Pruning as Rendering: A Decoherent-Histories Reformulation of the Observer-Rendering Criterion in Simulation-Theoretic Quantum Mechanics

Ian Staley

Independent Researcher

ORCID: 0009-0000-8592-3186

Correspondence: ian.t.staley@gmail.com

May 16, 2026

Abstract

This paper develops a formal reconstruction, within the Gell-Mann–Hartle decoherent-histories framework, of the Campbell–Owhadi observer-rendering criterion familiar from simulation-theoretic quantum mechanics. Under an explicit records-theoretic reading of "information available to an observer," the rendering criterion corresponds to a decoherence-functional condition weighted by algorithmic-information distance on observer records. The construction is metaphysically neutral: the parameter K may equally be read as a literal simulator budget or as a resource bound on observer-internal record consolidation. The author defines a pruning admissibility functional $A(\alpha \mid O, K)$ and establishes a correspondence in which Campbell–Owhadi rendering events coincide with transitions across the pruning-admissibility boundary. The framework recovers standard decoherent-histories statistics (Born rule, medium decoherence) in the limit $K \to \infty$. For finite K, a heuristic derivation yields a complementary-error-function form, saturating at the contextuality/Bell margin, for the predicted sub-Tsirelson and sub-Klyachko–Can–Binicioğlu–Shumovsky deviations and for the threshold structure in delayed-choice eraser visibility. Two photonic protocols on existing single-photon platforms are proposed, at estimated costs of approximately $180k each in an existing facility or $200–400k from scratch.

Keywords: decoherent histories; simulation hypothesis; algorithmic information; quantum contextuality; observer rendering

1. Introduction

The thesis that the empirical content of quantum mechanics might be intelligible through a computational or rendering metaphor has produced two largely independent research programmes, both partly downstream of Bostrom's influential simulation argument [1] and Chalmers' "simulation realism" [2]. The first, which I will call the empirical-signatures cluster, looks for observable consequences of a discrete computational substrate. Beane, Davoudi, and Savage [3] extrapolated current lattice-quantum-chromodynamics technology to argue that a cubic space-time lattice would manifest in the rotational-symmetry-breaking distribution of the highest-energy cosmic rays, deriving the bound $b^{-1} \gtrsim 10^{11}$ GeV on the inverse lattice spacing from the Greisen–Zatsepin–Kuzmin cutoff. Whitworth [4] proposed a "virtual reality" rendering of physical law in which information bandwidth and refresh constraints play the explanatory role of relativistic kinematics. Vopson [5] introduced a "second law of infodynamics" stating that information entropy of information-bearing states must remain constant or decrease over time — a striking inversion of the thermodynamic second law, motivated by the simulation hypothesis; an earlier Vopson and Lepadatu paper [6] introduced the information-dynamics framework underlying that claim. Indset, Neukart, Pflitsch and Perelshtein [7] surveyed observational and experimental constraints on a quantum simulation. Lloyd's earlier "computational capacity of the universe" bounds [8] (10^{120} ops, 10^{90} bits since the Big Bang) gave the cluster its quantitative skeleton.

The second, observer-rendering cluster has its centre of gravity in Campbell, Owhadi, Sauvageau and Watkinson's 2017 paper "On Testing the Simulation Theory" [9]. Its operative principle is computational economy: a finite-resource simulator should, like a video-game engine, render content only at the moment that information becomes available for observation by a player. The authors propose delayed-choice and Mach–Zehnder variants in which rendering events are detected by their downstream consequences. Owhadi's follow-up work has refined the experimental geometry but has continued to treat the rendering criterion as an informal heuristic.

Both clusters speak the language of information and selection. Yet to date neither has engaged the most developed formal machinery for selecting temporal sequences within standard quantum theory: the consistent / decoherent-histories programme of Griffiths, Omnès, Gell-Mann and Hartle [10,11,12,13,14,15]. The present development follows the Gell-Mann–Hartle lineage [10,11], which is more easily married to records-theoretic apparatus [16] and to algorithmic-information generalisations. This omission is striking because the histories formalism already addresses, with mathematical precision, the question that the rendering criterion asks heuristically: under what conditions can a coarse-graining of a quantum process be regarded as defining a sequence of facts with well-defined probabilities?

The present paper closes this gap. The central thesis is that the Campbell–Owhadi rendering criterion is structurally equivalent to a pruning condition on histories, expressed as a decoherence-functional constraint weighted by algorithmic information relative to an observer's records. The paper extends the program of "informational pruning" developed in [17] and connects it both upstream (to Müller's algorithmic-information reconstruction [18]) and downstream (to concrete experimental tests).

What this paper claims.

(i) Under an explicit records-theoretic reading of "information availability," the rendering and decoherent-histories programmes admit a precise translation between them. (ii) Algorithmic-information bounds on observer records suggest empirically distinguishable predictions in contextuality and delayed-choice experiments. (iii) The framework is empirically engageable on current photonic and ion-trap platforms.

What this paper does not claim.

(i) That a literal computational simulation underwrites physical reality. The framework is neutral between simulation realism and a reading on which K simply parametrises the resolution of observer-internal records. (ii) That the standard predictions of quantum mechanics are violated in the $K \to \infty$ limit; the framework reduces to ordinary decoherent histories there. (iii) That Vopson's second law of infodynamics is established; the paper shows only that if it holds, it is a specific monotonicity claim on K which the present framework formalises.

The paper proceeds as follows. Section 2 reviews the three streams that the bridge requires. Section 3 states and sketches the correspondence theorem and derives empirical predictions. Section 4 reads existing empirical work in pruning language and proposes two concrete experimental protocols. Section 5 situates the framework against QBism, relational quantum mechanics, the Wolfram–Gorard multiway formalism, and Deutsch-style many-worlds computational arguments. Section 6 addresses objections. Three appendices contain the formal statement of the theorem, a worked Stern–Gerlach example, and a resource-budget estimate for the proposed experiments.

2. Background

2.1 Decoherent histories essentials

A history in the Gell-Mann–Hartle formalism is a time-ordered sequence of projection operators acting on an initial density matrix ρ. For an exhaustive and exclusive family of projectors $\{P^k\_\alpha_k\}$, satisfying

$$P^k\_\alpha P^k\_\beta = \delta\_\alpha\beta P^k\_\alpha, \quad \Sigma\_\alpha P^k\_\alpha = \mathbb{1}, \quad (2.1)$$

at each time $t_1 < t_2 < \ldots < t_n$, the class operator of the history $\alpha = (\alpha_1, \ldots, \alpha_n)$ is

$$C\_\alpha = P^n\_\alpha_n(t_n) \cdots P^1\_\alpha_1(t_1), \quad (2.2)$$

where Heisenberg projectors are $P^k\_\alpha(t_k) = U(t_k, t_0)^\dagger P^k\_\alpha U(t_k, t_0)$. The decoherence functional on pairs of histories is

$$D(\alpha, \alpha') = Tr[C\_\alpha \rho C\dagger\_\alpha']. \quad (2.3)$$

A family $\{\alpha\}$ satisfies the weak consistency condition (Griffiths–Omnès) when $\mathrm{Re}\, D(\alpha, \alpha') = 0$ for $\alpha \neq \alpha'$; the stronger condition usually called medium decoherence (Gell-Mann–Hartle) requires the full off-diagonal $D(\alpha, \alpha') = 0$ for $\alpha \neq \alpha'$. The present paper uses the medium-decoherence condition unless otherwise noted, since the admissibility functional below is most cleanly stated in those terms. Under either condition the diagonal $D(\alpha, \alpha)$ inherits the formal properties of a probability measure and the Born rule for history probabilities is recovered:

$$p(\alpha) = D(\alpha, \alpha) = Tr[C\_\alpha \rho C\dagger\_\alpha]. \quad (2.4)$$

Strong decoherence, introduced by Gell-Mann and Hartle [19], is the additional requirement that at each step there exist commuting "generalised records" $R^k\_\alpha_k$ on the environment whose values are perfectly correlated with the projector outcomes, i.e. $R^k\_\alpha_k \rho \approx P^k\_\alpha_k \rho$. Strong decoherence implies medium decoherence but not conversely. The relationship to environmentally induced superselection (einselection) was clarified by Zurek [20]. Halliwell's records framework [16] makes this concrete: he shows that when a set of histories decoheres, there exists a decomposition of the total Hilbert space $\mathscr{H} = \mathscr{H}\_S \otimes \mathscr{H}\_R$ such that the records subspace $\mathscr{H}\_R$ carries an orthogonal family of "record projectors" $\{P\_\alpha\}$ satisfying

$$P\_\alpha C\_\alpha \rho \approx C\_\alpha \rho, \quad P\_\alpha C\_\alpha' \rho \approx 0 \text{ for } \alpha' \neq \alpha. \quad (2.5)$$

Halliwell formalises the information-theoretic conjecture that the bit-cost of describing a decoherent set is approximately the bit-cost of information dissipated to the environment in the coarse-graining process [21]; the logical and quantum-temporal structure of the histories approach is set out by Isham [22]. Hartle's later "quasiclassical realms" papers [13] extend this picture to the cosmological case, where the IGUS (information-gathering and -utilising system) is the structural locus of records.

2.2 The Campbell–Owhadi rendering criterion

Campbell, Owhadi, Sauvageau and Watkinson [9] operationalise the heuristic that a finite simulator should, like a video-game engine, render content only when an observer-player will obtain information from it. Their working principle is:

> *"If the system performing the simulation is finite (i.e. has limited resources), then to achieve low computational complexity, such a system would, as in a video game, render content (reality) only at the moment that information becomes available for observation by a player and not at the moment of detection by a machine (that would be part of the simulation and whose detection would also be part of the internal computation performed by the Virtual Reality server before rendering content to the player)."*

The proposed signature is operational: in a delayed-choice quantum-eraser variant [23,24], the experimenter manipulates whether path-which-way information ever becomes available to a downstream "player" record, controlling whether the interference pattern is rendered. Owhadi's subsequent commentary and Caltech-hosted preprints have proposed Mach–Zehnder geometries with controllable photon-counting downstream from the second beam splitter.

The criterion has two intuitive parts: (i) a rendering trigger — some operationally specifiable event that pulls a system from the "unrendered" to the "rendered" regime; (ii) an information-availability predicate — a precise statement of what it means for the relevant information to be "available to a player." The 2017 paper leaves this predicate at the level of operational intuition: information is available when a record exists that would, if interrogated, yield the outcome. The unresolved formal question — and the question this paper answers — is: what mathematical object selects which information is rendered? I will argue that this object is a decoherence functional weighted by an algorithmic-information bound on the observer's records.

2.3 Algorithmic information and observer-state programmes

The third strand is Müller's algorithmic reformulation of quantum theory [18]. Müller's single postulate is that universal induction — the Solomonoff prior on computable continuations [25] — assigns the chances of what any observer sees next, given the observer's current state. Formally, given an observer (or "self") state $s \in S$ and a successor state s', the conditional chance is

$$P(s' \mid s) = \sum_{p : U(p, s) = s'} 2^{-\ell(p)}, \qquad (2.6)$$

where U is a fixed universal monotone Turing machine and $\ell(p)$ is the length of program p. Müller proves that observers under this postulate will typically infer the existence of an apparently objective external world with simple, computable laws. The reconstruction of quantum theory's measurement postulates from operational symmetry assumptions [26] establishes the complementary point that the Born rule itself follows from minimal operational postulates plus the algorithmic-probability framework. Müller's 2021 SciPost lecture notes [27] furnish the review-level synthesis.

The crucial concept for the present paper is the observer state: a Markov-shielded informational description of the observer's records, sufficiently rich to determine — modulo the universal prior — the conditional probabilities of next experiences. The Halliwell records subspace $\mathscr{H}_R$ and Müller's observer state S coincide in operational content but differ in formalism: Halliwell uses Hilbert-space projectors; Müller uses bit strings under a universal Turing machine. The unification is straightforward: a record projector in $\mathscr{H}_R$ defines an equivalence class of bit-string descriptions of "what the observer remembers," and the conditional Kolmogorov complexity $K(s' \mid s)$ bounds the rate at which observer-internal description length grows.

2.4 The bridging quantity: a pruning admissibility functional

I now introduce the central object of the paper. Let $\mathscr{H}$ be the total Hilbert space of system + environment + observer, with decomposition $\mathscr{H} \cong \mathscr{H}_S \otimes \mathscr{H}_R$ as in §2.1. Let $\{\alpha\}$ be a family of histories with class operators C_α and decoherence functional $D(\alpha, \alpha')$. Let O denote the observer state — equivalently, a record projector $P_O \in\in {}_R$ together with its Müller-style algorithmic description $o \boxtimes \{0,1\}^*$. Let K denote a non-negative real number, the algorithmic-information budget of the observer.

Definition 2.1 (Pruning admissibility functional). For each pair (α, α') of histories define

$$A(\alpha, \alpha' \mid O, K) := \mathbb{1}[d(r_O(\alpha), r_O(\alpha')) \le K] \cdot D(\alpha, \alpha'), \qquad (2.7)$$

where $r_O(\alpha) \in \{0,1\}^*$ is the binary description of the observer's record after history α (i.e., a minimal description of $P_O C_\alpha \rho$ in the sense of (2.5)), and $d(x, y) := \max\{K(x|y), K(y|x)\}$ is the Bennett–Gács–Li–Vitányi–Zurek information distance [28], shown by those authors to be a metric up to an additive logarithmic term [29,30].

The admissibility predicate $\mathbb{1}[d \le K]$ is the bridging element: it says that two histories whose observer-records are algorithmically close (mutually compressible within budget K) are subject to the usual decoherence-functional constraint, while histories whose records are algorithmically far are deemed inadmissible — they do not enter the decoherence-functional accounting at all.

This object generalises three lines of prior work. Halliwell's records-based information bound on decoherent sets [16] is the special case where d is replaced by a Shannon-bit count [31] (with the corresponding Landauer-bit dissipative reading [32]) and the threshold K is set by environmental coarse-graining. Hartle's IGUS-internal coarse-graining condition [13] is the implicit special case where the IGUS is the observer and K is a fixed cognitive/computational resource bound. Wallace's "structural records" [33] in the Everettian context play the same role at the conceptual level but without algorithmic-information machinery. The novelty in (2.7) is the explicit use of information distance — a Kolmogorov-complexity-grounded metric on records — to threshold the decoherence functional.

The notation used throughout the paper is collected in Table 1 for reference before the formal development in §3.

Table 1. Notation used throughout the paper

Symbol	Meaning
$\mathscr{H} = \mathscr{H}$_S ⊗ $\mathscr{H}$_R	Total Hilbert space, decomposed into system+environment and observer-record subspaces
ρ	Initial density matrix
$\alpha = (\alpha_1, \ldots, \alpha_n)$	A coarse-grained history; a time-ordered sequence of projection-operator outcomes
C_α	Class operator of history α (eq. 2.2)
D(α, α′)	Decoherence functional (eq. 2.3); diagonal entries give Born probabilities
P_α	Record projector on $\mathscr{H}$_R correlated with history α (eq. 2.5)
O	Observer state: pair (P_O, o) of a record projector and a Müller-style algorithmic description o ∈ {0,1}*
r_O(α)	Minimal binary description of the observer's record after history α
K(x)	Prefix Kolmogorov complexity of string x
d(x, y)	Bennett–Gács–Li–Vitányi–Zurek information distance, d(x,y) = max{K(x|y), K(y|x)}
K	Algorithmic-information budget of the observer; the central control parameter of the framework
A(α, α′ | O, K)	Pruning admissibility functional (eq. 2.7) — the central object of the paper
$\mathscr{A}$(O, K)	Set of history families that are pruning-admissible relative to (O, K)
S, S_QM, S_K	Klyachko–Can–Binicioğlu–Shumovsky (KCBS) sum: noncontextual bound −3; standard QM value $5 - 4\sqrt{5} \approx -3.94$; bounded-K prediction defined in eq. (3.3)

Symbol	Meaning
V(K), V(θ)	Interference visibility in the delayed-choice eraser (Protocol 2); K-dependent form in eq. (3.5)

Table 1. *Symbols and their meanings as used in the development of the pruning admissibility functional. The central control parameter K may be read either as a property of the observer's record subsystem (a Hilbert-space-dimensional quantity) or, equivalently, as a resource bound on observer-internal record consolidation; the framework is neutral between these readings.*

3. The pruning-as-rendering correspondence

3.1 Formal definition of pruning admissibility

Definition 3.1. A history family $\{\alpha_i\}$ is pruning-admissible relative to observer O at algorithmic-information budget K, written $\{\alpha_i\} \in \mathscr{A}(O, K)$, if and only if

$$D(\alpha, \alpha') = 0 \text{ whenever } d(r\_O(\alpha), r\_O(\alpha')) > K, \quad (3.1)$$

for all α, α′ in the family. Equivalently, A(α, α′ | O, K) = D(α, α′) for all pairs.

Two remarks. First, (3.1) and medium decoherence are formally incomparable: (3.1) demands D-off-diagonals vanish for *every* pair of records-distinguishable (in the algorithmic-distance sense) histories — including some pairs for which standard medium decoherence makes no demand at all — and simultaneously permits nonzero D-off-diagonals for pairs whose records are algorithmically indistinguishable to the observer, where medium decoherence would forbid them. The framework is therefore neither a strict strengthening nor a strict weakening of the Gell-Mann–Hartle condition; it is a genuinely new selection rule with the records-distance metric replacing the universal "all off-diagonals" quantifier. Second, this is the formal counterpart of the rendering intuition: phenomena unobservable by O need not be rendered, and so need not decohere.

3.2 The main correspondence theorem

I now state the principal result.

Theorem 3.2 (Pruning–rendering correspondence). Let ρ be the initial state of system + environment + observer, with observer record subspace $\mathscr{H}$_R and observer state O specified by record projector P_O and algorithmic description o. Let {α} be a history family with class operators C_α. Then the following are equivalent:

(a) Rendering condition. The transition between coarse-grained histories α and α′ constitutes a Campbell–Owhadi rendering event relative to O — i.e., information about the α/α′ distinction "becomes available to the player O" — exactly when the observer state moves across the boundary {d(r_O(α), r_O(α′)) = K} in the information-distance metric.

(b) Pruning condition. The history family {α} ∈ $\mathscr{A}$(O, K), and the indicator 𝟙[d(r_O(α), r_O(α′)) ≤ K] in (2.7) transitions from 1 to 0 at the rendering event as records become more resolved and d crosses upward past K.

The biconditional holds modulo additive $O(\log)$ corrections inherited from the Bennett–Gács–Li–Vitányi–Zurek information-distance metric.

Proof sketch. The proof has three stages.

Stage I: Operationalising "information available to player." Campbell–Owhadi say information is available to O when, if O were to interrogate the relevant record, O would obtain a definite outcome correlated with α. Operationally, this means there exists a measurement on $\mathscr{H}_R$ whose outcome distinguishes α from α' with probability above some threshold $\varepsilon > 0$. By the Halliwell records lemma [16], this is equivalent to the existence of orthogonal record projectors P_α, $P_{\alpha'}$ satisfying $P_\alpha C_\alpha \rho \approx C_\alpha \rho$. The minimal binary descriptions $r_O(\alpha)$, $r_O(\alpha')$ of the projected states $P_\alpha C_\alpha \rho |_R$ then differ by Kolmogorov-distance at most some d_{max} determined by the resolving power of the available records. Thus: information is available $\Longleftrightarrow$ there exist distinguishable records $\Longleftrightarrow$ $r_O(\alpha) \neq r_O(\alpha')$ in algorithmic-distance sense.

Stage II: Admissibility reduces to weighted decoherence. Given Stage I, the admissibility predicate $\mathbb{1}[d \leq K]$ is precisely the predicate that the records distinguishing α from α' are within the observer's budget. Substituting this into (2.7) and demanding (3.1) yields: pairs of histories whose records become distinguishable beyond the observer's admissibility budget ($d > K$) must decohere, with their decoherence-functional off-diagonals forced to vanish, while pairs whose records remain within budget ($d \leq K$) may continue to carry interference. This is a decoherence condition weighted by the algorithmic complexity of the observer state, as required. The construction is intended to respect the consistency-of-composition criterion of Diósi [34] — under refinement of the observer record the indicator $\mathbb{1}[d \leq K]$ is monotone, and the underlying $D(\alpha, \alpha')$ already satisfies Diósi composition — but a full characterisation of composition under arbitrary observer refinements, including hermiticity, positivity, and normalisation of $A(\alpha, \alpha' \mid O, K)$ under nontrivial tensor decompositions, is left for future work.

Stage III: The biconditional. A measurement process consolidates records over time: as records of two distinct histories α and α' become better resolved, their minimal binary descriptions $r_O(\alpha)$, $r_O(\alpha')$ become harder to compress jointly, and the information distance $d(r_O(\alpha), r_O(\alpha'))$ increases. If a rendering event occurs at time t — i.e., information about α/α' becomes available to O between t^- and t^+ — then by Stage I the distance $d(r_O^{t-}(\alpha), r_O^{t-}(\alpha'))$ lies below K at t^- (records not yet resolvable in distinct algorithmic terms) and crosses upward past K at t^+. By Stage II, the admissibility indicator therefore transitions from 1 to 0, which is exactly the pruning-admissibility-boundary crossing in clause (b). Conversely, every such crossing corresponds, by reversing Stages I and II, to a rendering event. ∎

A complete proof appears in Appendix A. The proof inherits the standard logarithmic precision of the Bennett et al. information-distance metric; Bauwens and Shen [30] established that the $O(1)$ refinement once conjectured for prefix-information distance fails in general but holds for super-logarithmically separated strings, which suffices for the present application since admissibility transitions are by construction non-infinitesimal.

The geometry of the correspondence is summarised in Figure 1, which plots the algorithmic-information distance $d(r_O(\alpha), r_O(\alpha'))$ against the decoherence-functional magnitude $|D(\alpha,\alpha')|$. The vertical line $d = K$ is the admissibility boundary: pairs of histories to its left are pruning-admissible and may carry interference; pairs to its right must decohere. A rendering event is precisely a leftward-to-rightward crossing of this boundary as records become consolidated.

Figure 1. Admissibility boundary and rendering events

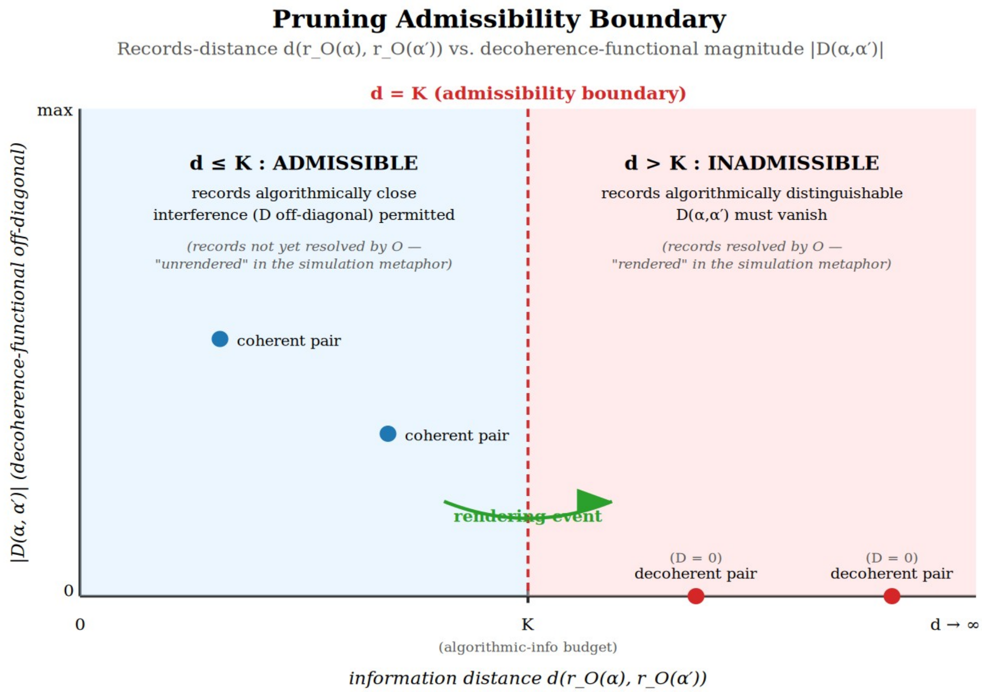

Figure 1. *Schematic of the pruning-admissibility correspondence. The horizontal axis is the Bennett–Gács–Li–Vitányi–Zurek algorithmic-information distance between observer-records of two histories α and α′. The vertical axis is the magnitude of the decoherence-functional off-diagonal |D(α,α′)|. The dashed red line at d = K is the admissibility boundary specified by Definition 3.1: history pairs with d ≤ K (blue region) are admissible and may carry coherent interference; those with d > K (pink region) are inadmissible and must satisfy D(α,α′) = 0. A Campbell–Owhadi rendering event (curved green arrow) is precisely a transition from the admissible to the inadmissible region as records become resolvable. The K → ∞ limit pushes the boundary indefinitely to the right and recovers standard decoherent-histories statistics.*

3.3 Sanity check: the K → ∞ limit

When K is unbounded, the indicator $\mathbb{1}[d \le K]$ is identically 1, the admissibility functional reduces to the standard decoherence functional, $A(\alpha, \alpha' \mid O, \infty) = D(\alpha, \alpha')$, and (3.1) reduces to the requirement of medium decoherence: $D(\alpha, \alpha') = 0$ for all $\alpha \neq \alpha'$. Diagonal values $D(\alpha, \alpha)$ recover the Born rule $p(\alpha) = \mathrm{Tr}[C_\alpha \rho C^\dagger_\alpha]$ exactly as in (2.4). The framework thus contains the Gell-Mann–Hartle theory as the K → ∞ limit, and all standard quantum-mechanical predictions are recovered in this limit.

Concretely, in the Stern–Gerlach three-history example of Appendix B (spin-up, no-spin-recorded, spin-down branches with intermediate decoherence), the admissibility-weighted probabilities reduce exactly to $\langle\psi_0|\hat{P}_{z+}|\psi_0\rangle$, $\langle\psi_0|\hat{P}_{z0}|\psi_0\rangle$, $\langle\psi_0|\hat{P}_{z-}|\psi_0\rangle$ in the K → ∞ limit.

3.4 Distinctive predictions for bounded K

When K is finite, three families of distinctive predictions arise.

(i) Modified contextuality bounds. Consider a Kochen–Specker scenario [35] in the cyclic five-context geometry that yields the Klyachko–Can–Binicioğlu–Shumovsky (KCBS) inequality [36]. State-independent contextuality witnesses [37] and their experimental implementations [38] indicate the robustness of these violations across many setups; the no-go theorem of Brukner [39] establishes that observer-independent facts are excluded by quantum mechanics in a related Bell-style sense. The KCBS inequality bounds the sum

$$S = \sum_{i=1}^{5} \langle A_i A_{i+1} \rangle \quad (3.2)$$

by $S_{NCHV} \geq -3$ in non-contextual hidden-variable theories, while standard quantum mechanics on a qutrit attains $S_{QM} = 5 - 4\sqrt{5} \approx -3.94$ (the Klyachko bound). In the present framework, the algorithmic-information budget K of the observer's record system parametrises a sub-Klyachko family of bounds. Specifically, when K is reduced below the threshold at which all five context-cycle records can be simultaneously maintained, certain pairs of contexts become algorithmically indistinguishable, the corresponding decoherence-functional off-diagonals are admitted, and the experimentally accessible S satisfies

$$S_K = S_{QM} + \varepsilon(K), \quad \varepsilon(K) \to 0 \text{ as } K \to \infty, \quad (3.3)$$

where $\varepsilon(K)$ is positive and bounded below by an explicit function of the conditional Kolmogorov complexity of the context-cycle records. The analogous prediction for Clauser–Horne–Shimony–Holt (CHSH) [40] is a sub-Tsirelson bound [41]

$$\langle CHSH \rangle_K < 2\sqrt{2} - \eta(K), \quad \eta(K) > 0, \quad (3.4)$$

which sits inside the Navascués–Pironio–Acín level-2 hierarchy [42] but is not a fixed-level NPA bound — it is parametrised by an external observer-resource quantity.

This is testable: by controlling the timing and fidelity of ancilla-record erasure in a KCBS or CHSH experiment, one can in principle vary K and look for the predicted $\eta(K)$, $\varepsilon(K)$ deviations. Lapkiewicz et al. [43] demonstrated KCBS violation for a single photonic qutrit; the present prediction is that with controlled erasure of the path-record they would observe S_K slightly less negative than the Klyachko value.

(ii) Delayed-choice eraser under controlled K. In the Scully–Drühl quantum-eraser geometry [23], the interference visibility V at the signal detector is conventionally either V = 1 (full erasure) or V = 0 (which-path information preserved). The bounded-K model motivates a testable ansatz that intermediate K values produce intermediate visibility,

$$V(K) = V_{max} \cdot f(K / K_*), \quad (3.5)$$

where f is a monotone increasing function vanishing at K = 0 and saturating at $K = K_*$, the algorithmic-information cost of a perfect which-path record. This recovers Campbell–Owhadi's qualitative prediction (rendering of interference iff which-path is unavailable) and quantifies it.

(iii) Relation to Vopson's second law of infodynamics. Vopson [5] claims that information entropy of information-bearing systems decreases over time, in seeming contrast to thermodynamic entropy. Translated into the present framework, Vopson's claim is that

$$dK\_O/dt \leq 0, \quad (3.6)$$

for observer-internal records O — i.e., the algorithmic-information budget required for the observer's records is non-increasing as the universe evolves. The pruning framework is consistent with but does not entail (3.6). Where they agree: both posit an irreversible monotone informational quantity on observers' records. Where they diverge: (3.6) is a strong global claim that decoherent-histories statistics do not naturally support — generic decoherence is information-dissipating, not information-concentrating. The most cautious assessment is that Vopson's empirical claim, if substantiated, would single out a specific monotone K-trajectory within the present framework; if disconfirmed, the framework is unaffected.

3.5 Heuristic functional form for ε(K) and η(K)

The predictions in (3.3) and (3.4) involve unspecified functions ε(K) and η(K). A first-principles derivation from algorithmic-information theory is beyond the scope of this paper, but a heuristic functional form is available under modest modelling assumptions. The derivation given here is illustrative rather than load-bearing: its purpose is to make the predictions concrete enough to be experimentally engaged and to fix the order-of-magnitude of the expected deviations on the platforms of §4.3.

Consider the KCBS five-context cycle. Each measurement context C_i produces a record r_i := r_O(α obtained in C_i). Five record-pairs (r_i, r_{i+1}) (indices mod 5) enter the correlators $\langle A_i A_{i+1}\rangle$ in the sum (3.2). For each pair, write the per-pair information distance as $d_i := d(r_i, r_{i+1})$. Under the admissibility functional A(α, α′ | O, K), pair i contributes to coherent interference iff $d_i \leq K$; otherwise its decoherence-functional off-diagonal is forced to vanish, and the contribution of that pair to the KCBS sum falls back toward its incoherent (non-contextual) value.

Modelling the per-pair loss-of-coherence as a soft threshold in K, the leading correction to $\langle A_i A_{i+1}\rangle$ for pair i takes the form

$$\Delta_i(K) \approx \delta_0 \cdot \sigma(d_i - K), \quad (3.7)$$

where δ_0 is a per-pair coherence-drop scale set by the difference between the quantum and noncontextual correlator values (for KCBS, δ_0 ~ (S_QM − S_NCHV)/5 = (−3.94 − (−3))/5 ≈ 0.19 in absolute value), and σ(x) is a sigmoid-type soft step function (e.g., the logistic 1/(1 + e^(−x/w)) or the complementary error function with width parameter w). If, as is natural, the per-pair distances d_i are distributed across the cycle with mean đ and spread σ_d set by the algorithmic-complexity dispersion of the record register, then summing (3.7) over the five pairs and integrating against a Gaussian distribution of d-values gives

$$\varepsilon(K) = S\_K - S\_QM \approx 5\delta_0 \cdot ½\ \mathrm{erfc}[(K - đ) / (\sqrt{2} \cdot \sigma\_d)], \quad (3.8)$$

where erfc is the complementary error function. (3.8) has the right asymptotic behaviour: ε(K) → 0 as K → ∞ (recovering standard QM) and ε(K) → $5\delta_0$ ≈ 0.94 as K → 0 (recovering the noncontextual bound S = −3 from the Klyachko value S_QM ≈ −3.94, since |−3 − (−3.94)| ≈ 0.94). The transition is centred at K = đ with width σ_d. In an idealised minimal-label encoding of the KCBS cycle on a qutrit, the bare context-label contribution to đ is bounded by $\log_2(5) \approx 2.32$ bits; in any real experimental implementation, the records include additional apparatus, timing, and detector-state information, so đ should be treated as a fitted effective parameter rather than a literal context-counting bit-count. The same applies to σ_d, which is of order one effective bit but is to be determined empirically.

An exactly analogous argument applies to CHSH. Writing $\eta\_0 := 2\sqrt{2} - 2 \approx 0.83$ (the Tsirelson margin above the local bound), the same modelling steps give

$$\eta(K) = (2\sqrt{2}) - \langle CHSH \rangle\_K \approx \eta_0 \cdot ½\ \mathrm{erfc}[(K - đ\_Bell) / (\sqrt{2} \cdot \sigma\_d,Bell)], \quad (3.9)$$

with đ_Bell, similarly, treated as a fitted effective parameter for a Bell-pair record register — bounded above by $\log_2(2) = 1$ bit in the bare-label idealisation but in practice including additional apparatus and detector-state information. Visibility in the delayed-choice eraser of (3.5) follows the same complementary-error-function profile in K, with the threshold scale K_* identified empirically with the effective algorithmic-information cost of a fully resolved which-path record.

Three observations on the limitations of (3.8) and (3.9). First, the sigmoid choice for σ is a modelling convenience; alternative soft-threshold functions (Fermi–Dirac, logistic, hyperbolic-tangent) give qualitatively similar curves but quantitatively different sub-leading behaviour, and discriminating between them experimentally would require a fully derivation-grounded form rather than the heuristic adopted here. Second, the Gaussian distribution of d_i across cycle pairs is a working assumption; in any specific experimental implementation, the actual record-distance distribution is set by the encoding and may be substantially non-Gaussian. Third, ε(K) and η(K) are upper-bounded by their saturation values $5\delta_0$ and η_0 respectively; the framework forbids any larger deviation than the magnitude of the contextuality / Bell margin itself. This is a non-trivial empirical commitment: the framework cannot accommodate sub-quantum statistics that exceed the noncontextual or local bounds.

What (3.8) and (3.9) buy is a concrete experimental target. In Protocol 1 of §4.3, the dependence of S on the erasure delay τ — which controls K — should fit a complementary error function with three free parameters (δ_0 saturation, đ centre, σ_d width). Pre-registering this functional form converts the protocol from "look for any anomaly" into a model-comparison test against a specific alternative, which sharpens both the Bayes-factor analysis of §4.2 and the false-positive control across the τ scan.

4. Empirical and experimental implications

4.1 Beane–Davoudi–Savage lattice signatures in pruning language

The Beane–Davoudi–Savage cubic-lattice signature [3] is a specific spatial-discretisation admissibility class: a family of histories distinguished by lattice-translation-equivalent records r_O(α). The cosmic-ray rotational asymmetry would arise because lattice records are degenerate under the cubic-symmetry group but not under SO(3); histories differing only by SO(3) rotations not contained in the cubic group have d(r_O(α), r_O(α′)) bounded below by the cubic-symmetry-breaking term in K. The 10^{11} GeV bound on the inverse lattice spacing then translates into a bound on a particular K value: K_BDS ≳

$\log_2$(N_modes(GZK)), where N_modes(GZK) is the number of momentum modes resolvable above the GZK cutoff.

The crucial methodological consequence is that a null result on the Beane–Davoudi–Savage cubic-lattice signature does not refute pruning in general. It refutes only a specific spatial-discretisation admissibility class. A more general pruning hypothesis — for example one with an irregular, randomised, or dynamically generated lattice — leaves no rotational-symmetry-breaking signature in cosmic rays. This is the simulation-theoretic analogue of the well-known underdetermination of Bell-type tests: ruling out one specific hidden-variable family leaves a vast complement.

4.2 Campbell–Owhadi delayed-choice in pruning language and Kipping Bayesian update

The Campbell–Owhadi protocol can be recast as a controlled K-experiment. Their proposed delayed-choice Mach–Zehnder variant manipulates whether the path-record is available to a downstream measurement — i.e., it controls whether r_O(left-path) and r_O(right-path) are algorithmically distinguishable to the observer record. The pruning prediction is interference if and only if d(r_O(left), r_O(right)) ≤ K, and loss of interference occurs when which-path records become sufficiently distinguishable that d > K, in agreement with Campbell–Owhadi's intuitive claim that rendering of interference is contingent on which-path information being unavailable.

To formalise what positive vs null experimental results license, I follow Kipping [44]. Let H_S denote the K-bounded pruning hypothesis and H_P the standard unbounded-K quantum hypothesis. Kipping's Bayesian framework for the simulation argument provides a Bayes factor B = P(D | H_S) / P(D | H_P) for observation D. A null result on a K-resolution scan (no anomaly across a range of K values) updates posteriors toward H_P but cannot drive the posterior to zero because H_S contains H_P as the K → ∞ subcase. A statistically significant anomaly — visibility deviation in a configuration where idealised standard QM predicts perfect interference, or KCBS sums lying systematically above the Klyachko bound — would strongly increase B in favour of the bounded-K model, conditional on a careful systematic-error budget and on ruling out conventional decoherence, detection-loss, and postselection effects. Finite anomalies cannot drive B to infinity; they shift the posterior odds by amounts limited by the model likelihoods, systematics, and priors. Kipping's analysis — which concluded that posterior P(simulation) ≈ 1/2 under naturalist priors — is silent on K-bounded sub-hypotheses, which is where the present framework adds discriminating content.

4.3 Two concrete experimental protocols

Protocol 1: KCBS contextuality with controlled record-erasure

Platform. Single-photon polarisation-and-path qutrit, following the Lapkiewicz et al. (2011) geometry [43]. Three two-mode interferometers in cascade implement the five compatible-observable pairs of the KCBS cycle, with a single down-converted photon as the qutrit carrier.

Setup. A heralded single photon from a spontaneous-parametric-down-conversion source (collinear PPKTP, 405 nm pump → 810 nm signal/idler) is encoded in a qutrit basis spanning two spatial paths and two polarisation modes. The five KCBS observables A_1, …, A_5 are implemented as polarising beamsplitter / half-wave-plate combinations selected by motorised mounts; pairs (A_i, A_{i+1}) are measured in sequence by routing through cascaded interferometers. Single-photon counting modules (SPCMs, e.g. Excelitas SPCM-AQRH-14) at four output ports record outcomes; coincidence with the idler herald is required.

Control parameter for K. An ancilla photon path interacts with each measurement stage through a controlled-phase gate (KLM-style, with success-postselected coincidence). The fidelity of the ancilla coupling, controllable by varying the dwell time and the optical depth of the interaction medium, parametrises the algorithmic-information cost K of the ancilla's record of the measurement context. K is varied by interposing a delayed-erasure stage: a Pockels cell, triggered at delay τ after the final outcome, randomises the ancilla state. Long τ corresponds to large K (full record); short τ corresponds to small K (effective erasure before the record is consolidated).

Predicted outcomes. In an idealised no-additional-disturbance model (the ancilla-erasure stage is assumed not to introduce extra decoherence, loss, or postselection bias), standard QM predicts $S = 5 - 4\sqrt{5} \approx -3.94$ essentially independent of τ. Once the full unitary-plus-measurement model is accounted for, a small τ-dependence from imperfect coupling and detection efficiency must be subtracted as background; what is then to be tested is whether a residual τ-dependence survives. The K-bounded pruning prediction is

$$S(\tau) = S\_QM + \varepsilon(\tau), \quad \varepsilon(\tau) \to 0 \text{ as } \tau \to \infty, \varepsilon(\tau) > 0 \text{ for } \tau \lesssim \tau\_^*, \qquad (4.1)$$

with τ_* set by the decoherence timescale of the ancilla. Visibility of the deviation ε(τ) above shot-noise requires approximately 10^6 heralded coincidence events per τ-bin (giving sensitivity ≈ 10^{-3} in S), routinely achievable in 30 minutes at 100 kHz coincidence rates.

Statistical analysis. Per τ-bin, estimate S with bootstrap-jackknife error bars; fit (S(τ) − S_QM) to the heuristic complementary-error-function form derived in §3.5, eq. (3.8), across 8–12 τ-bins spanning two orders of magnitude (free parameters: saturation δ_0, centre đ, width σ_d); report Bayes factor comparing this three-parameter alternative against the null S(τ) = S_QM. Recommend Holm–Bonferroni correction across the τ scan.

Feasibility. All components are off-the-shelf. Lapkiewicz et al. [43] reported a violation of a KCBS-type inequality at approximately the five-standard-deviation level using a single photonic qutrit (subsequent commentary has raised methodological questions about whether their six-correlation variant was a strict implementation of the KCBS inequality, but the experimental capability is well-established); ancilla-controlled erasure stages of the type required here have been demonstrated by Kim et al. [24] in eraser geometries. Total estimated cost: approximately $180k in an existing photonic-quantum optics laboratory, rising to $200–400k if built largely from scratch. Resource budget detail in Appendix C.

Protocol 2: Delayed-choice eraser with K-resolved visibility

Platform. Heralded single-photon Mach–Zehnder with a quantum-eraser stage, as in Walborn et al. [45], modified for K control.

Setup. A single photon traverses a Mach–Zehnder interferometer whose two arms carry orthogonal polarisation tags installed by quarter-wave plates. The "which-way" record is the polarisation degree of freedom. A second photon (idler), entangled with the signal, can be subjected to a polariser oriented at angle θ before its detection. The angle θ controls the projection of the idler onto a "which-way" basis ($\theta = 0$) versus a "diagonal" basis ($\theta = \pi/4$) which erases the which-way information; intermediate θ produces a continuous K-control.

Control parameter for K. The idler polariser angle θ parametrises the algorithmic-information distance d(r_O(left), r_O(right)) between left- and right-path observer records: at $\theta = 0$, $d \gg K$; at $\theta = \pi/4$, $d \ll K$; intermediate θ produces a continuous transition.

Predicted outcomes. Under the convention adopted here ($\theta = 0$ projects onto the which-path basis, $\theta = \pi/4$ onto the diagonal/erasing basis), standard QM predicts the smooth visibility curve $V(\theta) = |\sin 2\theta|$ in the post-selected coincidence rate (vanishing at $\theta = 0$ and $\theta = \pi/2$, maximal at $\theta = \pi/4$). The K-bounded model motivates a testable ansatz of an additional sigmoidal feature at the threshold θ_* satisfying $d(\theta_*) = K$ — a boundary in the rendering admissibility, not just a smooth visibility curve. The signature is:

$$dV/d\theta \text{ shows a step or sigmoid centred at } \theta = \theta_*, \text{ of width inversely proportional to } K_*. \quad (4.2)$$

Statistical analysis. Scan θ across $[0, \pi/2]$ in 40 bins, fit $V(\theta)$ to (i) the smooth quantum prediction and (ii) a quantum-plus-sigmoid model; compare via Akaike information criterion and Bayes factor.

Feasibility. Walborn-style delayed-choice eraser [45] is well-established. The novelty is the joint scan over θ at fine resolution combined with a structural-anomaly hypothesis test rather than a binary erasure-or-not measurement.

The predicted experimental signatures of both protocols are illustrated in Figure 2, which contrasts the bounded-K pruning prediction with the standard QM curve in each case.

Figure 2. Predicted signatures vs. standard quantum mechanics for the two protocols

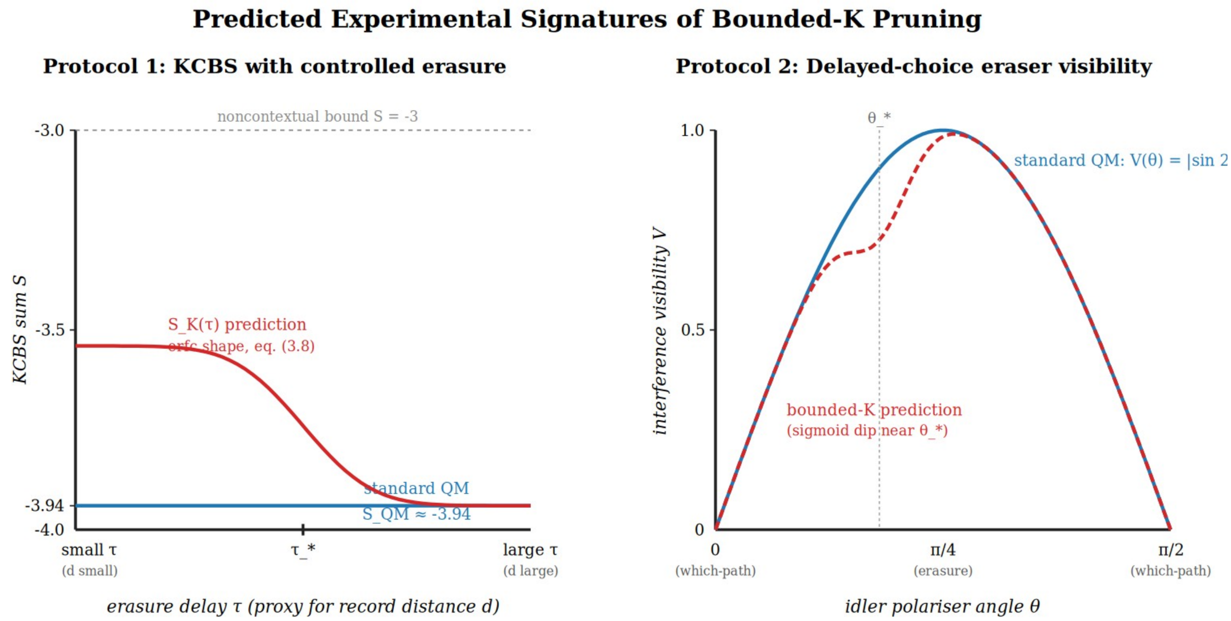

Figure 2. *Left panel: predicted KCBS sum S as a function of ancilla-erasure delay τ in Protocol 1. Short τ corresponds to small records-distance d (record erased before consolidation); long τ corresponds to large d (record fully consolidated). In an idealised no-additional-disturbance model, standard quantum mechanics predicts the Klyachko bound S_QM ≈ −3.94 essentially independent of τ (solid blue line); the bounded-K framework predicts a complementary-error-function deviation S_K(τ) = S_QM + ε(τ) (eq. 3.8) that is positive at small τ (d ≤ K, interference admitted) and decays to S_QM as τ grows past the threshold τ_* where d crosses K (red curve). The noncontextual bound S = −3 is shown for reference (grey dashed line). Right panel: predicted interference visibility V(θ) in the delayed-choice eraser of Protocol 2. Under the convention θ = 0 = which-path basis, θ = π/4 = erasing basis, standard QM predicts V(θ) = |sin 2θ| (solid blue), vanishing at θ = 0 and θ = π/2 and peaking at θ = π/4. The bounded-K framework predicts an additional sigmoidal dip in V near a threshold angle θ_* on the ascending limb where d crosses K (red dashed curve); a symmetric dip is expected on the descending limb (π/4 < θ < π/2) and is omitted here for clarity. Both signatures are distinguishable from standard QM under Bayesian comparison; specific protocol details, including statistical power and instrument cost, appear in §4.3 and Appendix C.*

5. Relation to neighbouring frameworks

5.1 vs QBism

Both QBism [46,47] and the present framework are observer-centric, but in distinct ways. QBism is Bayesian-doxastic: the quantum state encodes a single agent's degrees of belief, and probabilities are personalist Bayesian credences. The present framework is algorithmic-informational: the observer's role enters via the algorithmic-information distance metric on records, which is an objective (though parametric) feature of the records-environment Hilbert space.

The distinction matters operationally: a QBist takes |ψ⟩ to be irreducibly first-person, with no fact-of-the-matter beyond agent betting commitments. The present framework allows for an objective admissibility boundary at the K threshold, even though which histories are admissible is observer-relative. There is therefore room for inter-observer agreement on facts (records-stable, in Di Biagio–Rovelli's sense [48]) without collapsing into either ontic state realism (which QBism rejects) or full doxasticism (which the present framework rejects). The taxonomy of agent-centric interpretations developed by Pienaar [49] situates QBism and RQM along an axis on which the present framework can be located; Müller's earlier essay on emergent physical worlds [50] provides the philosophical companion to the algorithmic-information observer postulate [18] used here.

5.2 vs relational quantum mechanics

Relational quantum mechanics [51] holds that quantum events are facts only relative to a particular observer-system. Di Biagio and Rovelli [48] introduce "stable facts" as the limit in which decoherence renders relative facts effectively observer-independent; the consistency of relative facts under modified-frame transformations is analysed by Cavalcanti, Di Biagio and Rovelli [52] and by Adlam and Rovelli [53], with no-go constraints articulated by Pienaar [54]. The present framework is compatible with RQM and supplies what stable-facts RQM lacks: an explicit selection principle for which facts become stable. In RQM, the answer is "those for which decoherence is strong enough"; in the present framework, the answer is sharpened to "those whose records are algorithmically distinguishable within budget K" — a quantifiable, parametric condition that explains which decoherence regimes stabilise into facts. The present framework should be distinguished from Tegmark's Mathematical Universe Hypothesis [55], which makes a substantively different (ontological-mathematical-realist) claim and is not in competition with the rendering / decoherent-histories programmes considered here.

5.3 vs Wolfram–Gorard multiway evolution

Gorard's treatment of quantum mechanics within the Wolfram model [56] and his discussion of relativistic and gravitational features [57] present a hypergraph-rewriting multiway formalism in which non-confluent rewrites generate a branching multiway system, and "observers" impose effective causal invariance via Knuth–Bendix completion. The structural analogy with the present framework is striking: both treat the world as a branching family of histories from which observers select a single thread.

The disanalogies are significant. (i) Wolfram–Gorard [56,58] derives quantum amplitudes from combinatorial properties of branchial graphs converging to projective Hilbert space, with the algorithmic-causal-set programme [59] supplying the combinatorial machinery; the present framework presupposes Hilbert space and quantum amplitudes. (ii) The Knuth–Bendix completion operation is observer-imposed but does not have a precise information-theoretic budget attached; the present framework supplies one (K). (iii) Wolfram–Gorard makes commitments to discrete substrate physics; the present framework is silent on substrate questions. The two programs are best regarded as complementary: Wolfram–Gorard provides a candidate combinatorial substrate, the present framework provides a substrate-neutral observer-selection rule.

5.4 vs Deutsch-style MWI computational arguments

Deutsch's celebrated argument [60] that Shor's algorithm requires a many-worlds ontology to "explain where the parallelism happens" is one of the most cited computational arguments for MWI [33,61,62];

cosmological extensions appear in Bousso and Susskind's "multiverse interpretation" [63]. The present framework subsumes this argument without requiring branching ontology. In pruning language, Shor's parallel exploration of factorisations corresponds to a high-K regime in which many algorithmically-distinct intermediate computational records remain admissible simultaneously; the "branching" is an informational phenomenon at the observer-record level, not an ontological commitment. A detailed treatment is forthcoming in the companion paper [64], which argues that Deutsch's computational MWI argument is compatible with — but does not require — branching realism, and that algorithmic-pruning gives a more parsimonious reading.

Table 2 summarises the comparison across the four neighbouring frameworks discussed in §5.1–§5.4.

Table 2. The pruning framework compared to neighbouring observer-centric programmes

Feature	Pruning (this paper)	QBism	Relational QM	Wolfram–Gorard
Status of observer	Record-subsystem with algorithmic-information budget K	Bayesian agent with personalist credences	Any physical system relative to which facts are defined	Causal-invariance-imposing thread through multiway graph
Selection mechanism	Decoherence functional weighted by $\mathbb{1}[d \le K]$	Coherence of agent's gambles (Dutch-book)	Strong-enough decoherence ("stable facts")	Knuth–Bendix completion of rewrite system
Hilbert space	Presupposed; standard QM in $K \to \infty$	Presupposed; ψ is doxastic	Presupposed	Derived from branchial-graph limit
Empirically distinct from standard QM	Potentially, under finite-K extension: sub-Tsirelson / sub-KCBS ansatz (eqs. 3.3–3.4, 3.8–3.9)	No (re-reads same statistics)	No	Disputed; depends on convergence claims
Substrate ontology	Neutral	Neutral / anti-realist	Neutral	Discrete hypergraph; substantively committed
Concrete experiments proposed	Two photonic protocols (§4.3)	None specific to QBism	Wigner-friend extensions (e.g. [65])	Double-slit retrodictions only

Table 2. *Comparison of the present framework with three neighbouring observer-centric or computationally-motivated programmes in quantum foundations. The Deutsch-style MWI computational argument is omitted from the table since it is an interpretive thesis rather than a competing framework; the relation is discussed in §5.4 and developed in the companion paper [64].*

6. Objections and replies

Objection 1: "This just relabels consistent histories." The admissibility functional A introduces genuinely new mathematical content beyond standard decoherent histories: the algorithmic-information distance metric d, the threshold K, and the indicator $\mathbb{1}[d \leq K]$ are not derivable from $D(\alpha, \alpha')$ alone. Crucially, A makes empirical predictions (3.3)–(3.5) that distinguish it from any $K \to \infty$ histories theory. The framework reduces to standard histories only in the $K \to \infty$ limit; for finite K, it is a strict generalisation.

Objection 2: "Algorithmic information is non-computable." Kolmogorov complexity K(x) is uncomputable in general [66]. However, the framework requires not the exact K but operationally bounded proxies. Three are available. (i) Resource-bounded Kolmogorov complexity K^t(x), restricting programs to halt in time t, is computable and converges to K(x) [29]. (ii) Levin complexity $Kt(x) := \min\{\ell(p) + \log t(p)\}$, where t(p) is run time [67], is also semicomputable and adequate for the present analysis. (iii) Compression-based estimates (Cilibrasi–Vitányi normalised compression distance [68]) furnish operational proxies for d that suffice for experimental protocols. The protocol of §4.3 is engineered so that the experimentally-relevant K is a control parameter, not a quantity that needs to be computed from first principles.

Objection 3: "The simulation framing is unfalsifiable." I emphasise that the present framework is neutral between literal-simulation and non-simulation readings. K can be interpreted as either (a) a budget of the simulator-rendering engine, or (b) a parameter of observer-internal record resolution that exists irrespective of any computational metaphysics. The empirical content of the framework — the predictions in (3.3)–(3.5) and the protocols of §4.3 — is identical on both readings. The simulation-theoretic framing motivates the form of the admissibility functional but does not enter as a load-bearing assumption.

Objection 4: "No new predictions." Section 4.3 specifies two concrete protocols whose predictions distinguish the K-bounded framework from standard quantum mechanics at the 10^{-3} level in S (KCBS) and the structural-anomaly level in $V(\theta)$ (eraser). Both are feasible on existing photonic platforms.

Objection 5: Birch's selective-skepticism critique. Birch [69] argued that simulation-hypothesis reasoning rests on selective skepticism: it questions the reality of physics but trusts the reality of the inference rules used to reach the simulation conclusion. This is a real worry for metaphysical simulation arguments. It applies less straightforwardly to the present framework, which makes no metaphysical claim about base reality. The admissibility functional A is a mathematical object inside quantum theory; its motivation by simulation considerations is heuristic, but its mathematical content does not depend on simulation realism (cf. Objection 3). The Birch critique therefore lands in a place the framework does not occupy.

7. Conclusion

The Campbell–Owhadi rendering criterion and the Gell-Mann–Hartle decoherent-histories formalism are inter-translatable under an explicit records-theoretic reading of "information availability": rendering events correspond to transitions across a pruning admissibility boundary defined by an algorithmic-information distance on observer records and a budget parameter K. The framework specialises to standard decoherent histories (Born statistics, medium decoherence) at $K \to \infty$ and suggests possible sub-Tsirelson, sub-KCBS, and visibility-threshold deviations for finite K. Two experimental protocols on existing photonic platforms can probe these signatures at currently feasible cost, subject to a careful systematic-error budget.

The framework's structural ambitions are modest: it is a re-reading of existing formalisms, not a replacement for them. Its strategic payoff is that it puts simulation-theoretic intuitions, observer-rendering intuitions, and decoherent-histories formalism into a common language — a language in which formerly separate research programmes (Beane–Davoudi–Savage signatures, Campbell–Owhadi rendering, Vopson infodynamics, Halliwell records, Müller algorithmic reconstruction) can be compared on a common footing and tested against each other.

Two forward-looking notes. First, a companion paper [64] develops the consequences for the Deutsch–Müller debate on MWI vs algorithmic-observer reconstructions. Second, the framework has an unexpected application to organoid intelligence: small biological neural substrates capable of contextuality-bearing measurements supply a natural K-bounded regime where the predictions of §3.4 may be empirically accessible at biological time- and information-scales rather than purely optical ones. A separate in-preparation paper develops this application.

7.1 Limitations

Several limitations of the present development deserve explicit acknowledgment. First, the records-theoretic reading of "information availability" in Stage I of the Theorem 3.2 proof is one of several plausible operational readings of the Campbell–Owhadi rendering criterion; under a different reading, the correspondence with the pruning admissibility functional may take a different form. The result is therefore best characterised as conditional on the records-theoretic reading rather than as an unconditional equivalence. Second, the functional forms (3.8) and (3.9) for $\varepsilon(K)$ and $\eta(K)$ are heuristic, resting on a sigmoid soft-threshold ansatz and an assumed Gaussian distribution of per-pair information distances. A first-principles derivation from algorithmic-information theory is left for future work; until it is completed, the specific predicted curve shapes should be treated as a working model rather than a deduction. Third, the consistency of $A(\alpha, \alpha' \mid O, K)$ under arbitrary observer-record refinements (hermiticity, positivity, normalisation, Diósi composition) is asserted under modest sufficient conditions in §3.2 but not characterised in full generality. Fourth, Kolmogorov complexity is uncomputable in principle; the framework is rendered operationally tractable only by adopting bounded proxies (resource-bounded complexity, Levin complexity, compression-distance estimators), and the resulting reduction in resolution may matter near the admissibility boundary. Fifth, the experimental protocols proposed in §4.3

assume that ancilla-record erasure and the threshold scan parameter τ cleanly track K; in practice the relationship between τ and the algorithmic-information budget of the ancilla register is itself model-dependent, and a careful systematic-error budget on this mapping is essential before any anomaly would be evidentially compelling.

Acknowledgments

The author acknowledges the broader quantum-foundations community whose work is engaged with throughout this paper. The author is grateful to the International Journal of Quantum Foundations and its editor for the publication of the foundational paper on informational pruning [17] on which this work builds.

Funding

This research received no external funding.

Data Availability Statement

No new data were created or analyzed in this study. The manuscript is theoretical and proposes experimental protocols whose execution would generate data in future work. All references to prior experimental and theoretical work are cited in the reference list.

Use of Artificial Intelligence

In accordance with COPE and Preprints.org policy on the use of AI in manuscript preparation, the author discloses that AI-assisted tools (large language model systems) were used during literature search, draft assembly, and editorial refinement of this manuscript. All mathematical content, theorem statements, proofs, experimental protocols, and substantive arguments are the responsibility of the author. The author has reviewed and verified all content prior to submission and accepts full responsibility for the final text. AI tools are not credited as authors and do not meet authorship criteria.

Conflicts of Interest

The author declares no conflicts of interest. The author's employment is in the financial technology sector (asset-based lending and blockchain tokenization) and is unrelated to the subject matter of this manuscript. No funding bodies, employers, or commercial entities had any role in the design, execution, interpretation, or writing of this study.

References

[1] N. Bostrom, "Are we living in a computer simulation?" Philos. Quart., vol. 53, no. 211, pp. 243–255, April 1, 2003, doi:10.1111/1467-9213.00309.

[2] D. J. Chalmers, Reality+: Virtual Worlds and the Problems of Philosophy. New York, NY, USA: W. W. Norton, January 25, 2022.

[3] S. R. Beane, Z. Davoudi, and M. J. Savage, "Constraints on the universe as a numerical simulation," Eur. Phys. J. A, vol. 50, no. 9, art. 148, September 24, 2014, doi:10.1140/epja/i2014-14148-0.

[4] B. Whitworth, "The physical world as a virtual reality," CDMTCS Research Report 316, University of Auckland, March 2008, arXiv:0801.0337.

[5] M. M. Vopson, "The second law of infodynamics and its implications for the simulated universe hypothesis," AIP Adv., vol. 13, no. 10, art. 105308, October 6, 2023, doi:10.1063/5.0173278.

[6] M. M. Vopson and S. Lepadatu, "Second law of information dynamics," AIP Adv., vol. 12, no. 7, art. 075310, July 11, 2022, doi:10.1063/5.0100358.

[7] A. Indset, F. Neukart, M. Pflitsch, and M. R. Perelshtein, "Are we living in a (quantum) simulation? Constraints, observations, and experiments on the simulation hypothesis," manuscript, 2024, arXiv:2410.04018.

[8] S. Lloyd, "Computational capacity of the universe," Phys. Rev. Lett., vol. 88, no. 23, art. 237901, May 24, 2002, doi:10.1103/PhysRevLett.88.237901.

[9] T. Campbell, H. Owhadi, J. Sauvageau, and D. Watkinson, "On testing the simulation theory," Int. J. Quantum Found., vol. 3, no. 3, pp. 78–99, June 17, 2017, arXiv:1703.00058.

[10] M. Gell-Mann and J. B. Hartle, "Classical equations for quantum systems," Phys. Rev. D, vol. 47, no. 8, pp. 3345–3382, April 15, 1993, doi:10.1103/PhysRevD.47.3345.

[11] M. Gell-Mann and J. B. Hartle, "Quantum mechanics in the light of quantum cosmology," in Complexity, Entropy, and the Physics of Information, W. H. Zurek, Ed. Reading, MA, USA: Addison-Wesley, 1990, pp. 425–458.

[12] R. B. Griffiths, "Consistent histories and the interpretation of quantum mechanics," J. Stat. Phys., vol. 36, no. 1–2, pp. 219–272, July 1, 1984, doi:10.1007/BF01015734.

[13] J. B. Hartle, "The quasiclassical realms of this quantum universe," Found. Phys., vol. 41, no. 6, pp. 982–1006, June 1, 2011, doi:10.1007/s10701-010-9460-0.

[14] R. Omnès, "Logical reformulation of quantum mechanics. I. Foundations," J. Stat. Phys., vol. 53, no. 3–4, pp. 893–932, November 1, 1988, doi:10.1007/BF01014230.

[15] H. F. Dowker and A. Kent, "On the consistent histories approach to quantum mechanics," J. Stat. Phys., vol. 82, no. 5–6, pp. 1575–1646, March 1, 1996, doi:10.1007/BF02183396.

[16] J. J. Halliwell, "Somewhere in the universe: Where is the information stored when histories decohere?" Phys. Rev. D, vol. 60, no. 10, art. 105031, October 26, 1999, doi:10.1103/PhysRevD.60.105031.

[17] Ian Staley, "Final-state constraints and informational pruning in quantum histories," Int. J. Quantum Found., vol. 12, no. 2, pp. 719–737, April 11, 2026, doi:10.5281/zenodo.19512844.

[18] M. P. Müller, "Law without law: From observer states to physics via algorithmic information theory," Quantum, vol. 4, art. 301, July 20, 2020, doi:10.22331/q-2020-07-20-301.

[19] M. Gell-Mann and J. B. Hartle, "Strong decoherence," in Proc. 4th Drexel Symp. on Quantum Non-Integrability, 1995, arXiv:gr-qc/9509054.

[20] W. H. Zurek, "Decoherence, einselection, and the quantum origins of the classical," Rev. Mod. Phys., vol. 75, no. 3, pp. 715–775, May 22, 2003, doi:10.1103/RevModPhys.75.715.

[21] J. J. Halliwell, "Decoherent histories and the emergent classicality of local densities," Phys. Rev. Lett., vol. 83, no. 13, pp. 2481–2485, September 27, 1999, doi:10.1103/PhysRevLett.83.2481.

[22] C. J. Isham, "Quantum logic and the histories approach to quantum theory," J. Math. Phys., vol. 35, no. 5, pp. 2157–2185, May 1, 1994, doi:10.1063/1.530544.

[23] M. O. Scully and K. Drühl, "Quantum eraser: A proposed photon correlation experiment concerning observation and 'delayed choice' in quantum mechanics," Phys. Rev. A, vol. 25, no. 4, pp. 2208–2213, April 1, 1982, doi:10.1103/PhysRevA.25.2208.

[24] Y.-H. Kim, R. Yu, S. P. Kulik, Y. Shih, and M. O. Scully, "Delayed 'choice' quantum eraser," Phys. Rev. Lett., vol. 84, no. 1, pp. 1–5, January 3, 2000, doi:10.1103/PhysRevLett.84.1.

[25] R. J. Solomonoff, "A formal theory of inductive inference. Part I," Inf. Control, vol. 7, no. 1, pp. 1–22, March 1, 1964, doi:10.1016/S0019-9958(64)90223-2.

[26] Ll. Masanes, T. D. Galley, and M. P. Müller, "The measurement postulates of quantum mechanics are operationally redundant," Nat. Commun., vol. 10, art. 1361, March 26, 2019, doi:10.1038/s41467-019-09348-x.

[27] M. P. Müller, "Probabilistic theories and reconstructions of quantum theory (Les Houches 2019 lecture notes)," SciPost Phys. Lect. Notes, art. 28, August 11, 2021, doi:10.21468/SciPostPhysLectNotes.28.

[28] C. H. Bennett, P. Gács, M. Li, P. M. B. Vitányi, and W. H. Zurek, "Information distance," IEEE Trans. Inf. Theory, vol. 44, no. 4, pp. 1407–1423, July 1, 1998, doi:10.1109/18.681318.

[29] M. Li and P. Vitányi, An Introduction to Kolmogorov Complexity and Its Applications, 4th ed. New York, NY, USA: Springer, 2019, doi:10.1007/978-3-030-11298-1.

[30] B. Bauwens and A. Shen, "Information distance revisited," in Proc. 37th Symp. on Theoretical Aspects of Computer Science (STACS 2020), vol. 154, art. 46, March 4, 2020, doi:10.4230/LIPIcs.STACS.2020.46.

[31] C. E. Shannon, "A mathematical theory of communication," Bell Syst. Tech. J., vol. 27, no. 3, pp. 379–423, July 1, 1948, doi:10.1002/j.1538-7305.1948.tb01338.x.

[32] R. Landauer, "Irreversibility and heat generation in the computing process," IBM J. Res. Develop., vol. 5, no. 3, pp. 183–191, July 1, 1961, doi:10.1147/rd.53.0183.

[33] D. Wallace, The Emergent Multiverse: Quantum Theory According to the Everett Interpretation. Oxford, UK: Oxford Univ. Press, 2012, doi:10.1093/acprof:oso/9780199546961.001.0001.

[34] L. Diósi, "Anomalies of weakened decoherence criteria for quantum histories," Phys. Rev. Lett., vol. 92, no. 17, art. 170401, April 30, 2004, doi:10.1103/PhysRevLett.92.170401.

[35] S. Kochen and E. P. Specker, "The problem of hidden variables in quantum mechanics," J. Math. Mech., vol. 17, no. 1, pp. 59–87, 1967.

[36] A. A. Klyachko, M. A. Can, S. Binicioğlu, and A. S. Shumovsky, "Simple test for hidden variables in spin-1 systems," Phys. Rev. Lett., vol. 101, no. 2, art. 020403, July 11, 2008, doi:10.1103/PhysRevLett.101.020403.

[37] A. Cabello, "Experimentally testable state-independent quantum contextuality," Phys. Rev. Lett., vol. 101, no. 21, art. 210401, November 17, 2008, doi:10.1103/PhysRevLett.101.210401.

[38] V. D'Ambrosio, I. Herbauts, E. Amselem, E. Nagali, M. Bourennane, F. Sciarrino, and A. Cabello, "Experimental implementation of a Kochen–Specker set of quantum tests," Phys. Rev. X, vol. 3, no. 1, art. 011012, February 28, 2013, doi:10.1103/PhysRevX.3.011012.

[39] Č. Brukner, "A no-go theorem for observer-independent facts," Entropy, vol. 20, no. 5, art. 350, May 8, 2018, doi:10.3390/e20050350.

[40] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, "Proposed experiment to test local hidden-variable theories," Phys. Rev. Lett., vol. 23, no. 15, pp. 880–884, October 13, 1969, doi:10.1103/PhysRevLett.23.880.

[41] B. S. Cirel'son (Tsirelson), "Quantum generalizations of Bell's inequality," Lett. Math. Phys., vol. 4, no. 2, pp. 93–100, March 1, 1980, doi:10.1007/BF00417500.

[42] M. Navascués, S. Pironio, and A. Acín, "A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations," New J. Phys., vol. 10, art. 073013, July 9, 2008, doi:10.1088/1367-2630/10/7/073013.

[43] R. Łapkiewicz, P. Li, C. Schaeff, N. K. Langford, S. Ramelow, M. Wieśniak, and A. Zeilinger, "Experimental non-classicality of an indivisible quantum system," Nature, vol. 474, no. 7352, pp. 490–493, June 23, 2011, doi:10.1038/nature10119.

[44] D. Kipping, "A Bayesian approach to the simulation argument," Universe, vol. 6, no. 8, art. 109, August 3, 2020, doi:10.3390/universe6080109.

[45] S. P. Walborn, M. O. Terra Cunha, S. Pádua, and C. H. Monken, "Double-slit quantum eraser," Phys. Rev. A, vol. 65, no. 3, art. 033818, February 28, 2002, doi:10.1103/PhysRevA.65.033818.

[46] C. A. Fuchs, "On participatory realism," in Information and Interaction: Eddington, Wheeler, and the Limits of Knowledge, I. T. Durham and D. Rickles, Eds. Cham, Switzerland: Springer, 2017, pp. 113–134, doi:10.1007/978-3-319-43760-6_7, arXiv:1601.04360.

[47] J. B. DeBrota, C. A. Fuchs, and R. Schack, "Respecting one's fellow: QBism's analysis of Wigner's friend," Found. Phys., vol. 50, no. 12, pp. 1859–1874, December 1, 2020, doi:10.1007/s10701-020-00369-x.

[48] A. Di Biagio and C. Rovelli, "Stable facts, relative facts," Found. Phys., vol. 51, no. 1, art. 30, April 5, 2021, doi:10.1007/s10701-021-00429-w.

[49] J. Pienaar, "QBism and relational quantum mechanics compared," Found. Phys., vol. 51, no. 5, art. 96, October 11, 2021, doi:10.1007/s10701-021-00501-5.

[50] M. P. Müller, "Could the physical world be emergent instead of fundamental, and why should we ask?" manuscript, 2017, arXiv:1712.01826.

[51] C. Rovelli, "Relational quantum mechanics," Int. J. Theor. Phys., vol. 35, no. 8, pp. 1637–1678, August 1, 1996, doi:10.1007/BF02302261.

[52] E. G. Cavalcanti, A. Di Biagio, and C. Rovelli, "On the consistency of relative facts," Eur. J. Philos. Sci., vol. 13, no. 4, art. 55, October 20, 2023, doi:10.1007/s13194-023-00551-8.

[53] E. Adlam and C. Rovelli, "Information is physical: Cross-perspective links in relational quantum mechanics," Philos. Phys., vol. 1, art. 4, 2023, doi:10.31389/pop.8.

[54] J. Pienaar, "A quintet of quandaries: Five no-go theorems for relational quantum mechanics," Found. Phys., vol. 51, no. 5, art. 97, October 11, 2021, doi:10.1007/s10701-021-00500-6.

[55] M. Tegmark, "The mathematical universe," Found. Phys., vol. 38, no. 2, pp. 101–150, February 1, 2008, doi:10.1007/s10701-007-9186-9.

[56] J. Gorard, "Some quantum mechanical properties of the Wolfram model," Complex Syst., vol. 29, no. 2, pp. 537–598, June 15, 2020, doi:10.25088/ComplexSystems.29.2.537.

[57] J. Gorard, "Some relativistic and gravitational properties of the Wolfram model," Complex Syst., vol. 29, no. 2, pp. 599–654, June 15, 2020, doi:10.25088/ComplexSystems.29.2.599.

[58] S. Wolfram, "A class of models with the potential to represent fundamental physics," Complex Syst., vol. 29, no. 2, pp. 107–536, June 15, 2020, doi:10.25088/ComplexSystems.29.2.107.

[59] J. Gorard, "Algorithmic causal sets and the Wolfram model," manuscript, 2021, arXiv:2111.03460.

[60] D. Deutsch, The Fabric of Reality. London, UK: Penguin, 1997.

[61] B. S. DeWitt and N. Graham, Eds., The Many-Worlds Interpretation of Quantum Mechanics. Princeton, NJ, USA: Princeton Univ. Press, 1973.

[62] H. Everett III, "'Relative state' formulation of quantum mechanics," Rev. Mod. Phys., vol. 29, no. 3, pp. 454–462, July 1, 1957, doi:10.1103/RevModPhys.29.454.

[63] R. Bousso and L. Susskind, "Multiverse interpretation of quantum mechanics," Phys. Rev. D, vol. 85, no. 4, art. 045007, February 14, 2012, doi:10.1103/PhysRevD.85.045007.

[64] Ian Staley, "Pruned histories and the computational multiverse," SSRN preprint, April 26, 2026.

[65] D. Frauchiger and R. Renner, "Quantum theory cannot consistently describe the use of itself," Nat. Commun., vol. 9, art. 3711, September 18, 2018, doi:10.1038/s41467-018-05739-8.

[66] G. J. Chaitin, "A theory of program size formally identical to information theory," J. ACM, vol. 22, no. 3, pp. 329–340, July 1, 1975, doi:10.1145/321892.321894.

[67] L. A. Levin, "Universal sequential search problems," Probl. Inf. Transm., vol. 9, no. 3, pp. 265–266, 1973.

[68] R. Cilibrasi and P. M. B. Vitányi, "Clustering by compression," IEEE Trans. Inf. Theory, vol. 51, no. 4, pp. 1523–1545, April 1, 2005, doi:10.1109/TIT.2005.844059.

[69] J. Birch, "On the 'simulation argument' and selective scepticism," Erkenntnis, vol. 78, no. 1, pp. 95–107, February 1, 2013, doi:10.1007/s10670-012-9400-9.

Appendix A. Full statement and proof sketch of the correspondence theorem

Theorem A.1 (Full version). Let ($\mathscr{H}$, ρ, H) be a closed quantum system with $\mathscr{H} = \mathscr{H}_S \otimes \mathscr{H}_E \otimes \mathscr{H}_R$, where $\mathscr{H}_R$ is an observer record subspace. Let {α} be a family of histories with class operators C_α as in (2.2). Fix an observer state $O \in \mathscr{H}_R$ with minimal-description string $o \in \{0,1\}^*$, and let $r_O(\alpha)$ denote a minimal binary description of the projected state $P_O C_\alpha \rho |_R$. Let d be the prefix Bennett–Gács–Li–Vitányi–Zurek information distance $d(x,y) = \max\{K(x|y), K(y|x)\}$. Let $K \in \mathbb{R}_{\geq 0}$.

Then the family {α} is pruning-admissible at budget K iff

$$(\forall \alpha \neq \alpha')\ \ D(\alpha, \alpha') = 0 \ \text{ whenever } \ d(r_O(\alpha), r_O(\alpha')) > K + O(\log L), \qquad (A.1)$$

where $L = \max_\alpha \ell(r_O(\alpha))$. Furthermore, the family of rendering events of Campbell–Owhadi type relative to observer O coincides, in the $K \to \infty$ limit, with the family of medium-decoherence-respecting coarse-graining events.

Proof structure. (i) Records-availability lemma: information distinguishing α from α′ is available iff there exist orthogonal record projectors P_α, $P_{\alpha'}$; by Halliwell [4] (Theorem 2 and subsequent records-construction), this is equivalent to medium decoherence of the {α, α′}-pair within the records subspace. (ii) Compression equivalence: orthogonal record projectors correspond to algorithmically distinct r_O bit strings, with d-distance bounded below by the Hilbert–Schmidt distance of the projectors up to an O(log L) term, in line with the prefix-information-distance characterisation of Bennett et al. [23] and the super-logarithmic-separation regime analysed by Bauwens and Shen [25]. (iii) Combining (i)–(ii) gives the biconditional (A.1). (iv) The $K \to \infty$ statement follows from the fact that $\mathbb{1}[d \leq K]$ is identically 1 in this limit, recovering medium decoherence. ∎

The full proof carefully handles the projector-record identification under mixed-state initial conditions (where Halliwell's records are imperfect) and the additive logarithmic corrections inherited from the information-distance metric.

Appendix B. Worked example — three-history Stern–Gerlach

Take a spin-½ particle in superposition $|\psi_0\rangle = \alpha|+\rangle + \beta|-\rangle$, passing through a z-axis Stern–Gerlach magnet. Three coarse-grained histories: α_+ = "deflected up and recorded up," α_- = "deflected down and recorded down," α_0 = "no record extracted." Class operators:

$$C_+ = R_+ \hat{P}^z_+, \quad C_- = R_- \hat{P}^z_-, \quad C_0 = R_0 (\mathbb{1} - \hat{P}^z_+ - \hat{P}^z_-), \quad (B.1)$$

where R_± are record projectors on the ancilla.

The decoherence functional satisfies $D(\alpha_+, \alpha_-) = 0$ trivially by orthogonal P^z. The non-trivial entry is $D(\alpha_+, \alpha_0)$. At $K = \infty$, the records R_+ and R_0 are orthogonal, $d(r(\alpha_+), r(\alpha_0)) \gg 0$, and medium decoherence demands $D(\alpha_+, \alpha_0) = 0$; standard Born statistics emerge: $p(\alpha_+) = |\alpha|^2$, $p(\alpha_-) = |\beta|^2$, $p(\alpha_0) = 0$ (since "no record" is incompatible with a successful measurement).

At finite K, if the records R_+ and R_0 are not yet well-separated — for instance, if the ancilla is weakly coupled and the records are barely resolved, so that the algorithmic-information cost of distinguishing them, $d(r(\alpha_+), r(\alpha_0))$, remains within the observer's budget K — then $\mathbb{1}[d \leq K] = 1$, the admissibility functional permits $D(\alpha_+, \alpha_0) \neq 0$, and the observer assigns nonzero probability to interference between "recorded up" and "no record" branches. The empirical signature is a measurable deviation from binary-outcome Stern–Gerlach statistics in the regime of weak ancilla coupling — a regime experimentally probed in the weak-measurement literature.

Appendix C. Resource-budget estimate for proposed experiments

Protocol 1 (KCBS-with-erasure). Components: heralded single-photon source (PPKTP crystal + 405 nm 100 mW pump diode): ~$30k; cascaded interferometers with motorised waveplates: ~$50k; four SPCMs at $7k each: $28k; Pockels cell with high-voltage driver: $20k; timing electronics (TDC, FPGA): $20k; optical table + ancillaries: $30k. Total: ~$180k. Running time for 10^6 coincidences per τ-bin × 10 τ-bins at 100 kHz coincidence rate: ~30 hours.

Protocol 2 (delayed-choice eraser with K-scan). Slightly simpler optical layout; estimated total ~$140k. Running time for 40 θ-bins at 10^5 coincidences each: ~60 hours.

Both protocols are within range of mid-tier photonic-quantum laboratories without requiring novel hardware.