

The Sheaf-Theoretic Structure of Definite Causality

Stefano Gogioso
Hashberg Quantum
quantum(at)hashberg(dot)io

Nicola Pinzani
University of Oxford
nicola.pinzani@cs.ox.ac.uk

We fill a gap in the study of contextuality by extending the sheaf-theoretic framework for non-locality by Abramsky and Brandenburger to deal with operational scenarios in the presence of arbitrary definite causal orders.

1 Introduction

Understanding and classifying correlations that can be realised in quantum theory is of crucial interest to quantum foundations and quantum computation. Even someone as agnostic towards interpretational problems as Paul Dirac was puzzled by the irreducible importance that young quantum theory was assigning to correlations between sets of choices and sets of outcomes:

Also in the quantum theory one starts from certain numbers from which one deduces other numbers. [...] The perturbations which an observer inflicts on a system in order to observe it are directly subject to his control and are acts of his free will. It is exclusively the numbers which describe these acts of free choice which can be taken as initial numbers for a calculation in quantum theory. [48]

It is this connection between the perturbations inflicted by an observer and the permanent registration of the outcomes—the results of observations—which frames our understanding of what John Bell famously saw as violations of local causality.

In this work, we put that very connection front and centre instead: when modelling the interaction of localised processes (e.g. “operations” or “experiments”) in spacetime, the only information relevant to our understanding of causality is the observed distribution of outcomes (from local observations) corresponding to various choices of local perturbations. Our characterisation is therefore independent of both the underlying physical theory and the specific realisation of such processes: information cannot travel against the flow of causality, regardless of how it was obtained.

Several works before ours aimed at describing correlations with definite or indefinite causality, mostly consisting of various generalisations of either Bell scenario (e.g. [38]) or classical Bayesian networks (e.g. [46]). In this paper, we instead build on previous work by Abramsky and Brandenburger [9], which provided a unified framework for non-locality and contextuality in terms of sheaf theory.

The conceptual unification of causality and contextuality touches upon subtle considerations of the nature of contextuality itself. In a Bell scenario, it is not too controversial to believe that, when performing a choice of measurement on one of the entangled subsystems, the observer ends up changing the very conditions defining the possible predictions that can be made about the entire system. In fact, one can trace this belief all the way back to Bohr’s original reply to EPR, relating to Bohr’s own understanding of the notion of ‘complementarity’¹:

¹A discussion on the strong connections between Bohr’s notion of complementarity, contextuality, and non-locality is explained by Arkady Plotnitsky in [66].

Of course there is in a case like that just considered no question of a mechanical disturbance of the system under investigation during the last critical stage of the measuring procedure. But even at this stage there is essentially the question of an influence on the very conditions which define the possible types of predictions regarding the future behaviour of the system. [20]

This line of reasoning resonates with the approaches to contextuality described in [9] and in the Contextuality-by-Default framework [34], which consider the presence of contextual phenomena to be a property related to compatibility conditions of intrinsically different and incomparable probabilistic distribution of outcomes.

Recent works on causality, stemming from the seminal intuition on the “quantum switch” and frameworks of process theories [27, 62], have resulted in attempts to provide an experimental realisation [45] and verification [68] of indefinite causal structures. Facing the need to certify the realisability of indefinite causal structures, we are confronted with a common limitation of all descriptive approaches to date: they are all a generalisation of quantum theory, from which causality is hard to extricate as an independent and observable phenomenon. What do we really mean when we talk about indefinite causality? Are the realisable operational contexts implicit in the aforementioned experimental realisations compatible with a definite causal structure, or are we in the presence of genuine indefinite causality?

The sheaf-theoretic language already provided a stand-alone description of non-locality and contextuality, independent of physical theory or experimental realisation. Today, we extend that language to answering questions about causality, moving one step closer to an understanding of Bohr’s cryptic claim that contextuality and complementarity provide a quantum generalisation of the notion of causality itself.

Literature review

Over the past decade, the sheaf-theoretic framework has evolved from its seminal form [9] in a number of directions. The core framework itself has grown through the addition of extended operational semantics [10], All-vs-Nothing arguments [5], logical characterisation of no-signalling polytopes [6], contextual fraction [7], an extension to continuous variable models [15], and the development of a categorical structure on empirical models [4, 8, 49, 50]. Deep connections have been discovered with cohomology [1, 11, 14, 22, 23], logic [3, 13, 52, 53, 56, 59, 60] and other frameworks, including database theory [2], Spekkens’s contextuality [70], effect algebras [69] and non-commutative geometry [33].

Development of the sheaf-theoretic framework is both recent and limited to a relatively small community: this allows for a relatively exhaustive review. The modern study of causality, in contrast, is much older and broader, so we restrict our attention to a sample of the literature most relevant to our work. Literature on quantum correlations in time [37, 39, 40] is certainly relevant to the aspects of this work generalising aspects of non-locality to arbitrary causal structures. However, the main driver for this work—and its main source of models—is the development of causal frameworks in the context of process theories [28, 29, 32, 43, 54, 55, 64] and operational probabilistic theories [24, 25, 35]. Recent years have also seen the development of different approaches to causality, including quantum generalisations [16, 17] of Reichenbach’s Principle [67] and causal generalisations of the Bell inequalities [21, 63].

We could not find any actual attempts at a full causal generalisation in the existing literature on the sheaf-theoretic framework. However, some published works took steps in that direction by dealing with the case of sequential protocols. In [41, 44], a broad generalisation of the HBB hybrid quantum-classical secret sharing protocol is discussed, connecting device-independent security to contextuality. In [71], sequential communication protocols are discussed, connecting success probability to contextuality. In [61], a (non sheaf-theoretic) notion of sequential-transformation contextuality is introduced, with application to information retrieval tasks discussed in the beautifully illustrated [36].

A proposed extension of the sheaf-theoretic approach can be found in some presentations [57,58]. The proposal focuses on a single concrete example, but already points at an approach very different from our own. The underlying sheaf-theoretic machinery is not discussed explicitly, but private conversation with the author reveals it to be the same as in the original sheaf-theoretic framework: a set X of measurements, its powerset $\mathcal{P}(X)$ as locale of contexts, the usual sheaf of sections and presheaf of distributions, the usual definition of locality/non-contextuality as existence of a global section. Measurement contexts are modified: they are obtained as the maximal elements under a certain partial order on $\mathcal{P}(X)$ induced by a definite causal order imposed on the set of measurements X . Empirical models are generalised to any family of probability distributions on the measurement contexts and their subsets, regardless of no-signalling or causality conditions. Causal empirical models are defined to be the ones respecting causality equations, analogous to those found in previous literature on causality (e.g. see [37, 39]): for these models, the distributions on subsets of the maximal contexts are exactly those obtained by restriction in the presheaf of distributions.

This proposed approach is not applicable to the causal scenarios we consider in this work: causal order is defined between the events, but there is no canonical notion of causal order between alternative input choices at an event. The only permutation invariant way to impose a causal order on the different measurement choices at a given event is to consider them unrelated in the causal order, but this results in all measurements for an event appearing together in the maximal contexts. Indeed, taking the maximal contexts in the non-locality case (on the discrete partial order, as described at the end of the presentations) yields the top element X as the only member of the measurement cover, allowing only local/non-contextual scenarios to be specified.

Furthermore, this proposed approach does not result in causal classical functions always being local/non-contextual. Because the full powerset $\mathcal{P}(X)$ is considered, the presheaf of distributions includes restrictions to sets of measurements which are not downward closed with respect to the causal order: for functions where outcomes in the future depend on measurement choices in the past, some of the restrictions from the top element X would not actually be well-defined (in the sheaf of sections), implying that those functions cannot be local/non-contextual. This is not what we want to achieve: causal classical functions and their mixtures should correspond exactly to non-contextual empirical models.

At this point, it is worth remarking that our formulation, via a locale of inputs over lower-sets and a sheaf of causal sections, was not actually obtained by direct modification of the original sheaf-theoretic framework. In particular, it was not obtained using the ideas from [57, 58]. Rather, it was developed starting from an earlier observation about marginalisation in probabilistic theories and no-signalling models [41, 43]. This observation was then extended to the more general framework for causal diagrams developed by the authors in [64, 65], leading to the sheaf-theoretic formulation of causality presented here. (Indeed, there is a tight correspondence between the two lines of work, cf. Proposition 22.)

2 Definite Causal Scenarios

We give the following operational interpretation to definite causal scenarios. We think of the causal set Ω as defining a discrete subset of events in some spacetime. To each event ω corresponds the local operation of a black-box device, or gadget: an input value for the device is freely chosen from the finite set I_ω and and output value from the finite set O_ω is generated in response by the device.

Definition 1. A **definite causal scenario** is a triple $\Sigma := (\Omega, \underline{I}, \underline{O})$ where:

- Ω is a finite poset, the **causal set**

- $\underline{I} = (I_\omega)_{\omega \in \Omega}$ is a family of non-empty finite sets
- $\underline{O} = (O_\omega)_{\omega \in \Omega}$ is a family of non-empty finite sets

We refer to the points $\omega \in \Omega$ as **events**, to the set I_ω as the set of **inputs** at event ω and to the set O_ω as the set of **outputs** at event ω . We refer to $\underline{I}_\Sigma := \prod_{\omega \in \Omega} I_\omega$ as the set of **joint inputs** and to $\underline{O}_\Sigma := \prod_{\omega \in \Omega} O_\omega$ as the set of **joint outputs**.

A definite causal scenario describes the causal structure of events and the possible inputs and outputs, but not the observed behaviour of the devices: that is encoded in an **empirical model**. An empirical model is an assignment of a probability distribution on joint outputs given any (allowed) choice of joint inputs for all events, in such a way as to respect the causal structure. Just as was originally done in the sheaf-theoretic framework for non-locality, we will not restrict ourselves to probabilities—modelled by the commutative semiring \mathbb{R}^+ —but rather allow any notion of non-determinism modelled by some commutative semiring R . Examples of other semirings of interest include the real field \mathbb{R} (for signed probabilities), the Boolean semiring \mathbb{B} (for possibilities), the p -adic fields (for p -adic probabilities [42, 51]).

3 The original sheaf-theoretic framework for non-locality

The scenarios considered in the original sheaf-theoretic framework for non-locality [9] correspond to choices of causal set Ω which are discrete partially ordered sets, i.e. those where each event in Ω is causally unrelated to every other event.

Definition 2. Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario. We say that Σ is a **non-locality scenario** if Ω is a discrete partial order, i.e. one such that $x \not\leq y$ for all $x, y \in \Omega$ such that $x \neq y$.

The set of individual inputs is defined as the disjoint union $X := \coprod_{\omega \in \Omega} I_\omega = \{(i, \omega) \mid i \in I_\omega, \omega \in \Omega\}$ and the powerset $\mathcal{P}(X)$ is considered. The powerset is naturally a locale, corresponding to X endowed with the discrete topology. The **sheaf of sections** \mathcal{E}_Σ for one such scenario Σ is defined by mapping each $U \in \mathcal{P}(X)$ to the set of **sections over U** :

$$\mathcal{E}_\Sigma(U) := \prod_{(i, \omega) \in U} O_\omega \quad (1)$$

and by mapping each inclusion $V \subseteq U$ of subsets in $\mathcal{P}(X)$ to the following **restriction map**:

$$\mathcal{E}_\Sigma(V \subseteq U) : \begin{array}{ccc} \mathcal{E}_\Sigma(U) & \longrightarrow & \mathcal{E}_\Sigma(V) \\ s \in \prod_{(i, \omega) \in U} O_\omega & \mapsto & s|_V \in \prod_{(i, \omega) \in V} O_\omega \end{array} \quad (2)$$

The **presheaf of distributions** $\mathcal{D}_R \mathcal{E}_\Sigma$ is obtained by applying the the distribution monad \mathcal{D}_R . Each $U \in \mathcal{P}(X)$ is now sent to the set of R -valued distributions on sections over U :

$$\mathcal{D}_R \mathcal{E}_\Sigma(U) := \left\{ d : \mathcal{E}_\Sigma(U) \rightarrow R \mid \sum_{s \in \mathcal{E}_\Sigma(U)} d(s) = 1 \right\} \quad (3)$$

Each inclusion $V \subseteq U$ in $\mathcal{P}(X)$ is now sent to the marginalisation of R -distributions:

$$\mathcal{D}_R \mathcal{E}_\Sigma(V \subseteq U) : \begin{array}{ccc} \mathcal{D}_R \mathcal{E}_\Sigma(U) & \longrightarrow & \mathcal{D}_R \mathcal{E}_\Sigma(V) \\ d & \mapsto & d|_V := \left(s \mapsto \sum_{s'|_V=s} d(s') \right) \end{array} \quad (4)$$

We say that a subset $\mathcal{J} \subseteq \prod_{\omega \in \Omega} I_\omega$ of the joint inputs is a **global cover** if it includes all local inputs, i.e. if $I_\omega = \{i_\omega \mid \underline{i} \in \mathcal{J}\}$ for all $\omega \in \Omega$. Let $U \subseteq \prod_{\omega \in \Omega} I_\omega$ then each $\underline{i} \in U$ is identified with a subset $\varphi(\underline{i}) := \{(i_\omega, \omega) \mid \omega \in \Omega\} \subseteq X$. The subset \mathcal{J} being a global cover is equivalent to the statement that $\bigcup_{\underline{i} \in \mathcal{J}} \varphi(\underline{i}) = X$ (closer to the formulation given in the original framework).

An **empirical model** for a scenario Σ is a **compatible family** for some global cover \mathcal{J} , i.e. an assignment $e \in \prod_{\underline{i} \in \mathcal{J}} \mathcal{D}_R \mathcal{E}_\Sigma(\varphi(\underline{i}))$ satisfying the following condition:

$$(e_{\underline{i}})|_{\underline{k}} = (e_{\underline{j}})|_{\underline{k}} \text{ for all } \underline{i}, \underline{j} \in \mathcal{J} \text{ and all } \underline{k} \subseteq \underline{i} \cap \underline{j} \quad (5)$$

where $\underline{i} \cap \underline{j}$ is the intersection of \underline{i} and \underline{j} as partial functions (the same as $\varphi^{-1}(\varphi(\underline{i}) \cap \varphi(\underline{j}))$). An empirical model e is said to be **local** if there is some global section $g \in \mathcal{D}_R \mathcal{E}_\Sigma(X)$ for the presheaf of distributions such that $g|_{\varphi(\underline{i})} = e_{\underline{i}}$ for all $\underline{i} \in \mathcal{J}$; the empirical model e is said to be **non-local** if no such global section exists.

4 Sheaf of Sections

The particular construction adopted in the original framework for the presheaf of distributions is due to an interpretation of inputs as measurement choices and outputs as measurement outcomes. The global cover \mathcal{J} is interpreted to define sets of measurements which are jointly compatible, i.e. ones for which a well-defined distribution over joint measurement outcomes can be defined. In this context, defining the presheaf of distributions over the powerset of X signals that we are interested in knowing which subsets $U \subseteq X$ are composed of compatible measurements.

When a non-trivial causal structure is imposed on events, it no longer makes sense to consider *all* possible subsets of X : when $x, y \in \Omega$ are causally related as $x < y$, the outcome for a choice of input $i_y \in I_y$ at event y is (in principle) allowed to depend on some choice of input $i_x \in I_x$ for the preceding event x . Given a generic causal set Ω , we consider the poset $\Lambda(\Omega)$ formed by its **lowersets**—that is, sets closed below w.r.t the order relation \leq of Ω —ordered by subset inclusion:

$$\Lambda(\Omega) := (\{\lambda \subseteq \Omega \mid \forall \omega \in \lambda. \forall \omega' \in \Omega. (\omega' \leq \omega \Rightarrow \omega' \in \lambda)\}, \subseteq) \quad (6)$$

From a spacetime perspective, a set of events $\lambda \subseteq \Omega$ is a lowerset if it contains all points $\omega' \in \Omega$ which lie in the past of any event $\omega \in \lambda$. In this more general causal context, we no longer define $X := \prod_{\omega \in \Omega} I_\omega$ and work within the locale $\mathcal{P}(X)$: instead, we define a new locale \mathcal{L}_Σ which takes the necessary causality constraints into account.

Definition 3. Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario. The associated **locale of inputs** \mathcal{L}_Σ is defined as follows:

$$\mathcal{L}_\Sigma := \coprod_{\lambda \in \Lambda(\Omega)} \prod_{\omega \in \lambda} (\mathcal{P}(I_\omega) \setminus \{\emptyset\}) \quad (7)$$

The generic element $U \in \mathcal{L}_\Sigma$ is a family of non-empty subsets $U_\omega \subseteq I_\omega$ of inputs for all events in some lowerset λ_U of Ω :

$$U = (U_\omega)_{\omega \in \lambda_U} \in \prod_{\omega \in \lambda_U} (\mathcal{P}(I_\omega) \setminus \{\emptyset\}) \quad (8)$$

The partial order on \mathcal{L}_Σ is then defined as follows:

$$V \leq U \Leftrightarrow \lambda_V \subseteq \lambda_U \text{ and } \forall \omega \in \lambda_V. V_\omega \subseteq U_\omega \quad (9)$$

When convenient, we can silently extend elements $U \in \mathcal{L}_\Sigma$ to families defined for all $\omega \in \Omega$ by canonically setting $U_\omega := \emptyset$ for all $\omega \notin \lambda_U$ (so that $\lambda_U = \{\omega \in \Omega \mid U_\omega \neq \emptyset\}$).

Just as in the previous section we identified elements of \mathcal{J} with elements of $\mathcal{P}(X)$, this can be done for all elements of \mathcal{L}_Σ . This turns out to be extremely convenient, so we define the map φ formally.

Definition 4. We will denote by $\varphi : \mathcal{L}_\Sigma \rightarrow \mathcal{P}(X)$ the injective function defined as follows:

$$\varphi(U) := \bigcup_{\omega \in \lambda_U} U_\omega \times \{\omega\} = \{(i, \omega) \mid i \in U_\omega, \omega \in \lambda_U\} \quad (10)$$

The function φ thus defined is clearly order-preserving.

Proposition 5. Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario. The poset \mathcal{L}_Σ from Definition 3 is a locale with the following join and meet:

$$U \cap V := (U_\omega \cap V_\omega)_{\omega \in \lambda_{U \cap V}} \quad U \cup V := (U_\omega \cup V_\omega)_{\omega \in \lambda_{U \cup V}} \quad (11)$$

where:

$$\lambda_{U \cap V} := \{\omega \in \lambda_U \cap \lambda_V \mid U_\omega \cap V_\omega \neq \emptyset\} \quad \lambda_{U \cup V} := \lambda_U \cup \lambda_V \quad (12)$$

The top element is the family of all inputs for all events $\underline{I} = (I_\omega)_{\omega \in \Omega}$ and the bottom element is the empty family (which is indexed by the empty lower set \emptyset and which we itself denote by \emptyset). The covers on an object U are exactly those sets $\mathfrak{R} \subseteq U \downarrow$ of elements $V \leq U$ such that:

$$U_\omega = \bigcup_{V \in \mathfrak{R}} V_\omega \quad \text{for all } \omega \in \Omega \quad (13)$$

Above, we have written $U \downarrow$ for the **downset** $U \downarrow := \{V \in \mathcal{L}_\Sigma \mid V \leq U\}$.

In a locale, a **global cover** is a cover for the top object. In the case of \mathcal{L}_Σ , the top object is $\underline{I} = (I_\omega)_{\omega \in \Omega}$ and a global cover is a subset $\mathfrak{R} \subseteq \mathcal{L}_\Sigma$ such that all local inputs are covered:

$$I_\omega = \bigcup_{V \in \mathfrak{R}} V_\omega \quad \text{for all } \omega \in \Omega \quad (14)$$

A special subclass of covers in \mathcal{L}_Σ is obtained from all possible non-empty subsets $\mathcal{K} \subseteq \prod_{\omega \in \lambda} I_\omega$, for any choice of lower set $\lambda \in \Lambda(\Omega)$. Indeed, each such subset \mathcal{K} corresponds to a cover \mathfrak{R} for an object U defined as follows:

$$\mathfrak{R} := \{(\{i_\omega\})_{\omega \in \lambda} \mid \underline{i} \in \mathcal{K}\} \quad U := (\{i_\omega \mid \underline{i} \in \mathcal{K}\})_{\omega \in \lambda} \quad (15)$$

We say that \mathfrak{R} is the cover **induced by** \mathcal{K} and that U is the object **covered by** \mathcal{K} . In particular, Proposition 5 states that the cover induced by $\mathcal{I}_\Sigma = \prod_{\omega \in \Omega} I_\omega$ is a global cover, which we henceforth denote by \mathfrak{I}_Σ .

In the original framework, the sheaf of sections was defined to send a subset $U \in \mathcal{P}(\prod_{\omega \in \Omega} I_\omega)$ to the set $\mathcal{E}_\Sigma(U) := \prod_{(i, \omega) \in U} O_\omega$: each $f \in \mathcal{E}_\Sigma(U)$ is a function, mapping each input (i, ω) at an event ω to some output $f(i, \omega) \in O_\omega$ at the same event ω . This formulation is special to the non-locality case—where no two distinct events are causally related—and does not generalise directly. Instead, our sheaf of sections will explicitly incorporate the requirement of causality.

Definition 6. Let Ω be a partial order and let $f : \prod_{\omega \in \lambda} A_\omega \rightarrow \prod_{\omega \in \lambda} B_\omega$ be a function for some lower set $\lambda \in \Lambda(\Omega)$ and some families $(A_\omega)_{\omega \in \lambda}$ and $(B_\omega)_{\omega \in \lambda}$ of non-empty sets. We say that f is **causal** if for

the output value $f(\underline{i})_\omega \in O_\omega$ at any event ω depends only on the input values $i_{\omega'}$ at events ω' causally preceding ω (i.e. events $\omega' \in \omega_\downarrow := \{\omega'' \in \Omega \mid \omega'' \leq \omega\}$ in its downset):

$$\forall \underline{i}, \underline{j} \in \prod_{\omega \in \lambda} A_\omega. \left[\underline{i}|_{\omega_\downarrow} = \underline{j}|_{\omega_\downarrow} \Rightarrow f(\underline{i})_\omega = f(\underline{j})_\omega \right] \quad (16)$$

Definition 7. Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario. The **sheaf of sections** \mathcal{E}_Σ is the functor $\mathcal{L}_\Sigma \rightarrow \text{Set}^{\text{op}}$ defined as follows. To each object U in \mathcal{L}_Σ , the functor \mathcal{E}_Σ assigns the following set of functions:

$$\mathcal{E}_\Sigma(U) := \left\{ f : \prod_{\omega \in \lambda_U} U_\omega \rightarrow \prod_{\omega \in \lambda_U} O_\omega \mid f \text{ causal for } \Omega \right\} \quad (17)$$

The definition of the functor \mathcal{E}_Σ on arrows $V \leq U$ in \mathcal{L}_Σ is by restriction:

$$\mathcal{E}_\Sigma(V \leq U) := f \mapsto f|_V \quad (18)$$

where the restricted function $f|_V$ is defined on all $\underline{i} \in V$ as follows:

$$f|_V(\underline{i}) := \left(f(\underline{j}) \right)|_V \text{ for any } \underline{j} \in U \text{ such that } \forall \omega \in \lambda_V. j_\omega = i_\omega \quad (19)$$

Proposition 8. Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario. The restricted function $f|_V$ from Definition 7 is well-defined and the functor $\mathcal{E}_\Sigma : \mathcal{L}_\Sigma \rightarrow \text{Set}^{\text{op}}$ is a sheaf over the locale \mathcal{L}_Σ .

At first, Definition 7 for the sheaf of sections appears unrelated to the one from the original sheaf-theoretic framework for non-locality. However, this turns out to be merely a matter of presentation. In the original framework, the no-signalling sections are defined **explicitly**, in terms of a factorisation into independent functions sending local input to local output at each individual event:

$$\prod_{(i, \omega) \in \varphi(U)} O_\omega \cong \prod_{\omega \in \lambda_U} (U_\omega \rightarrow O_\omega) \quad (20)$$

In this work, the sections are instead defined **implicitly**, as functions from global inputs to global outputs satisfying a causality condition. But the causality condition itself can be used to obtain an equivalent explicit formulation, as families of functions where the output at any event $\omega \in \lambda_U$ can depend on the input at all events $\omega' \leq \omega$:

$$\left\{ f : \prod_{\omega \in \lambda_U} U_\omega \rightarrow \prod_{\omega \in \lambda_U} O_\omega \mid f \text{ causal for } \Omega \right\} \cong \prod_{\omega \in \lambda_U} \left(\left(\prod_{\omega' \leq \omega} U_{\omega'} \right) \rightarrow O_\omega \right) \quad (21)$$

In the non-locality case, where distinct events are causally unrelated, the two explicit formulations coincide. This leads to the following result.

Proposition 9. Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a non-locality scenario and write $X := \coprod_{\omega \in \Omega} I_\omega$. We have that $\mathcal{P}(X)$ and \mathcal{L}_Σ are isomorphic as locales. Furthermore, we have that the sheaf of sections \mathcal{E}_Σ in the original framework for non-locality—defined over the locale $\mathcal{P}(X)$ by Equations 1 and 2—is isomorphic to the sheaf of sections \mathcal{E}_Σ introduced in this work—defined over the locale \mathcal{L}_Σ by Definition 7.

5 Example of a sheaf of sections

Consider a definite causal scenario Σ with four events $\Omega = \{A, B, C, D\}$, arranged into a “diamond” where $C \leq A \leq D$, $C \leq B \leq D$ and the events A and B are causally unrelated. The causal order Ω has the following six lowersets:

$$\Lambda(\Omega) = \{\emptyset, \{C\}, \{C, A\}, \{C, B\}, \{C, A, B\}, \Omega\} \quad (22)$$

We now provide an explicit formulation for $\mathcal{E}_\Sigma(U)$, where U is a generic element of \mathcal{L}_Σ with fixed λ_U , for all possible choices of $\lambda_U \in \Lambda(\Omega)$. When $\lambda_U = \emptyset$, U is the empty function \emptyset and $\mathcal{E}_\Sigma(U) \cong \{\emptyset\}$ contains only the empty family. When $\lambda_U = \{C\}$, functions in $\mathcal{E}_\Sigma(U)$ are simply functions mapping local input at C to local output at C :

$$\mathcal{E}_\Sigma(U) \cong I_C \rightarrow O_C \quad (23)$$

When $\lambda_U = \{C, A\}$, functions in $\mathcal{E}_\Sigma(U)$ can be explicitly seen as pairs of individual functions for the two events, where the output at C can only depend on the input at C while the output at A can depend jointly on the input at both C and A :

$$\mathcal{E}_\Sigma(U) \cong (I_C \rightarrow O_C) \times (I_C \times I_A \rightarrow O_A) \quad (24)$$

The case $\lambda_U = \{C, B\}$ is analogous. When $\lambda_U = \{C, A, B\}$, functions in $\mathcal{E}_\Sigma(U)$ can be explicitly seen as triples of individual functions for the three events, where the output at C can only depend on the input at C while the output at A can depend jointly on the input at both C and A , and analogously for the output at B :

$$\mathcal{E}_\Sigma(U) \cong (I_C \rightarrow O_C) \times (I_C \times I_A \rightarrow O_A) \times (I_C \times I_B \rightarrow O_B) \quad (25)$$

Finally, when $\lambda_U = \Omega = \{C, A, B, D\}$, functions in $\mathcal{E}_\Sigma(U)$ can be explicitly seen as quadruples of individual functions for the four events. The output at C can only depend on the input at C , the output at A can depend jointly on the input at both C and A , the output at B can depend jointly on the input at both C and B , and the output at D can depend on the joint input at all four events:

$$\mathcal{E}_\Sigma(U) \cong (I_C \rightarrow O_C) \times (I_C \times I_A \rightarrow O_A) \times (I_C \times I_B \rightarrow O_B) \times (I_C \times I_A \times I_B \times I_D \rightarrow O_D) \quad (26)$$

6 Empirical Models

The additional constraints introduced by causality are entirely captured by the definition of the sheaf of sections: given that, the definition of the presheaf of distributions and empirical models parallels that of the original framework.

Definition 10. *Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario and let R be a commutative semiring. The presheaf of R -distributions for Σ is the presheaf $\mathcal{D}_R \mathcal{E}_\Sigma$ obtained by composing the R -distribution monad \mathcal{D}_R with the sheaf of sections \mathcal{E}_Σ .*

Explicitly, each $U \in \mathcal{L}_\Sigma$ is sent to the set of R -valued distributions on $\mathcal{E}_\Sigma(U)$:

$$\mathcal{D}_R \mathcal{E}_\Sigma(U) := \left\{ d : \mathcal{E}_\Sigma(U) \rightarrow R \mid \sum_{f \in \mathcal{E}_\Sigma(U)} d(f) = 1 \right\} \quad (27)$$

Each inclusion $V \leq U$ in \mathcal{L}_Σ is sent to the marginalisation of R -distributions:

$$\begin{aligned} \mathcal{D}_R \mathcal{E}_\Sigma(V \leq U) : \mathcal{D}_R \mathcal{E}_\Sigma(U) &\longrightarrow \mathcal{D}_R \mathcal{E}_\Sigma(V) \\ d &\longmapsto d|_V := \left(f \mapsto \sum_{f'|_V=f} d(f') \right) \end{aligned} \quad (28)$$

Definition 11. Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario and let R be a commutative semiring. An (R -valued) **empirical model** is a compatible family $e \in \prod_{\underline{i} \in \mathcal{I}_\Sigma} \mathcal{D}_R \mathcal{E}_\Sigma(\underline{i})$ for the global cover \mathfrak{F}_Σ induced by the set of joint inputs $\mathcal{I}_\Sigma = \prod_{\omega \in \Omega} I_\omega$.

Remark 12. If R is not specified, we take $R = \mathbb{R}^+$ by default: in this case, we drop the “ R -valued” qualification for the empirical model and we refer to $\mathcal{D}_{\mathbb{R}^+} \mathcal{E}_\Sigma$ as the “presheaf of distributions”, sometimes also denoted by $\mathcal{D} \mathcal{E}_\Sigma$.

The definition of the empirical model is sheaf-theoretic, as a compatible family for the global cover \mathfrak{F}_Σ . In practice, however, it is easier to work with the following equivalent formulation as a causal conditional distribution.

Definition 13. Let Ω be a partial order and let R be a commutative semiring. An (R -valued) **conditional distribution** is a function $d : \prod_{\omega \in \lambda} A_\omega \rightarrow \mathcal{D}_R(\prod_{\omega \in \lambda} B_\omega)$ for some lower set $\lambda \in \Lambda(\Omega)$ and some families $(A_\omega)_{\omega \in \lambda}$ and $(B_\omega)_{\omega \in \lambda}$ of non-empty sets. We adopt the following notation, in analogy with probability theory:

$$d(\underline{o} | \underline{i}) := d_{\underline{i}}(\underline{o}) \quad (29)$$

for $\underline{i} \in \prod A_\omega$, $\underline{o} \in \prod B_\omega$.

Definition 14. Let Ω be a partial order, let R be a commutative semiring and consider an R -valued conditional distribution $d : \prod_{\omega \in \lambda} A_\omega \rightarrow \mathcal{D}_R(\prod_{\omega \in \lambda} B_\omega)$. We define restriction $d|_\mu$ of d to μ by marginalisation as follows:

$$d|_\mu((o_\omega)_{\omega \in \mu} | \underline{i}) := \sum_{\substack{\underline{o}' \in \prod_{\omega \in \lambda} B_\omega \\ \text{s.t. } \underline{o}'|_\mu = \underline{o}|_\mu}} d(\underline{o}' | \underline{i}) \quad \text{for all } \underline{i} \in \prod_{\omega \in \lambda} A_\omega \quad (30)$$

We say that d is **causal** if for every sub-lower set $\mu \subseteq \lambda$ in $\Lambda(\Omega)$ we have that the restriction $d|_\mu$ is independent of the value of inputs $(i_\omega)_{\omega \in \lambda \setminus \mu}$ outside of μ :

$$\underline{i}|_\mu = \underline{j}|_\mu \implies \forall \underline{o} \in \prod_{\omega \in \mu} B_\omega. \left[d|_\mu(\underline{o} | \underline{i}) = d|_\mu(\underline{o} | \underline{j}) \right] \quad (31)$$

In this case, we freely identify the restriction $d|_\mu$ with the equivalent distribution supported on μ , i.e. we deem the following to be well-typed:

$$d|_\mu : \prod_{\omega \in \mu} A_\omega \rightarrow \mathcal{D}_R \left(\prod_{\omega \in \mu} B_\omega \right) \quad (32)$$

Proposition 15. Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario. The data defining an (R -valued) empirical model $e \in \prod_{\underline{i} \in \mathcal{I}_\Sigma} \mathcal{D}_R \mathcal{E}_\Sigma(\underline{i})$ is exactly the data defining a causal (R -valued) conditional distribution $\hat{e} : \mathcal{I}_\Sigma \rightarrow \mathcal{D}_R(\underline{O}_\Sigma)$:

$$\hat{e}(\underline{o} | \underline{i}) := e_{\underline{i}}(\underline{i} \mapsto \underline{o}) \quad (33)$$

Above, we denoted by $\underline{i} \mapsto \underline{o}$ a generic function in $\mathcal{E}_\Sigma(\{(i_\omega)_{\omega \in \Omega}\})$, with singleton domain $\{\underline{i}\}$.

Conditional distributions are familiar objects in the study of non-locality. One possible way of representing them is in tabular form, with rows indexed by the joint inputs in $\mathcal{I}_\Sigma = \prod_{\omega \in \Omega} I_\omega$, columns indexed by the joint outputs in $\mathcal{O}_\Sigma = \prod_{\omega \in \Omega} O_\omega$, and cells containing the R values defining the distribution. The tabular representation makes it clear that—were it not for the requirement of causality—conditional distributions would form a polytope, namely the \mathcal{I}_Σ -fold product (i.e. the product over the rows) of a simplex (specifying the distribution value for each row):

$$\left(\Delta^{O_\Sigma} \right)^{\times \mathcal{I}_\Sigma} \subset \mathbb{R}^{O_\Sigma \times \mathcal{I}_\Sigma} \quad (34)$$

The question, then, is whether the introduction of the causality requirement changes things substantially. The (negative) answer lies in the following result.

Proposition 16. *Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario. The probabilistic empirical models (i.e. the \mathbb{R}^+ -valued ones) form a polytope, obtained by intersecting the product polytope from 34 with the \mathbb{R} -linear subspace defined by the following **causality equations**:*

$$\left(\sum_{\substack{\underline{o} \in O_\Sigma \text{ s.t.} \\ \underline{o}|_\lambda = \underline{o}'}} (d(\underline{o}|\underline{i}) - d(\underline{o}|\underline{j})) = 0 \right) \quad \text{for all} \quad \begin{cases} \lambda \in \Lambda(\Omega) \\ \underline{o}' \in \prod_{\omega \in \lambda} O_\omega \\ \underline{i}, \underline{j} \in \mathcal{I}_\Sigma \text{ s.t. } \underline{i}|_\lambda = \underline{j}|_\lambda \end{cases} \quad (35)$$

Note that the unknown in the system of equations above is the distribution d , which can be seen as a vector in $\mathbb{R}^{O_\Sigma \times \mathcal{I}_\Sigma}$ (e.g. from its tabular representation).

Remark 17. The causality equations above are highly redundant. Instead of taking all $\underline{i}, \underline{i}' \in \mathcal{I}_\Sigma$, it is enough to impose some total order \leq on \mathcal{I}_Σ —e.g. lexicographic order, together with some choice of total orders on the sets I_ω for all $\omega \in \Omega$ —and consider all $\underline{i} < \underline{i}'$ instead.

Definition 18. *Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario and let R be a commutative semiring. An (R -valued) empirical model $e \in \prod_{\underline{i} \in \mathcal{I}_\Sigma} \mathcal{D}_R \mathcal{E}_\Sigma(\underline{i})$ is said to be **local** if there is some global section $g \in \mathcal{D}_R \mathcal{E}_\Sigma(\underline{I})$ for the presheaf of distributions such that $g|_{\{(i_\omega)\}_{\omega \in \Omega}} = e_{\underline{i}}$ for all $\underline{i} \in \mathcal{I}_\Sigma$. The empirical model e is said to be **non-local** if no such global section exists.*

The definition of locality for an empirical model does not change from the original sheaf-theoretic framework. What's more interesting is the physical meaning of locality: what does it mean, operationally, to have a global section in a generic definite causal scenario? The answer is provided by Proposition 21 below: having a global section means that there exists a classical hidden variable model in terms of causal functions (or, more precisely, in terms of “deterministic” distributions).

Definition 19. *Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario and let R be a commutative semiring. The **delta distributions** are the following R -valued causal conditional distributions, indexed by all causal functions $f : \mathcal{I}_\Sigma \rightarrow O_\Sigma$:*

$$\delta_f(\underline{o}|\underline{i}) = \begin{cases} 1 & \text{if } f(\underline{i}) = \underline{o} \\ 0 & \text{otherwise} \end{cases} \quad (36)$$

Remark 20. The delta distributions are nothing but the causal functions themselves, seen as boolean conditional distributions; these are valid distributions for all choices of R , because every commutative semiring R has the elements 0 and 1. We also refer to these as the **deterministic** distributions.

Proposition 21. *Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario and let R be a commutative semiring. Let $e \in \prod_{i \in \mathcal{I}_\Sigma} \mathcal{D}_R \mathcal{E}_\Sigma(i)$ be an R -valued empirical model. The empirical model e is local if and only if there is a family F of causal functions $f : \mathcal{I}_\Sigma \rightarrow \mathcal{O}_\Sigma$ and an R -valued distribution $p \in \mathcal{D}_R(F)$ such that:*

$$\hat{e} = \sum_{f \in F} p(f) \delta_f \quad (37)$$

where $\hat{e} : \mathcal{I}_\Sigma \rightarrow \mathcal{D}_R(\mathcal{O}_\Sigma)$ is the causal conditional distribution corresponding to e .

7 Realisability

Empirical models are not, by themselves, operational: rather, they paint a distributional and phenomenological picture of operational reality. They are concerned with the results of experiments—e.g. as the observed distribution of joint outputs conditional on joint inputs—but they say nothing about how the experiments could be realised within some operational theory of physics. The study of realisability is not—and shouldn't be—a question to be tackled from a sheaf-theoretic perspective. Instead, the question should be addressed by connecting the sheaf-theoretic frameworks to other frameworks, already designed to deal with the operational perspective [24, 26, 30, 31, 47].

The basic building blocks of operational theories are *instruments*: local transformations of physical systems controlled by some local classical input and returning some local classical output. The instruments can be combined by wiring their input/output physical systems together in various ways—including provision of initial states and discarding of systems no longer needed—and the resulting circuit/diagram defines a distribution on joint instrument outputs conditional on joint instrument inputs. When talking about a *realisation* of an empirical model in such an operational theory, we mean exactly one such diagram, having the empirical model as its distribution and where the wiring of physical systems respects the given causal order.

Here, we choose to work within the framework of probabilistic theories by [43], an operational framework born from the universe of categorical quantum mechanics [12, 30, 31]. For the special case of non-locality scenarios, it was already shown in the original [43] that all diagrams of a certain shape in R -probabilistic theories yield R -valued empirical models according to the original framework for non-locality. We now extend that result to the sheaf-theoretic framework for (definite) causality: we show that diagrams over definite causal scenarios in R -probabilistic theories, as defined by the authors in the recent [64], always yield R -valued empirical models for the new sheaf-theoretic framework.

Proposition 22. *Let $\Theta = (\Gamma, \underline{I}, \underline{O})$ be a definite causal scenario as defined in [64], where Γ has no boundary nodes. Let Ω be the definite causal order underlying the acyclic framed multigraph Γ , so that $\Sigma = (\Omega, \underline{I}, \underline{O})$ is a definite causal scenario as defined in this work. If Δ is a diagram over Θ in a R -probabilistic theory \mathcal{C} , as defined in [64], and Δ involves normalised processes at all events, then Δ yields an R -valued empirical model e for Σ . We refer to such a Δ as a **realisation** of e in \mathcal{C} .*

8 Example of a causal empirical model

Consider the “diamond” definite causal scenario Σ from our previous example, on four events $\Omega = A, B, C, D$. Set boolean ($\mathbb{B} = \{0, 1\}$) inputs and outputs at all events, i.e. $I_\omega = O_\omega = \mathbb{B}$. We now construct a simple empirical model for Σ “from the ground up”, by providing a representation in quantum theory (which is an example of a probabilistic theory, i.e. one with $R = \mathbb{R}^+$). At event C , an instrument outputs one of the 4

9 Beyond Definite Causal Scenarios

Because our operational interpretation of definite causal scenarios is in terms of *local* black-box devices, the same interpretation can be extended, essentially unchanged, to indefinite causal scenarios. What changes in the transition to indefinite causality is that the classical inputs are now allowed to influence not only the classical outputs at events after them, but also the causal order of the events themselves. Within our framework, this can be implemented by changing the way in which causal functions are defined: instead of being causal with respect to a fixed partial order Ω , they will be patched together from partial functions which are individually causal for different partial orders Ω' , all compatible with the same *pre-order* Ω (where equivalence classes represent events in mutually indefinite causal order).

The sheaf-theoretic construction itself is too long and complicated to find adequate exposition in the few remaining pages of this work, and will be presented fully in an upcoming publication. However, the transition from definite to indefinite causal scenarios is easy to describe, and we can already use it to reason about an unusual construction by Baumeler, Feix and Wolf [18, 19].

Definition 23. A **causal scenario** is the same as a definite causal scenario $\Sigma := (\Omega, \underline{I}, \underline{O})$, except that Ω is now allowed to be a finite pre-order (instead of just a partial order). We write $x \simeq y$ for events $x, y \in \Omega$ such that $x \leq y$ and $y \leq x$: when $x \neq y$, we say that two such x and y are in **indefinite causal order** (we write $x \not\leq y$ when $x \neq y$ and at most one of $x \leq y$ or $y \leq x$ holds). We say that a causal scenario is **definite** if $x \neq y$ implies $x \not\leq y$ for all $x, y \in \Omega$ —in which case they coincide with the definition of Section 2—and we say that it is **indefinite** otherwise.

Remark 24. We have four possible causal relationships for two events $x \neq y$ in a pre-order:

- if $x \not\leq y$ and $x \not\geq y$, then x and y are **causally unrelated**
- if $x < y$, then x **causally precedes** y
- if $x > y$, then x **causally succeeds** y
- if $x \simeq y$, then x and y are in **indefinite causal order**

If $x \neq y$, we sometime say that x and y are in **definite causal order**, and this includes the possibility that they are causally unrelated.

For the purposes of discussing the indefinite causal scenario below, it is enough to understand the polytope structure of indefinite causality, i.e. to describe how causal equations extend from partial orders to pre-orders. As it turns out, there is no work to be done here: Proposition 16 already describes the equations for both definite and indefinite causal scenarios. What changes, in practice, is that pre-orders have fewer lowersets, because all events in mutually indefinite causal order must be discarded together.

10 Example of an indefinite causal empirical model

Baumeler, Feix and Wolf [18, 19] describe an empirical model e^{BFW} for a certain classical scenario which cannot be realised under the assumption of a definite causal order:

ABC	000	001	010	011	100	101	110	111
000	1/2	0	0	0	0	0	0	1/2
001	0	0	0	1/2	1/2	0	0	0
010	0	1/2	0	0	0	0	1/2	0
011	0	0	1/2	0	0	1/2	0	0
100	0	0	1/2	0	0	1/2	0	0
101	0	1/2	0	0	0	0	1/2	0
110	0	0	0	1/2	1/2	0	0	0
111	1/2	0	0	0	0	0	0	1/2

(42)

To come to this conclusion, it is enough to compute the Ω -causal fraction for the model e^{BFW} above with respect to all definite causal orders Ω on the three events $\{A, B, C\}$: this is the maximum $\lambda \in [0, 1]$ such that $\lambda \cdot e \leq e^{BFW}$ for some empirical model e which is causal for Ω , where the order \leq on conditional distributions is defined componentwise as:

$$e \leq f \Leftrightarrow \forall i, \underline{o}. e(\underline{o}|i) \leq f(\underline{o}|i) \quad (43)$$

For the empirical model e^{BFW} above, the Ω -causal fraction turns out to be 0 for all partial orders Ω : no part of the observed correlations can be explained by definite causality. Moving to indefinite causality, we find the Ω -causal fraction to be 1 for exactly four pre-orders: the indiscrete pre-order on $\{A, B, C\}$, where all three events are in mutually indefinite causal order, and the three pre-orders obtained by putting each of the three events definitely before the other two, which remain in indefinite causal order. All other pre-orders have causal fraction 0.

The results above (which will be presented in full in the aforementioned upcoming publication) are compatible with the original description of cyclicity for the Baumeler-Feix-Wolf construction: the correlations can be explained *fully* by letting *any* one of the three parties act first, with the remaining two in indefinite causal order. What's even more interesting, and certainly deserving of further investigation, is that fixing the input of the first acting party breaks the indefinite causality entirely, resulting in a no-signalling empirical model for the remaining two parties. For example, considering the pre-order where C is before $\{A, B\}$. Fixing the input of C to 0 and marginalising over its output results in the following no-signalling empirical model for A and B :

AB	00	01	10	11
00	1/2	0	0	1/2
01	1/2	0	0	1/2
10	0	1/2	1/2	0
11	0	1/2	1/2	0

(44)

The model obtained by fixing the input of C to 1 is analogous.

Acknowledgements

SG would like to acknowledge support from Hashberg Ltd. NP would like to acknowledge financial support from EPSRC and the Pirie-Reid Scholarship. The authors would like to thank Shane Mansfield for

constructive discussion. This publication was made possible through the support of the grant 61466 ‘The Quantum Information Structure of Spacetime (QISS)’ (qiss.fr) from the John Templeton Foundation. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

References

- [1] Sivert Aasnæss (2020): *Cohomology and the algebraic structure of contextuality in measurement based quantum computation*. *EPTCS* 318, pp. 242–253, doi:10.4204/EPTCS.318.15.
- [2] Samson Abramsky (2013): *Relational databases and Bell's theorem*. In: *In search of elegance in the theory and practice of computation*, Springer, pp. 13–35, doi:10.1007/978-3-642-41660-6_2.
- [3] Samson Abramsky & Rui Soares Barbosa (2021): *The logic of contextuality*. *Leibniz International Proceedings in Informatics (LIPIcs)* 183: 5:1–5:18, doi:10.4230/LIPIcs.CSL.2021.5.
- [4] Samson Abramsky, Rui Soares Barbosa, Martti Karvonen & Shane Mansfield (2019): *A comonadic view of simulation and quantum resources*. In: *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, IEEE, pp. 1–12, doi:10.1109/LICS.2019.8785677.
- [5] Samson Abramsky, Rui Soares Barbosa, Kohei Kishida, Raymond Lal & Shane Mansfield (2015): *Contextuality, cohomology and paradox*. *Leibniz International Proceedings in Informatics (LIPIcs)* 41:211–228, doi:10.4230/LIPIcs.CSL.2015.211.
- [6] Samson Abramsky, Rui Soares Barbosa, Kohei Kishida, Raymond Lal & Shane Mansfield (2016): *Possibilities determine the combinatorial structure of probability polytopes*. *Journal of Mathematical Psychology* 74, pp. 58–65, doi:10.1016/j.jmp.2016.03.006.
- [7] Samson Abramsky, Rui Soares Barbosa & Shane Mansfield (2017): *Contextual fraction as a measure of contextuality*. *Physical review letters* 119(5), p. 050504, doi:10.1103/PhysRevLett.119.050504.
- [8] Samson Abramsky, Rui Soares Barbosa, Nadish de Silva & Octavio Zapata (2017): *The quantum monad on relational structures*. *Leibniz International Proceedings in Informatics (LIPIcs)* 83:35:1–35:19, doi:10.4230/LIPIcs.MFCS.2017.35.
- [9] Samson Abramsky & Adam Brandenburger (2011): *The sheaf-theoretic structure of non-locality and contextuality*. *New Journal of Physics* 13(11), p. 113036, doi:10.1088/1367-2630/13/11/113036.
- [10] Samson Abramsky & Adam Brandenburger (2014): *An operational interpretation of negative probabilities and no-signalling models*. In: *Horizons of the mind. A tribute to Prakash Panangaden*, Springer, pp. 59–75, doi:10.1007/978-3-319-06880-0_3.
- [11] Samson Abramsky & Giovanni Carù (2019): *Non-locality, contextuality and valuation algebras: a general theory of disagreement*. *Philosophical Transactions of the Royal Society A* 377(2157), p. 20190036, doi:10.1098/rsta.2019.0036.
- [12] Samson Abramsky & Bob Coecke (2009): *Categorical quantum mechanics*. *Handbook of quantum logic and quantum structures* 2, pp. 261–325, doi:10.1016/B978-0-444-52869-8.50010-4.
- [13] Samson Abramsky & Lucien Hardy (2012): *Logical bell inequalities*. *Physical Review A* 85(6), p. 062114, doi:10.1103/PhysRevA.85.062114.
- [14] Samson Abramsky, Shane Mansfield & Rui Soares Barbosa (2012): *The cohomology of non-locality and contextuality*. *EPTCS* 95, doi:10.4204/EPTCS.95.1.
- [15] Rui Soares Barbosa, Tom Douce, Pierre-Emmanuel Emeriau, Elham Kashefi & Shane Mansfield (2019): *Continuous-variable nonlocality and contextuality*. *arXiv preprint arXiv:1905.08267*.
- [16] Jonathan Barrett, Robin Lorenz & Ognjan Oreshkov (2019): *Quantum causal models*. *arXiv preprint arXiv:1906.10726*.
- [17] Jonathan Barrett, Robin Lorenz & Ognjan Oreshkov (2021): *Cyclic quantum causal models*. *Nature communications* 12(1), pp. 1–15, doi:10.1038/s41467-020-20456-x.
- [18] Ämin Baumeler, Adrien Feix & Stefan Wolf (2014): *Maximal incompatibility of locally classical behavior and global causal order in multiparty scenarios*. *Phys. Rev. A* 90, p. 042106, doi:10.1103/PhysRevA.90.042106.
- [19] Ämin Baumeler & Stefan Wolf (2014): *Perfect signaling among three parties violating predefined causal order*. In: *2014 IEEE International Symposium on Information Theory*, IEEE, pp. 526–530, doi:10.1109/ISIT.2014.6874888.

- [20] N. Bohr (1935): *Can Quantum-Mechanical Description of Physical Reality be Considered Complete?* *Phys. Rev.* 48, pp. 696–702, doi:10.1103/PhysRev.48.696. Available at <https://link.aps.org/doi/10.1103/PhysRev.48.696>.
- [21] Cyril Branciard, Mateus Araújo, Adrien Feix, Fabio Costa & Ľaslav Brukner (2015): *The simplest causal inequalities and their violation*. *New Journal of Physics* 18(1), p. 013008, doi:10.1088/1367-2630/18/1/013008.
- [22] Giovanni Caru (2017): *On the cohomology of contextuality*. *EPTCS* 266, doi:10.4204/EPTCS.236.2.
- [23] Giovanni Caru (2018): *Towards a complete cohomology invariant for non-locality and contextuality*. *arXiv preprint arXiv:1807.04203*.
- [24] Giulio Chiribella, Giacomo Mauro D’Ariano & Paolo Perinotti (2010): *Probabilistic theories with purification*. *Physical Review A* 81(6), p. 062348, doi:10.1103/PhysRevA.81.062348.
- [25] Giulio Chiribella, Giacomo Mauro D’Ariano & Paolo Perinotti (2011): *Informational derivation of quantum theory*. *Physical Review A* 84(1), p. 012311, doi:10.1103/PhysRevA.84.012311.
- [26] Giulio Chiribella, Giacomo Mauro D’Ariano & Paolo Perinotti (2016): *Quantum from principles*. In: *Quantum theory: informational foundations and foils*, Springer, pp. 171–221, doi:10.1007/978-94-017-7303-4_6.
- [27] 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.
- [28] Bob Coecke (2014): *Terminality implies non-signalling*. *EPTCS* 172, pp. 27–35, doi:10.4204/EPTCS.172.3.
- [29] Bob Coecke (2016): *Terminality implies no-signalling... and much more than that*. *New Generation Computing* 34(1-2), pp. 69–85, doi:10.1007/s00354-016-0201-6.
- [30] 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.
- [31] Bob Coecke & Aleks Kissinger (2017): *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning*. Cambridge University Press, doi:10.1017/9781316219317.
- [32] Bob Coecke & Raymond Lal (2012): *Time asymmetry of probabilities versus relativistic causal structure: An arrow of time*. *Physical review letters* 108(20), p. 200403, doi:10.1103/PhysRevLett.108.200403.
- [33] Nadish De Silva & Rui Soares Barbosa (2019): *Contextuality and noncommutative geometry in quantum mechanics*. *Communications in Mathematical Physics* 365(2), pp. 375–429, doi:10.1007/s00220-018-3222-9.
- [34] Ehtibar N. Dzhafarov, Janne V. Kujala & Victor H. Cervantes (2016): *Contextuality-by-Default: A Brief Overview of Ideas, Concepts, and Terminology*. In Harald Atmanspacher, Thomas Filk & Emmanuel Pothos, editors: *Quantum Interaction*, Springer International Publishing, Cham, pp. 12–23, doi:10.1007/978-3-319-28675-4_2.
- [35] 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.
- [36] Pierre-Emmanuel Emeriau, Mark Howard & Shane Mansfield (2020): *Quantum Advantage in Information Retrieval*. *arXiv preprint arXiv:2007.15643*.
- [37] Tobias Fritz (2010): *Quantum correlations in the temporal Clauser–Horne–Shimony–Holt (CHSH) scenario*. *New Journal of Physics* 12(8), p. 083055, doi:10.1088/1367-2630/12/8/083055.
- [38] Tobias Fritz (2012): *Beyond Bell’s theorem: correlation scenarios*. *New Journal of Physics* 14(10), p. 103001, doi:10.1088/1367-2630/14/10/103001.
- [39] Tobias Fritz (2016): *Beyond Bell’s theorem II: Scenarios with arbitrary causal structure*. *Communications in Mathematical Physics* 341(2), pp. 391–434, doi:10.1007/s00220-015-2495-5.
- [40] Rodrigo Gallego, Lars Erik Würflinger, Rafael Chaves, Antonio Acín & Miguel Navascués (2014): *Non-locality in sequential correlation scenarios*. *New Journal of Physics* 16(3), p. 033037, doi:10.1088/1367-2630/16/3/033037/.
- [41] Stefano Gogioso (2017): *Categorical Quantum Dynamics*. Ph.D. thesis, University of Oxford.

- [42] Stefano Gogioso (2017): *Fantastic quantum theories and where to find them*. To appear in *Compositionality (ACT 2019)*. Available at <https://arxiv.org/abs/1703.10576>.
- [43] Stefano Gogioso & Carlo Maria Scandolo (2018): *Categorical probabilistic theories*. *EPTCS* 266, pp. 367–385, doi:10.4204/EPTCS.266.23.
- [44] Stefano Gogioso & William Zeng (2019): *Generalised Mermin-type non-locality arguments*. *Logical Methods in Computer Science* 15, doi:10.23638/LMCS-15(2:3)2019.
- [45] Kaumudibikash Goswami, Christina Giarmatzi, Michael Kewming, Fabio Costa, Cyril Branciard, Jacqueline Romero & Andrew G. White (2018): *Indefinite Causal Order in a Quantum Switch*. *Physical Review Letters* 121(9):090503, doi:10.1103/PhysRevLett.121.090503.
- [46] Joe Henson, Raymond Lal & Matthew F Pusey (2014): *Theory-independent limits on correlations from generalised Bayesian networks*. *New Journal of Physics* 16(11), p. 113043, doi:10.1088/1367-2630/16/11/113043.
- [47] Christiaan Heunen & Jamie Vicary (2019): *Categories for quantum theory: an introduction*. Oxford University Press, United Kingdom, doi:10.1093/oso/9780198739623.001.0001.
- [48] Jørgen Kalckar (1985): *General Discussion At The Fifth Solvay Conference: Unpublished Manuscript From Folder Labelled Notes From Solvay Meeting(1927)*. In: *Foundations of Quantum Physics I (1926–1932), Niels Bohr Collected Works* 6, Elsevier, pp. 99 – 106, doi:10.1016/S1876-0503(08)70332-3.
- [49] Martti Karvonen (2018): *Categories of empirical models*. *EPTCS* 287, doi:10.4204/eptcs.287.14.
- [50] Martti Karvonen (2021): *Neither contextuality nor non-locality admit catalysts*. *arXiv preprint arXiv:2102.07637*.
- [51] Andrei Khrennikov (1993): *p-Adic probability theory and its applications. The principle of statistical stabilization of frequencies*. *Theoretical and Mathematical Physics* 97(3), pp. 1340–1348, doi:10.1007/BF01015763.
- [52] Kohei Kishida (2014): *Stochastic Relational Presheaves and Dynamic Logic for Contextuality*. *Electronic Proceedings in Theoretical Computer Science* 172, p. 115–132, doi:10.4204/eptcs.172.9.
- [53] Kohei Kishida (2016): *Logic of local inference for contextuality in quantum physics and beyond*. *Forty-Third International Colloquium on Automata, Languages, and Programming (ICALP 2016)* 113:1–113:14, doi:10.4230/LIPIcs.ICALP.2016.113.
- [54] Aleks Kissinger, Matty Hoban & Bob Coecke (2017): *Equivalence of relativistic causal structure and process terminality*. *arXiv preprint arXiv:1708.04118*.
- [55] Aleks Kissinger & Sander Uijlen (2017): *A categorical semantics for causal structure*. In: *2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, IEEE, pp. 1–12, doi:10.23638/LMCS-15(3:15)2019.
- [56] Shane Mansfield (2017): *Consequences and applications of the completeness of Hardy’s nonlocality*. *Physical Review A* 95(2), p. 022122, doi:10.1103/PhysRevA.95.022122.
- [57] Shane Mansfield (2017): *Contextuality in Causal scenarios*. <https://simons.berkeley.edu/sites/default/files/docs/8998/leachtsimons.pdf> (accessed 14 Mar 2021).
- [58] Shane Mansfield (2017): *A unified approach to contextuality and violations of macrorealism*. https://www1.psych.purdue.edu/~ehtibar/workshop/files/leacht_prague.pdf (accessed 14 Mar 2021).
- [59] Shane Mansfield & Rui Soares Barbosa (2014): *Extendability in the sheaf-theoretic approach: Construction of Bell models from Kochen-Specker models*. *arXiv preprint arXiv:1402.4827*.
- [60] Shane Mansfield & Tobias Fritz (2012): *Hardy’s non-locality paradox and possibilistic conditions for non-locality*. *Foundations of Physics* 42(5), pp. 709–719, doi:10.1007/s10701-012-9640-1.
- [61] Shane Mansfield & Elham Kashefi (2018): *Quantum advantage from sequential-transformation contextuality*. *Physical review letters* 121(23), p. 230401, doi:10.1103/PhysRevLett.121.230401.
- [62] Ognjan Oreshkov, Fabio Costa & Časlav Brukner (2012): *Quantum correlations with no causal order*. *Nature Communications* 3:1092, doi:10.1038/ncomms2076.

- [63] Ognjan Oreshkov & Christina Giarmatzi (2016): *Causal and causally separable processes*. *New Journal of Physics* 18(9), p. 093020, doi:10.1088/1367-2630/18/9/093020/.
- [64] Nicola Pinzani & Stefano Gogioso (2020): *Giving Operational Meaning to the Superposition of Causal Orders*. To appear in *EPTCS (QPL 2020)*. Available at <https://arxiv.org/abs/2003.13306>.
- [65] 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, pp. 1–20, doi:10.1109/LICS.2019.8785664.
- [66] Arkady Plotnitsky (2019): *Spooky predictions at a distance: reality, complementarity and contextuality in quantum theory*. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 377(2157), p. 20190089, doi:10.1098/rsta.2019.0089.
- [67] Hans Reichenbach (1956): *The direction of time*. University of California Berkeley Press, doi:10.1063/1.3059791.
- [68] Giulia Rubino, Lee A. Rozema, Adrien Feix, Mateus Araújo, Jonas M. Zeuner, Lorenzo M. Procopio, Časlav Brukner & Philip Walther (2017): *Experimental verification of an indefinite causal order*. *Science Advances* 3(3), doi:10.1126/sciadv.1602589.
- [69] Sam Staton & Sander Uijlen (2018): *Effect algebras, presheaves, non-locality and contextuality*. *Information and Computation* 261, pp. 336–354, doi:10.1016/j.ic.2018.02.012.
- [70] Linde Wester (2018): *Almost Equivalent Paradigms of Contextuality*. *Electronic Proceedings in Theoretical Computer Science* 266, p. 1–22, doi:10.4204/eptcs.266.1.
- [71] Linde Wester (2018): *Classical and Quantum Structures in Computation*. Ph.D. thesis, University of Oxford.

A Proofs

Proof of Proposition 5

Proposition 5. *Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario. The poset \mathcal{L}_Σ from Definition 3 is a locale with the following join and meet:*

$$\begin{aligned} U \cap V &:= (U_\omega \cap V_\omega)_{\omega \in \lambda_{U \cap V}} & U \cup V &:= (U_\omega \cup V_\omega)_{\omega \in \lambda_{U \cup V}} \\ \lambda_{U \cap V} &:= \{\omega \in \lambda_U \cap \lambda_V \mid U_\omega \cap V_\omega \neq \emptyset\} & \lambda_{U \cup V} &:= \lambda_U \cup \lambda_V \end{aligned}$$

The top element is the family of all inputs for all events $\underline{I} = (I_\omega)_{\omega \in \Omega}$ and the bottom element is the empty family (which is indexed by the empty lower set \emptyset and which we itself denote by \emptyset). The covers on an object U are exactly those sets $\mathfrak{R} \subseteq U \downarrow$ of elements $V \leq U$ such that:

$$U_\omega = \bigcup_{V \in \mathfrak{R}} V_\omega \quad \text{for all } \omega \in \Omega$$

Above, we have written $U \downarrow$ for the **downset** $U \downarrow := \{V \in \mathcal{L}_\Sigma \mid V \leq U\}$.

Proof. By definition, $U \cap V$ and $U \cup V$ are elements of \mathcal{L}_Σ : the only subtlety is in the definition of $\lambda_{U \cap V}$, where we must take care to exclude the ω where $U_\omega \cap V_\omega = \emptyset$. Because $\lambda_{U \cap V} \subseteq \lambda_U, \lambda_V$ and $U_\omega \cap V_\omega \subseteq U_\omega, V_\omega$, we have that $U \cap V \leq U, V$. Analogously, $U \cup V \geq U, V$. If $W \leq U, V$, then $W_\omega \subseteq U_\omega, V_\omega$ and $\lambda_W \subseteq \lambda_U, \lambda_V$. Furthermore, if $\omega \in \lambda_W$ then $W_\omega \neq \emptyset$, so that $U_\omega \cap V_\omega \neq \emptyset$ and hence $\omega \in \lambda_{U \cap V}$. We conclude that $W \leq U \cap V$. Analogously, if $W \geq U, V$, then $W \geq U \cup V$. This proves that $U \cap V$ and $U \cup V$ define meets and joins on \mathcal{L}_Σ , turning it into a locale (because it is finite).

The family $\underline{I} = (I_\omega)_{\omega \in \Omega}$ is clearly an element of \mathcal{L}_Σ , because $\emptyset \neq I_\omega \subseteq I_\omega$ for all ω . If $U \in \mathcal{L}_\Sigma$, then $U_\omega \subseteq I_\omega$ for all ω and $\lambda_U \subseteq \Omega$, proving $U \leq \underline{I}$. Hence \underline{I} is the top element of \mathcal{L}_Σ . The empty family $\emptyset := (\emptyset)_{\omega \in \emptyset}$ is also an element of \mathcal{L}_Σ : it is vacuously true that $\forall \omega \in \emptyset$ we have $\emptyset \neq \emptyset \subseteq I_\omega$.

The covers on an object U are the subsets $\mathfrak{R} \subseteq \mathcal{L}_\Sigma$ such that $U = \bigvee \mathfrak{R}$. This last statement is equivalent to $\lambda_U = \bigcup_{V \in \mathfrak{R}} \lambda_V$ and $U_\omega = \bigcup_{V \in \mathfrak{R}} V_\omega$ for all $\omega \in \lambda_U$, with the convention that $V_\omega := \emptyset$ if $\omega \notin \lambda_V$. Because $\lambda_V := \{\omega \in \Omega \mid V_\omega \neq \emptyset\}$, the second statement actually implies the first, i.e. it is enough to ask for $U_\omega = \bigcup_{V \in \mathfrak{R}} V_\omega$. This concludes our proof. \square

Proof of Proposition 8

Proposition 8. *Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario. The restricted function $f|_V$ from Definition 7 is well-defined and the functor $\mathcal{E}_\Sigma : \mathcal{L}_\Sigma \rightarrow \text{Set}^{\text{op}}$ is a sheaf over the locale \mathcal{L}_Σ .*

Proof. The restricted function $f|_V$ is defined as follows by Definition 7:

$$f|_V(\underline{i}) := \left(f(\underline{j}) \right) \Big|_V \text{ for any } \underline{j} \in U \text{ such that } \forall \omega \in \lambda_V. j_\omega = i_\omega$$

This is well-defined because the causality requirement on f , together with the fact that $V \leq U$, makes the value of $f|_V(\underline{i})$ independent of the specific choice of \underline{j} : if $j_\omega = j'_\omega$ for all $\omega \in V$, then by causality we have that $f(\underline{j})_\omega = f(\underline{j}')_\omega$ for all $\omega \in V$.

For \mathcal{E}_Σ to be a sheaf, it is enough to show that given any two objects $U, V \in \mathcal{L}_\Sigma$ and any two functions $f \in \mathcal{E}_\Sigma(U)$ and $g \in \mathcal{E}_\Sigma(V)$, if we have that $f|_{U \cap V} = g|_{U \cap V}$, i.e. if the two functions are compatible, then there is a unique function $h \in \mathcal{E}_\Sigma(U \cup V)$ such that $h|_U = f$ and $h|_V = g$. For $\underline{i} \in U \cup V$ and $\lambda \subseteq \lambda_U \cup \lambda_V$, write $i|_\lambda := (i_{\omega'})_{\omega' \in \lambda}$. The requirements that $h|_U = f$ and $h|_V = g$ already provide a unique candidate for the gluing $h \in \mathcal{E}_\Sigma(U \cup V)$:

$$h(\underline{i})_\omega := \begin{cases} f(i|_{\lambda_U})_\omega & \text{if } \omega \in \lambda_U \text{ and } i|_{\lambda_U} \in U \\ g(i|_{\lambda_V})_\omega & \text{if } \omega \in \lambda_V \text{ and } i|_{\lambda_V} \in V \end{cases}$$

The only thing to establish is that the above yields a well-defined function $h \in \mathcal{E}_\Sigma(U \cup V)$. The only place where h can be ill-defined is on the overlap between the f clause and the g clause above: $\omega \in \lambda_U \cap \lambda_V$, $i|_{\lambda_U} \in U$ and $i|_{\lambda_V} \in V$. The latter two conditions, in particular, imply that $U_\omega \cap V_\omega \neq \emptyset$ for all $\omega \in \lambda_U \cap \lambda_V$, i.e. that $\lambda_{U \cap V} = \lambda_U \cap \lambda_V$. Under these conditions, the causality requirements for f and g together with their compatibility imply that h is indeed well-defined:

$$f(i|_{\lambda_U})_\omega = f|_{U \cap V}(i|_{\lambda_{U \cap V}})_\omega = g|_{U \cap V}(i|_{\lambda_{U \cap V}})_\omega = g(i|_{\lambda_V})_\omega$$

This concludes our proof. \square

Proof of Proposition 9

Proposition 9. *Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a non-locality scenario and write $X := \coprod_{\omega \in \Omega} I_\omega$. We have that $\mathcal{P}(X)$ and \mathcal{L}_Σ are isomorphic as locales. Furthermore, we have that the sheaf of sections \mathcal{E}_Σ in the original framework for non-locality—defined over the locale $\mathcal{P}(X)$ by Equations 1 and 2—is isomorphic to the sheaf of sections \mathcal{E}_Σ introduced in this work—defined over the locale \mathcal{L}_Σ by Definition 7.*

Proof. Firstly, we show that in the case of non-locality scenarios the locale \mathcal{L}_Σ is isomorphic to the locale $\mathcal{P}(X)$. More specifically, we will show that the map φ is a bijection. To see why, observe that for non-locality scenarios the poset of lowersets $\Lambda(\Omega)$ coincides with the poset of all subsets of Ω , i.e. the powerset $\mathcal{P}(\Omega)$. This means that \mathcal{L}_Σ can be written as $\coprod_{\lambda \subseteq \Omega} \prod_{\omega \in \lambda} (\mathcal{P}(I_\omega) \setminus \{\emptyset\})$. If $A \in \mathcal{P}(X)$, then we can set $\lambda_U := \{\omega \in \Omega \mid A \cap (I_\omega \times \{\omega\}) \neq \emptyset\}$ and $U_\omega := A \cap (I_\omega \times \{\omega\})$ to obtain a $U \in \mathcal{L}_\Sigma$ such that $\varphi(U) = A$, making φ surjective. The inverse φ^{-1} is also clearly order-preserving, so that φ is an isomorphism of locales. In our definitions, we have used the fact that X has elements in the form $(i, \omega) \in I_\omega \times \{\omega\}$ for all $\omega \in \Omega$.

We now need to define a natural isomorphism between the sets of functions associated to the objects by the original sheaf of sections (left) and the one from Definition 7 (right):

$$\prod_{(i, \omega) \in \varphi(U)} \mathcal{O}_\omega \cong \left\{ f : \prod_{\omega \in \lambda_U} U_\omega \rightarrow \prod_{\omega \in \lambda_U} \mathcal{O}_\omega \mid f \text{ causal for } \Omega \right\}$$

On the left hand side, note that $\varphi(U)$ is partitioned as $\varphi(U) = \coprod_{\omega \in \lambda_U} U_\omega$ across different sectors of the disjoint union $X := \coprod_{\omega \in \Omega} I_\omega$, with $U_\omega \neq \emptyset$ for all $\omega \in \lambda_U$. This yields the following equivalent expression:

$$\prod_{(i, \omega) \in \varphi(U)} \mathcal{O}_\omega \cong \prod_{\omega \in \lambda_U} (U_\omega \rightarrow \mathcal{O}_\omega)$$

On the right hand side, note that the condition of f being causal (for Ω discrete) becomes the statement that $f(i)_y$ is independent of i_x for all $x \neq y$, i.e $f(i)_y$ only depends on the value i_y . This means that f factors as $f = \prod_{\omega \in \lambda_U} s_\omega$ for a suitable family $s_\omega : U_\omega \rightarrow \mathcal{O}_\omega$, yielding our desired isomorphism:

$$\left\{ f : \prod_{\omega \in \lambda_U} U_\omega \rightarrow \prod_{\omega \in \lambda_U} \mathcal{O}_\omega \mid f \text{ causal for } \Omega \right\} \cong \prod_{\omega \in \lambda_U} (U_\omega \rightarrow \mathcal{O}_\omega)$$

To see that this isomorphism is in fact natural, let's look at function restriction for the original framework:

$$\left(s \in \prod_{(i, \omega) \in \varphi(U)} \mathcal{O}_\omega \right) \mapsto \left(s|_{\varphi(V)} \in \prod_{(i, \omega) \in \varphi(V)} \mathcal{O}_\omega \right)$$

We can re-arrange the above into the following equivalent formulation:

$$\left(s \in \prod_{\omega \in \lambda_U} (U_\omega \rightarrow \mathcal{O}_\omega) \right) \mapsto \left(\prod_{\omega \in \lambda_V} (s_\omega)|_{V_\omega} \in \prod_{\omega \in \lambda_V} (V_\omega \rightarrow \mathcal{O}_\omega) \right)$$

Let's compare the above to the function restriction from Definition 7:

$$\left(f : \prod_{\omega \in \lambda_U} U_\omega \rightarrow \prod_{\omega \in \lambda_U} \mathcal{O}_\omega \right) \mapsto \left(f|_V : \prod_{\omega \in \lambda_V} V_\omega \rightarrow \prod_{\omega \in \lambda_V} \mathcal{O}_\omega \right)$$

Because f factors as $f = \prod_{\omega \in \lambda_U} s_\omega$, the two restrictions are actually the same:

$$f|_V = \left(\prod_{\omega \in \lambda_U} s_\omega \right) \Big|_V = \prod_{\omega \in \lambda_V} (s_\omega)|_{V_\omega}$$

This concludes our proof. □

Proof of Proposition 15

Proposition 15. *Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario. The data defining an (R -valued) empirical model $e \in \prod_{\underline{i} \in \mathcal{I}_\Sigma} \mathcal{D}_R \mathcal{E}_\Sigma(\underline{i})$ is exactly the data defining a causal (R -valued) conditional distribution $\hat{e} : \mathcal{I}_\Sigma \rightarrow \mathcal{D}_R(\mathcal{O}_\Sigma)$:*

$$\hat{e}(\underline{o}|\underline{i}) := e_{\underline{i}}(\underline{i} \mapsto \underline{o})$$

Above, we denoted by $\underline{i} \mapsto \underline{o}$ a generic function in $\mathcal{E}_\Sigma(\{(i_\omega)\}_{\omega \in \Omega})$, with singleton domain $\{\underline{i}\}$.

Proof. The above certainly gives a well-defined (R -valued) conditional distribution $\hat{e} : \mathcal{I}_\Sigma \rightarrow \mathcal{D}_R(\mathcal{O}_\Sigma)$. The definition of causality for $\hat{e}(\underline{o}|\underline{i})$ reads:

$$\underline{i}|_\mu = \underline{j}|_\mu \Rightarrow \forall \underline{o} \in \prod_{\omega \in \mu} \mathcal{O}_\omega. \left[\hat{e}|_\mu(\underline{o}|\underline{i}) = \hat{e}|_\mu(\underline{o}|\underline{j}) \right]$$

The definition of the restriction $\hat{e}|_\mu$ reads:

$$\hat{e}|_\mu((o_\omega)_{\omega \in \mu} | \underline{i}) := \sum_{\substack{\underline{o}' \in \mathcal{O}_\Sigma \\ \text{s.t. } \underline{o}'|_\mu = \underline{o}|_\mu}} \hat{e}(\underline{o}'|\underline{i}) \quad \text{for all } \underline{i} \in \mathcal{I}_\Sigma \quad (45)$$

We rephrase this in terms of the original empirical model:

$$\underline{i}|_\mu = \underline{j}|_\mu \Rightarrow \forall \underline{o} \in \prod_{\omega \in \mu} \mathcal{O}_\omega. \left[\sum_{\substack{\underline{o}' \in \mathcal{O}_\Sigma \\ \text{s.t. } \underline{o}'|_\mu = \underline{o}|_\mu}} e_{\underline{i}}(\underline{i} \mapsto \underline{o}') = \sum_{\substack{\underline{o}' \in \mathcal{O}_\Sigma \\ \text{s.t. } \underline{o}'|_\mu = \underline{o}|_\mu}} e_{\underline{j}}(\underline{j} \mapsto \underline{o}') \right]$$

Because the function $\underline{i} \mapsto \underline{o}$ (resp. $\underline{j} \mapsto \underline{o}$) is defined on a singleton, asking for the $\underline{o}' \in \mathcal{O}_\Sigma$ such that $\underline{o}'|_\mu = \underline{o}|_\mu$ is the same as asking for the $f : \prod_{\omega \in \Omega} \{i_\omega\} \rightarrow \mathcal{O}_\Sigma$ (resp. $f : \prod_{\omega \in \Omega} \{j_\omega\} \rightarrow \mathcal{O}_\Sigma$) such that $f|_{(\{i_\omega\})_{\omega \in \mu}} = \underline{i} \mapsto \underline{o}|_{(\{i_\omega\})_{\omega \in \mu}} = \underline{i}|_\mu \mapsto \underline{o}|_\mu$. This makes the equation above the same as the definition of a compatible family for the global cover \mathfrak{F}_Σ :

$$\underline{i}|_\mu = \underline{j}|_\mu \Rightarrow \forall \underline{o} \in \prod_{\omega \in \mu} \mathcal{O}_\omega. \left[\sum_{\substack{f|_{(\{i_\omega\})_{\omega \in \mu}} \\ = (\underline{i}|_\mu \mapsto \underline{o}|_\mu)}} e_{\underline{i}}(f) = \sum_{\substack{f|_{(\{j_\omega\})_{\omega \in \mu}} \\ = (\underline{i}|_\mu \mapsto \underline{o}|_\mu)}} e_{\underline{j}}(f) \right]$$

This shows that the definition of causality for conditional distributions and the compatibility condition for empirical models are equivalent, so that causal conditional distributions correspond exactly to empirical models. \square

Proof of Proposition 16

Proposition 16. *Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario. The probabilistic empirical models (i.e. the \mathbb{R}^+ -valued ones) form a polytope, obtained by intersecting the product polytope from 34 with the*

\mathbb{R} -linear subspace defined by the following **causality equations**:

$$\left(\sum_{\substack{\underline{o} \in \mathcal{O}_\Sigma \text{ s.t.} \\ \underline{o}|_\lambda = \underline{o}'}} (d(\underline{o}|\underline{i}) - d(\underline{o}|\underline{j})) = 0 \right) \quad \text{for all} \quad \begin{cases} \lambda \in \Lambda(\Omega) \\ \underline{o}' \in \prod_{\omega \in \lambda} \mathcal{O}_\omega \\ \underline{i}, \underline{j} \in \mathcal{I}_\Sigma \text{ s.t. } \underline{i}|_\lambda = \underline{j}|_\lambda \end{cases}$$

Note that the unknown in the system of equations above is the distribution d , which can be seen as a vector in $\mathbb{R}^{\mathcal{O}_\Sigma \times \mathcal{I}_\Sigma}$ (e.g. from its tabular representation).

Proof. The requirement that $d(\underline{o}|\underline{i})$ be a \mathbb{R}^+ -valued distribution is exactly the requirement that $d(\underline{o}|\underline{i})$ be a point in a simplex for each fixed \underline{i} :

$$d \in \Delta^{\mathcal{O}_\Sigma} \subset \mathbb{R}^{\mathcal{O}_\Sigma}$$

Putting the values of $d(\underline{o}|\underline{i})$ for all $\underline{i} \in \mathcal{I}_\Sigma$ together into a single vector results in the requirement that d be a point in a product of simplexes (a polytope):

$$d \in \left(\Delta^{\mathcal{O}_\Sigma} \right)^{\times \mathcal{I}_\Sigma} \subset \mathbb{R}^{\mathcal{O}_\Sigma \times \mathcal{I}_\Sigma}$$

The causality equations are simply the causality constraints for d turned into linear equations:

$$\sum_{\substack{\underline{o} \in \mathcal{O}_\Sigma \text{ s.t.} \\ \underline{o}|_\lambda = \underline{o}'}} (d(\underline{o}|\underline{i}) - d(\underline{o}|\underline{j})) = 0 \quad \iff \quad \sum_{\substack{\underline{o} \in \mathcal{O}_\Sigma \text{ s.t.} \\ \underline{o}|_\lambda = \underline{o}'}} d(\underline{o}|\underline{i}) = \sum_{\substack{\underline{o} \in \mathcal{O}_\Sigma \text{ s.t.} \\ \underline{o}|_\lambda = \underline{o}'}} d(\underline{o}|\underline{j})$$

Imposing these equations is the same as intersecting the polytope $(\Delta^{\mathcal{O}_\Sigma})^{\times \mathcal{I}_\Sigma}$ with the corresponding linear subspaces, yielding another polytope (because there are finitely many equations). \square

Proof of Proposition 21

Proposition 21. *Let $\Sigma = (\Omega, \underline{I}, \underline{O})$ be a definite causal scenario and let R be a commutative semiring. Let $e \in \prod_{\underline{i} \in \mathcal{I}_\Sigma} \mathcal{D}_R \mathcal{E}_\Sigma(\underline{i})$ be an R -valued empirical model. The empirical model e is local if and only if there is a family $(f^{(\xi)})_{\xi \in \Xi}$ of causal functions $f^{(\xi)} : \mathcal{I}_\Sigma \rightarrow \mathcal{O}_\Sigma$ and an R -valued distribution $p \in \mathcal{D}_R(\Xi)$ such that:*

$$\hat{e} = \sum_{\xi \in \Xi} p(\xi) \delta_{f^{(\xi)}}$$

where $\hat{e} : \mathcal{I}_\Sigma \rightarrow \mathcal{D}_R(\mathcal{O}_\Sigma)$ is the causal conditional distribution corresponding to e .

Proof. The empirical model is local if and only if there is a global section $g \in \mathcal{D}_R \mathcal{E}_\Sigma(\underline{I})$ such that $g|_{(\{i_\omega\})_{\omega \in \Omega}} = e_{\underline{i}}$ for all $\underline{i} \in \mathcal{I}_\Sigma$. This immediately implies a decomposition for \hat{e} in the required form:

$$\hat{e} = \sum_{f \in \mathcal{E}_\Sigma(\underline{I})} g(f) \delta_f$$

Conversely, if $\hat{e} = \sum_{\xi \in \Xi} p(\xi) \delta_{f^{(\xi)}}$ for some $p \in \mathcal{D}_R(\Xi)$ and some family $(f^{(\xi)})_{\xi \in \Xi} \in (\mathcal{E}_\Sigma(\underline{I}))^\Xi$ then we can define our desired global section $g \in \mathcal{D}_R \mathcal{E}_\Sigma(\underline{I})$ as follows:

$$g(f) := \sum \left\{ p(\xi) \mid \xi \in \Xi \text{ s.t. } f^{(\xi)} = f \right\}$$

This completes our proof. \square

Proof of Proposition 22

Proposition 22. *Let $\Theta = (\Gamma, \underline{I}, \underline{O})$ be a definite causal scenario as defined in [64], where Γ has no boundary nodes. Let Ω be the definite causal order underlying the acyclic framed multigraph Γ , so that $\Sigma = (\Omega, \underline{I}, \underline{O})$ is a definite causal scenario as defined in this work. If Δ is a diagram over Θ in a R -probabilistic theory C , as defined in [64], and Δ involves normalised processes at all events, then Δ yields an R -valued empirical model e for Σ . We refer to such a Δ as a **realisation** of e in C .*

Proof. The discarding map τ_{O_ω} for the classical output system O_ω at any event ω can be written as $\tau_{O_\omega} := \sum_{o \in O_\omega} \langle o |$, so that diagrammatic discarding of classical systems as defined in [64] is exactly the same as marginalisation of distributions as defined in this work. Because the framed multigraph Γ has no boundary nodes, any diagram Δ over Θ in any R -probabilistic theory C only has classical inputs and outputs, so it defines a conditional distribution $d(\underline{o} | \underline{i})$ (normalisation of d follows from normalisation of the processes in Δ when discarding all classical outputs). Now consider discarding the classical output wire for the process at an event $\omega \in \Omega$ which is maximal in the causal order. Because ω is maximal, the process at ω has no non-classical output. Because of normalisation, then, we obtain a diagram on $\Omega \setminus \{\omega\}$, disconnected from the classical input wire at ω . We can repeat this process to obtain restrictions of the original diagram to any lower set $\lambda \in \Lambda(\Omega)$, independently of the classical inputs at $\omega \notin \lambda$. This means that the conditional distribution $d(\underline{o} | \underline{i})$ is causal, i.e. that it corresponds to an R -valued empirical model. \square