

Causal Order as Representation Data: Process Matrices, Causal Inequalities, and Locally Finite Posets as Reconstructions of Causality

Matthew Long

The YonedaAI Collaboration, YonedaAI Research Collective

Chicago, IL

matthew@yonedaai.com · <https://yonedaai.com>

6 July 2026

Abstract

We treat causal order not as a fixed backdrop for physics but as a piece of representation data to be reconstructed from operational or combinatorial invariants, in the same spirit in which a metric is reconstructed from a spectrum or a bulk from a boundary. Two formalisms carry the load. The process-matrix formalism of Oreshkov, Costa, and Brukner encodes every correlation compatible with local quantum mechanics in a single positive operator W that presupposes no global order between the local laboratories; a causal inequality is a linear bound satisfied by every process built from a definite (possibly classically randomized) order, and its violation certifies that no such order exists. The causal-set proposal of Bombelli, Lee, Meyer, and Sorkin encodes spacetime as a locally finite partial order $(C, \preceq)$ from which conformal geometry is read off directly and volume is recovered by counting. We give a self-contained account of both, state the central facts as labeled propositions carrying the project's S/H/P epistemic tags, and work three examples in full: the quantum switch and its causal nonseparability, the bipartite Oreshkov–Costa–Brukner causal inequality with its bound $p \leq 3/4$ and quantum value $(2 + \sqrt{2})/4$, and a small sprinkled causal diamond whose order relation faithfully reproduces the continuum light-cone structure. That process matrices form a well-defined convex set, that the causal bound holds for all definite-order strategies, that the switch is causally nonseparable, and that order determines conformal structure (Malament) are standard mathematics ([S]). That laboratory indefinite order is a model of quantum-gravitational causal indefiniteness, and that causal-set order plus counting recovers a Lorentzian metric, are heuristic ([H]). The claim that gravity actually produces causally nonseparable processes, and the proposed identification of the two formalisms under a single causal-order invariant, are speculative ([P]). We flag throughout that this topic — indefinite causal order — is a *proposed* extension (QG-X) of the library's roadmap, with no prior dictionary entry for process matrices; we propose

one, at status [H/P]. A companion Haskell development checks the poset axioms, the causal bound, and the switch’s order dependence.

Contents

1	Introduction	2
1.1	The reconstruction thesis for causal order	2
1.2	An honesty note on status	3
1.3	What is and is not claimed	3
1.4	Organization	3
2	Mathematical framework: reconstruction from representation data	4
2.1	The representation stack and the S/H/P calculus	4
2.2	The dictionary rows that bear on this paper, and the one we propose	4
2.3	The proposed module QG-X	5
2.4	What “reconstruction” means for order specifically	5
3	The process-matrix formalism	5
3.1	Local laboratories and local operations	5
3.2	The Choi–Jamiołkowski representation	6
3.3	The process matrix	6
4	Causal, causally separable, and causally nonseparable processes	7
4.1	Definite order and its convex closure	7
4.2	Causal inequalities	8
4.3	Causal witnesses	9
5	The quantum switch	10
5.1	Construction	10
5.2	Why no fixed circuit reproduces it	10
5.3	The switch as a process matrix, and its nonseparability	11
6	Causal sets as discrete causal order	12
6.1	The poset axioms	12
6.2	Order plus number equals geometry	13
6.3	Dynamics: sequential growth	14
7	Order as representation data	14
8	Worked examples	15
8.1	Example 1: the quantum switch on Pauli operations	15
8.2	Example 2: the OCB causal inequality, bound and violation	16
8.3	Example 3: a sprinkled causal diamond	16

9	Relation to the other regimes and the S/H/P calculus	18
9.1	Lorentzian geometry (Topic 1)	18
9.2	Noncommutative geometry (Topic 11)	18
9.3	Spin foams and histories (Topic 4)	18
9.4	The weakest-link summary	18
10	A Haskell model of causal order	19
11	Limitations and open problems	20
11.1	The gravitational reading is speculative	20
11.2	No dynamics for indefinite order	20
11.3	The synthesis of the two halves is unproven	20
11.4	Interpretational status of “indefinite”	21
11.5	The Hauptvermutung is open	21
11.6	Roadmap and dictionary status	21
12	Discussion	21
13	Conclusion	22
A	The core Haskell module	22

1 Introduction

The governing perspective of this library is that quantum gravity is not primarily the quantization of material objects moving in spacetime, but the reconstruction of spacetime, matter, causality, and measurement from invariant representation structures. Most of the twelve regimes in the library reconstruct *geometry*: a metric from a spectrum, a bulk from a boundary code, a semiclassical three-slice from spin-network representation labels. This paper takes up the one item on the governing perspective’s own short list that is not geometry at all. It is *causality*: the relation “before and after” between physical events, ordinarily assumed as a precondition for even writing down a quantum operation, here treated as reconstructed — and, at the quantum level, possibly indefinite — structure.

Two independent research programs supply the mathematics. The first is the causal-set program of Bombelli, Lee, Meyer, and Sorkin [1]. Its thesis is that spacetime is fundamentally a *locally finite partially ordered set* $(C, \preceq)$: the order is causal precedence, and local finiteness — every causal interval $[p, q] = \{r : p \preceq r \preceq q\}$ is a finite set — makes the substrate discrete without a separately imposed lattice scale. The program’s slogan, “order plus number equals geometry,” is a literal reconstruction claim: the order recovers the conformal (light-cone) structure of a Lorentzian manifold, and counting elements recovers the volume element, and the two together fix the metric up to nothing. The second is the process-matrix program of Oreshkov, Costa, and Brukner [7]. Its thesis is more radical: even the order itself can fail to be definite. A process matrix W is an operator, formally analogous to a density matrix but for a background causal structure rather than a state, that encodes all the statistics available to a collection of local laboratories — each obeying ordinary quantum mechanics internally

— *without* presupposing any global order among them. Some process matrices are causally nonseparable: no definite order between the parties, not even a classically randomized one, reproduces their statistics.

1.1 The reconstruction thesis for causal order

We state the thesis this paper defends. In the orthodox account, one fixes a causal structure — a Lorentzian manifold, a circuit diagram, a tensor-product factorization into “systems” with a time axis — and *then* places quantum operations on it. The causal structure is input. In the reconstruction account, the causal structure is output. One is handed operational data (a family of local instruments and the joint probabilities they generate) or combinatorial data (a finite set with a partial order), and one reads the causal structure — including the answer to the question “is there a definite order at all?” — off that data as an invariant.

The word *invariant* is doing real work. In the causal-set case the invariant is the isomorphism class of the poset $(C, \preceq)$: two labelings of the same order are the same causal set, and the physically meaningful object is the order relation, not any embedding of it into a manifold. In the process-matrix case the invariant is the operator W together with the convex-geometric question of whether it lies in the subset of causally separable processes. Neither invariant presupposes what it reconstructs. The causal set does not presuppose a manifold; the manifold, when it exists, is recovered as an approximation. The process matrix does not presuppose an order; the order, when it exists, is recovered as a convex decomposition.

1.2 An honesty note on status

This is the twelfth and epistemically weakest of the library’s topics, and we say so at the outset rather than burying it. Three facts about its standing should be stated plainly.

First, the library’s declared roadmap names eight modules, QG-I through QG-VIII. Indefinite causal order is *not among them*. Following the knowledge base, we present it as a *proposed* ninth-and-tenth-style extension, “QG-X: causal-order representation,” and we mark it as proposed every time it appears. It is a suggestion for where this material would sit, not an adopted part of the project.

Second, the library’s dictionary of twenty-six mathematics-to-physics entries contains a row for *causal set / locally finite poset* at status [H], but *no row at all* for process matrices, the quantum switch, or causal inequalities. We propose adding one, and we propose its status: [H/P]. The *H* covers the mathematically well-defined process-matrix formalism and the finite linear algebra of causal inequalities; the *P* covers the ontological claim that quantum gravity actually realizes indefinite order, which remains an extrapolation.

Third, the sharpest results in this paper are honest [S] mathematics — the convex structure of process matrices, the causal bound for definite-order strategies, the causal nonseparability of the switch, Malament’s theorem that order determines conformal structure — but the *physical* readings that motivate the topic are [H] at best and, where they concern gravity, [P]. We are careful to attach the tag to the claim, not to the section.

1.3 What is and is not claimed

We do *not* claim that the process-matrix formalism proves that quantum gravity produces indefinite causal order. The mathematical consistency of causally nonseparable processes is one thing; their gravitational realizability is another, and the laboratory quantum switch is realized on an ordinary fixed background spacetime [15, 16]. The gravitational reading rests on the gedanken-arguments of Hardy [13] and of Zych, Costa, Pikovski, and Brukner [14], which we present as motivation, labeled [P], not as demonstration. We also do not claim that causal sets and process matrices are two descriptions of the same thing; we propose, as an organizing conjecture at status [P], that they are two reconstructions of a single “causal-order” invariant, and we are explicit that the literature does not establish this.

1.4 Organization

Section 2 sets up the representation-stack language and the S/H/P calculus, and records the proposed QG-X dictionary entry. Section 3 develops the process-matrix formalism: local laboratories, the Choi–Jamiołkowski representation, and the characterization of valid W by positivity and a trace-projection constraint. Section 4 defines causal, causally separable, and causally nonseparable processes and proves the bipartite causal bound. Section 5 constructs the quantum switch and proves its causal nonseparability. Section 6 develops causal sets as discrete causal order, states Malament’s theorem and the sprinkling reconstruction, and records the Hauptvermutung. Section 7 collects the paper’s viewpoint as labeled results on order as representation data. Section 8 works the three central examples in full. Section 9 places the topic against the other regimes and runs the weakest-link calculus. Section 10 describes the companion Haskell model. Section 11 is an unusually long limitations section, as befits the weakest topic in the library. Section 12 and Section 13 close.

2 Mathematical framework: reconstruction from representation data

2.1 The representation stack and the S/H/P calculus

We reuse the library’s common scaffolding. A representation entry is a tuple $E = (M, P, \tau, \sigma)$ of a mathematical object M , a physical target P , a translation datum τ , and an epistemic status σ . The statuses form an ordered set of three warrants,

$$S < H < P, \tag{1}$$

read: S = standard mathematical or mathematical-physics correspondence; H = strong heuristic physical representation; P = speculative ontological extension. A composite status is a nonempty tuple rendered [S], [H], [P], or a composite such as [S/H] or [H/P]. The single structural fact we use about statuses is that composition is monotone and worst-component-wins:

$$\sigma(\bigcirc_i E_i) = \max_i \sigma(E_i), \tag{2}$$

where max is taken in the order (1). Chaining regimes can never raise a composite above its weakest link. This is the precise, checkable sense in which the library is modular rather than unified, and it is what keeps the honest-status discipline of this paper from being mere rhetoric: a claim that touches the gravitational reading of indefinite order inherits a P no matter how solid its other ingredients are.

2.2 The dictionary rows that bear on this paper, and the one we propose

Of the twenty-six dictionary rows, exactly one bears directly on this topic:

Status	Math	Physical representation	Translation
[H]	Causal set / locally finite poset	causal-discrete spacetime substrate candidate	causal-order representation

No row exists for the modern content of the topic — process matrices, the quantum switch, causal inequalities. This is a genuine gap in the library, and we propose to close it with the following row, flagged as *proposed*, not adopted:

Status	Math	Physical representation	Translation
[H/P]	Process matrix / indefinite causal order	operationally certifiable order (in)definiteness	causal-order-as-operator representation

The status is composite. The H is warranted: the process-matrix formalism is mathematically well posed, and causal-inequality reasoning is finite convex geometry (both established below at [S] internally, carried at [H] as a *physical* representation because the claim that this is the right physics of causal structure remains a modeling judgment). The P is the extrapolation that quantum gravity realizes such processes. Under (2) the row composes as [H/P].

2.3 The proposed module QG-X

The knowledge base proposes, and we adopt as the paper’s framing, a roadmap extension:

QG-X (proposed) — Causal-order representation. Mathematics: causal sets / locally finite posets; process matrices; causal inequalities. Physics: causal-discrete spacetime substrate candidates; operationally certifiable indefinite causal order. Core invariant: the order relation $(C, \preceq)$ (causal sets) or the process matrix W (quantum causal structures).

We stress once more: QG-X is not in the declared roadmap. It is a natural place to put this material, proposed here.

2.4 What “reconstruction” means for order specifically

For geometry the reconstruction schema is a map algebraic/order data $\rightarrow$ geometry. For order the schema is subtler because there are two levels. At the first level, order data reconstruct *geometry*: this is the causal-set claim, order plus number gives a Lorentzian metric. At the second level, operational data reconstruct *order itself*, including the possibility that there is no definite order to reconstruct: this is the process-matrix claim. The two levels are not the same reconstruction, and a recurring discipline of this paper is to keep them apart. Causal sets assume a definite (if discrete) order and reconstruct geometry from it. Process matrices do not assume a definite order and reconstruct the order — or its absence — from statistics.

3 The process-matrix formalism

3.1 Local laboratories and local operations

Fix a finite collection of parties, each operating a closed laboratory. Party A has an input Hilbert space $\mathcal{H}^{A_I}$ (a system enters) and an output Hilbert space $\mathcal{H}^{A_O}$ (a system leaves); write $d_{A_I} = \dim \mathcal{H}^{A_I}$ and so on. Inside the laboratory, ordinary quantum mechanics holds: the party implements a *quantum instrument*, a collection $\{\mathcal{M}_a^A\}_a$ of completely positive (CP) maps $\mathcal{M}_a^A : \mathcal{B}(\mathcal{H}^{A_I}) \rightarrow \mathcal{B}(\mathcal{H}^{A_O})$ summing to a completely positive trace-preserving (CPTP) map $\sum_a \mathcal{M}_a^A$. The index a is the classical outcome. Everything the party can do — prepare, measure, transform, correlate with a memory — is a choice of instrument.

The formalism makes one assumption and refuses another. It assumes local quantum mechanics: each laboratory is an ordinary quantum system with a well-defined in/out structure. It refuses to assume a global order: nothing is said, a priori, about whether A ’s output feeds B ’s input, or the reverse, or neither, or a superposition.

3.2 The Choi–Jamiołkowski representation

The technical device that makes “no global order” expressible is the Choi–Jamiołkowski (CJ) isomorphism, which turns a map into an operator. For a CP map $\mathcal{M} : \mathcal{B}(\mathcal{H}^I) \rightarrow \mathcal{B}(\mathcal{H}^O)$, its (standard, Choi) operator is

$$M = (\mathcal{I} \otimes \mathcal{M}) |\phi^+\rangle\langle\phi^+|, \quad |\phi^+\rangle = \sum_{j=1}^{d_I} |jj\rangle \in \mathcal{H}^I \otimes \mathcal{H}^I, \quad (3)$$

an operator on $\mathcal{H}^I \otimes \mathcal{H}^O$ (with the non-normalized $|\phi^+\rangle$). CP corresponds to $M \geq 0$; trace preservation corresponds to $\text{Tr}_O M = \mathbb{1}_I$. An instrument $\{\mathcal{M}_a\}$ becomes a set of positive operators $\{M_a\}$ with $\sum_a M_a$ satisfying the trace-preservation constraint. Conventions differ by at most a partial transpose on the input leg: the process-matrix literature [7] places that partial transpose so that the generalized Born rule (5) is a plain trace pairing of W with the local CJ operators; nothing below depends on the choice beyond this bookkeeping, and we keep the standard Choi form for definiteness.

3.3 The process matrix

The central object encodes the “background” against which all local operations are composed.

Definition 3.1 (Process matrix, bipartite case, [S]). Let parties A, B have spaces $\mathcal{H}^{A_I}, \mathcal{H}^{A_O}, \mathcal{H}^{B_I}, \mathcal{H}^{B_O}$. A *process matrix* is an operator

$$W \in \mathcal{B}(\mathcal{H}^{A_I} \otimes \mathcal{H}^{A_O} \otimes \mathcal{H}^{B_I} \otimes \mathcal{H}^{B_O}) \quad (4)$$

such that, for every pair of local instruments $\{M_a^A\}, \{M_b^B\}$, the numbers

$$P(a, b \mid x, y) = \text{Tr}[W (M_{a|x}^A \otimes M_{b|y}^B)] \quad (5)$$

form a valid probability distribution (nonnegative, normalized) for all settings x, y .

Here $M_{a|x}^A$ is the CJ operator of the outcome- a branch of the instrument chosen by setting x . Formula (5) is the generalized Born rule for an indefinite background: it looks exactly like $\text{Tr}[\rho E]$, but W plays the role of the state *and* the causal glue, and the local operators carry both input and output legs.

The requirement in theorem 3.1 is an infinite family of constraints (one for every instrument). It is a standard result that they collapse to a finite, checkable characterization.

Proposition 3.2 (Validity characterization, [S]). *For a subsystem label x and an operator X on the total space, write the trace-and-replace map*

$${}_x X := \frac{1}{d_x} \mathbb{1}_x \otimes \text{Tr}_x X. \quad (6)$$

An operator W on the space (4) is a bipartite process matrix if and only if

$$W \geq 0, \quad \text{Tr } W = d_{A_O} d_{B_O}, \quad (7)$$

$${}_{B_I B_O} W = {}_{A_O B_I B_O} W, \quad {}_{A_I A_O} W = {}_{A_I A_O B_O} W, \quad (8)$$

$$W = {}_{A_O} W + {}_{B_O} W - {}_{A_O B_O} W. \quad (9)$$

Proof sketch. Positivity of (5) for all product instruments forces $W \geq 0$ (choose rank-one CJ operators and vary). Normalization for all CPTP local maps forces the affine constraints: expanding W in a Hilbert–Schmidt basis of Pauli-type operators on each leg, the requirement that $\text{Tr}[W(M^A \otimes M^B)] = 1$ whenever M^A, M^B are trace-preserving is equivalent to the vanishing of every basis coefficient of “forbidden” type — terms that would let a party signal to its own past or create a closed signalling loop. Collecting the surviving terms is exactly the projection onto the subspace $\mathcal{L}_V$ cut out by (8)–(9); the trace value $d_{A_O} d_{B_O}$ is fixed by normalization on the maximally mixed instruments. The converse is a direct check that any W obeying (7)–(9) yields valid probabilities. Full details are in Oreshkov–Costa–Brukner [7] and in the operational reformulation of Araújo et al. [10]. $\square$

The content of (8) is “no party signals into its own past”: ${}_{B_I B_O} W = {}_{A_O B_I B_O} W$ says that tracing out Bob’s lab leaves something independent of Alice’s *output*, i.e. Alice cannot signal to a Bob who has already been fully traced away. Condition (9) is the affine “no closed loop” constraint, forbidding the self-referential terms that would make (5) ill defined. The set of valid W is a convex subset of a real affine subspace; it is a compact convex body, and its extreme points are the vertices we care about below.

Remark 3.3 (Why this is not just a state, [S]). A density matrix ρ is the special case in which the parties are ordered and non-interacting: $W = \rho^{A_I B_I} \otimes \mathbb{1}^{A_O B_O}$ describes two parties who each receive a share of a state and whose outputs feed nothing. A quantum channel from A to B is the case $W = \rho^{A_I} \otimes C^{A_O B_I} \otimes \mathbb{1}^{B_O}$ with C the CJ operator of the channel: here A precedes B . The process matrix is strictly larger than either, because (8)–(9) admit solutions that are *neither* of these forms nor any mixture of them. Those solutions are the subject of the next section.

4 Causal, causally separable, and causally nonseparable processes

4.1 Definite order and its convex closure

Definition 4.1 (Causally ordered and causally separable processes, [S]). A bipartite process is *ordered* $A \prec B$ if it has the channel form $W^{A \prec B} = \rho^{A_I} \otimes C^{A_O B_I} \otimes \mathbb{1}^{B_O}$ (Alice’s output may influence Bob’s input; not the reverse), and *ordered* $B \prec A$ symmetrically. A process is *causally separable* if it is a convex combination

$$W^{\text{sep}} = q W^{A \prec B} + (1 - q) W^{B \prec A}, \quad q \in [0, 1], \quad (10)$$

of an $A \prec B$ process and a $B \prec A$ process. A process is *causally nonseparable* if it admits no such decomposition.

Causal separability is exactly “there is a definite order, possibly chosen by a classical coin.” The coin models a shared classical record that fixes, run by run, who is before whom; the party never learns anything a definite order would forbid. Nonseparability is the failure of even this randomized-order account.

Remark 4.2 (Three, not two, levels, [S]). For more than two parties the picture refines: Oreshkov and Giarmatzi [12] distinguish causally separable from the strictly stronger *extensibly* causally separable processes, where the order decomposition must survive the addition of ancillary correlated parties. The switch of section 5 is causally nonseparable already at the bipartite-with-control level; we do not need the multipartite refinement for the results here, but we flag it because it is where the convex geometry becomes genuinely subtle.

4.2 Causal inequalities

A causal inequality is a device-independent certificate: a linear inequality on the observable probabilities $P(a, b \mid x, y)$ that holds for every causally separable process and can be violated by some valid W .

We use the bipartite game of [7]. Alice receives a uniformly random bit $x \in \{0, 1\}$ and outputs a bit a ; she does *not* receive the task selector b' , so her strategy cannot depend on which of the two guessing tasks is in force. Bob receives a uniformly random bit $y \in \{0, 1\}$ and, additionally, a uniformly random bit $b' \in \{0, 1\}$ that selects his task; he outputs a bit b . The task is asymmetric in b' :

- if $b' = 0$, Bob must output $b = x$ (guess Alice's bit);
- if $b' = 1$, Alice must output $a = y$ (guess Bob's bit).

The figure of merit is the average success probability

$$p = \frac{1}{2} \left[P(b = x \mid b' = 0) + P(a = y \mid b' = 1) \right]. \quad (11)$$

Proposition 4.3 (Causal bound, [S]). *Every causally separable process satisfies $p \leq 3/4$ in the game (11).*

Proof. Because p is an affine functional of the process (it is a fixed linear combination of Born-rule probabilities, each linear in W), and the causal separable set is the convex hull of the two ordered sets, it suffices to bound p on a definite order, say $A \prec B$. In an $A \prec B$ process Alice's output a is generated with no access to Bob's data; in CJ terms $W^{A \prec B}$ has trivial B_O dependence and Alice's marginal is ρ^{A_I} , so $P(a \mid x, y, b') = P(a \mid x)$. Hence when $b' = 1$ Alice must guess a uniformly random bit y she cannot see, and $P(a = y \mid b' = 1) \leq \frac{1}{2}$. Bob, on the other hand, may receive Alice's full data through the channel $C^{A_O B_I}$, so $P(b = x \mid b' = 0) \leq 1$, saturated by Alice broadcasting $a = x$ and Bob copying. Thus $p \leq \frac{1}{2}(1 + \frac{1}{2}) = \frac{3}{4}$. The order $B \prec A$ gives the same bound by symmetry. Any convex mixture of the two inherits the bound because p is affine. $\square$

Proposition 4.4 (Quantum violation, [S]). *There is a valid bipartite process matrix W_{OCB} , together with local qubit instruments, achieving*

$$p = \frac{2 + \sqrt{2}}{4} \approx 0.8536 > \frac{3}{4}. \quad (12)$$

Proof sketch. Take all four local spaces to be qubits. The Oreshkov–Costa–Brukner process [7] is

$$W_{\text{OCB}} = \frac{1}{4} \left[\mathbb{1} + \frac{1}{\sqrt{2}} \left(\underbrace{\sigma_z^{A_O} \sigma_z^{B_I}}_{T_1} + \underbrace{\sigma_z^{A_I} \sigma_x^{B_I} \sigma_z^{B_O}}_{T_2} \right) \right], \quad (13)$$

with identities on the unwritten legs and σ the Pauli matrices. Its two off-identity terms are of the two allowed signalling types, T_1 an “ A signals to B ” term and T_2 a “ B signals to A ” term, one for each order, which is the hallmark of causal nonseparability. Positivity is the delicate point and is worth spelling out. The two Pauli strings T_1, T_2 overlap on the single leg B_I , where T_1 carries σ_z and T_2 carries σ_x ; since σ_z and σ_x anticommute and the strings are otherwise disjoint, T_1 and T_2 *anticommute*. Hence $S := \frac{1}{\sqrt{2}}(T_1 + T_2)$ obeys $S^2 = \frac{1}{2}(T_1^2 + T_2^2 + \{T_1, T_2\}) = \frac{1}{2}(\mathbb{1} + \mathbb{1} + 0) = \mathbb{1}$, so S has eigenvalues ± 1 and $W_{\text{OCB}} = \frac{1}{4}(\mathbb{1} + S)$ has eigenvalues in $\{\frac{1}{2}, 0\}$, confirming $W_{\text{OCB}} \geq 0$. (Had both terms carried σ_z on B_I they would *commute*, S would have an eigenvalue $-\sqrt{2}$, and $\frac{1}{4}(1 - \sqrt{2}) < 0$ would appear in the spectrum, so the anticommuting placement of $\sigma_x^{B_I}$ is what keeps W_{OCB} a genuine process matrix.) The trace and validity constraints (7)–(9) hold by inspection. Choosing the local instruments of [7] — Alice measures σ_z or re-prepares, Bob likewise conditioned on b' — yields the two conditional success probabilities $\frac{1}{2}(1 + \frac{1}{\sqrt{2}})$ each, whose average is (12), matching [7]. The classical bound half of this computation is reproduced exactly in the companion code (section 10). $\square$

Corollary 4.5 (Certified indefiniteness, [S/H]). *Any process achieving $p > 3/4$ in the game (11) is causally nonseparable: no definite or classically randomized order reproduces its statistics. The mathematics is [S]; reading the violation as evidence that a physical causal order is genuinely absent (rather than merely hidden) is [H].*

The tag split in theorem 4.5 is the crux of the topic’s honesty problem. That the OCB process violates the bound is a matrix fact. That the violation means there is *no* causal order — as opposed to a causal order we have failed to model — is an interpretive commitment, exactly as the violation of a Bell inequality is a matrix fact whose reading as “no local hidden variables” is a physical stance. We carry the interpretive step at [H].

4.3 Causal witnesses

Where causal inequalities are device-independent but hard to violate with physically motivated processes, a *causal witness* is a device-dependent but flexible certificate. Since the causally separable processes form a closed convex set $\mathcal{S}$, any nonseparable $W \notin \mathcal{S}$ can be separated from $\mathcal{S}$ by a hyperplane: there is a Hermitian S with $\text{Tr}[S W'] \geq 0$ for all separable W' and $\text{Tr}[S W] < 0$.

Proposition 4.6 (Existence of a causal witness, [S]). *For every causally nonseparable process W there is a Hermitian operator S (a causal witness) such that $\text{Tr}[S W] < 0 \leq \text{Tr}[S W']$ for all causally separable W' . Membership of W in the separable set, and the optimal witness, are computable by semidefinite programming.*

Proof sketch. The separable set is a spectrahedron: it is the projection of a set defined by positive-semidefinite and affine constraints (positivity of the two ordered components and the mixing equation (10)). Convexity plus closedness gives the separating hyperplane by Hahn–Banach in finite dimension; the constructive form is the dual of the semidefinite feasibility program deciding $W \in \mathcal{S}$. This is the witness construction of Araújo et al. [10] and underlies the experimental certifications of [16]. $\square$

This is where the topic touches the executable spirit of the library’s formalization module: deciding causal separability is a finite convex program, and the causal bound of theorem 4.3 is a finite enumeration over deterministic strategies. Both are the kind of “small, checkable” statement that the roadmap’s formalization item asks for, and both appear in the companion code.

5 The quantum switch

5.1 Construction

The quantum switch of Chiribella, D’Ariano, Perinotti, and Valiron [8] is the cleanest physically motivated causally nonseparable process. Two black-box operations, with CJ unitaries A and B acting on a target system $\mathcal{H}_t$, are to be composed — but a control qubit $\mathcal{H}_c$ decides the order coherently. Define the switch unitary on $\mathcal{H}_c \otimes \mathcal{H}_t$ by its action conditioned on the control:

$$S_{\text{sw}}(A, B) = |0\rangle\langle 0|_c \otimes (B A) + |1\rangle\langle 1|_c \otimes (A B). \quad (14)$$

On control $|0\rangle$ the target sees A then B ; on control $|1\rangle$ it sees B then A . On a control superposition $\alpha|0\rangle + \beta|1\rangle$ the two orders are coherently superposed. Crucially, each of A and B is applied *exactly once*: the switch is a higher-order operation, a map on operations, not a circuit with A and B wired in a fixed place.

5.2 Why no fixed circuit reproduces it

The defining feature is that no circuit inserting A and B each once in a fixed order produces S_{sw} . A fixed order gives either BA or AB on the target for *every* control value; the switch gives BA on one control branch and AB on the other, coherently. When A and B do not commute the two branches differ, and the coherence between them — detectable by measuring the control in the $|\pm\rangle$ basis — has no fixed-order origin.

Proposition 5.1 (Order dependence of the switch, [S]). *Let A, B be unitaries with $[A, B] \neq 0$. Then there is a target state $|\psi\rangle$ for which $S_{\text{sw}}(A, B)(|+\rangle_c \otimes |\psi\rangle_t)$ is not of the spectator form $|+\rangle_c \otimes U|\psi\rangle_t$ for any fixed order $U \in \{AB, BA\}$; equivalently, no fixed-order circuit that leaves the control an untouched spectator reproduces the switch’s action on all inputs. In the generic case the switch output is control–target entangled; in special cases (such as anticommuting A, B) it factors, but only at the cost of an observable rotation of the control away from $|+\rangle$, which a fixed order cannot produce.*

Proof. With $|+\rangle_c = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$,

$$S_{\text{sw}}(A, B)(|+\rangle_c \otimes |\psi\rangle) = \frac{1}{\sqrt{2}}(|0\rangle_c \otimes BA|\psi\rangle + |1\rangle_c \otimes AB|\psi\rangle). \quad (15)$$

A fixed order with the control a spectator produces $|+\rangle_c \otimes U|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle_c + |1\rangle_c) \otimes U|\psi\rangle$. Matching this to (15) term by term requires $BA|\psi\rangle = U|\psi\rangle$ and $AB|\psi\rangle = U|\psi\rangle$, hence $BA|\psi\rangle = AB|\psi\rangle$, i.e. $[A, B]|\psi\rangle = 0$. If $[A, B] \neq 0$, choose $|\psi\rangle$ with $[A, B]|\psi\rangle \neq 0$; then $BA|\psi\rangle \neq AB|\psi\rangle$ and no such spectator U exists. Two sub-cases realize the failure. If $BA|\psi\rangle$ and $AB|\psi\rangle$ are linearly independent, (15) is a genuinely control–target entangled state. If instead $BA|\psi\rangle = \lambda AB|\psi\rangle$ with $|\lambda| = 1$ but $\lambda \neq 1$ (the anticommuting case gives $\lambda = -1$), the state factors as $\frac{1}{\sqrt{2}}(\lambda|0\rangle_c + |1\rangle_c) \otimes AB|\psi\rangle$, a product state whose *control* has been rotated to $\frac{1}{\sqrt{2}}(\lambda|0\rangle_c + |1\rangle_c) \neq |+\rangle_c$. Either way the control’s response distinguishes the switch from every fixed order. It is this observable control response, not entanglement per se, that certifies the absence of a fixed order. $\square$

Remark 5.2 (Consistency with section 8.1, [S]). The correction matters: the worked section 8.1 takes anticommuting Paulis and finds the *product* state $i|-\rangle_c \otimes \sigma_y|\psi\rangle$. This is not a counterexample to theorem 5.1 but exactly its $\lambda = -1$ sub-case: the state factors, yet the control has swung from $|+\rangle$ to $|-\rangle$, an observable change no spectator fixed order produces.

5.3 The switch as a process matrix, and its nonseparability

Promoting the two calls of A and B to two local laboratories (each party implements one call on the target as it passes through) and treating the control plus target as the process “background,” the switch becomes a process matrix W_{sw} on the parties’ in/out spaces together

with a global future laboratory that reads the control. The following is the switch’s headline property.

Theorem 5.3 (Causal nonseparability of the switch, [S/H]). *The quantum switch W_{sw} is a valid process matrix that is causally nonseparable: there is no decomposition $W_{\text{sw}} = q W^{A \prec B} + (1 - q) W^{B \prec A}$ into definite-order processes. Consequently there is a causal witness S with $\text{Tr}[S W_{\text{sw}}] < 0$. The nonseparability is [S]; its reading as a laboratory realization of “indefinite causal order between the two operations” is [H].*

Proof sketch. Validity is checked as in theorem 3.2: W_{sw} is the CJ operator of the isometry (14) viewed as a map from the two parties’ outputs to their inputs and the final control readout, and it obeys (7)–(9) with the control leg included. Nonseparability follows from theorem 5.1: a decomposition into $W^{A \prec B}$ and $W^{B \prec A}$ would make the target evolution, conditioned on the (hidden) order variable, a fixed order on each branch, hence a classical mixture of AB and BA with no control-target coherence; but the switch produces coherence (15) detectable by a $|\pm\rangle$ control measurement, which a classical order mixture cannot reproduce. The separating witness exists by theorem 4.6. The S/H split is as stated: the operator computation is [S]; “the two orders are physically superposed” is a physical reading, [H], and one that holds on a fixed background spacetime, so it does not by itself say anything about gravity. $\square$

Remark 5.4 (Experimental status, [H]). The switch has been realized photonically (Procopio et al. [15]; Rubino et al. [16]) with causal-witness certification, and device-independently certified under a no-superluminal-signalling assumption by van der Lugt, Barrett, and Chiribella [17]. These realizations are on ordinary spacetime: they demonstrate that the *operational* content of indefinite order is physical, not that gravity sources it. That distinction is the entire content of the H -versus- P split for this topic.

6 Causal sets as discrete causal order

We now turn to the other, older, and epistemically firmer half of the topic. A causal set reconstructs *geometry* from *definite* order; it does not touch indefiniteness. Its role in this paper is twofold: it is the library’s existing dictionary entry, and it is the cleanest example of the first-level reconstruction (order $\rightarrow$ geometry) against which the process matrix’s more radical second-level reconstruction (data $\rightarrow$ order) is measured.

6.1 The poset axioms

Definition 6.1 (Causal set, [S]). A *causal set* is a pair $(C, \preceq)$ where C is a set and $\preceq$ a relation that is

- (i) *reflexive*: $x \preceq x$ (or, in the strict convention $\prec$, irreflexive $x \not\prec x$);
- (ii) *antisymmetric*: $x \preceq y$ and $y \preceq x$ imply $x = y$ (equivalently, $\prec$ has no cycles);
- (iii) *transitive*: $x \preceq y$ and $y \preceq z$ imply $x \preceq z$;
- (iv) *locally finite*: every order interval $[x, z] = \{y : x \preceq y \preceq z\}$ is a finite set.

Axioms (i)–(iii) make $(C, \preceq)$ a partial order; (iv) makes it locally finite and is what supplies discreteness.

Axioms (i)–(iv) are the entire ontology of the causal-set program: no manifold, no coordinates, no metric, only a set and a locally finite order. The order is read as causal precedence: $x \prec y$ means x is in the causal past of y . Antisymmetry is the discrete statement of no closed causal loops. Local finiteness is the discrete statement that between any two causally related events only finitely many events intervene — the causal analogue of a finite proper-time separation containing finitely much “stuff.”

6.2 Order plus number equals geometry

The reconstruction claim has two ingredients, and they are of different epistemic strength. The first, that order determines conformal (light-cone) structure, is a theorem about the continuum.

Theorem 6.2 (Order determines conformal structure; Malament, [S]). *Let (M_1, g_1) and (M_2, g_2) be distinguishing Lorentzian manifolds of the same dimension $d \geq 3$, and $f : M_1 \rightarrow M_2$ a bijection preserving the chronological order in both directions ($x \ll y \iff f(x) \ll f(y)$). Then f is a smooth conformal isometry: $f^*g_2 = \Omega^2 g_1$ for some positive function Ω . Hence the causal order of such a manifold determines its metric up to a local conformal factor.*

This is Malament’s 1977 theorem [6], building on Hawking, King, and McCarthy; the “distinguishing” hypothesis (points are separated by their chronological pasts and futures) is a mild causality condition. The dimension restriction $d \geq 3$ is essential: in $1+1$ dimensions the causal order factorizes through the two null coordinates, and an independent monotone (not necessarily smooth) reparametrization of each null coordinate preserves the order without being a conformal isometry, so order alone does not fix the conformal class in two dimensions. For $d \geq 3$, Theorem 6.2 is the precise sense in which “order is almost geometry”: everything about the metric except one scalar function per point is fixed by the causal order alone.

The missing scalar is the conformal factor, equivalently the volume element, and it is recovered by counting.

Proposition 6.3 (Number recovers volume; order plus number equals geometry, [S/H]). *Fix a Lorentzian manifold (M, g) and generate a causal set by a Poisson sprinkling at density ρ : place points by a Poisson process with expected count $\rho \text{Vol}_g(R)$ in each region R , and induce the order from g ’s causal structure. Then the expected number of sprinkled points in a region equals ρ times its volume, so counting elements estimates volume; combined with theorem 6.2 (order gives conformal structure), order and number together determine the metric g up to the overall density scale. The combinatorial-to-continuum direction (that a given abstract causal set arises from such a sprinkling) is the content of the *Hauptvermutung* below and is [H].*

The *S/H* split is the whole subtlety of causal-set kinematics. That sprinkling *into* a manifold reproduces its volume and causal structure is straightforward ([S], it is the definition of the Poisson process plus Malament). That an abstract causal set handed to you without a manifold *came from* such a sprinkling — that it is “manifoldlike” — is the hard, still-open direction.

Conjecture 6.4 (Hauptvermutung of causal-set theory, [H]). *Two Lorentzian manifolds into which a given causal set can be faithfully embedded (as a sprinkling at the same density) are approximately isometric: the manifold, if it exists, is essentially unique. Equivalently, “manifoldlikeness” of a causal set is a well-defined property picking out an approximate geometry.*

The Hauptvermutung is proven only in restricted settings and remains the central open problem of the reconstruction direction of causal-set theory; Surya’s review [3] is the current comprehensive account. That Poisson sprinkling breaks no local Lorentz symmetry on average (Bombelli, Henson, Sorkin [5]) is a supporting [S] result: the discreteness is Lorentz-invariant, unlike a fixed lattice, which is why order-plus-number can hope to reconstruct a Lorentzian rather than a merely Euclidean geometry.

6.3 Dynamics: sequential growth

Causal sets come with a candidate dynamics, the classical sequential growth models of Rideout and Sorkin [4]: a causal set is built one element at a time, each new element added to the future of some subset of the existing ones, with transition probabilities constrained by discrete general covariance (Bell causality and label independence). This is a genuinely *definite-order* stochastic process: it grows a fixed partial order. We flag it because it is the sharpest contrast with the process-matrix side. The causal-set dynamics grows a definite order stochastically; the process matrix contemplates order that is not definite at all. Whether the former is a special (classical, definite) case of a suitable “quantum sequential growth” living in the process-matrix world is, to our knowledge, open, and we return to it as a proposed problem in section 11.

7 Order as representation data

We now state the paper’s viewpoint as labeled results. The theme is that in both formalisms the causal structure is an invariant read off representation data, and that the two formalisms sit at two different levels of the reconstruction.

Proposition 7.1 (Causal order is the isomorphism invariant of the poset, [S]). *For a causal set the physical content is the isomorphism class of $(C, \preceq)$ in the category of posets: two order-isomorphic causal sets are the same causal set, and any manifold embedding is auxiliary data, not part of the invariant. The automorphism group $\text{Aut}(C, \preceq)$ is the discrete residual symmetry, playing the role that the isometry (Killing) group plays for a Lorentzian manifold (Topic 1).*

Proof. Immediate from theorem 6.1: the axioms are stated purely in terms of the relation $\preceq$, so any order isomorphism preserves all of them and all derived structure (intervals, chains, antichains, the counting measure). An embedding into a manifold is a map to auxiliary data that is not referenced by the axioms. $\square$

This makes causal sets a clean instance of the library’s recurring structural fact, “physical content is a groupoid/isomorphism class, not a coarse representative.” The poset is the object; its automorphisms are retained symmetry, not gauge to be quotiented away.

Proposition 7.2 (Causal (in)definiteness is a convex-geometric invariant of W , $[S]$). *For a process matrix the analogous invariant is the pair (the operator W up to local unitary relabeling; its membership class in the convex separable set $\mathcal{S}$). Causal definiteness is the property $W \in \mathcal{S}$; indefiniteness is $W \notin \mathcal{S}$; the degree of indefiniteness is measured by the optimal causal-witness value or by the random-robustness (the least mixing with noise restoring separability). All of these are invariants of W under local operations, computable by semidefinite programming.*

Proof. Local unitary relabeling is a congruence $W \mapsto UWU^\dagger$ with U a product of local unitaries on the in/out legs; it preserves positivity, the validity constraints (8)–(9), and the separable set $\mathcal{S}$ (which is defined by the same constraints on its ordered pieces). Hence membership in $\mathcal{S}$, the witness value $\min_S \text{Tr}[SW]$ over normalized witnesses, and the robustness are all invariant. Computability is theorem 4.6. $\square$

Theorems 7.1 and 7.2 are the two halves of the topic’s reconstruction thesis stated as invariance statements: in the causal-set case the invariant is order-isomorphism type; in the process-matrix case it is the local-operation class together with convex-position data. Neither invariant mentions a background spacetime.

Theorem 7.3 (Two levels of causal reconstruction, $[P/H]$). *The two formalisms reconstruct at different levels:*

- (A) (Order $\rightarrow$ geometry, $[S/H]$.) *A causal set reconstructs a Lorentzian geometry from a definite order via theorem 6.2 and theorem 6.3, subject to the Hauptvermutung.*
- (B) (Data $\rightarrow$ order, $[S/H]$.) *A process matrix reconstructs the causal order — or certifies its absence — from operational statistics via theorems 3.2, 4.3 and 7.2.*

The proposal that (A) and (B) are two faces of a single “causal-order representation” invariant — that causal-set order data and process-matrix indefiniteness are special and general cases of one reconstruction — is $[P]$: it is not established in the literature and is offered here as an organizing conjecture, not a theorem.

Discussion of the tag. Parts (A) and (B) are each supported by the cited $[S]/[H]$ results and carry $[S/H]$. The synthesis claim is unproven. A precise version would ask for a single mathematical object — perhaps a poset-valued or “causal-structure-valued” generalization of the process matrix — specializing to a causal set in a classical, definite-order limit and to a process matrix in an operational, possibly-indefinite limit. No such object is constructed here or, to our knowledge, in the literature (checked against 2023–2026 surveys including the “map of indefinite causal order” [19]); hence $[P]$. We keep the composite label $[P/H]$ to record that the two ingredient reconstructions are $[S/H]$ while their proposed unification is $[P]$, and worst-component-wins gives P . $\square$

8 Worked examples

8.1 Example 1: the quantum switch on Pauli operations

Take the target a qubit and the two operations the Pauli unitaries $A = \sigma_x$, $B = \sigma_z$, which anticommute: $\sigma_x \sigma_z = -\sigma_z \sigma_x$, so $[A, B] = 2\sigma_x \sigma_z \neq 0$. The switch acts on control $\otimes$ target by

(14):

$$\begin{aligned} S_{\text{sw}}(\sigma_x, \sigma_z) &= |0\rangle\langle 0|_c \otimes (\sigma_z \sigma_x) + |1\rangle\langle 1|_c \otimes (\sigma_x \sigma_z) \\ &= |0\rangle\langle 0|_c \otimes (i\sigma_y) + |1\rangle\langle 1|_c \otimes (-i\sigma_y), \end{aligned} \quad (16)$$

using $\sigma_z \sigma_x = i\sigma_y$ and $\sigma_x \sigma_z = -i\sigma_y$. On control $|+\rangle$ and any target $|\psi\rangle$,

$$\begin{aligned} S_{\text{sw}}(|+\rangle_c \otimes |\psi\rangle) &= \frac{1}{\sqrt{2}}(|0\rangle_c \otimes (i\sigma_y|\psi\rangle) + |1\rangle_c \otimes (-i\sigma_y|\psi\rangle)) \\ &= i|-\rangle_c \otimes (\sigma_y|\psi\rangle), \end{aligned} \quad (17)$$

so measuring the control in the $|\pm\rangle$ basis is deterministic (outcome $|-\rangle$): the coherence between the two orders is total and visible. Had the two orders been classically mixed rather than coherently superposed, the control would be maximally mixed in the $|\pm\rangle$ basis, outcome probability $1/2$ each. The gap between “deterministic $|-\rangle$ ” and “coin-flip” is the operational signature of causal nonseparability, and it is exactly the quantity the companion code checks by comparing BA and AB on the target and confirming they differ (theorem 5.1). For commuting operations, e.g. $A = B = \sigma_z$, the two orders agree, the control ends in $|+\rangle$, and the switch degenerates to a fixed order — no nonseparability, as it must be.


8.2 Example 2: the OCB causal inequality, bound and violation

We carry out the bound of theorem 4.3 by explicit enumeration, which is also what the companion code does. Fix the order $A \prec B$ and restrict to deterministic strategies (the extreme points of the strategy polytope; convex mixing cannot beat them for an affine objective). Alice’s output is a function $a = f(x)$; Bob’s is a function $b = g(a, y, b')$ of everything he can see, including Alice’s message a . The objective (11) averages over uniform x, y, b' :

$$p = \frac{1}{2} \left[\frac{1}{4} \sum_{x,y} [g(f(x), y, 0) = x] + \frac{1}{4} \sum_{x,y} [f(x) = y] \right]. \quad (18)$$

The second bracket is $\frac{1}{4} \sum_{x,y} [f(x) = y] = \frac{1}{2}$ for any f (for each x , exactly one of the two y values matches). The first bracket is maximized by $f = \text{id}$ (Alice broadcasts x) and $g(a, y, 0) = a$ (Bob copies), giving 1. Hence $p_{\text{max}}^{A \prec B} = \frac{1}{2}(1 + \frac{1}{2}) = \frac{3}{4}$, and by symmetry $p_{\text{max}}^{B \prec A} = \frac{3}{4}$; convex mixtures do no better. This is the causal bound, obtained by finite enumeration over the $2^2 \times 2^{2 \cdot 2}$ deterministic strategies — a computation the code performs exhaustively, confirming the maximum is exactly $3/4$.

The OCB process (13) reaches $p = (2 + \sqrt{2})/4 \approx 0.8536$ (theorem 4.4), a violation of about 0.1036 above the bound. The violation is not large — causal inequalities are famously hard to violate by much, and whether any physically motivated process (as opposed to the abstract W_{OCB}) violates one at all is itself a subtle question — but it is nonzero and certifies causal nonseparability device-independently.




8.3 Example 3: a sprinkled causal diamond

Take the 1+1 Minkowski causal diamond, the set of points causally between an origin event o and a top event t , in light-cone coordinates $u = t+x$, $v = t-x$ with $u, v \in [0, 1]$. Two points $p = (u_p, v_p)$, $q = (u_q, v_q)$ satisfy $p \preceq q$ iff $u_p \leq u_q$ and $v_p \leq v_q$: the (reflexive) causal order is the product order on the two null coordinates, with $p \prec q$ its strict part. Sprinkle a handful of points and read off the order.

Concretely, sprinkle five points (in addition to $o = (0, 0)$ and $t = (1, 1)$):

point	u	v
o	0.00	0.00
p_1	0.20	0.55
p_2	0.35	0.15
p_3	0.55	0.60
p_4	0.60	0.30
p_5	0.80	0.75
t	1.00	1.00

The induced order (product order on (u, v)) makes o the minimum and t the maximum; $p_2 \prec p_4 \prec p_5$ is a chain (each later in both coordinates); $p_1 \prec p_3 \prec p_5$ is another; p_1 and p_2 are incomparable (spacelike: $u_{p_1} < u_{p_2}$ but $v_{p_1} > v_{p_2}$), as are p_3 and p_4 . The Hasse diagram of covering relations is:



Every order interval is finite (the whole set has seven elements), so this is a causal set in the sense of theorem 6.1. Its order faithfully reflects the continuum diamond’s causal structure: the chains are exactly the continuum timelike-related families, and the antichains ($\{p_1, p_2\}$, $\{p_3, p_4\}$) are the continuum spacelike-related ones. The number of elements, seven, estimates the diamond’s volume at the sprinkling density used; refining the density (more points) sharpens the volume estimate while preserving the conformal order structure — “order plus number” in miniature. The companion code builds this poset, verifies the four axioms of theorem 6.1 (irreflexivity, antisymmetry, transitivity of the induced relation, and finiteness of every interval), and confirms that the order it computes from the coordinates agrees with the product-order causal relation, i.e. the discrete order is faithful to the continuum one on this

sample. We use $1 + 1$ dimensions only because the poset is then small enough to draw; the example illustrates the order/counting structure (chains, antichains, interval finiteness, and counting-as-volume), not the full conformal reconstruction of theorem 6.2, whose hypotheses require $d \geq 3$. The order-plus-number recovery of a metric is a $d \geq 3$ statement; the $1 + 1$ diagram is a visual aid for the combinatorics.

9 Relation to the other regimes and the S/H/P calculus

This topic sits at the boundary of the library and touches three other regimes. We record the composition each time via the worst-component rule (2).

9.1 Lorentzian geometry (Topic 1)

Topic 1 supplies the continuum causal structure that both halves of this topic reconstruct or dissolve. Causal sets reconstruct it: theorem 6.2 is the bridge, and the composite claim “a manifoldlike causal set reconstructs a Topic-1 Lorentzian geometry” composes $[S/H]$ (Topic 1) with the Hauptvermutung $[H]$ to give

$$\max([S/H], [H]) = [H]. \quad (19)$$

Process matrices dissolve it: the claim “a causally nonseparable process has no Topic-1 causal structure” is the sharpest possible departure from the fixed background of Topic 1, and reading it as a statement about *gravitational* causal structure composes Topic 1’s $[S/H]$ with this topic’s P to give $[P]$.

9.2 Noncommutative geometry (Topic 11)

Topic 11 (the sibling proposed extension, QG-IX) reconstructs metric geometry from spectral data, and its Lorentzian/causal generalization — Krein-space spectral triples with an algebraic notion of causal order (Franco–Eckstein and collaborators) — is the natural bridge to this topic. The composite “an algebraic reconstruction of causal order from spectral data” pairs two proposed-extension topics, both carrying an H or worse, and composes to $[H]$ at best, $[P]$ once the indefinite-order reading is added. We flag this bridge as the most promising unexplored link for either topic but claim nothing proven.

9.3 Spin foams and histories (Topic 4)

Topic 4’s covariant histories are sums over discrete geometries with a fixed combinatorial causal skeleton (the two-complex). A causal set is a candidate for that skeleton made primary rather than auxiliary; “causal-set-histories quantum gravity” (the quantum sequential growth / decoherence-functional program) is the point of contact. The composite is $[H]$ (Topic 4 is bare H) degrading to $[P]$ if the histories are allowed genuinely indefinite order. We do not develop it.

9.4 The weakest-link summary

The table records the composite status of each bridge under (2).

Composite claim	Ingredients	Status
Manifoldlike causal set reconstructs Lorentzian geometry	Topic 1 [S/H] $\circ$ Hauptvermutung [H]	[H]
Causally nonseparable process $\Rightarrow$ no gravitational order	Topic 1 [S/H] $\circ$ this topic [P]	[P]
Algebraic causal order from spectral data	Topic 11 [H] $\circ$ causal reading [P]	[P]
Causal-set histories with indefinite order	Topic 4 [H] $\circ$ indefiniteness [P]	[P]
Process-matrix mathematics alone	theorems 3.2 and 4.3 [S]	[S]

The pattern is the intended one: the internal mathematics is [S], the physical representations are [H], and every bridge that reaches toward gravity lands at [P]. This topic is where the library’s weakest-link calculus does its most important work, because it is the topic most tempting to overstate.

10 A Haskell model of causal order

The companion code (in `src/quantum-causal-structures/`) realizes the three checkable fragments of the paper. It is written against `base` only and compiles with `ghc`.

Core.hs: definite orders and the causal bound. A definite bipartite strategy is a pair of functions, Alice’s $f : \{0, 1\} \rightarrow \{0, 1\}$ and Bob’s $g : \{0, 1\}^3 \rightarrow \{0, 1\}$ (his inputs being Alice’s message, his own bit y , and the task bit b'), constrained to the order $A \prec B$ (Bob may read Alice’s message; Alice reads nothing of Bob’s). The module enumerates all such deterministic strategies, evaluates the OCB objective (11) exactly as a rational number, and exposes the maximum. The checkable property is that this maximum equals $3/4$ over every definite-order strategy in both orders — a computational proof of theorem 4.3.

Poset.hs: finite causal sets. A causal set is a finite carrier with a relation given as a predicate. The module checks the four axioms of theorem 6.1 (irreflexivity of the strict order, antisymmetry, transitivity, and finiteness of every interval, the last automatic for a finite carrier and verified by computing interval cardinalities), builds the sprinkled causal diamond of section 8.3 from its coordinate table, and confirms that the induced product order is a valid strict partial order faithful to the continuum causal relation on the sample.

Switch.hs: order dependence of the switch. Operations are modeled as 2×2 complex unitaries (the Pauli matrices, via a small `Data.Complex`-based matrix layer). The module computes both orders BA and AB on the target and reports whether they agree; for non-commuting operations they differ, which is the discrete witness of theorem 5.1 that no fixed order reproduces both control branches. For commuting operations they agree, recovering the degenerate fixed-order case.

Properties.hs and Main.hs. `Properties.hs` collects the checkable claims — the causal bound equals $3/4$; the sprinkled diamond satisfies the poset axioms and is order-faithful; the switch’s two orders disagree for anticommuting Paulis and agree for commuting ones — and `Main.hs` runs the demonstrations and the property suite, exiting nonzero if any property fails.

The code proves nothing about gravity, and is not meant to. It discharges the finite, [S]-level content of the paper: the causal bound, the poset axioms, and the order dependence of the switch. That is exactly the scope the honest status of the topic allows.

11 Limitations and open problems

This is the weakest topic in the library, and its limitations section is correspondingly the longest. We separate what is solid from what is not.

11.1 The gravitational reading is speculative

The entire motivation for caring about indefinite causal order as *quantum gravity* rests on the gedanken-arguments of Hardy [13] and Zych et al. [14]: a mass in spatial superposition sources a superposed gravitational field, hence a superposed light-cone structure, hence indefinite temporal order between events. This is a physical argument, not a theorem, and it has never been realized gravitationally. The laboratory quantum switch is implemented on ordinary spacetime with photonic or interferometric control [15, 16, 17]; it demonstrates that the *operational* content of indefinite order is real, not that gravity is its source. Everything in this paper that connects process matrices to gravity is [P], and we have tried never to let a P masquerade as an H .

11.2 No dynamics for indefinite order

Causal sets have a candidate dynamics (Rideout–Sorkin sequential growth [4]). Process matrices have none: the formalism is kinematical, a description of correlations given a fixed background operator W , with no equation of motion for W and no account of how an indefinite causal order would evolve, form, or decohere into a definite one. The quantum-to-classical transition that OCB [7] gestures at (“in the classical limit causal order always arises”) is not made dynamical. This is the single largest gap: a theory of indefinite causal order without a dynamics of causal order is a theory of a snapshot.

11.3 The synthesis of the two halves is unproven

Theorem 7.3’s proposal — that causal sets and process matrices are two faces of one causal-order invariant — is [P] and, as far as we have found, absent from the literature (the recent survey [19] organizes the process-matrix side without touching causal sets). A precise formulation would require a mathematical object interpolating between a locally finite poset and a process matrix; we do not have one. Proposed problem: construct a “causal-structure-valued” generalization of the process matrix whose classical, definite-order limit is a causal set and whose sequential-growth analogue recovers Rideout–Sorkin, thereby making the analogy of section 6 into a theorem. We list this as the topic’s flagship open problem.

11.4 Interpretational status of “indefinite”

Whether a causal-inequality violation shows that causal order is *absent* (ontologically indefinite) or merely *unmodeled* (epistemically hidden) is the causal analogue of the Bell-inequality interpretation debate, and it is not settled by the mathematics. Theorem 4.5 carries this step at [H] deliberately. A committed operationalist reads the violation as “no order exists”; a committed realist may seek a hidden order at the cost of superluminal signalling or retro-causality (the loopholes the device-independent result of [17] must explicitly assume away). We take no side; we record that the interpretive step is not [S].

11.5 The Hauptvermutung is open

The causal-set half’s reconstruction direction (theorem 6.4) is proven only in restricted settings. Without it, “manifoldlikeness” is not known to pick out a unique approximate geometry, and order-plus-number reconstruction of the metric is secure only in the sprinkling-into-a-known-manifold direction, not the recover-the-manifold-from-the-order direction that the reconstruction thesis actually needs. This is why the causal-set entry is H , not S .

11.6 Roadmap and dictionary status

Finally, the structural limitation flagged from the first page: this topic has no assigned roadmap module and, for its modern content, no dictionary entry. QG-X is proposed here, not adopted; the process-matrix dictionary row is proposed here, not adopted. A reader should treat both as suggestions for the library’s future, carrying the composite status [H/P], and should not cite them as established parts of the project.

12 Discussion

The two formalisms this paper collects share a stance and differ in radicalism. Both refuse to take a background causal structure as given and instead read causal structure off other data. Causal sets read it off a discrete order and recover geometry; they are conservative, keeping a definite order and merely making it discrete and primary. Process matrices read it off operational statistics and are willing to conclude that there is no definite order at all;

they are radical, dissolving the one structure — “before and after” — that even the causal set retains.

The library’s governing perspective names four things to be reconstructed: spacetime, matter, causality, and measurement. Most of the twelve topics reconstruct spacetime. This one is the library’s only sustained attempt at causality per se, and its lesson is cautionary. The mathematics is clean and, in its finite fragments, fully checkable: process matrices form a convex body, causal inequalities have provable bounds, the switch is provably nonseparable, order provably determines conformal structure. But the physics — and above all the *gravitational* physics — that would make this a reconstruction of causality in nature rather than in the laboratory is not in hand. The honest report is that we can reconstruct causal order from representation data in two mathematically precise senses, that laboratory systems realize the operational content of the more radical sense, and that the leap from “the laboratory switch is causally nonseparable” to “quantum gravity produces indefinite causal order” is a leap we can motivate but not make.

That is why the topic’s status is [H/P] and why it is the weakest link in the library’s chain. Under the composition calculus, any claim that depends on this topic’s gravitational reading caps the whole chain at [P] — which is the precise, checkable sense in which the reconstruction of causality remains, for now, a program rather than a result.

13 Conclusion

We have presented indefinite causal order as a reconstruction of causality from representation data, in two formalisms. The process-matrix formalism encodes all correlations compatible with local quantum mechanics in a positive operator W that presupposes no global order; causal separability is a convex condition, causal inequalities are provable linear bounds with the value $p \leq 3/4$ for definite order, and the quantum switch and the OCB process violate separability provably. The causal-set formalism encodes spacetime as a locally finite partial order from which conformal geometry follows by Malament’s theorem and volume by counting. We stated the central facts as labeled propositions, worked the switch, the causal inequality, and a sprinkled causal diamond in full, and backed the finite content with a Haskell development that checks the poset axioms, the causal bound, and the switch’s order dependence.

We have also been unusually explicit about what is not established. This topic is a proposed extension of the library’s roadmap (QG-X), with a proposed dictionary entry at status [H/P]; its internal mathematics is [S], its physical readings are [H], and its gravitational readings are [P]. The synthesis of its two halves under a single causal-order invariant is a conjecture, and the dynamics of indefinite order is absent. Under the weakest-link calculus this makes causality the epistemically softest of the four things the governing perspective sets out to reconstruct. Reconstructing it securely — giving indefinite causal order a dynamics, proving the Hauptvermutung, and building the object that would fuse posets with process matrices — is the work that would move this topic from proposed extension to established module.

A The core Haskell module

For completeness we reproduce the signature-level content of the causal-order model; the full sources are in `src/quantum-causal-structures/`. The strategy types and the causal-bound evaluator (module `Core`) are:

```
type Bit = Int
data Order = ABefore | BBefore
ocbValueAB :: (Bit -> Bit) -> ((Bit, Bit, Bit) -> Bit) -> Rational
ocbValueBA :: ((Bit, Bit) -> Bit) -> ((Bit, Bit) -> Bit) -> Rational
causalBound :: Rational -- = 3/4, the max over all definite strategies
```

and the poset layer (module `Poset`) is:

```
data Causet a = Causet { carrier :: [a], leq :: a -> a -> Bool }
isStrictPartialOrder :: Eq a => Causet a -> Bool
allIntervalsFinite :: Eq a => Causet a -> Bool
diamondFaithful :: Bool
```

with `Switch.hs` providing `ordersAgree :: Mat2 -> Mat2 -> Bool` over 2×2 complex unitaries. The property suite in `Properties.hs` asserts `causalBound == 3 % 4`, the poset axioms for the diamond, and `not (ordersAgree sx sz)` for the anticommuting Paulis.

References

- [1] L. Bombelli, J. Lee, D. Meyer, and R. D. Sorkin, “Space-time as a causal set,” *Phys. Rev. Lett.* **59** (1987) 521–524; doi:10.1103/PhysRevLett.59.521.
- [2] R. D. Sorkin, “Causal sets: discrete gravity,” in *Lectures on Quantum Gravity* (A. Gomberoff and D. Marolf, eds.), Springer, 2005, pp. 305–327; arXiv:gr-qc/0309009.
- [3] S. Surya, “The causal set approach to quantum gravity,” *Living Rev. Relativ.* **22** (2019) 5; doi:10.1007/s41114-019-0023-1; arXiv:1903.11544.
- [4] D. P. Rideout and R. D. Sorkin, “Classical sequential growth dynamics for causal sets,” *Phys. Rev. D* **61** (2000) 024002; arXiv:gr-qc/9904062.
- [5] L. Bombelli, J. Henson, and R. D. Sorkin, “Discreteness without symmetry breaking: a theorem,” *Mod. Phys. Lett. A* **24** (2009) 2579–2587; arXiv:gr-qc/0605006.
- [6] D. B. Malament, “The class of continuous timelike curves determines the topology of spacetime,” *J. Math. Phys.* **18** (1977) 1399–1404; doi:10.1063/1.523436.
- [7] O. Oreshkov, F. Costa, and Č. Brukner, “Quantum correlations with no causal order,” *Nat. Commun.* **3** (2012) 1092; doi:10.1038/ncomms2076; arXiv:1105.4464.
- [8] G. Chiribella, G. M. D’Ariano, P. Perinotti, and B. Valiron, “Quantum computations without definite causal structure,” *Phys. Rev. A* **88** (2013) 022318; doi:10.1103/PhysRevA.88.022318; arXiv:0912.0195.

- [9] G. Chiribella, G. M. D’Ariano, and P. Perinotti, “Theoretical framework for quantum networks,” *Phys. Rev. A* **80** (2009) 022339; arXiv:0904.4483.
- [10] M. Araújo, C. Branciard, F. Costa, A. Feix, C. Giarmatzi, and Č. Brukner, “Witnessing causal nonseparability,” *New J. Phys.* **17** (2015) 102001; arXiv:1506.03776.
- [11] C. Branciard, M. Araújo, A. Feix, F. Costa, and Č. Brukner, “The simplest causal inequalities and their violation,” *New J. Phys.* **18** (2016) 013008; doi:10.1088/1367-2630/18/1/013008; arXiv:1508.01704.
- [12] O. Oreshkov and C. Giarmatzi, “Causal and causally separable processes,” *New J. Phys.* **18** (2016) 093020; arXiv:1506.05449.
- [13] L. Hardy, “Probability theories with dynamic causal structure: a new framework for quantum gravity,” arXiv:gr-qc/0509120 (2005).
- [14] M. Zych, F. Costa, I. Pikovski, and Č. Brukner, “Bell’s theorem for temporal order,” *Nat. Commun.* **10** (2019) 3772; arXiv:1708.00248.
- [15] L. M. Procopio, A. Moqanaki, M. Araújo, F. Costa, I. A. Calafell, E. G. Dowd, D. R. Hamel, L. A. Rozema, Č. Brukner, and P. Walther, “Experimental superposition of orders of quantum gates,” *Nat. Commun.* **6** (2015) 7913; arXiv:1412.4006.
- [16] G. Rubino, L. A. Rozema, A. Feix, M. Araújo, J. M. Zeuner, L. M. Procopio, Č. Brukner, and P. Walther, “Experimental verification of an indefinite causal order,” *Sci. Adv.* **3** (2017) e1602589; arXiv:1608.01683.
- [17] T. van der Lugt, J. Barrett, and G. Chiribella, “Device-independent certification of indefinite causal order in the quantum switch,” *Nat. Commun.* **14** (2023) 5811; doi:10.1038/s41467-023-40162-8.
- [18] J.-M. A. Allen, J. Barrett, D. C. Horsman, C. M. Lee, and R. W. Spekkens, “Quantum common causes and quantum causal models,” *Phys. Rev. X* **7** (2017) 031021; arXiv:1609.09487.
- [19] “A map of indefinite causal order,” arXiv:2506.04607 (2025); a survey of the process-matrix classification of indefinite causal order, consulted during the July 2026 research pass.