When Does a Quantum Transition Become a Record?

Instruments, export, and cadence in finite quantum dynamics

David Elliman

Neuro-Symbolic Ltd

ORCID: 0009-0001-9792-4036

dave@neusym.ai

Abstract

Quantum measurement theory distinguishes outcome probabilities from state transformations, while open-system theory distinguishes a reduced channel from its environmental realisation. In finite models, however, a transition table is still easily read as a record, and a record as evidence that the same measurement occurred after every elementary step. We formalise a six-layer separation between transition support, effects, instruments, environmental export, invocation cadence, and autonomous dissipative dynamics. Exact calibration of a finite record observable fixes its effects to the spectral projectors, but does not fix the post-measurement instrument. A continuous family of positive-semidefinite environmental Gram matrices then realises the same basis transitions and populations with no, partial, or projective record export. Endpoint observables do not determine cadence, and a separately lawful Lindblad term has an independent rate. These distinctions are exhibited in an eight-register example. Its truthful transfer observable has nine sectors, whereas its physical exit code has seven live addresses. One positive invocation supplies 104 exit transitions; the complementary sector supplies 120; and their polarity-complete union supplies 224. Three finite tests separate record export, the complementary dissipator, and polarity-resolved unravellings. The result is not a new interpretation of quantum mechanics. It is a criterion for what additional microscopic information is needed before a finite transition grammar can be promoted to a record-producing physical law.

Keywords: quantum measurement; quantum instruments; measurement records; open quantum systems; environmental decoherence; quantum trajectories; process tensors; finite-state quantum models

1 Introduction

A detector record is more than a transition which could have occurred. It is also more than an effect operator that assigns the correct probabilities. Something must carry the outcome away from the system, and the time at which that happens is part of the physical model. These statements are familiar in open-system language, but they are often compressed when a finite transition grammar is used as a microscopic description. The compression can be consequential. A transition may be reversible while its record is not; the same effects may accompany different state changes; and a jump operator that is algebraically available need not occur at the rate, or on the schedule, suggested by another part of the model.

The standard theory already provides the correct vocabulary. A quantum instrument combines outcome statistics and conditional state transformations [1–4]. Stinespring dilation separates a completely positive map from an environment which realises it [5]; quantum dynamical semigroups separate an allowed jump operator from the coefficient with which it enters a generator [6, 7]. Decoherence and quantum Darwinism explain how some system observables become stable, redundantly accessible environmental records [8, 9]. Collision models make the supply of fresh environmental units explicit [10], and process tensors show why multitime dynamics cannot in general be inferred from endpoint maps [11, 12].

The purpose of this paper is not to replace those theories. It is to put six objects side by side and then keep them apart:

layer	object	question answered
1	transition support	Which basis transitions are possible?
2	effects or POVM	Which visible outcomes have which probabilities?
3	instrument	What conditional state transformation follows an outcome?
4	environmental export	Has distinguishable outcome information left the system?
5	invocation cadence	At which microscopic boundaries is that export attempted?
6	dissipative generator	Which channels occur autonomously, and at what rates?

Several implications are immediate once this hierarchy is made explicit, but the useful point is that they can be proved by small finite witnesses. We first give general finite-dimensional statements. We then apply them to an eight-register model whose transition and record alphabets can be exhaustively enumerated. The example is intentionally placed after the general results. It is evidence that the distinctions bite in a nontrivial construction, not a premise the reader must accept in order to understand them.

Each separation in the hierarchy is known in its appropriate formal setting. The contribution here is to make their succession checkable by exhaustive finite support counts, and to isolate cadence as an independent multitime datum which cannot be read from either a one-step channel or its fixed points.

The principal findings are as follows.

1. Exact calibration fixes the visible effects of a finite record observable, but it does not select a unique instrument.

2. A transition table does not decide whether an environment receives no record, a partial record, or an orthogonal projective record.

3. Single-time or endpoint data do not determine whether the same record coupling was invoked at each intermediate transition.

4. An algebraically constructed jump family may enter an autonomous dissipator with an independent, undetermined rate.

5. Even a complete union of jump edges does not decide whether the environment resolves the constituent branches.

In the finite example these statements lead to three different experimental questions: an inverse-interference test for record export, a complementary-only exit-rate test for an independent dissipator, and a conditioned-coherence test for the jump unravelling. A single count cannot answer all three.

2 Relation to established measurement theory

It is important to delimit the contribution before giving the formalism. Quantum instruments were introduced precisely because an observable alone does not specify the state change caused by measurement [1]. Completely positive instruments and indirect measurement models were subsequently connected in a general form [3], and modern accounts treat effects, operations, instruments, and dilations as distinct objects from the outset [4]. We do not claim this distinction as new.

Nor is environmental nonuniqueness new. Stinespring representations, quantum trajectories and stochastic unravellings all make clear that reduced dynamics can admit more than one environmental or conditioned description [5, 13, 14]. Recent work has made the freedom in jump unravellings particularly explicit and has studied which trajectory-level quantities can distinguish them [15, 16]. The connection between distinguishable environment states and loss of interference is likewise standard [8, 17].

The present contribution is narrower and combinatorial. It packages the standard separations into a finite implication hierarchy, supplies elementary proof certificates for each failed implication, and shows in one model that the numerical supports differ substantially: 104, 120, and 224 transitions answer three differently typed questions. It also isolates cadence as a multitime datum, rather than treating a fixed point or a one-step transition rule as a clock. This complements the operational process-tensor viewpoint, in which a sequence of interventions is part of the object being reconstructed [11, 12].

Thus the novelty is not a new definition of instrument, dilation, or decoherence. It is a finite proof architecture in which each failed arrow has its own witness and the later failures cannot be hidden inside the familiar earlier ones.

The finite ledger language and the eight-register model used later arise from the author's finite-QEC substrate research programme. The record–response discipline was developed in Elliman [18], while the programme's code, crystallisation, ledger, and audit framework is set out in Elliman [19]. This provenance motivates the example; it is not a hidden premise of the results. The model is restated here in conventional operator language, and every count used in the argument is reproduced independently. Programme-level hypotheses about service clocks, cosmology, or gravitational normalisation are therefore listed explicitly as non-claims in Appendix B. That list records the boundary of this paper rather than disguising the origin of the construction.

3 Six layers of a finite record process

Let Z be a finite event set and let $\mathcal{H}_Z = \mathrm{span}\{|z\rangle : z \in Z\}$. An event may include both a system basis state and the address of an attempted operation. Let $r : Z \to \mathcal{R}$ be a proposed record map. Its spectral projectors are

$$Q_a = \sum_{z:r(z)=a} |z\rangle\langle z|, \qquad a \in \mathcal{R}, \qquad \sum_a Q_a = \mathbb{I}. \tag{1}$$

The map r says how event states are grouped into visible outcomes. It does not yet say that a measurement occurs.

Definition 1 (Effect and instrument)**.** A finite quantum instrument is a family of completely positive, trace-nonincreasing maps $\{\mathcal{I}_a\}_{a\in\mathcal{R}}$ such that $\sum_a \mathcal{I}_a$ is trace preserving. Its effects are $E_a = \mathcal{I}_a^*(\mathbb{I})$, so that

$$p(a|\rho) = \mathrm{Tr}[\mathcal{I}_a(\rho)] = \mathrm{Tr}(E_a\rho). \tag{2}$$

The effects answer the probability question. The maps $\mathcal{I}_a$ add the conditional output state. Different instruments can have the same effects.

Definition 2 (Exported record)**.** An outcome is *exported* when an indirect measurement realisation writes its value into an environmental pointer. For an event-basis transition T, write

$$V\,|z\rangle|0\rangle_E = |Tz\rangle|\eta_{r(z)}\rangle_E. \tag{3}$$

The export is projective at the record level when environmental states belonging to distinct records are orthogonal.

This definition does not require the entire unknown system state to be copied. Only an already selected commuting pointer variable is copied. That distinction is essential: arbitrary pure states cannot be cloned [20], and noncommuting mixed states cannot be broadcast [21]. Pointer records evade these no-go results because the record alphabet is orthogonal and classical.

Definition 3 (Cadence)**.** A cadence is a specification of the times, operation boundaries, or causal interfaces at which a record coupling such as Eq. (3) is invoked. A one-step channel does not by itself specify this schedule.

Finally, a continuously acting open-system model introduces a generator. We use

$$\mathcal{D}[A](\rho) = A\rho A^\dagger - \frac{1}{2}\{A^\dagger A, \rho\} \tag{4}$$

for a Lindblad dissipator. Writing down A does not fix its coefficient in the generator, still less show that the corresponding environmental monitoring is the one realised in nature.

4 Exact finite statements

4.1 Calibration fixes effects, not instruments

Theorem 4 (Exact finite calibration)**.** *Let E_a be an effect, $0 \leq E_a \leq \mathbb{I}$, calibrated on the event basis so that*

$$\langle z|E_a|z\rangle = \begin{cases} 1, & r(z) = a, \\ 0, & r(z) \neq a. \end{cases} \tag{5}$$

Then $E_a = Q_a$.

Proof. For a positive matrix E_a, every 2×2 principal minor is positive, and therefore

$$|(E_a)_{uv}|^2 \leq (E_a)_{uu}(E_a)_{vv}. \tag{6}$$

Every matrix element touching a zero diagonal entry consequently vanishes. Apply the same inequality to the positive matrix $\mathbb{I} - E_a$. On the block where E_a has unit diagonal, $\mathbb{I} - E_a$ has zero diagonal, so every off-diagonal entry in that block also vanishes. The only surviving entries are the unit diagonal entries selected by $r(z) = a$, which is precisely Eq. (1). □

Repeatability often supplies the physical reason for demanding exact calibration, but it is not needed in the algebraic proof once Eq. (5) is given. The theorem selects the visible observable. It says nothing about the conditional state within a record sector.

Corollary 5 (Approximate finite calibration)**.** *Let $d = |Z|$, $m_a = \operatorname{rank} Q_a$, and $n_a = d - m_a$. Suppose $0 \leq \epsilon \leq 1$ and the diagonal calibration errors obey*

$$1 - \epsilon \leq \langle z|E_a|z\rangle \leq 1 \quad \textit{when } r(z) = a, \qquad 0 \leq \langle z|E_a|z\rangle \leq \epsilon \quad \textit{otherwise.}$$

Then

$$\|E_a - Q_a\| \leq \|E_a - Q_a\|_{\mathrm{F}} \leq \sqrt{(m_a^2 + n_a^2)\epsilon^2 + 2m_a n_a \epsilon} \leq d\sqrt{\epsilon}. \tag{7}$$

Proof. Within the Q_a block, positivity of $\mathbb{I} - E_a$ bounds every entry of $E_a - Q_a$ by ϵ in magnitude. Within its complement, positivity of E_a gives the same bound. A cross-block entry has squared magnitude at most ϵ, again by positivity of E_a. Counting the two square blocks and the two cross blocks gives the Frobenius bound; the last inequality uses $\epsilon \le 1$. □

The square-root scaling cannot generally be replaced by a dimension-free linear bound. In two dimensions, take $Q = |0\rangle\langle 0|$ and $|\psi_\epsilon\rangle = \sqrt{1-\epsilon}|0\rangle + \sqrt{\epsilon}|1\rangle$. The effect $E = |\psi_\epsilon\rangle\langle\psi_\epsilon|$ has diagonal calibration error ϵ, but $\|E - Q\| = \sqrt{\epsilon}$.

Proposition 6 (Instrument nonuniqueness)**.** *Unless every Q_a has rank one and further state-update conditions are imposed, the effects $E_a = Q_a$ admit distinct instruments.*

Proof. The Lüders instrument

$$\mathcal{I}_a(\rho) = Q_a \rho Q_a \tag{8}$$

has effect Q_a. For any unitary U_a acting inside $Q_a\mathcal{H}_Z$, the instrument

$$\widetilde{\mathcal{I}}_a(\rho) = U_a Q_a \rho Q_a U_a^\dagger \tag{9}$$

has the same effect and generally a different post-measurement state. A dephasing channel, or any trace-preserving channel confined to the same record block, gives further examples. □

More generally, instruments with fixed effects form a convex set: convex mixtures preserve complete positivity, normalization, and the effects [4].

This familiar freedom is included because it prevents the first common over-reading: a uniquely calibrated outcome map is not a uniquely selected measurement dynamics.

4.2 Transition support does not fix environmental export

Suppose T is a reversible permutation of the event basis. For q record labels, choose normalised environment states with Gram matrix

$$G_{ab}(\mu) = (1-\mu)\delta_{ab} + \mu, \qquad 0 \le \mu \le 1. \tag{10}$$

The eigenvalues are $1-\mu$, with multiplicity $q-1$, and $1+(q-1)\mu$, with multiplicity one. Hence the matrix is positive semidefinite throughout the stated interval and such pointer states exist.

Theorem 7 (Export non-identifiability)**.** *The basis transition $z \mapsto Tz$, its transition support, and all event-basis output populations are compatible with a continuous family of record exports ranging from no which-record information to an orthogonal projective record.*

Proof. Use Eq. (3) with environmental inner products given by Eq. (10). The system basis transition is T for every value of μ, and basis populations are unchanged. At $\mu = 1$, all record states coincide and no which-record information is exported. For $0 < \mu < 1$, the environment carries partial information. At $\mu = 0$, distinct record states are orthogonal and the export is projective. Thus the transition table and its classical population data do not select one member of the family. □

The corresponding reduced channels differ on coherence. This is not a loophole; it provides the discriminator. Take two event states $|z_0\rangle$ and $|z_1\rangle$ carrying different records and prepare $|\psi\rangle = (|z_0\rangle + |z_1\rangle)/\sqrt{2}$. Apply the transition and possible record coupling, undo the system transition, and project onto $|\psi\rangle$. Direct calculation gives

$$P_{\text{return}} = \frac{1 + \operatorname{Re}\mu}{2}. \tag{11}$$

The values are 1, 3/4, and 1/2 for $\mu = 1, 1/2, 0$, respectively. This is the finite analogue of reading which-path information through visibility [17].

4.3 Endpoint data do not fix cadence

Theorem 8 (Endpoint cadence boundary)**.** *Suppose all currently accessible effects factor through a final state x_T, and two microscopic histories have the same x_T. No statistic built only from those effects can decide how many intermediate record couplings occurred.*

Proof. For every accessible effect F, both histories assign the same probability $\mathrm{Tr}(F\rho_T)$ because they produce the same final state. The number and location of intermediate interventions are therefore not identifiable from the stated data. They become identifiable only when a multitime intervention or a retained history register is added to the access class. □

The theorem is elementary, but it rules out a surprisingly common inference: the fixed point of a channel is rate independent, so it cannot determine how often a monitor was invoked. A cadence claim is a claim about a quantum stochastic process, not merely its endpoint channel. Process tensors provide the general operational setting for precisely this distinction [11, 12].

In a finite XOR monitor the boundary is especially sharp. If a pointer stores an increment r_t by XOR, two identical increments return one reused pointer to blank. Two fresh pointer factors retain two non-idle records. Both descriptions can have the same final system state and the same final net syndrome.

4.4 A lawful operator does not fix an autonomous rate

Let J_p be the jump support associated with one declared invocation and let L_p be a larger, independently lawful jump operator. The family

$$\mathcal{G}_{\gamma,\kappa} = \gamma \sum_p \mathcal{D}[J_p] + \kappa \sum_p \mathcal{D}[L_p], \qquad \gamma, \kappa \geq 0, \tag{12}$$

is a valid GKSL generator for every nonnegative pair [6, 7].

Proposition 9 (Independent dissipator coefficient)**.** *The transition table of J_p does not determine whether $\kappa = 0$, or the value of κ when it is nonzero.*

Proof. Varying κ in Eq. (12) leaves the definition and transition support of every J_p unchanged. Complete positivity requires only nonnegativity. An additional action, microscopic bath model, or experimental rate is therefore needed to select κ. □

If a basis source $|s_K\rangle$ satisfies $J_p|s_K\rangle = 0$ but $L_p|s_K\rangle = |t_K\rangle$, then

$$\frac{\mathrm{d}}{\mathrm{d}t}\langle t_K|\rho(t)|t_K\rangle\bigg|_{t=0} = \kappa \tag{13}$$

for the corresponding normalised trial. A complementary-only source thus measures the independent coefficient directly.

4.5 Edge support does not fix the unravelling

Suppose $L = J + K$ and J, K have orthogonal source supports. The one-jump description L and the two-jump description $\{J, K\}$ have the same union of basis edges. They nevertheless act differently on coherence.

Proposition 10 (Polarity record discriminator)**.** *Let $J|s_J\rangle = |t_J\rangle$, $K|s_K\rangle = |t_K\rangle$, with the crossed actions zero. For the input $(|s_J\rangle + |s_K\rangle)/\sqrt{2}$, conditioning on an unresolved L jump leaves the output cross term with coefficient $1/2$. Separately recorded J and K jumps leave coefficient zero after the polarity label is discarded.*

Proof. The unresolved jump produces

$$L|\psi\rangle\langle\psi|L^\dagger = \frac{1}{2}(|t_J\rangle + |t_K\rangle)(\langle t_J| + \langle t_K|), \tag{14}$$

which contains both cross terms. The resolved operation produces

$$J|\psi\rangle\langle\psi|J^\dagger + K|\psi\rangle\langle\psi|K^\dagger = \frac{1}{2}|t_J\rangle\langle t_J| + \frac{1}{2}|t_K\rangle\langle t_K|, \tag{15}$$

which contains none. □

This is a finite, linear discriminator between two unravellings of the same edge union. It is consistent with the broader trajectory literature, where the average channel does not in general exhaust the operational content of a monitored process [13, 15, 16].

5 An eight-register worked example

We now apply the hierarchy to the eight-register local cell used in the author's finite-QEC substrate programme [18, 19]. The construction was not invented solely as a toy for this paper, and its programme provenance is part of the record. It is not assumed here to be a complete microscopic theory. Its value for the present question is that all sets are finite and every support claim can be checked exhaustively.

The construction was selected because one finite object exhibits all of the later separations simultaneously. Its transfer alphabet has a nontrivial mismatch with its physical exit alphabet; one transfer address is exit-silent; and positive-control, complementary-control, and polarity-complete supports are all different. The dual cube record complex labels all eight transfer addresses while retaining a transparent parity-incidence census. The example is therefore diagnostic by design, rather than an arbitrary enumeration.

5.1 Register, code, and transition family

Let the register basis be the 256 bit strings

$$x = (G_0, G_1, L_Q, C_0, C_1, I_3, \chi, W) \in \mathbb{F}_2^8. \tag{16}$$

The physical subset P is defined by

$$\neg(G_0 G_1), \qquad W = \chi, \qquad (L_Q = 0) \Longleftrightarrow (C_0 = 0 \text{ and } C_1 = 0). \tag{17}$$

Exactly 48 strings satisfy Eq. (17); the complement Q has 208 strings. Let Π_P and Π_Q be their projectors.

For target address $p \in \mathbb{Z}_8$, let X_p flip target bit p, and let D_p project onto control bit

$$c(p) = p - 3 \pmod 8 \tag{18}$$

having value one. The addressed positive-control operation is

$$C_p = (\mathbb{I} - D_p) + D_p X_p. \tag{19}$$

The eight target addresses are covered bijectively. On the 48-state code, the family has 188 active moves: 104 leave P, while 84 remain internal.

5.2 The truthful transfer record

The record complex is the cube graph Q_3, used here as a dual face-record complex rather than as a claim about a literal cubic matter cell. Its eight vertices label the eight register addresses, and its twelve edges are parity checks. If $\sigma(x) \in \mathbb{F}_2^{12}$ lists the edge parities, the record of an addressed event (x, p) is

$$r(x, p) = \sigma(C_p x) \oplus \sigma(x). \tag{20}$$

The $2048 = 256 \times 8$ event states have exactly nine record values. There are 1024 idle events and 128 moved events at each of eight addresses. Every moved event carries the weight-three incidence word of the target vertex. The record projectors therefore have ranks

$$\operatorname{rank} Q_{\text{idle}} = 1024, \qquad \operatorname{rank} Q_p = 128 \quad (p = 0, \dots, 7). \tag{21}$$

Theorem 4 now has a concrete consequence. Any exactly calibrated effect which claims to measure the syndrome increment in Eq. (20) must equal the corresponding Q_r. A hidden increment, a spurious emitted record, or an address permutation fails on exactly 1024 moved or idle event states. Yet Proposition 6 still applies: the visible effects do not select a unique post-state channel.

The distinction between the transfer alphabet and the exit alphabet matters. All eight addresses occur in Eq. (20). The validity condition Eq. (17), however, does not depend on I_3. Flipping I_3 can therefore never take a physical string out of P. It is a live transfer address but a silent exit address.

5.3 The $J/K/L$ split

Define the polarity-complete exit operator

$$L_p = \Pi_{\mathsf{Q}} X_p \Pi_{\mathsf{P}}. \tag{22}$$

Its positive- and negative-control parts are

$$J_p = L_p D_p, \qquad K_p = L_p(\mathbb{I} - D_p), \qquad L_p = J_p + K_p. \tag{23}$$

The source supports of J_p and K_p are disjoint. Exhaustive counts are shown in Table 1.

Table 1: Exit-transition census on the 48-state physical code. The control column gives the bit controlling the positive branch whose target is the address in the first column. The zero I_3 row is structural because the validity predicate is independent of I_3.

target	control	J_p	K_p	L_p
G_0	I_3	8	8	16
G_1	χ	8	8	16
L_Q	W	24	24	48
C_0	G_0	8	16	24
C_1	G_1	8	16	24
I_3	L_Q	0	0	0
χ	C_0	24	24	48
W	C_1	24	24	48
total		104	120	224

These are not three estimates of one quantity. They are supports of three operators. One literal positive invocation of C_p can produce only the J_p exits, hence 104 transitions. The remaining 120 exits require the complementary control polarity. The full bit-flip exit family contains their union.

The latch location makes the distinction operational. Define the complementary controlled operation

$$\overline{C}_p = D_p + (\mathbb{I} - D_p)X_p. \tag{24}$$

Because the control and target bits are different,

$$\overline{C}_p C_p = C_p \overline{C}_p = X_p. \tag{25}$$

Exactly one of the two controlled substeps moves any basis state. A latch immediately after C_p therefore records the 104-edge J support. A latch only after both polarities records the 224-edge L support. The latter is a two-test macrocycle, not one invocation of the declared positive operator. This is a typing statement, independent of circuit depth.

5.4 An exact cadence witness

The code contains the two-step trajectory

$$64 \longrightarrow 72 \longrightarrow 64. \tag{26}$$

Both transfers carry the same non-idle address record. The state types are $\mathsf{P} \to \mathsf{Q} \to \mathsf{P}$. Coupling the two transfers to fresh environmental pointer factors E_1, E_2 retains two non-idle entries. A single reused XOR pointer E returns to its blank state, and a delayed final syndrome read reports idle. The declared exit instrument sees one $\mathsf{P} \to \mathsf{Q}$ jump rather than two transfer records.

This witness prevents three identifications at once:

$$\text{transfer count} \neq \text{net syndrome count} \neq \text{exit-jump count}. \tag{27}$$

No fixed-point theorem can repair the mismatch. Fixed points constrain the long-time state conditional on a channel; they do not decide where along a microscopic history a separate environment was coupled. Results on fixed algebras and stationary states [22, 23] remain fully applicable once the channel and its cadence have been selected.

In process-tensor language, the fresh-pointer and reused-pointer constructions are distinct two-time processes. An intervention which reads or swaps the first pointer produces different statistics, so the corresponding multitime process-tensor contraction is nonzero. Contracting with no intermediate intervention gives the same endpoint system state. The trajectory in Eq. (26) is therefore an explicit finite witness of multitime information which an endpoint map cannot reconstruct [11, 12].

5.5 Three separate discriminators

The example realises the general tests without introducing fitted parameters.

1. **Record export.** At address G_0, use the idle source 0 and the crossing source $68 \to 196$. Reinterference after the inverse positive operation obeys Eq. (11). It separates coherent, partial, and projective export.
2. **Independent L dynamics.** The state 64 is a complementary-only G_0 source, with $X_{G_0}64 = 192$. The initial population rate of 192 is κ in Eq. (12). It vanishes for J-only dynamics.

3. **Polarity resolution.** A coherent superposition of one J_p-source and one K_p-source, at fixed p, retains output cross coherence $1/2$ under one unresolved L_p jump and zero under separately recorded J_p, K_p jumps.

The tests answer different questions. Observing a K-only exit would not show that every J-capable invocation exports a record. Observing complete edge support would not show that J/K polarity is erased. And a projective export at one tested boundary would not prove a universal one-record-per-step cadence.

6 A sufficient open-system completion, and its price

The preceding results identify what is missing. They also identify a simple sufficient completion. For each elementary invocation t, supply a fresh record factor R_t, initially blank, and apply

$$|z_t\rangle|0\rangle_{R_t} \longmapsto |T_t z_t\rangle|r_t\rangle_{R_t}. \tag{28}$$

Require distinct non-idle records to be orthogonal and require R_t, together with any purifying side information, to leave the algebra of future system operations before the next invocation. Tracing it out gives

$$\rho \longmapsto \sum_a Q_a \rho Q_a. \tag{29}$$

This is a collision-model construction with one fresh ancilla per event [10]. It is sufficient to prevent coherent uncopying and to retain multiplicity through time. It is not forced by the system transition operator.

A local history-tape action makes the same point. Let $h_{t,n}$ hold the record at history depth $n \geq 0$, and introduce multipliers $\lambda_{t,n}$. The first-order constraint action

$$S_+ = \sum_t \lambda_{t+1,0} \cdot (h_{t+1,0} - j_t) + \sum_{t,n\geq 0} \lambda_{t+1,n+1} \cdot (h_{t+1,n+1} - h_{t,n}) \tag{30}$$

injects the supplied record j_t and shifts all earlier records away from the system. Within the binary, translation-invariant, radius-one linear class, outward shift is the unique one-step rule with this property. But spatial reflection produces an equally local incoming action S_-. The action transports whatever source sequence it is given; it does not decide whether j_t is supplied after every positive invocation, after a polarity-complete macrocycle, or only after a delayed read.

There is also a finite-capacity obstruction. A source alphabet with q distinguishable records has q^T formal histories of length T. No fixed finite carrier can preserve all of them indefinitely without returning, overwriting, or exporting information to a larger environment. The need for a fresh mode or entropy-carrying bath is therefore real. Its physical origin, chirality and coupling remain additional microscopic inputs.

The completion has to specify five things, not merely one:

1. the environment couples to the calibrated record r, rather than an arbitrary pointer;
2. distinct records generate orthogonal environmental states;
3. a fresh factor is supplied at each selected invocation boundary;
4. that factor cannot coherently return before record commitment; and
5. any additional L dissipator, including its rate and its treatment of J/K polarity, is fixed independently.

Without those clauses Eq. (28) is an implementation proposal, not a deduction from the closed transition grammar.

7 Discussion

7.1 What the finite hierarchy adds

The hierarchy is useful because each arrow can fail for a different reason:

$$\text{support} \not\Rightarrow \text{effect} \not\Rightarrow \text{instrument} \not\Rightarrow \text{export} \not\Rightarrow \text{cadence} \not\Rightarrow \text{autonomous rate}. \tag{31}$$

The first two failures are standard quantum measurement theory. The worked example makes the later failures concrete. It has a uniquely calibrated nine-sector transfer observable, yet the exit instrument has only seven live addresses. Its positive invocation and polarity-complete macrocycle differ by 120 exit transitions. None of these facts decides the environmental overlap μ, the independent rate κ, or whether polarity is resolved.

This is why the term *record* should be reserved for a stronger statement than *transition*. A record is relational: it specifies a correlation between system alternatives and a degree of freedom which remains available outside the subsequent coherent system algebra. Decoherence supplies the loss of local interference; redundant environmental proliferation supplies public accessibility [8, 9]. The finite hierarchy does not compete with that account. It prevents a discrete model from claiming its conclusion merely because it contains a labelled edge.

7.2 A restrained epistemological reading

One may describe a physical theory as a grammar of questions and admissible responses. In that language, an instrument is the rule by which an intervention returns both an outcome and a changed state. A preparation acts like a constraint on the responses which may subsequently be written; a record is the part of the response that has become stable and externally available.

The formal results do not settle whether an unrecorded quantum state is an underlying reality, information, disposition, or merely part of a predictive calculus. They establish a more modest point. The structure before record export can affect later interference even though it is not yet a public record, and different environmental couplings can turn the same basis transition into different observable processes. The hierarchy is therefore compatible with more than one interpretation of quantum mechanics. It is a constraint on model construction, not an ontological conclusion.

7.3 Limitations

Five limitations should be explicit.

First, all theorems are finite-dimensional. This is sufficient for the worked example and for finite experimental records, but domains and continuity require additional care in infinite dimension.

Second, exact calibration is an idealisation. Corollary 5 supplies a finite robustness bound, but its rank dependence and $\sqrt{\epsilon}$ scaling show that diagonal calibration alone does not give a dimension-free linear estimate.

Third, the constant-overlap family in Eq. (10) is a witness family, not a classification of all environmental realisations.

Fourth, the history action Eq. (30) is conditional. No substrate-native principle in the worked example selects its outgoing chirality or supplies a fresh mode at each invocation.

Fifth, the eight-register construction does not derive that nature implements the independent L dissipator, nor that every positive invocation exports a projective record. Those are exactly the two physical questions the paper leaves open.

8 Conclusion

The phrase "a transition occurred" is ambiguous unless the object and the access class are named. It may refer to an allowed basis edge, a visible outcome, a conditional state update, a distinguishable environmental record, an invocation in a multitime history, or a term in an autonomous dissipative generator. Standard quantum theory has separate mathematical objects for all of them. The finite hierarchy developed here turns that vocabulary into a sequence of exact implication tests.

In the worked register model, the tests expose a clean structure. The truthful transfer observable has idle plus eight address sectors. The physical exit code has seven live addresses because I_3 is exit-silent. One positive invocation has 104 exit transitions, its complementary control sector has 120, and their polarity-complete union has 224. The environmental record, the invocation cadence, the independent dissipator rate, and the polarity unravelling remain separate choices. Return visibility, a complementary-only rate, and conditioned coherence test those choices separately.

The result does not solve the measurement problem and does not select a new interpretation. It does something more local: it says what a finite model must add before its transition grammar becomes a physical record law. That is a smaller claim, but one that can be checked exactly.

A Independent finite census

For completeness, the central support counts can be reproduced without any matrix diagonalisation. Enumerate all $x \in \{0, \ldots, 255\}$, decode bits most-significant first in the order $(G_0, G_1, L_Q, C_0, C_1, I_3, \chi, W)$, and retain x precisely when Eq. (17) holds. For each target p, set $c = p - 3 \pmod 8$ and $y = x \oplus 2^{7-p}$. If $y \notin \mathsf{P}$, count the edge in J_p when bit c of x is one and in K_p otherwise. The union gives L_p.

The resulting tuples $(|J_p|, |K_p|, |L_p|)$, in the address order used in Table 1, are

$$(8,8,16),\ (8,8,16),\ (24,24,48),\ (8,16,24),\ (8,16,24),\ (0,0,0),\ (24,24,48),\ (24,24,48). \quad (32)$$

They sum to $(104, 120, 224)$. Because the validity predicate contains no I_3, the zero in the sixth entry is structural rather than accidental.

B Reproducibility and claim boundary

The finite calculations were implemented as exact enumeration gates with named negative controls. The public archive accompanying this report includes independent checks of: the 48/208 code partition; all 2048 addressed event states; the nine record-sector ranks; the $J/K/L$ supports; the polarity-latch identity; the cadence witness; positivity of the environmental Gram family; and the three discriminator formulae. The portable entry point is `verify_transition_to_record.py`. The archive also contains its machine-readable JSON summary and CSV check ledger, so the finite verification surface is deposited with the paper rather than promised for a later date.

The executable checks establish finite algebraic statements. They do not establish that the environment of the worked model is realised in nature. In particular, the following claims are not made:

- one projective record is exported after every elementary invocation;
- the polarity-complete macrocycle is physically elementary;
- the independent L_p dissipator has nonzero rate;

- a universal service clock or 1/28 scheduler follows from the register operator; or
- any downstream gravitational or cosmological normalisation is thereby derived.

Data and code availability

No empirical data were analysed. All numerical statements are finite exact enumerations or rational/algebraic evaluations. The public Zenodo record for this report, doi:10.5281/zenodo.22151372, contains the PDF, LaTeX source, bibliography, the standard-library verifier `verify_transition_to_record.py`, and machine-readable JSON and CSV verification reports. The deposited files are therefore sufficient to rerun the finite census without access to private data or software.

Funding

This research received no external funding.

Competing interests

The author declares no competing interests.

Author contributions

D.E. conceived the study, developed and checked the finite constructions, interpreted the results, and wrote the manuscript.

Acknowledgements

The author used OpenAI Codex as an interactive aid within a consistency framework that tracked claims, sources, executable checks, negative controls, and explicit non-claims. The system assisted with exploratory code, adversarial test design, literature-search support, and language editing. The author selected the claims, inspected the derivations and cited sources, reran the calculations, resolved inconsistencies, and accepts responsibility for the content of the manuscript.

References

[1] E. B. Davies and J. T. Lewis. An operational approach to quantum probability. *Communications in Mathematical Physics*, 17(3):239–260, 1970. doi: 10.1007/BF01647093.

[2] Karl Kraus. *States, Effects, and Operations: Fundamental Notions of Quantum Theory*, volume 190 of *Lecture Notes in Physics*. Springer, Berlin, 1983. doi: 10.1007/3-540-12732-1.

[3] Masanao Ozawa. Quantum measuring processes of continuous observables. *Journal of Mathematical Physics*, 25(1):79–87, 1984. doi: 10.1063/1.526000.

[4] Paul Busch, Pekka Lahti, Juha-Pekka Pellonpää, and Kari Ylinen. *Quantum Measurement*. Theoretical and Mathematical Physics. Springer, Cham, 2016. doi: 10.1007/978-3-319-433 89-9.

[5] W. Forrest Stinespring. Positive functions on C^*-algebras. *Proceedings of the American Mathematical Society*, 6(2):211–216, 1955. doi: 10.1090/S0002-9939-1955-0069403-4.

[6] Vittorio Gorini, Andrzej Kossakowski, and E. C. G. Sudarshan. Completely positive dynamical semigroups of N-level systems. *Journal of Mathematical Physics*, 17(5):821–825, 1976. doi: 10.1063/1.522979.

[7] Göran Lindblad. On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48(2):119–130, 1976. doi: 10.1007/BF01608499.

[8] Wojciech H. Zurek. Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75:715–775, 2003. doi: 10.1103/RevModPhys.75.715.

[9] Harold Ollivier, David Poulin, and Wojciech H. Zurek. Environment as a witness: Selective proliferation of information and emergence of objectivity in a quantum universe. *Physical Review A*, 72:042113, 2005. doi: 10.1103/PhysRevA.72.042113.

[10] Francesco Ciccarello, Salvatore Lorenzo, Vittorio Giovannetti, and G. Massimo Palma. Quantum collision models: Open system dynamics from repeated interactions. *Physics Reports*, 954:1–70, 2022. doi: 10.1016/j.physrep.2022.01.001.

[11] Felix A. Pollock, César Rodríguez-Rosario, Thomas Frauenheim, Mauro Paternostro, and Kavan Modi. Non-Markovian quantum processes: Complete framework and efficient characterization. *Physical Review A*, 97:012127, 2018. doi: 10.1103/PhysRevA.97.012127.

[12] Simon Milz and Kavan Modi. Quantum stochastic processes and quantum non-Markovian phenomena. *PRX Quantum*, 2:030201, 2021. doi: 10.1103/PRXQuantum.2.030201.

[13] Howard M. Wiseman and Gerard J. Milburn. *Quantum Measurement and Control*. Cambridge University Press, Cambridge, 2010. doi: 10.1017/CBO9780511813948.

[14] Alberto Barchielli and Matteo Gregoratti. *Quantum Trajectories and Measurements in Continuous Time: The Diffusive Case*, volume 782 of *Lecture Notes in Physics*. Springer, Berlin, 2009. doi: 10.1007/978-3-642-01298-3.

[15] Dariusz Chruściński, Kimmo Luoma, Jyrki Piilo, and Andrea Smirne. How to design quantum-jump trajectories via distinct master equation representations. *Quantum*, 6:835, 2022. doi: 10.22331/q-2022-10-13-835.

[16] Eloy Piñol, Themis K. Mavrogordatos, Dustin Keys, Romain Veyron, Piotr Sierant, Miguel Angel García-March, Samuele Grandi, Morgan W. Mitchell, Jan Wehr, and Maciej Lewenstein. Telling different unravelings apart via nonlinear quantum-trajectory averages. *Physical Review Research*, 6:L032057, 2024. doi: 10.1103/PhysRevResearch.6.L032057.

[17] Berthold-Georg Englert. Fringe visibility and which-way information: An inequality. *Physical Review Letters*, 77:2154–2157, 1996. doi: 10.1103/PhysRevLett.77.2154.

[18] David Elliman. Records and responses: Dressing-blind theorems for monitored quantum instruments. Zenodo preprint, 2026. First public version, 4 July 2026.

[19] David Elliman. Foundations and methodology for a finite-QEC substrate: Code, crystallisation, ledgers, and audit protocol. Zenodo preprint, 2026. Canon-synchronised version dated 1 July 2026.

[20] William K. Wootters and Wojciech H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982. doi: 10.1038/299802a0.

[21] Howard Barnum, Carlton M. Caves, Christopher A. Fuchs, Richard Jozsa, and Benjamin Schumacher. Noncommuting mixed states cannot be broadcast. *Physical Review Letters*, 76:2818–2821, 1996. doi: 10.1103/PhysRevLett.76.2818.

[22] David E. Evans. Irreducible quantum dynamical semigroups. *Communications in Mathematical Physics*, 54(3):293–297, 1977. doi: 10.1007/BF01614091.

[23] Alberto Frigerio. Stationary states of quantum dynamical semigroups. *Communications in Mathematical Physics*, 63(3):269–276, 1978. doi: 10.1007/BF01196936.