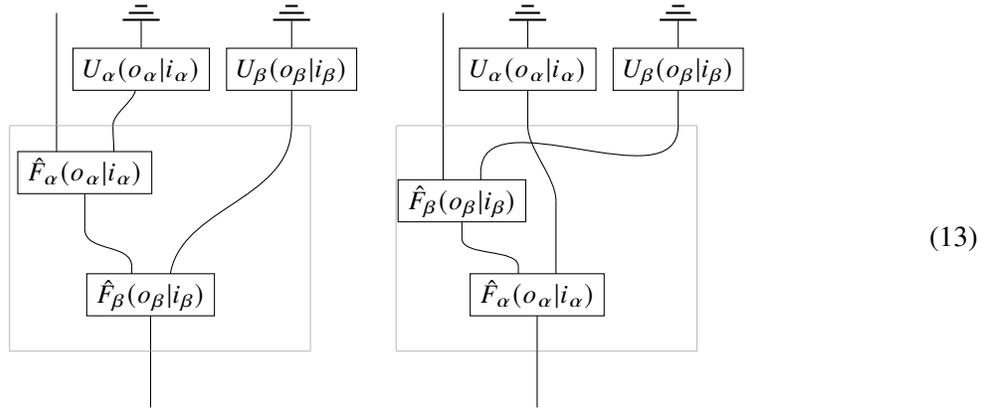
where we introduced the short-hand $\text{dbl}[U]$ for the CP map $\text{dbl}[U] := UU^\dagger$ corresponding to a linear map U and we have freely confused the pure CP maps $\hat{F}_\omega(o_\omega | i_\omega)$ with the corresponding linear maps.

Proposition 6 tells us that, in general, a controlled process such as (11) will depend on our choices of purification \hat{F}_ω . However, this turns out not to be the case for the switch. Indeed, we can pull the two environments to the boundary of the scenario, keeping them throughout the superposition of causal orders

and only discarding afterwards:



Any alternative choice of purification can be obtained by applying a suitable unitary to each environment. However, the same map \hat{F}_ω appears at the bottom of the environment E_ω in all branches of the superposition, and hence the same unitary appears applied to the environment. This means that we can pull the unitaries themselves out of the scenario:



The unitaries will then be cancelled by the discarding maps, leading us to conclude that the controlled process we constructed for the switch was actually independent of our choices of purification \hat{F}_ω and is therefore a function of the original quantum instruments F_ω . This argument is not unique to the switch, but instead generalises to coherent control for all diagrams over indefinite causal scenarios in quantum theory, as dictated by our final result below. As a consequence, our framework can be used to give well-defined semantics to superposition of causal orders in quantum theory.

Definition 19. Let $\Phi = (\Omega, \Xi, I, O, \Sigma^{in}, \Sigma^{out}, \Pi^{in}, \Pi^{out})$ be an indefinite causal scenario and Δ be a diagram over Φ in quantum theory. The **purification** of Φ is the indefinite causal scenario obtained from Φ by adding fresh symbols ε_ω to Ξ and Π^{out} for all $\omega \in \Omega$ (in some chosen order). A **purification of Δ with environments** $(E_\omega)_{\omega \in \Omega}$ is the diagram over the purification of Φ obtained by considering purifications of all CP maps in Δ , with fixed environment E_ω for each event $\omega \in \Omega$.⁴

The purification of a diagram—and the necessary corresponding “purification” of the underlying indefinite causal scenario—generalise the process seen in (12) and (13) above, where the environment wires were “pulled to the boundary” of the diagram. We use purifications of a diagram Δ to define its coherent control.

⁴The purifications considered here are for each fixed choice of classical input and output values, i.e. they are indexed families of purifications $\hat{F}_\omega(o_\omega|i_\omega)$ for the CP maps $F_\omega(o_\omega|i_\omega)$ corresponding to each given choice of $i_\omega \in I_\omega$ and $o_\omega \in O_\omega$.

Definition 20. Let Φ be an indefinite causal scenario and Δ be a diagram over Φ in quantum theory. Let $P(\varphi) := \sum_{\Theta \in \text{DCaus}(\Phi)} e^{i\varphi\Theta} |\Theta\rangle \langle \Theta|$ be a phase gate for the computational basis of $\mathbb{C}^{\text{DCaus}(\Phi)}$. The **coherent control of Δ with phase $P(\varphi)$** is defined to be the coherent control—with phase $P(\varphi)$ and control system $\mathbb{C}^{\text{DCaus}(\Phi)}$ —of the following family of processes, where Δ^{pure} is a purification of Δ with environments $(E_\omega)_{\omega \in \Omega}$ and we have defined the global environment $E := \bigotimes_{\omega \in \Omega} E_\omega$:

$$\left(\text{Tr}_E \llbracket \Delta^{\text{pure}} |_{\Theta} \rrbracket \right)_{\Theta \in \text{DCaus}(\Phi)} \quad (14)$$

Note that $\text{Tr}_E \llbracket \Delta^{\text{pure}} |_{\Theta} \rrbracket$ is a diagram over Θ for all $\Theta \in \text{DCaus}(\Phi)$.

Proposition 21. Let Φ be an indefinite causal scenario and Δ be a diagram over Φ in quantum theory. For each possible choice of phase gate $P(\varphi)$ for the computational basis of $\mathbb{C}^{\text{DCaus}(\Phi)}$, the coherent control of Δ with phase $P(\varphi)$ is well-defined independently of the choice of purification for the processes in Δ .

Having shown that the coherent control with phase for diagrams over indefinite causal scenarios is a well-defined concept, we conclude by providing the following definition for sake of clarity and future use.

Definition 22. Let Φ be an indefinite causal scenario and Δ be a diagram over Φ in quantum theory. By a **superposition of causal orders for Δ with phase $P(\varphi)$** we mean the quantum instrument obtained by:

- (i) considering the coherent control of Δ with phase $P(\varphi)$;
- (ii) pre-composing the control system with the uniform superposition state $|+\rangle$ below:

$$|+\rangle := \frac{1}{\sqrt{|\text{DCaus}(\Phi)|}} \sum_{\Theta \in \text{DCaus}(\Phi)} |\Theta\rangle \quad (15)$$

- (iii) post-composing the control system with some measurement which is unbiased with respect to the computational basis (e.g. one in the Fourier basis for some group structure on $\text{DCaus}(\Phi)$).

Note that the choice of fixed input state $|+\rangle$ was made to avoid redundancy in the relative phase between the different causal orders, which is already controlled by $P(\varphi)$.

Acknowledgements

NP would like to acknowledge financial support from EPSRC and the Pirie-Reid Scholarship. SG would like to acknowledge support from Hashberg Ltd.

References

- [1] Alastair A. Abbott, Julian Wechs, Dominic Horsman, Mehdi Mhalla & Cyril Branciard (2018): *Communication through coherent control of quantum channels*. Available at <https://arxiv.org/abs/1810.09826>.
- [2] Mateus Araújo, Adrien Feix, Fabio Costa & Časlav Brukner (2014): *Quantum circuits cannot control unknown operations*. *New Journal of Physics* 16(9), p. 093026, doi:10.1088/1367-2630/16/9/093026.
- [3] Giulio Chiribella, Giacomo Mauro D’Ariano & Paolo Perinotti (2010): *Probabilistic theories with purification*. *Physical Review A* 81, p. 062348, doi:10.1103/PhysRevA.81.062348.
- [4] Giulio Chiribella, Giacomo Mauro D’Ariano, Paolo Perinotti & Benoit Valiron (2013): *Quantum computations without definite causal structure*. *Physical Review A* 88(2), doi:10.1103/physreva.88.022318.
- [5] Giulio Chiribella & Hlér Kristjánsson (2019): *Quantum Shannon theory with superpositions of trajectories*. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 475(2225), p. 20180903, doi:10.1098/rspa.2018.0903.
- [6] Marios Christodoulou & Carlo Rovelli (2019): *On the possibility of laboratory evidence for quantum superposition of geometries*. *Physics Letters B* 792, pp. 64 – 68, doi:10.1016/j.physletb.2019.03.015.
- [7] Bob Coecke & Ross Duncan (2011): *Interacting quantum observables: categorical algebra and diagrammatics*. *New Journal of Physics* 13, p. 043016, doi:10.1088/1367-2630/13/4/043016.
- [8] Bob Coecke, Chris Heunen & Aleks Kissinger (2016): *Categories of quantum and classical channels*. *Quantum Information Processing* 15(12), pp. 5179–5209, doi:10.1007/s11128-014-0837-4.
- [9] Bob Coecke & Raymond Lal (2012): *Causal Categories: Relativistically Interacting Processes*. *Foundations of Physics* 43(4), p. 458–501, doi:10.1007/s10701-012-9646-8.
- [10] Giacomo Mauro D’Ariano, Giulio Chiribella & Paolo Perinotti (2017): *Quantum Theory from First Principles: An Informational Approach*. Cambridge University Press, doi:10.1017/9781107338340.
- [11] Andreas Döring & Chris J. Isham (2010): “What is a Thing?": *Topos Theory in the Foundations of Physics*. *Lecture Notes in Physics*, p. 753–937, doi:10.1007/978-3-642-12821-913.
- [12] Giacomo Mauro D’Ariano (2018): *Causality re-established*. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376(2123), p. 20170313, doi:10.1098/rsta.2017.0313.
- [13] Stefano Gogioso (2017): *Fantastic Quantum Theories and Where to Find Them*. Available at <https://arxiv.org/abs/1703.10576>.
- [14] Stefano Gogioso & Carlo Maria Scandolo (2018): *Categorical Probabilistic Theories*. *Electronic Proceedings in Theoretical Computer Science* 266, p. 367–385, doi:10.4204/eptcs.266.23.
- [15] Chris J. Isham (2011): *Topos Methods in the Foundations of Physics*, p. 187–206. Cambridge University Press, doi:10.1017/CBO9780511976971.005.
- [16] David B. Malament (1977): *The class of continuous timelike curves determines the topology of spacetime*. *Journal of Mathematical Physics* 18(7), pp. 1399–1404, doi:10.1063/1.523436.
- [17] Chiara Marletto & Vladko Vedral (2017): *Gravitationally Induced Entanglement between Two Massive Particles is Sufficient Evidence of Quantum Effects in Gravity*. *Physical Review Letters* 119(24):240402, doi:10.1103/PhysRevLett.119.240402.
- [18] Chiara Marletto & Vlatko Vedral (2019): *Answers to a few questions regarding the BMV experiment*. Available at <https://arxiv.org/abs/1907.08994>.
- [19] Daniel K. L. Oi (2003): *Interference of Quantum Channels*. *Physical Review Letters* 91(6), doi:10.1103/physrevlett.91.067902.
- [20] Ognjan Oreshkov (2019): *Time-delocalized quantum subsystems and operations: on the existence of processes with indefinite causal structure in quantum mechanics*. *Quantum* 3, p. 206, doi:10.22331/q-2019-12-02-206.
- [21] Ognjan Oreshkov, Fabio Costa & Časlav Brukner (2012): *Quantum correlations with no causal order*. *Nature Communications* 3(1), p. 1092, doi:10.1038/ncomms2076.

- [22] Nikola Paunkovic & Marko Vojinovic (2019): *Causal orders, quantum circuits and spacetime: distinguishing between definite and superposed causal orders*. Available at <https://arxiv.org/abs/1905.09682>.
- [23] Nicola Pinzani, Stefano Gogioso & Bob Coecke (2019): *Categorical Semantics for Time Travel*. In: 2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), IEEE Computer Society, Los Alamitos, CA, USA, pp. 1–20, doi:10.1109/LICS.2019.8785664.
- [24] Sina Salek, Daniel Ebler & Giulio Chiribella (2018): *Quantum communication in a superposition of causal orders*. Available at <https://arxiv.org/abs/1809.06655>.
- [25] Jayne Thompson, Kavan Modi, Vlatko Vedral & Mile Gu (2018): *Quantum plug n' play: modular computation in the quantum regime*. *New Journal of Physics* 20(1), p. 013004, doi:10.1088/1367-2630/aa99b3.
- [26] Magdalena Zych, Fabio Costa, Igor Pikovski & Časlav Brukner (2019): *Bell's theorem for temporal order*. *Nature Communications* 10(1), p. 3772, doi:10.1038/s41467-019-11579-x.

A Categorical Probabilistic Theories

Quantum theory can be formulated using different mathematical formalisms, each highlighting different conceptual aspects of the theory. In the purely categorical formulation of process theories [8], classical systems emerge as a compositional part of the theory itself, but the probabilistic structure is not explicitly axiomatised. In the formulation of *operational probabilistic theories* (OPT) by [3, 10], on the other hand, the probabilistic structure is explicitly axiomatised, but classical systems are *not* treated compositionally as part of the theory, with the quantum-classical interface relegated to a mere labelling of quantum processes. The formulation of probabilistic theories by [14] tries to take the best of both worlds: the compositional treatment of classical system from categorical formulations and the explicit axiomatisation of probabilistic structure from OPTs. This leads to the following definition of a probabilistic theory.

Definition 23. A *probabilistic theory* is a (strict) symmetric monoidal category (SMC) C which satisfies the following requirements:

- there is a full sub-SMC of C , denoted by C_K , which is equivalent to the SMC \mathbb{R}^+ -Mat modelling **classical theory** (itself a probabilistic theory);
- the SMC C is enriched in commutative monoids and the induced enrichment on C_K coincides with the one given by the linear structure of \mathbb{R}^+ -Mat;
- the SMC C comes with an environment structure, i.e with a family of effects $\top_A : A \rightarrow 1$ which satisfy the following requirements:

$$\begin{array}{c} \overline{\top} \\ | \\ \mathcal{H} \otimes \mathcal{K} \end{array} = \begin{array}{c} \overline{\top} \\ | \\ \mathcal{H} \end{array} \begin{array}{c} \overline{\top} \\ | \\ \mathcal{K} \end{array} \quad \begin{array}{c} \overline{\top} \\ | \\ 1 \end{array} = \boxed{} \quad (16)$$

The environment structure induced on C_K coincides with the one given by marginalisation in \mathbb{R}^+ -Mat.

The requirements above imply that processes $A \rightarrow B$ in C have the structure of a convex cone, i.e. they are \mathbb{R}^+ -modules. The effects $\top_A : A \rightarrow 1$ are known as **discarding maps**. The systems in the full sub-category C_K are known as **classical systems** and the processes between them as **classical processes**.

The probabilistic theory most relevant to this work is *quantum theory*, defined by taking CP maps together with classical theory and introducing the quantum-classical interface by linearity from families of quantum processes (cf. resolution of the classical identity below).

From a diagrammatic perspective, dashed wires are used to denote systems which are guaranteed to be classical, while solid wires denote generic systems. It is convenient to assume the following general form for processes in probabilistic theories, with distinguished classical input and output systems (finite sets X and Y respectively):



The linear structure of \mathbb{R}^+ -Mat can be used to explicitly perform resolutions of the classical identity:



The resolution of the classical identity can be used to freely confuse between the following two perspectives, establishing a direct link to the OPT formalism:

- processes $F : A \otimes X \rightarrow B \otimes Y$, with X and Y classical systems (i.e finite sets);
- families $F(y|x) : A \rightarrow B$ of processes indexed by the classical values $x \in X$ and $y \in Y$.

The discarding maps axiomatise the notion of marginalisation (aka *partial trace* in the context of quantum systems). They can also be used to define a sub-SMC of *normalised* processes:

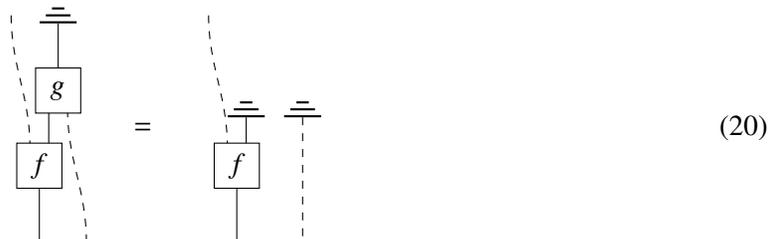
Definition 24. A process $f : A \rightarrow B$ is said to be **normalised** if it satisfies the following equation:



A process $f : A \rightarrow B$ is said to be **sub-normalised** if there exists some $g : A \rightarrow B$ such that $f + g$ is normalised (in which case g is also sub-normalised).

In particular, the normalised states on a classical system X are the probability distributions on X , the normalised processes $X \rightarrow Y$ are the Y -by- X stochastic matrices, and discarding on a classical output of a classical process is the same as marginalisation.

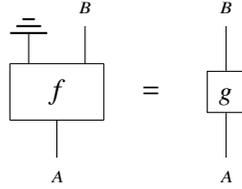
The axiomatisation of causality as “no-signalling from the future” [10] is embodied in probabilistic theories by the following observation about normalised processes:



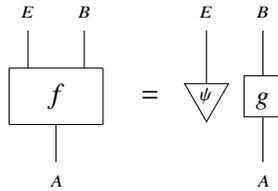
In the above, we see that the classical outcome of a test f cannot be influenced by a controlled process g in its future (i.e. it is independent of the classical input used to control g).

The framework also adopts a notion of *purity*, defined as the lack of non-trivial interaction with a discarded environment.

Definition 25. A process $g : A \rightarrow B$ is said to be **pure** if whenever we can find a system E and a process $f : A \rightarrow E \otimes B$ such that the following equality holds:



then there exists a normalised state $\psi : 1 \rightarrow E$, dependent on f , such that the following equality also holds:



Note that neither g nor f are required to be normalised as part of this definition.

Finally, we recall the definition of a sharp preparation-observation pair, capturing the idea of perfect encoding/decoding of classical information in arbitrary systems.

Definition 26. A **sharp preparation-observation pair** (SPO pair) is a pair (p, m) of a **preparation** process $p : X \rightarrow H$ and an **observation** process $m : H \rightarrow X$ on some classical system X and some arbitrary system H , such that the following equation holds:



B Proofs

Proof of Proposition 4

Proposition 4. Let $(F_x)_{x \in X}$ be a family of pure CP maps in quantum theory, not necessarily normalised (i.e. not necessarily trace-preserving). Assume that (G, p, m) is a controlled process for the family with control system \mathbb{C}^X , where (p, m) is the SPO for the canonical basis $(|x\rangle)_{x \in X}$ and where G is itself a pure CP map. Then G takes form (4) for some phase α .

Proof. In quantum theory, the normalised SPO pairs connecting a classical system X and the quantum system \mathbb{C}^X are exactly the preparations and measurements in some orthonormal basis of \mathbb{C}^X . Without loss of generality, assume that the SPO pair is for the canonical basis $(|x\rangle)_{x \in X}$ of system \mathbb{C}^X :

$$p = \begin{array}{c} | \\ \circ \\ | \end{array} \quad m = \begin{array}{c} | \\ \circ \\ | \end{array} \quad (22)$$

Now assume that (G, p, m) is a controlled process for the family $(F_x)_{x \in X}$. In particular, it satisfies the following equalities for all $x \in X$, where on the right we have used the coherent control of Example 3:

$$\begin{array}{c} B \\ | \\ \triangle x \\ | \\ \boxed{G} \\ | \quad | \\ \mathbb{C}^X \quad A \end{array} = \begin{array}{c} B \\ | \\ \triangle x \\ | \\ \boxed{F_x} \\ | \\ \mathbb{C}^X \quad A \end{array} = \begin{array}{c} B \\ | \\ \triangle x \\ | \\ \circ \\ | \quad | \\ \mathbb{C}^X \quad A \end{array} \quad \boxed{F} \quad (23)$$

Because G is pure, the equality of CP maps above is an equality of the corresponding linear maps up to a global phase α_x dependent on $x \in X$. Using resolution of the identity $\text{id}_{\mathbb{C}^X} = \sum_{x \in X} |x\rangle \langle x|$ and putting all the phases together in a phase gate for the canonical basis we obtain the following equality:

$$\begin{array}{c} \mathbb{C}^X \quad B \\ | \quad | \\ \boxed{G} \\ | \quad | \\ \mathbb{C}^X \quad A \end{array} = \begin{array}{c} \mathbb{C}^X \quad B \\ | \quad | \\ \circ \alpha \\ | \quad | \\ \circ \quad \boxed{F} \\ | \quad | \\ \mathbb{C}^X \quad A \end{array} = \begin{array}{c} \mathbb{C}^X \quad B \\ | \quad | \\ \circ \alpha \\ | \quad | \\ \boxed{F} \\ | \quad | \\ \mathbb{C}^X \quad A \end{array} \quad (24)$$

This concludes our proof. □

Proof of Proposition 6

Proposition 6. *Let $(F_x)_{x \in X}$ be a family of CP maps in quantum theory. It is not generally possible to construct a coherent control of the family which is a function of the family $(F_x)_{x \in X}$ alone, i.e. one which is independent of a choice of purification for the CP maps in the family. This is the same as the statement that it is not generally possible to construct a coherent control $(\text{Tr}_E(G'), p, m)$ of the family $(F_x)_{x \in X}$ in such a way that the following equation holds for all choices of unitaries $U_x : E \rightarrow E$:*

$$\begin{array}{c} H \quad B \\ | \quad | \\ \boxed{G'} \\ | \quad | \\ H \quad A \end{array} = \begin{array}{c} H \quad B \\ | \quad | \\ \circ \\ | \quad | \\ \boxed{G'} \\ | \quad | \\ H \quad A \end{array} \begin{array}{c} \overline{\overline{E}} \\ | \\ \boxed{U} \\ | \\ E \end{array} \quad (5)$$

Note that the unitaries $(U_x)_{x \in X}$ correspond to all possible choices of purification for the CP maps $(F_x)_{x \in X}$.

Proof. Consider a coherent control $(\text{Tr}_E(G'), p, m)$ of a family $(F_x)_{x \in X}$ of CP maps, where G' is pure. In particular, the following equality is satisfied:

$$\text{Diagram (25): } G' = \sum_{x \in X} F_x$$

If \hat{F}_x is a chosen purification of F_x for each $x \in X$ —without loss of generality all with environment E —then we can create the coherent control of the family $(\hat{F}_x)_{x \in X}$ and obtain the following equality:

$$\text{Diagram (26): } G' = \hat{F}_x = \text{coherent control of } \hat{F}_x$$

By essential uniqueness of purification, the following equality of pure CP maps must hold for some choice of unitary $V_x : E \rightarrow E$, dependent on each specific value of $x \in X$:

$$\text{Diagram (27): } G' = \text{coherent control of } \hat{F}_x \text{ with } V_x$$

The equality above is equivalently an equality of linear maps up to a global phase φ_x , also dependent on each specific value of $x \in X$. We can therefore put all $x \in X$ together and obtain the following equality of pure CP maps:

$$\text{Diagram (28): } G' = \text{coherent control of } \hat{F}_x \text{ with } \varphi \text{ and } V$$

Any alternative choice of purification for each CP maps F_x can be obtained by applying some unitary $W_x : E \rightarrow E$ to the environment of our current choice of purification \hat{F}_x . If the coherent control is to be

invariant under this choice of purification, the following equality must hold for all possible choices of unitaries $(W_x)_{x \in X}$:

$$(29)$$

For each $x \in X$, we can define a unitary $U_x := V_x W_x V_x^\dagger$ such that $U_x V_x = V_x W_x$, so that the equality above for all possible choices of unitaries $(W_x)_{x \in X}$ can be equivalently recast as the equality below for all possible choices of unitaries $(U_x)_{x \in X}$:

$$(30)$$

The above is equivalent to Equation (5) in the statement of this Proposition holding for all $(U_x)_{x \in X}$:

$$(5)$$

For a general such G' —i.e. for a general choice of $(\hat{F}_x)_{x \in X}$ —this equation cannot always be made to hold for all $(U_x)_{x \in X}$, leading to the statement of the proposition. □

Proof of Proposition 21

Proposition 21. *Let Φ be an indefinite causal scenario and Δ be a diagram over Φ in quantum theory. For each possible choice of phase gate $P(\varphi)$ for the computational basis of $\mathbb{C}^{\text{DCaus}(\Phi)}$, the coherent control of Δ with phase $P(\varphi)$ is well-defined independently of the choice of purification for the processes in Δ .*

Proof. The trick to circumventing Proposition 6 is the following: the choice of purification is done locally at the level of the processes associated to each event, not globally at the level of the processes associated to the diagrams. As a consequence, the unitary maps implementing the change of purification always

factor over the individual events. Furthermore, the same CP map appears attached to the same event in all diagrams, marking a difference with the situation discussed in [19].

More specifically, assume that the purification Δ^{pure} used in constructing the coherent control of Δ uses purifications $\hat{F}_\omega(o_\omega|i_\omega)$ for the CP maps $F_\omega(o_\omega|i_\omega)$ appearing in Δ . Another choice of purification corresponds to chosen unitaries $U_\omega(o_\omega|i_\omega)$, yielding as the following purifications instead:

$$(\text{id} \otimes U_\omega(o_\omega|i_\omega)) \circ \hat{F}_\omega(o_\omega|i_\omega) \quad (31)$$

We now fix global classical inputs and output $\underline{i} \in \underline{I}$ and $\underline{o} \in \underline{O}$ and pull all the unitaries to the boundary of the diagrams $\llbracket \Delta|_\Theta \rrbracket$ for the definite scenarios $\Theta \in \text{DCaus}(\Phi)$. The same unitary $U(\underline{o}|\underline{i}) : E \rightarrow E$ then appears on the global environment in all diagrams $\llbracket \Delta|_\Theta \rrbracket$:

$$U(\underline{o}|\underline{i}) := \bigotimes_{\omega \in \Omega} U_\omega(o_\omega|i_\omega) \quad (32)$$

The coherent control over $\mathbb{C}^{\text{DCaus}(\Phi)}$ is thus trivial, so that the unitary $U(\underline{o}|\underline{i})$ can be cancelled by the discarding map on the global environment E . This means that the coherent control of Δ is actually independent of the specific choice of purifications \hat{F}_ω , i.e. that it is well-defined as a function of Δ alone. \square