

The Thermodynamic Cost of a Final Boundary Condition: A Resource Theory of Informational Pruning and a Rank-Indexed Contextuality Conjecture

Ian Staley

Independent Researcher

ORCID: 0009-0000-8592-3186

Correspondence: ian.t.staley@gmail.com

Abstract

In prior work I modelled the macroscopic development of a closed quantum system as a directed graph of coarse-grained histories on which a final boundary condition induces an admissible subgraph, and I graded terminal conditions into a trivial, structured, and computational typology. That programme left two structural gaps: it imported its formal backbones rather than proving results of its own, and it never joined its pruning half to its agency-and-contextuality half or assigned pruning a thermodynamic cost. This paper addresses both. I construct a resource theory whose objects are two-boundary-conditioned decoherent-history ensembles, whose free objects are those in which the terminal condition leaves the history distribution unchanged, and whose free operations are boundary-compatible coarse-grainings. *Theorem 1* establishes that the relative entropy of the conditioned against the unconditioned history distribution is a faithful monotone under every such coarse-graining, jointly convex in the pair of history distributions; I also exhibit a counterexample showing that the more obvious counting-based deficit is *not* a monotone, which is why the relative-entropy form is the correct object. *Theorem 2* establishes that a register satisfying the Reeb-Wolf hypotheses, carried from the unconditioned to the conditioned distribution, deposits heat at least $k_B T \ln 2$ times the entropy difference, and, combined with the Craig rank-product bound, at least $k_B T \ln[n_0 / (r_\alpha r_\omega)]$ for equidistributed history weights. Lower terminal rank therefore raises the derived cost floor under the stated assumptions. I am explicit that this is a floor and not a total ordering: rank does not by itself order terminal conditions by resource content, and full rank does not imply a trivial condition. I then advance, explicitly as a conjecture rather than a theorem, a bridge to the contextuality half of the programme: along a nested family of normalized projectors terminating at the identity, within the paradoxical class, and against a noncontextual baseline, the Contextuality-by-Default measure of pre- and post-selected cyclic statistics should be non-increasing in the family index. I do not claim that the universe physically erases histories; the more limited claim is that a register implementing admissibility filtering pays a rank-graded cost. Conjecture C is, in the first instance, a mathematical conjecture about quantum statistics, settleable by computation, not a proposed departure from quantum mechanics.

Keywords: decoherent histories; final boundary conditions; quantum resource theories; Landauer's principle; quantum contextuality; Contextuality-by-Default; algorithmic information; quantum foundations

1. Introduction

A final boundary condition prunes. This is not controversial: in the two-boundary formulation of generalized quantum mechanics developed by Gell-Mann and Hartle [3], [4], imposing a terminal density operator ρ_f alongside the initial ρ_i reduces the set of histories carrying non-zero probability, and Craig [5] proved that in any decohering set the number of such histories is bounded above by the product of the ranks, $r_\alpha r_\omega$.

The literature has asked what this pruning *permits* — which retrodictions become valid, which paradoxes arise, which sets are consistent. It has not asked what it *costs*. That omission is not a failure of the field; it is a consequence of how the question has been framed. Pruning has been treated as a feature of the formalism rather than as an operation performed on a physical register, and operations on registers are the things thermodynamics knows how to price. The moment one asks who or what implements the filter,

Landauer's principle [21] becomes unavoidable, because filtering is logically irreversible: it maps a broad distribution over histories onto a narrow one, and entropy that leaves the register must appear as heat in a bath.

This paper asks the cost question and answers it. In doing so it addresses a second gap. My previous work has run on two tracks that never met. One concerns pruning, final-state constraints, and the informational grading of terminal conditions [1], [65], [66]. The other concerns agency, substrate-flexibility, and the operational signatures by which one would recognize an epistemic system at all [63], [64]. The first speaks the language of histories and description length; the second the language of contextuality and generalized-noncontextual ontological models. They have shared a vocabulary of constraint without sharing a formalism. A resource theory supplies the missing joint, because a resource theory is exactly the structure in which one asks what a constraint costs and what it buys.

The paper is organized around five questions the existing literature has not asked.

- Q1. What is the thermodynamic cost of a final boundary condition? If imposing an admissible subgraph is a selection process, what heat does it dissipate?
- Q2. Is the trivial/structured/computational typology of terminal conditions a monotone ordering? It was introduced informally in [1]; does a resource-theoretic order recover it, and how much of it?
- Q3. Does the rank of the terminal condition control contextuality? Post-selection has been linked to contextuality through anomalous weak values [44], [45] and through quasiprobability negativity [31], [32]; nobody has asked whether terminal rank is itself a controlling variable.
- Q4. Is there a complexity-contextuality interface? Kolmogorov complexity and contextuality have never been formally related. Does terminal-condition logical depth set a floor on achievable contextual signatures?
- Q5. Can boundary-induced contextuality distinguish substrates? Does dual-boundary conditioning yield signatures on engineered biological substrates unavailable to simplex-embeddable classical learners?

Sections 4 and 5 answer Q1 and, partially, Q2 with theorems. Section 7 addresses Q3 with a conjecture, stated as such. Q4 and Q5 are scoped in Section 11 as open, and I do not pretend otherwise. I follow the claim discipline of my earlier papers throughout: formal, informational, and interpretive claims are separated, and a non-claims table appears in Section 10.

2. Background: assembling a permission structure

2.1 Two-boundary decoherent histories

The decoherent-histories programme originates with Griffiths [2] and reaches its standard form in Gell-Mann and Hartle [3]. A history is a time-ordered sequence of coarse-grained propositions, represented by a class operator C_α formed as a product of Heisenberg-picture projectors drawn at each time from an exhaustive orthogonal set. Hartle's time-neutral generalization [4] admits a terminal as well as an initial condition, giving the two-boundary decoherence functional

$$D(\alpha, \alpha') = N^{-1} \text{Tr}[\rho f C_\alpha \rho_i C_{\alpha'}^\dagger], \quad N = \sum \alpha \text{Tr}[\rho f C_\alpha \rho_i C_\alpha^\dagger], \quad (1)$$

with medium decoherence the condition that off-diagonal entries vanish, under which the diagonal is a probability distribution. Craig [5] proved the rank-product bound: in any decohering set the number of

histories of non-zero probability is at most $r_\alpha r_\omega$. The pure-state corollary — at most two histories when both boundaries are pure — appears in his published cosmological paper [6]. The provenance matters enough to state precisely: the two-boundary functional is due to Gell-Mann and Hartle [3], [4]; the counting bound is Craig's, and its primary source [5] appears never to have been formally published.

The two-boundary setting has an older probabilistic ancestor in the Aharonov-Bergmann-Lebowitz rule [12], which assigns probabilities to intermediate outcomes conditioned on both a preparation and a post-selection. I use it only as motivation, and flag at the outset the standard caution that its counterfactual reading is not valid in general [13], and that its relation to consistent-histories reasoning requires care [14]. Nothing in Sections 3 to 5 depends on a counterfactual construal of ABL; the objects there are ordinary conditional distributions over a decoherent family. Griffiths has more recently developed a consistent-quantum-causes framework [69], against which I positioned an earlier paper [66]. The set-selection problem — that mutually incompatible consistent sets exist and the formalism does not by itself choose among them — was pressed by Dowker and Kent [7], [8] and remains live, though Strasberg and collaborators [11] argue it substantially dissolves for slow, coarse observables.

Halliwell's records programme [9], [10] supplies the complementary intuition that decoherence is accompanied by information deposited in the environment, which is the intuition this paper monetizes. I use it as motivation and not as a lemma, for a specific reason: the correspondence between medium decoherence and the existence of record projectors is an equivalence only for *pure* initial states. Halliwell shows that for mixed initial states records may still exist but are necessarily imperfect, and that where decoherence relies on the impurity of the initial state there are no records perfectly correlated with the past alternatives [9]. Boundary-conditioned ensembles are generically mixed, so this is precisely the regime in which the correspondence weakens. Nothing in Theorems 1 or 2 depends on it: the register in Section 5 is posited as part of the setting, not derived from a records theorem.

2.2 Resource theories

A quantum resource theory in the sense of Chitambar and Gour [15] is a pair consisting of a set of free objects and a set of free operations, together with monotones — functions non-increasing under the free operations, vanishing on free objects, and typically convex. Recent work has pushed resource theories beyond states to processes: Milz, Bavaresco and Chiribella construct a resource theory of causal connection [16]; Berk and collaborators construct resource theories of multi-time processes [17], [18]; Zambon and Adesso treat quantum processes as thermodynamic resources [19]. All take process matrices or quantum combs [20] as their objects.

None takes history families as its objects. The nearest recent work in spirit is De Saegher's treatment of decoherent-histories contextuality [70], which argues that Bell-Kochen-Specker contextuality persists into the classical limit of decoherent histories; that is adjacent to the bridge attempted in Section 7 but is not a resource theory, has no free operations or monotone, and does not index anything by terminal rank. I searched deliberately for a resource theory over decoherent-history sets or boundary-conditioned ensembles and found none. The construction in Section 3 therefore appears to be new, and I state that with the appropriate hedge: to my knowledge, no prior work defines a resource theory whose objects are boundary-conditioned decoherent-history ensembles.

2.3 Landauer's principle and the thermodynamics of post-selection

Landauer's principle [21] in its modern rigorous form is due to Reeb and Wolf [22], who prove an equality-form statement: for a system coupled to a thermal reservoir by a global unitary, with system and reservoir

initially uncorrelated and the reservoir in a Gibbs state, the heat deposited in the reservoir equals the system's entropy decrease plus a mutual-information term and a relative-entropy term, both non-negative. Dropping them recovers the familiar inequality; retaining them gives finite-reservoir corrections. Related statements appear in Faist and collaborators [23], Anderson on conditional operations [24], Sagawa and Ueda on the cost of measurement [25], and Chiribella, Yang and Renner [26]. Hsieh [27] has recently formulated a dynamical Landauer principle for transmission.

Post-selection has its own thermodynamics. Purves and Short [28] analyze channels and measurements with post-selection; Boes and collaborators [29] show how post-selection can bypass fluctuation theorems; Arvidsson-Shukur and collaborators [30] demonstrate metrological advantage from post-selection and trace it to negativity of the Kirkwood-Dirac distribution [33], [34]. Lostaglio [31], [32] shows that negativity of the work quasiprobability certifies contextuality. Two cautions are essential and I take them seriously throughout. First, Kirkwood-Dirac negativity is not necessary for anomalous quantum Fisher information, as Das and collaborators argue [35] — a claim contested and defended in a subsequent exchange. Second, Kirkwood-Dirac negativity does not by itself certify generalized noncontextuality, as Schmid and collaborators establish [53]. I therefore route every contextuality claim through a bona fide contextuality measure and never through negativity as a proxy.

Critically for Section 7: throughout this literature the controlling variable is the post-selection *probability*, or the quasiprobability negativity, and never the *rank* of the post-selecting operator. The one place rank appears prominently is Xu, Yu and Kleinmann [48], and there it is the rank of the *measurement* projectors in a state-independent Kochen-Specker construction, a different object entirely.

2.4 Contextuality

I use two frameworks. Spekkens' generalized noncontextuality [36], [37] and its GPT characterization by Schmid and collaborators [38], [39] identify classicality with simplex-embeddability of the accessible GPT fragment. Contextuality-by-Default [40], [41] provides an operational measure, CNT_2 , for cyclic systems of n binary measurements, of which the KCBS scenario [42] is the $n = 5$ case; the Kochen-Specker landscape is reviewed by Budroni and collaborators [43]. Contextuality is itself a resource [49], [50], [51], [52]. On the post-selection side, Leifer and Spekkens [46] and Pusey [44] establish that pre- and post-selection paradoxes are contextual, with robustness analysed by Kunjwal, Lostaglio and Pusey [45]. Singh, Bhati and Arvind [47] prove the crucial complementary result: there is no contextual advantage in *non-paradoxical* scenarios of the two-state vector formalism. Any claim relating terminal conditions to contextuality must therefore be conditional on scenario class, and mine is.

3. The formal model

3.1 Objects

Fix a finite-dimensional Hilbert space of dimension d . A **boundary-conditioned history ensemble** is a tuple

$$E = (\mathcal{A}, \{C\alpha\}_{\alpha \in \mathcal{A}}, \rho_i, \rho_f) \quad (2)$$

where the index set is finite, the class operators sum to the identity, ρ_i is a density operator of rank r_ω and ρ_f is a positive operator of rank r_ω . Each ensemble carries two distributions: the *unconditioned* distribution, obtained by setting the terminal condition to the identity,

$$p_0(\alpha) = \text{Tr}[C\alpha \rho_i C\alpha^\dagger] / \sum_{\alpha'} \text{Tr}[C\alpha' \rho_i C\alpha'^\dagger], \quad (3)$$

and the conditioned distribution obtained from the full two-boundary functional (1),

$$pf(\alpha) = \text{Tr}[\rho_f C \alpha \rho_i C \alpha^\dagger] / \sum \alpha' \text{Tr}[\rho_f C \alpha' \rho_i C \alpha'^\dagger]. \quad (4)$$

Figure 1 illustrates the construction. I call the ensemble **admissible** when the family is medium-decoherent with respect to both the trivial and the actual terminal condition, so that both distributions are genuine probability distributions rather than merely diagonal entries of a non-decohering functional. All results below are stated for admissible ensembles. This is a real restriction and I flag it as such: on non-decoherent families the objects of the theory are not defined, and the theory is silent.




Figure 1. A boundary-conditioned history ensemble. Above, the terminal condition is trivial and every history of non-zero unconditioned weight survives. Below, a terminal condition of rank r_ω renders most histories inadmissible (dashed). Craig's bound caps the surviving set at $r_\alpha r_\omega$, and the pruning content measures how far the conditioned distribution has been driven from the unconditioned one.

3.2 A support lemma

Lemma 1 (support inclusion). For any boundary-conditioned history ensemble, $\text{supp } pf \subseteq \text{supp } p_0$.

Proof. Suppose $p_0(\alpha) = 0$, so that $\text{Tr}[C \alpha \rho_i C \alpha^\dagger] = 0$. The operator $C \alpha \rho_i C \alpha^\dagger$ is positive semi-definite, and a positive semi-definite operator of vanishing trace is the zero operator. Hence $C \alpha \rho_i C \alpha^\dagger = 0$, so $\text{Tr}[\rho_f C \alpha \rho_i C \alpha^\dagger] = 0$ and $pf(\alpha) = 0$. ■

Lemma 1 is elementary but it licenses reading pf as a pruning of p_0 rather than as an unrelated distribution: a terminal condition can only remove histories from the support, never add them.

3.3 Free objects and free operations

Free objects. An admissible ensemble is *free* when its terminal condition induces no change in the history distribution, that is, when $pf = p_0$. The definition is operational rather than representational, and the distinction is not cosmetic. Every ensemble with ρ_f proportional to the identity is free, but the converse fails: a terminal operator not proportional to the identity may act proportionally to it on the subspace actually explored by a given history family, in which case $pf = p_0$ although ρ_f is nontrivial. Freedom is thus a property of the pair (terminal condition, history family), not of the terminal operator alone. I adopt the operational definition because it is what makes the monotone of Section 4 faithful rather than merely vanishing on a subclass, and because resources are properly individuated by their operational effects. The trivial class of the typology in [1] is contained in the free class, and coincides with it only when the history family is rich enough to detect any nontriviality in the terminal condition.

Free operations. A *boundary-compatible coarse-graining* is determined by a partition of the index set, with merged class operators formed by summing within each block, subject to the requirement that the coarse-grained family remain medium-decoherent with respect to both terminal conditions and that both boundary operators be left untouched. These are the only free operations for which I prove anything below. Appending an ancilla carrying a free ensemble, and discarding a subsystem, are natural further candidates, but defining them properly requires a tensor product of boundary-conditioned ensembles — composite index sets, composite class operators, composite boundaries, and admissibility of the product — which I do not develop here. Theorem 1 is therefore stated for boundary-compatible coarse-grainings alone, with composition and marginalization recorded as expected extensions rather than established components of the theory.


Lemma 2 (coarse-graining is marginalization). On an admissible ensemble, a boundary-compatible coarse-graining acts on both distributions as the same deterministic marginalization map, summing weights within each block.

Proof. By bilinearity of the decoherence functional, the diagonal entry of a merged block is the double sum of the functional over its constituents. Medium decoherence kills every off-diagonal term, leaving the sum of the constituent diagonal entries. The same computation applies with the terminal operator replaced by the identity. ■

4. Theorem 1: the pruning monotone

The obvious candidate monotone is a counting quantity — the number of histories the terminal condition kills, or the shortfall of the admissible set against the Craig bound. I record at once that this candidate fails, because the failure is instructive and because a reader who has followed the programme this far will reach for it first.

Remark 1 (the counting deficit is not a monotone). Define the counting deficit as the base-two logarithm of the ratio of the unconditioned to the conditioned support size. This is not non-increasing under boundary-compatible coarse-graining. Let the unconditioned support be $\{1,2,3,4\}$ and the conditioned support $\{1,2\}$, so the deficit is $\log_2(4/2) = 1$. Coarse-grain by the partition $\{1,2\}, \{3\}, \{4\}$. The merged block $\{1,2\}$ meets the conditioned support, and the blocks $\{3\}$ and $\{4\}$ each meet the unconditioned support, so the coarse-grained value is $\log_2(3/1) \approx 1.585 > 1$. Merging surviving histories while keeping the pruned ones distinct inflates the deficit. ■



The deficit increased under a free operation so Φ_{count} is not a monotone.
Merging surviving histories while leaving pruned ones separate inflates the count

By contrast $\Phi = D(p^f \parallel p_0)$ is non-increasing under every such map, by data processing

Filled = admissible (supp p^f) · Dashed = pruned

Figure 2. The counterexample of Remark 1. Coarse-graining that merges two surviving histories while leaving the pruned ones separate raises the counting deficit from 1 to $\log_2 3$, so the counting quantity increases under a free operation and cannot be a monotone. The relative-entropy pruning content is immune to this failure because it tracks weight rather than support, and weight is conserved by coarse-graining.

The failure is displayed in Figure 2. The lesson is that a monotone must be sensitive to weights, not merely to supports, because coarse-graining redistributes weight but preserves it. The correct object is a divergence.

Definition (pruning content). For an admissible ensemble, the *pruning content* is the Kullback-Leibler divergence of the conditioned against the unconditioned history distribution,

$$\mathcal{P}(E) := D(p_f \parallel p_0) = \sum \alpha p_f(\alpha) \log_2 [p_f(\alpha) / p_0(\alpha)]. \quad (5)$$

The expression is well defined by Lemma 1, which guarantees p_0 is strictly positive wherever p_f is, so no term diverges.

Theorem 1. The pruning content is a resource monotone in the sense of Chitambar and Gour [15], with respect to the free objects and boundary-compatible coarse-grainings of Section 3.3:

- (i) Non-negativity and faithfulness. The pruning content is non-negative, with equality if and only if $p_f = p_0$, that is, if and only if the ensemble is free.
- (ii) Monotonicity. For every boundary-compatible coarse-graining, the pruning content is non-increasing.
- (iii) Convexity. It is jointly convex in the associated pair of history distributions. Where a probabilistic mixture of admissible ensembles induces the corresponding mixture of those distributions and remains admissible, the monotone is convex under that mixture; I do not construct ensemble mixing as part of the theory, so the claim is stated at the level of the distribution pair.

Proof. (i) is Gibbs' inequality together with the observation that the free objects were *defined* as those with $p_f = p_0$; faithfulness is therefore genuine rather than one-directional. Had freedom instead been defined by the terminal operator being proportional to the identity, only the implication $\text{free} \Rightarrow \text{zero pruning content}$ would follow, since a nontrivial terminal operator can leave the history distribution unchanged. For (ii), Lemma 2 shows that a coarse-graining acts on both distributions as one and the same deterministic stochastic map, namely the marginalization induced by the partition. Relative entropy is monotone under every stochastic map — the data-processing inequality [67] — which gives the result. (iii) is the standard joint convexity of relative entropy. ■

Two corollaries make the monotone concrete.

Corollary 1 (pruning content as surprisal). If the terminal condition acts by pure conditioning — if p_f is p_0 restricted to a subset and renormalized — then the pruning content equals $-\log_2 q$, where q is the unconditioned weight of the admissible set. The pruning content is exactly the surprisal of the terminal condition.

Corollary 2 (rank floor). In general, by the log-sum inequality, the pruning content is at least the negative base-two logarithm of the unconditioned weight of the conditioned support. If the unconditioned weights are equidistributed over n_0 histories, then since the conditioned support is bounded by $r\alpha$ $r\omega$ by Craig's bound [5],

$$\mathcal{P}(E) \geq \log_2 [n_0 / (r\alpha r\omega)]. \quad (6)$$

Corollary 2 bears on Q2, though it answers it only partially, and the limits are worth stating carefully because an earlier version of this work overstated them. The pruning content is bounded below by a quantity decreasing in the terminal rank, so under equidistribution a lower-rank terminal condition carries a higher *guaranteed minimum* of pruning content. That is a statement about a floor, not about actual values. A high-rank but strongly nonuniform terminal operator may induce a larger divergence than a low-rank one, so rank does not totally order terminal conditions by resource content. Nor does full rank imply triviality: a terminal operator has full rank whenever it is invertible, and an invertible operator such as $\text{diag}(0.8, 0.15, 0.05)$ is far from proportional to the identity and will reweight histories substantially. Full rank is necessary but not sufficient for freedom, and the free class is characterized operationally, as in Section 3.3, rather than by rank. What survives is the modest claim: the free class sits at the bottom of the order, and the Craig bound supplies a rank-dependent floor above it. I do not claim the typology of [1] is recovered as a total ordering — the order induced by the pruning content is a preorder and, as with causal non-separability [16], incomparable pairs are to be expected.

5. Theorem 2: the cost of filtering

Theorem 1 prices pruning in bits. Theorem 2 prices it in joules. The setting is deliberately modest: I do not model the universe as a computer performing selection. I model a register that undergoes the corresponding entropy reduction, and ask what thermodynamics requires of it.

Setting. Let R be a register whose Hamiltonian is degenerate on the logical states labelled by the history index set, prepared in the distribution p_0 . Let B be a reservoir prepared in the Gibbs state of its Hamiltonian at inverse temperature $\beta = 1/k_B T$, uncorrelated with R initially, and let the joint evolution be a unitary on $R \otimes B$. Suppose it brings R to the conditioned distribution p_f , and let Q be the heat deposited in the reservoir, defined as the change in the reservoir's mean energy. These are exactly the hypotheses of Reeb and Wolf [22]: product initial state, thermal reservoir, global unitary. Nothing further is assumed about how the process is realized.

Theorem 2 (entropy-reduction cost for an admissibility register). Under the setting above,

$$Q \geq k_B T \ln 2 \cdot [H(p_0) - H(p_f)], \quad (7)$$


where H is the Shannon entropy in bits. Combining with Craig's bound [5], and assuming the unconditioned weights equidistributed over n_0 histories,

$$Q \geq k_B T \ln [n_0 / (r \alpha r \omega)], \quad (8)$$

which is non-vacuous precisely when n_0 exceeds the rank product.

Proof. Equation (7) is the Reeb-Wolf inequality [22] applied to R : their equality-form result expresses βQ as the register's entropy decrease plus a mutual-information term and a relative-entropy term, both non-negative; discarding them yields the inequality, and for a register with degenerate Hamiltonian the von Neumann entropy of the diagonal state is the Shannon entropy of the distribution. For (8), note that $H(p_f)$ is at most the base-two logarithm of the conditioned support size, which is at most the logarithm of the rank product by Craig's bound; the first step is the maximum-entropy bound on a distribution of given support. With $H(p_0) = \log_2 n_0$ under equidistribution, substitution into (7) and conversion of logarithm base gives the result. ■

Theorem 2: cost floor as a function of terminal rank



$Q \geq k_B T \ln[n_0 / (r\alpha r\omega)]$, shown for $n_0 = 64$, $r\alpha = 2$

A rank-dependent floor, not a total ordering: a high-rank but strongly nonuniform pf may still yield larger pruning content than a low-rank one.

Figure 3. The cost floor of equation (8) as a function of terminal rank, for $n_0 = 64$ and $r\alpha = 2$. Lower terminal rank raises the minimum heat that a register satisfying the stated hypotheses must deposit. The bound degrades to vacuity once the rank product reaches n_0 . The curve is a floor, not a resource ordering: a high-rank but strongly nonuniform terminal operator may still induce larger pruning content than a low-rank one.

The content of (8) is the answer to Q1, and it is rank-indexed (Figure 3): the cost floor rises as terminal rank falls, with the exchange rate set by the Craig bound. Three remarks on scope. First, the finite-reservoir corrections of Reeb and Wolf [22] tighten (7) whenever the reservoir is finite, so (7) is conservative. Second, the equidistribution assumption in (8) is a convenience; without it one retains (7) with the unconditioned Shannon entropy in place of $\log_2 n_0$, which is the honest general statement. Third, and most important, the theorem is about a particular implementation: a register carried from p_0 to pf under the Reeb-Wolf hypotheses. It is not a claim about every physical realization of conditional sampling. A laboratory post-selection protocol that retains successful trials and discards the rest has a different and more involved accounting, involving success probability, flag registers, preparation costs, and discarded runs; equation (7) does not automatically apply to it. What the theorem bounds is the entropy reduction of an admissibility register, no more.

I want to be exact about what Theorem 2 does not say, because the temptation to over-read it is strong and I have seen the analogous over-reading in the simulation-theoretic literature. Theorem 2 does not say the universe expends work to prune histories. Pruning in the two-boundary formalism is not a dynamical process, and nothing here makes it one. It says that *if* a physical system carries a record of which histories are admissible, and *if* that record is narrowed as stipulated, *then* the narrowing deposits at least the heat in (7). The bound applies to observers, apparatus, and simulators. It does not apply to the boundary condition itself, which is not the kind of thing that has a register.

6. The informational reading, and its limits

My earlier work [1] stated the central relation algorithmically, as an inequality between the Kolmogorov complexity of the full history graph and that of the admissible subgraph together with the terminal condition. The connection to Theorem 1 runs through Levin's coding theorem, by which complexity is approximated by the negative logarithm of the universal semimeasure up to an additive constant [56]. If the unconditioned weights are algorithmically typical, Corollary 1 gives

$$\mathcal{P}(E) = -\log_2 q \approx K(\varphi) + O(1), \quad (9)$$

identifying the pruning content with the description length of the terminal condition. This is a satisfying identification and I think it is the right one, but I label it heuristic rather than theorem, for three reasons I do not think careful writing can repair. Kolmogorov complexity is uncomputable, so (9) can never be evaluated. The invariance theorem fixes it only up to an additive constant depending on the universal machine, and that constant is not small in any absolute sense. And algorithmic typicality is an assumption about the physics, not a mathematical fact; Müller's programme [68] shows how much structure is needed to make such assumptions bite.

The division of labour I propose is therefore: the pruning content is the operational monotone, computable in principle and provably monotone; the algorithmic quantity is the interpretive ceiling, in the sense that Zurek's physical entropy [54] is an interpretive ceiling over Shannon entropy. Bennett's logical depth [55] then supplies the resource distinguishing structured from computational terminal conditions: a computational endpoint has small description length but large depth, so the pruning content is cheap to specify and expensive to unpack. Making that trade-off precise — a depth-indexed refinement of Theorem 2 — is open, and I do not attempt it. Related algorithmic-thermodynamic machinery appears in Janzing and collaborators [57] and, recently, in Edmonds [58].

7. The contextuality bridge, as a conjecture

Sections 4 to 6 concern one half of the programme. This section attempts the joint with the other half, and it is the part of the paper I hold most loosely.

The motivating observation is structural. Corollary 2 shows that lowering the terminal rank raises the floor on pruning content, and pruning content measures how far the conditioned history distribution has been driven from the unconditioned one. Driving a distribution away from its Born-rule baseline is precisely what generates the anomalies — negative quasiprobabilities, anomalous weak values, pre- and post-selection paradoxes — that Leifer and Spekkens [46] and Pusey [44] showed to be contextual. It is therefore natural to ask whether the terminal *rank* is a controlling variable for contextuality, and, as Section 2.3 recorded, nobody appears to have asked.

Before stating the conjecture I must dispose of an under-specification that an earlier version of this work did not handle. Rank alone does not determine a terminal operator. At each rank there are many choices of support subspace, spectrum, and orientation relative to the measurement projectors, and these will generally give different statistics; fixing the post-selection probability narrows but does not remove the freedom. A conjecture phrased simply as decreasing in the terminal rank therefore has no unique referent. I state it instead over a specified family.

Conjecture C (rank-indexed contextuality). Fix a cyclic system of n binary measurements — the KCBS scenario for $n = 5$ [42] — an initial state, and a target post-selection probability q_0 that is achievable at every rank $r = 1, \dots, d$. For each r let $\mathcal{R}(r, q_0)$ be the set of terminal operators of rank r whose induced post-selection probability equals q_0 , and define $C(r) = \sup \{ \text{CNT}_2(\rho f) : \rho f \in \mathcal{R}(r, q_0) \}$, the supremum taken over ensembles in the paradoxical class of Singh, Bhati and Arvind [47] with an unconditioned baseline that is itself noncontextual. Then C is non-increasing in r , and $C(d) = 0$.

Two points about this statement deserve emphasis, because the obvious alternative formulation is defective. One might instead fix a nested sequence $\Pi_1 \subset \dots \subset \Pi_d = \mathbb{1}$ and set $\rho f^{(r)} = \Pi_r / r$, which is attractively concrete and terminates at the maximally mixed operator. But one cannot then also stipulate a fixed post-selection probability. Along a generic nested family the post-selection probability varies with

r , because the terminal event's probability depends on the overlap between the evolving state and Πr , and enlarging the projector generally enlarges that overlap. Holding q fixed and fixing the family are therefore competing requirements, not compatible ones. I state the conjecture over the constrained set $\mathcal{R}(r, q_0)$ precisely to avoid that conflict: rank and post-selection probability are both pinned, and the orientation and spectrum of the terminal operator are what vary. The nested-family picture remains a legitimate special case, but only for families along which q happens to be constant in r , and I do not claim such families exist in any given scenario. Second, the supremum is over a set that must be non-empty at every rank, which is why achievability of q_0 at each r is a hypothesis of the conjecture rather than an afterthought.

Two endpoints are secure and are the reason I think the conjecture is worth stating.

- (a) At $r = d$ the only rank- d terminal operator that leaves the statistics unconditioned is a multiple of the identity, for which $p_f = p_0$ and, given a noncontextual baseline, CNT_2 vanishes. The endpoint $C(d) = 0$ is therefore a claim about that operator specifically, and it is secure for it. Note carefully what is not being said: full rank does not by itself imply proportionality to the identity, and an earlier formulation of this work wrongly leaned on that implication. A rank- d operator such as $\text{diag}(0.8, 0.15, 0.05)$ reweights histories substantially, so $C(d) = 0$ should be read as pinned by the trivial operator rather than as a consequence of rank alone. Where the constraint set at $r = d$ contains nontrivial operators, the endpoint is a conjecture like the rest of the curve.
- (b) At $r = 1$ the terminal condition is a normalized rank-one projector, and with paradoxical pre- and post-selection the logical pre- and post-selection paradoxes are contextual by Leifer and Spekkens [46] and by Pusey's result on anomalous weak values [44], with robustness quantified by Kunjwal, Lostaglio and Pusey [45]. So a strictly positive measure is attainable at the sharpest terminal condition.

A weaker variant replaces the supremum by an average over a specified measure on $\mathcal{R}(r, q_0)$, which is easier to estimate numerically but requires committing to a measure. Either is well posed; a bare appeal to rank is not, and neither is a nested family combined with fixed q . What remains conjectural in all of them is the *monotone interpolation* between the endpoints. I offer, in the register of the heuristic functional forms of [64] and with the same epistemic status, the ansatz

$$CNT_2(r) \approx C \cdot [\text{erfc}(\kappa(r-1)/\sqrt{r\alpha}) - \varepsilon d] / (1 - \varepsilon d), \quad \varepsilon d := \text{erfc}(\kappa(d-1)/\sqrt{r\alpha}), \quad (10)$$

with C and κ platform-dependent. The normalization by εd is not cosmetic: a bare complementary error function is asymptotic and never reaches zero, so it would fail to reproduce endpoint (a); subtracting its value at $r = d$ and renormalizing enforces a vanishing measure at the trivial endpoint exactly, while preserving the value C at $r = 1$ and monotonicity throughout. Equation (10) is nonetheless an ansatz. It is not derived. Its virtues are that it reproduces both secure endpoints exactly, that it has the complementary-error-function threshold structure I predicted on independent grounds for sub-Tsirelson and sub-KCBS deviations in [64], and that it is falsifiable.

One further caveat belongs here rather than in a footnote, because it is a live assumption and not a technicality. Contextuality-by-Default was constructed for systems that may be *inconsistently connected*, meaning single-variable marginals differ across contexts, and its maximal-coupling machinery accommodates exactly such differences; that is in its favour here, since post-selection generically shifts marginals. But Dzhamalov, Kujala and Cervantes [40] do not treat post-selection, conditioning on terminal outcomes, or pre- and post-selected subensembles anywhere in their development. No theorem there

licenses applying CNT_2 to a post-selected subensemble. I therefore do not claim that measuring CNT_2 on post-selected trials is a sanctioned use of CbD; the more limited claim is that it is a natural extension carrying an assumption which must be discharged rather than assumed — that the post-selected marginals still support couplings of the kind CbD requires. I flag it as the weakest joint in the proposal. A second consequence: CNT_2 as defined in [40] is an unnormalized L_1 distance, and the framework supplies no canonical normalization to the unit interval, so the constant C in (10) is mine to fix and is not inherited.




Figure 4. Conjecture C over the nested family. The two filled endpoints are secure: at $r = d$ the terminal operator is maximally mixed, so $\text{pf} = p_0$ and the measure vanishes given a noncontextual baseline, while at $r = 1$ with paradoxical pre- and post-selection the statistics are known to be contextual. The dashed interpolation is the ansatz of equation (10) and is not derived. The curve is drawn to be falsified: a numerical sweep at fixed post-selection probability showing a flat measure would retire the conjecture without disturbing Theorems 1 and 2.

Figure 4 plots the conjecture, drawn so the secure endpoints and the conjectural interpolation are visually distinguishable. The restriction to the paradoxical class is not a hedge added after the fact; it is forced. Singh, Bhati and Arvind [47] prove that non-paradoxical two-state-vector scenarios exhibit no contextual advantage, so any rank dependence must vanish identically outside that class. This meshes with the typology of [1] in a way I find suggestive rather than conclusive: operationally trivial terminal conditions are free and non-contextual; structured and computational terminal conditions are those capable of placing the ensemble in the paradoxical class.

One methodological point deserves emphasis because it is easy to get wrong and because getting it wrong would misrepresent the paper. Conjecture C is not a proposed departure from quantum mechanics. Everything in Sections 3 to 7 is standard quantum theory with two boundary conditions. Conjecture C is therefore, in the first instance, a *mathematical* conjecture about the behaviour of quantum statistics under rank-varied post-selection, and it is settleable by computation — a semidefinite or linear-programming sweep over the nested family in a qutrit KCBS scenario would decide it, and that is the cheapest next step for anyone who wants to test this paper. The experimental protocols in Section 8 are not tests of quantum mechanics against an alternative. They test whether the predicted dependence survives realistic noise, and whether it is measurable on substrates of independent interest.

8. Falsifiable predictions and protocols

The discriminating statistic is the **rank-resolved contextuality measure at fixed post-selection probability**. The existing literature relating post-selection to nonclassicality varies the post-selection probability [30] or the quasiprobability negativity [31], [32]. The protocol here holds that probability fixed and varies only the family index. Prediction P1: the measure decreases monotonically along the nested family within the paradoxical class, with the threshold structure of (10). Prediction P2: outside the paradoxical class it is flat and consistent with zero, per [47]. P2 is a control; its failure would falsify the framing rather than the conjecture.

8.1 Photonic implementation

A KCBS cyclic scenario on a qutrit encoded in path or orbital-angular-momentum degrees of freedom, with pre-selection on a fixed initial state and post-selection implemented by a programmable multiport — a spatial light modulator or reconfigurable integrated interferometer — realizing the nested family at $r \in \{1, 2, 3\}$. For a qutrit the accessible range is narrow, and I regard that as the principal weakness of the photonic route; a ququart or two-qutrit encoding widens it at the cost of interferometric stability. Holding the post-selection probability fixed while r varies is the essential control, and it is worth being concrete about how, since a vague appeal to compensation would leave the protocol untestable. The matching is done by choosing, at each rank, the orientation of the terminal projector so that its overlap with the evolving state yields the target probability; orientation is a free parameter within the constraint set, and adjusting it changes the physical post-selection probability, which is the quantity that must be matched. Two tempting alternatives do not work. Subsampling already-accepted trials equalizes sample counts but leaves the physical post-selection probability untouched, so it does not implement the constraint. Attenuating the beam or altering detection rates changes the count rate rather than the conditional statistics, and risks introducing a second varying factor alongside the rank. Where no orientation at a given rank achieves the target probability, that rank lies outside the accessible range and should be reported as such rather than matched by other means. This control is also what distinguishes the protocol from a conventional weak-value measurement. The analysis pipeline is the CbD estimator of [40] with bootstrapped confidence intervals; the robustness methodology of [45] transfers directly.

On cost: my earlier estimate of roughly US\$180,000 per protocol in an existing quantum-optics facility, and US\$200,000 to 400,000 to instrument from scratch [64], carries over, since the optical requirements are those of a standard pre- and post-selection contextuality experiment plus a programmable projector. I flag that this figure is my own order-of-magnitude estimate rather than a quotation.

8.2 Engineered cortical wetware

The second protocol extends the substrate-flexibility argument of [63]. A CL1-class culture is driven with a five-context cyclic stimulation protocol in the CbD paradigm, electrophysiological output supplying the binary outcomes, and the terminal condition imposed by post-selecting trials on a terminal population-activity pattern whose *effective* rank — the number of admissible terminal microstates — is tunable by the coarseness of the pattern classifier. The four operational signatures specified in the companion work [63] are then evaluated as functions of that effective rank.

I am obliged to be careful here in two directions. First, identifying a coarse-grained neural activity class with a rank- r terminal operator is an analogy, not a derivation; the culture is not a qutrit and I do not claim it is. What the protocol tests is whether an operational analogue of rank dependence appears, which would be evidence the framework describes something substrate-independent, and whose absence would be uninformative rather than falsifying. Second, no contextuality measurement on a biological neural culture

has been published, so this is a first-of-kind measurement with the attendant risk. Cortical Labs' CL1 [59], [60], the organoid-intelligence programme [61], and the FinalSpark Neuroplatform [62] supply the platforms. Published CL1 specifications are manufacturer-claimed and not peer-reviewed; the widely reported figures are approximately 800,000 human iPSC-derived cortical neurons on a 59-electrode array with sub-millisecond bidirectional latency and viability up to six months, at approximately US\$35,000 per unit. I cite these as manufacturer claims and not as verified specifications.

9. Objections and replies

Objection 1. This is a relabelling of an existing resource theory.

The free operations here — coarse-grainings preserving a terminal condition and both decoherence properties — are not the free operations of any theory I have found. Causal-connection free adapters [16], comb concatenations [17], [20], and contextuality wirings [50] all act on different objects. I do not claim a new mathematical genus; the more limited claim is a new instance within the Chitambar-Gour framework [15] whose objects are two-boundary history ensembles.

Objection 2. Kolmogorov complexity is uncomputable, so the informational apparatus is vacuous.

I agree, and Section 6 concedes it explicitly. The operational monotone is a relative entropy, computable in principle and provably monotone; the algorithmic quantity is an interpretive ceiling only, in the sense of Zurek [54]. Nothing in Theorem 1 or Theorem 2 depends on the algorithmic reading. If Section 6 were deleted the results would stand.

Objection 3. The thermodynamics of post-selection is already worked out.

The existing results concern work extraction and fluctuation theorems under post-selection [28], [29], [30] and are controlled by the post-selection probability. I do not claim to improve them. The more limited claim is a lower bound on the entropy reduction of an admissibility register, indexed by the rank of the terminal condition — a quantity those papers do not compute and have no reason to.

Objection 4. Post-selection does not generate contextuality; Singh, Bhati and Arvind proved as much.

Correct in the non-paradoxical class, and I concede it without reservation. Conjecture C is stated only for the paradoxical class, and Prediction P2 makes the concession an experimental control rather than a caveat. Rank dependence appearing in the non-paradoxical class would falsify my framing, not confirm it.

Objection 5. Kirkwood-Dirac negativity already captures the phenomenon.

Negativity is neither necessary for the metrological advantage [35] nor sufficient for generalized noncontextuality [53]. That is precisely why every contextuality claim here is routed through CbD or generalized-noncontextuality measures and never through negativity. The rank argument is orthogonal to that literature and stands or falls independently of how the dispute resolves.

Objection 6. Landauer's principle governs erasure, and pruning is selection, not erasure.

Agreed, and Section 5 states the scope accordingly. The bound applies to a register carrying admissibility records, because narrowing such a record is logically irreversible whether one calls it selection or erasure. I do not claim the universe erases; the more limited claim is that a register undergoing the stipulated transformation pays.

Objection 7. Conjecture C is not an empirical prediction, since the framework is just quantum mechanics.

This is right and I have said so in Section 7 rather than waiting to be told. Conjecture C is a mathematical conjecture about quantum statistics, and the cheapest decisive test is numerical, not optical. I regard that as a strength: the paper contains a tractable open problem rather than an untestable posture. The experiments test noise-robustness and substrate-independence, which are separate questions worth answering.

Objection 8. Rank does not specify a terminal condition, so the conjecture is ill-posed.

It would be, stated baldly, and Section 7 now says so before an objector has to. Rank fixes neither support subspace nor spectrum nor orientation, and fixing the post-selection probability does not close the gap. The conjecture is therefore stated over a nested normalized-projector family, or equivalently as an optimization over rank- r terminal operators at fixed post-selection probability. Both are well posed; the bare rank formulation is not, and I have withdrawn it.

Objection 9. Full rank is not the same as a trivial terminal condition, so the free end of the order is misidentified.

This is correct and an earlier version of this work got it wrong. An invertible terminal operator has full rank and may still reweight histories heavily. I have therefore defined freedom operationally, by the condition that the history distribution is unchanged, and I have restated the secure endpoint of Conjecture C in terms of a family that terminates at a multiple of the identity rather than in terms of rank. Corollary 2 is correspondingly weakened to a floor rather than an ordering.

10. Explicit non-claims

Following the practice of my earlier papers, I tabulate what is not being asserted.

I do not claim	The more limited claim is
That the universe physically erases or expends work to prune histories.	A register satisfying the Reeb-Wolf hypotheses and carried from p_0 to p_f deposits heat at least $k_B T \ln 2$ times the entropy difference.
That Theorem 2 bounds every physical implementation of post-selection.	It bounds the stipulated register transformation. Trial-discarding laboratory protocols have a different accounting.
That a trivial terminal condition is the same thing as a full-rank one.	Full rank is necessary but not sufficient for freedom; freedom is defined operationally by $p_f = p_0$.
That rank totally orders terminal conditions by resource content.	The Craig bound supplies a rank-dependent floor under equidistribution; the induced order is a preorder with incomparable pairs.
That the typology of [1] is fully recovered.	Operationally trivial conditions constitute the free class; the rest of the typology is not recovered as an ordering.
That the pruning content is computable in the algorithmic form of [1].	The relative-entropy form is computable in principle and provably monotone; the algorithmic form is an interpretive ceiling.
That Conjecture C is well posed in terms of rank alone.	It is stated over a nested normalized-projector family, or as an optimization at fixed post-selection probability.
That post-selection generates contextuality in general.	Rank dependence is conjectured only within the paradoxical class of [47], and only against a noncontextual baseline.

I do not claim	The more limited claim is
That Conjecture C is a theorem, or that (10) is derived.	Both endpoints are secure; the interpolation and the erfc form are an ansatz awaiting numerical or experimental test.
That CbD sanctions application of CNT_2 to post-selected data.	Dzhafarov, Kujala and Cervantes do not treat post-selection; this is a novel extension carrying an undischarged assumption.
That ancilla-appending and subsystem-discarding are established free operations.	Only boundary-compatible coarse-grainings are defined and proved; the others are expected extensions.
That ensemble mixing is constructed, or that convexity is established at the level of ensembles.	Joint convexity is a property of the associated distribution pairs; ensemble mixing is not developed here.
That a nested projector family can be combined with a fixed post-selection probability.	Along a generic nested family the post-selection probability varies with rank; the conjecture is stated over rank-and-probability-constrained sets instead.
That CL1-class wetware is conscious, sentient, or an agent.	It is a candidate substrate whose input-output statistics can be tested for rank-graded contextuality signatures.

Table 1. Explicit non-claims.

11. Limitations

The theory is defined only on medium-decoherent families. Where decoherence fails, the two distributions are not distributions, Lemma 2 fails, and the monotone is undefined. This is a genuine restriction inheriting the whole set-selection difficulty [7], [8], mitigated but not eliminated by results of the kind in [11].

The free operations are developed only for coarse-grainings. A complete resource theory would require a tensor product of boundary-conditioned ensembles and the associated composition rules; I have not constructed one, and Theorem 1 is scoped accordingly.

Theorem 2 bounds a stipulated register transformation, not the boundary condition and not every implementation of conditional sampling. The gap between these is where most of the interpretive risk in this programme lives, and I have tried to keep it visible rather than paper over it.

Q4 and Q5 remain open. I have not related logical depth to contextuality, and the complexity-contextuality interface identified in Section 1 is untouched — I could find no prior work relating Kolmogorov complexity to any contextuality measure, which makes it attractive and also unsupported. Q5 is addressed only by a protocol proposal, not a result.

The primary source for the rank-product bound [5] appears never to have been formally published, though the pure-state corollary appears in the peer-reviewed [6]. Anyone building on Corollary 2 or equation (8) should know the load-bearing citation is a preprint.

Finally, the ansatz (10) has no derivation. If a numerical sweep shows the measure flat along the family at fixed post-selection probability, Conjecture C should be withdrawn and Sections 4 through 6 retained; they do not depend on it.

12. Conclusion

A final boundary condition prunes, and pruning has a price. This paper has fixed the price in two currencies. In bits, the relative entropy of the conditioned against the unconditioned history distribution is a faithful monotone under boundary-compatible coarse-graining, jointly convex in the pair of distributions, and the naive counting alternative is not — a distinction that matters because the counting version is the one a reader of my earlier work would expect. In joules, a register carried from the unconditioned to the conditioned distribution under the Reeb-Wolf hypotheses deposits at least $k_B T \ln[n_0 / (r_\alpha r_\omega)]$, so that sharper terminal conditions carry a higher cost floor, with the exchange rate set by Craig's rank-product bound.

The typology I introduced informally in [1] is only partly recovered. Its trivial class is contained in the operationally defined free class, and coincides with it only for history families capable of detecting every nontrivial terminal reweighting; the Craig bound then places a rank-dependent floor above that class. But rank does not order terminal conditions by resource content, and I have withdrawn the stronger claim. The bridge to contextuality remains a conjecture, now stated over a specified family rather than in terms of rank alone, and I have been explicit that it is a mathematical conjecture before it is an experimental one. The next step is not an interferometer; it is a computation.

Author Contributions

I. Staley is the sole author of this work and is responsible for its conceptualization, formal analysis, methodology, investigation, writing of the original draft, and review and editing. The author has read and agreed to the posted version of the manuscript.

Funding

This research received no external funding.

Data Availability Statement

No new data were created or analyzed in this study, and data sharing is therefore not applicable to this article. The numerical counterexample supporting Remark 1 and the verification of Lemmas 1 and 2 are fully reproducible from the definitions given in the text. The code generating the numerical checks and Figures 1 through 4 is provided as a supplementary archive accompanying this preprint, together with dependency information sufficient to reproduce the figures as they appear here.

Use of Artificial Intelligence

During preparation of this manuscript the author used a large language model (Claude, Anthropic) for literature discovery, assistance with bibliographic verification, prose drafting and organization, numerical checking, and programmatic figure generation. The numerical assistance confirmed the counterexample reported in Remark 1 and tested the inequalities underlying Theorem 1 prior to their analytical proof. The figures were produced programmatically with such assistance from the equations and parameter values stated in their captions; no figure is a generative image, Figures 1 and 2 being schematic diagrams drawn to specification and Figures 3 and 4 being plotted directly from equations (8) and (10).

All mathematical statements, definitions, theorems, proofs, interpretive claims, and figure content were formulated, checked, and are vouched for by the author, who accepts full responsibility for the entire content of the manuscript including any error introduced by such tools. No artificial-intelligence system

is listed as an author, consistent with the COPE position statement on authorship and AI tools and with Preprints.org policy. As this is an analytical paper without a Methods section, the disclosure is made here.

Conflicts of Interest

The author declares no conflicts of interest. The author is employed in a role unrelated to the subject matter of this paper; that employment played no part in the conception, design, analysis, or conclusions of this work, and no funder or employer had any role in the decision to publish.

References

- [1] I. Staley, "Final-state constraints and informational pruning in quantum histories," *Int. J. Quantum Found.*, vol. 12, no. 2, pp. 719-737, Apr. 2026, doi: 10.5281/zenodo.19512844.
- [2] R. B. Griffiths, "Consistent histories and the interpretation of quantum mechanics," *J. Stat. Phys.*, vol. 36, no. 1-2, pp. 219-272, Jul. 1984, doi: 10.1007/BF01015734.
- [3] M. Gell-Mann and J. B. Hartle, "Classical equations for quantum systems," *Phys. Rev. D*, vol. 47, no. 8, pp. 3345-3382, Apr. 1993, doi: 10.1103/PhysRevD.47.3345.
- [4] J. B. Hartle, "Spacetime quantum mechanics and the quantum mechanics of spacetime," in *Gravitation and Quantizations: Proc. 1992 Les Houches Summer School, Session LVII*, B. Julia and J. Zinn-Justin, Eds. Amsterdam, The Netherlands: North-Holland, 1995. arXiv:gr-qc/9304006.
- [5] D. A. Craig, "The geometry of consistency: Decohering histories in generalized quantum theory," arXiv:quant-ph/9704031, Apr. 1997. [Preprint].
- [6] D. A. Craig, "Observation of the final boundary condition: Extragalactic background radiation and the time symmetry of the universe," *Ann. Phys.*, vol. 251, no. 2, pp. 384-425, Nov. 1996, doi: 10.1006/aphy.1996.0118.
- [7] F. Dowker and A. Kent, "On the consistent histories approach to quantum mechanics," *J. Stat. Phys.*, vol. 82, no. 5-6, pp. 1575-1646, Mar. 1996, doi: 10.1007/BF02183396.
- [8] F. Dowker and A. Kent, "Properties of consistent histories," *Phys. Rev. Lett.*, vol. 75, no. 17, pp. 3038-3041, Oct. 1995, doi: 10.1103/PhysRevLett.75.3038.
- [9] J. J. Halliwell, "Somewhere in the universe: Where is the information stored when histories decohere?," *Phys. Rev. D*, vol. 60, no. 10, art. 105031, Oct. 1999, doi: 10.1103/PhysRevD.60.105031.
- [10] J. J. Halliwell, "Incompatible multiple consistent sets of histories and measures of quantumness," *Phys. Rev. A*, vol. 96, no. 1, art. 012123, Jul. 2017, doi: 10.1103/PhysRevA.96.012123.
- [11] P. Strasberg, T. E. Reinhard, and J. C. Schindler, "First principles numerical demonstration of emergent decoherent histories," *Phys. Rev. X*, vol. 14, no. 4, art. 041027, Oct. 2024, doi: 10.1103/PhysRevX.14.041027.
- [12] Y. Aharonov, P. G. Bergmann, and J. L. Lebowitz, "Time symmetry in the quantum process of measurement," *Phys. Rev.*, vol. 134, no. 6B, pp. B1410-B1416, Jun. 1964, doi: 10.1103/PhysRev.134.B1410.
- [13] L. Vaidman, "Validity of the Aharonov-Bergmann-Lebowitz rule," *Phys. Rev. A*, vol. 57, no. 3, pp. 2251-2253, Mar. 1998, doi: 10.1103/PhysRevA.57.2251.
- [14] O. Cohen, "Pre- and postselected quantum systems, counterfactual measurements, and consistent histories," *Phys. Rev. A*, vol. 51, no. 6, pp. 4373-4380, Jun. 1995, doi: 10.1103/PhysRevA.51.4373.
- [15] E. Chitambar and G. Gour, "Quantum resource theories," *Rev. Mod. Phys.*, vol. 91, no. 2, art. 025001, Apr. 2019, doi: 10.1103/RevModPhys.91.025001.
- [16] S. Milz, J. Bavaresco, and G. Chiribella, "Resource theory of causal connection," *Quantum*, vol. 6, art. 788, Aug. 2022, doi: 10.22331/q-2022-08-25-788.
- [17] G. D. Berk, A. J. P. Garner, B. Yadin, K. Modi, and F. A. Pollock, "Resource theories of multi-time processes: A window into quantum non-Markovianity," *Quantum*, vol. 5, art. 435, Apr. 2021, doi: 10.22331/q-2021-04-20-435.
- [18] G. D. Berk, S. Milz, F. A. Pollock, and K. Modi, "Extracting quantum dynamical resources: Consumption of non-Markovianity for noise reduction," *npj Quantum Inf.*, vol. 9, art. 104, Nov. 2023, doi: 10.1038/s41534-023-00774-w.
- [19] G. Zambon and G. Adesso, "Quantum processes as thermodynamic resources: The role of non-Markovianity," *Phys. Rev. Lett.*, vol. 134, no. 20, art. 200401, May 2025, doi: 10.1103/PhysRevLett.134.200401.
- [20] G. Chiribella, G. M. D'Ariano, and P. Perinotti, "Theoretical framework for quantum networks," *Phys. Rev. A*, vol. 80, no. 2, art. 022339, Aug. 2009, doi: 10.1103/PhysRevA.80.022339.
- [21] R. Landauer, "Irreversibility and heat generation in the computing process," *IBM J. Res. Develop.*, vol. 5, no. 3, pp. 183-191, Jul. 1961, doi: 10.1147/rd.53.0183.
- [22] D. Reeb and M. M. Wolf, "An improved Landauer principle with finite-size corrections," *New J. Phys.*, vol. 16, art. 103011, Oct. 2014, doi: 10.1088/1367-2630/16/10/103011.

- [23] P. Faist, F. Dupuis, J. Oppenheim, and R. Renner, "The minimal work cost of information processing," *Nat. Commun.*, vol. 6, art. 7669, Jul. 2015, doi: 10.1038/ncomms8669.
- [24] N. G. Anderson, "Generalized Landauer bound for information processing: Proof and applications," *Entropy*, vol. 24, no. 11, art. 1568, Oct. 2022, doi: 10.3390/e24111568.
- [25] T. Sagawa and M. Ueda, "Minimal energy cost for thermodynamic information processing: Measurement and information erasure," *Phys. Rev. Lett.*, vol. 102, no. 25, art. 250602, Jun. 2009, doi: 10.1103/PhysRevLett.102.250602.
- [26] G. Chiribella, Y. Yang, and R. Renner, "Fundamental energy requirement of reversible quantum operations," *Phys. Rev. X*, vol. 11, no. 2, art. 021014, Apr. 2021, doi: 10.1103/PhysRevX.11.021014.
- [27] C.-Y. Hsieh, "Dynamical Landauer principle: Quantifying information transmission by thermodynamics," *Phys. Rev. Lett.*, vol. 134, no. 5, art. 050404, Feb. 2025, doi: 10.1103/PhysRevLett.134.050404.
- [28] T. Purves and A. J. Short, "Channels, measurements, and postselection in quantum thermodynamics," *Phys. Rev. E*, vol. 104, no. 1, art. 014111, Jul. 2021, doi: 10.1103/PhysRevE.104.014111.
- [29] P. Boes, R. Gallego, N. H. Y. Ng, J. Eisert, and H. Wilming, "By-passing fluctuation theorems," *Quantum*, vol. 4, art. 231, Feb. 2020, doi: 10.22331/q-2020-02-20-231.
- [30] D. R. M. Arvidsson-Shukur, N. Yunger Halpern, H. V. Lepage, A. A. Lasek, C. H. W. Barnes, and S. Lloyd, "Quantum advantage in postselected metrology," *Nat. Commun.*, vol. 11, art. 3775, Jul. 2020, doi: 10.1038/s41467-020-17559-w.
- [31] M. Lostaglio, "Quantum fluctuation theorems, contextuality, and work quasiprobabilities," *Phys. Rev. Lett.*, vol. 120, no. 4, art. 040602, Jan. 2018, doi: 10.1103/PhysRevLett.120.040602.
- [32] M. Lostaglio, "Certifying quantum signatures in thermodynamics and metrology via contextuality of quantum linear response," *Phys. Rev. Lett.*, vol. 125, no. 23, art. 230603, Dec. 2020, doi: 10.1103/PhysRevLett.125.230603.
- [33] D. R. M. Arvidsson-Shukur et al., "Properties and applications of the Kirkwood-Dirac distribution," *New J. Phys.*, vol. 26, no. 12, art. 121201, Dec. 2024, doi: 10.1088/1367-2630/ada05d.
- [34] S. Gherardini and G. De Chiara, "Quasiprobabilities in quantum thermodynamics and many-body systems," *PRX Quantum*, vol. 5, no. 3, art. 030201, Jul. 2024, doi: 10.1103/PRXQuantum.5.030201.
- [35] S. Das, S. Modak, and M. N. Bera, "Saturating quantum advantages in postselected metrology with the positive Kirkwood-Dirac distribution," *Phys. Rev. A*, vol. 107, no. 4, art. 042413, Apr. 2023, doi: 10.1103/PhysRevA.107.042413. See also the Comment, *Phys. Rev. A*, vol. 109, art. 066401, 2024, and Reply, *Phys. Rev. A*, vol. 109, art. 066402, 2024.
- [36] R. W. Spekkens, "Contextuality for preparations, transformations, and unsharp measurements," *Phys. Rev. A*, vol. 71, no. 5, art. 052108, May 2005, doi: 10.1103/PhysRevA.71.052108.
- [37] R. W. Spekkens, "Negativity and contextuality are equivalent notions of nonclassicality," *Phys. Rev. Lett.*, vol. 101, no. 2, art. 020401, Jul. 2008, doi: 10.1103/PhysRevLett.101.020401.
- [38] D. Schmid, J. H. Selby, E. Wolfe, R. Kunjwal, and R. W. Spekkens, "Characterization of noncontextuality in the framework of generalized probabilistic theories," *PRX Quantum*, vol. 2, no. 1, art. 010331, Feb. 2021, doi: 10.1103/PRXQuantum.2.010331.
- [39] D. Schmid, J. H. Selby, M. F. Pusey, and R. W. Spekkens, "A structure theorem for generalized-noncontextual ontological models," *Quantum*, vol. 8, art. 1283, Apr. 2024, doi: 10.22331/q-2024-04-30-1283.
- [40] E. N. Dzhafarov, J. V. Kujala, and V. H. Cervantes, "Contextuality and noncontextuality measures and generalized Bell inequalities for cyclic systems," *Phys. Rev. A*, vol. 101, no. 4, art. 042119, Apr. 2020, doi: 10.1103/PhysRevA.101.042119.
- [41] E. N. Dzhafarov, J. V. Kujala, and V. H. Cervantes, "Contextuality-by-default: A brief overview of ideas, concepts, and terminology," in *Quantum Interaction (QI 2015)*, *Lecture Notes in Computer Science*, vol. 9535, Cham, Switzerland: Springer, 2016, pp. 12-23, doi: 10.1007/978-3-319-28675-4_2.
- [42] 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, Jul. 2008, doi: 10.1103/PhysRevLett.101.020403.

- [43] C. Budroni, A. Cabello, O. Gühne, M. Kleinmann, and J.-Å. Larsson, "Kochen-Specker contextuality," *Rev. Mod. Phys.*, vol. 94, no. 4, art. 045007, Dec. 2022, doi: 10.1103/RevModPhys.94.045007.
- [44] M. F. Pusey, "Anomalous weak values are proofs of contextuality," *Phys. Rev. Lett.*, vol. 113, no. 20, art. 200401, Nov. 2014, doi: 10.1103/PhysRevLett.113.200401.
- [45] R. Kunjwal, M. Lostaglio, and M. F. Pusey, "Anomalous weak values and contextuality: Robustness, tightness, and imaginary parts," *Phys. Rev. A*, vol. 100, no. 4, art. 042116, Oct. 2019, doi: 10.1103/PhysRevA.100.042116.
- [46] M. S. Leifer and R. W. Spekkens, "Pre- and post-selection paradoxes and contextuality in quantum mechanics," *Phys. Rev. Lett.*, vol. 95, no. 20, art. 200405, Nov. 2005, doi: 10.1103/PhysRevLett.95.200405.
- [47] J. Singh, R. S. Bhati, and Arvind, "No contextual advantage in nonparadoxical scenarios of the two-state vector formalism," *Phys. Rev. A*, vol. 107, no. 1, art. 012206, Jan. 2023, doi: 10.1103/PhysRevA.107.012206.
- [48] Z.-P. Xu, X.-D. Yu, and M. Kleinmann, "State-independent quantum contextuality with projectors of nonunit rank," *New J. Phys.*, vol. 23, no. 4, art. 043025, Apr. 2021, doi: 10.1088/1367-2630/abe6e3.
- [49] S. Abramsky, R. S. Barbosa, and S. Mansfield, "Contextual fraction as a measure of contextuality," *Phys. Rev. Lett.*, vol. 119, no. 5, art. 050504, Aug. 2017, doi: 10.1103/PhysRevLett.119.050504.
- [50] C. Duarte and B. Amaral, "Resource theory of contextuality for arbitrary prepare-and-measure experiments," *J. Math. Phys.*, vol. 59, no. 6, art. 062202, Jun. 2018, doi: 10.1063/1.5018582.
- [51] B. Amaral, "Resource theory of contextuality," *Phil. Trans. R. Soc. A*, vol. 377, no. 2157, art. 20190010, Nov. 2019, doi: 10.1098/rsta.2019.0010.
- [52] R. Wagner, R. D. Baldijão, A. Tezzin, and B. Amaral, "Using a resource theoretic perspective to witness and engineer quantum generalized contextuality for prepare-and-measure scenarios," *J. Phys. A: Math. Theor.*, vol. 56, no. 50, art. 505303, Dec. 2023, doi: 10.1088/1751-8121/ad0bcc.
- [53] D. Schmid, R. D. Baldijão, Y. Ying, R. Wagner, and J. H. Selby, "Kirkwood-Dirac representations beyond quantum states and their relation to noncontextuality," *Phys. Rev. A*, vol. 110, no. 5, art. 052206, Nov. 2024, doi: 10.1103/PhysRevA.110.052206.
- [54] W. H. Zurek, "Algorithmic randomness and physical entropy," *Phys. Rev. A*, vol. 40, no. 8, pp. 4731-4751, Oct. 1989, doi: 10.1103/PhysRevA.40.4731.
- [55] C. H. Bennett, "Logical depth and physical complexity," in *The Universal Turing Machine: A Half-Century Survey*, R. Herken, Ed. Oxford, U.K.: Oxford Univ. Press, 1988, pp. 227-257.
- [56] M. Li and P. Vitányi, *An Introduction to Kolmogorov Complexity and Its Applications*, 4th ed. Cham, Switzerland: Springer, 2019, doi: 10.1007/978-3-030-11298-1.
- [57] D. Janzing, R. Chaves, and B. Schölkopf, "Algorithmic independence of initial condition and dynamical law in thermodynamics and causal inference," *New J. Phys.*, vol. 18, no. 9, art. 093052, Sep. 2016, doi: 10.1088/1367-2630/18/9/093052.
- [58] J. Edmonds, "A physical framework for algorithmic entropy," *Entropy*, vol. 28, no. 1, art. 61, Jan. 2026, doi: 10.3390/e28010061.
- [59] B. J. Kagan, A. C. Kitchen, N. T. Tran, F. Habibollahi, M. Khajehnejad, B. J. Parker, A. Bhat, B. Rollo, A. Razi, and K. J. Friston, "In vitro neurons learn and exhibit sentience when embodied in a simulated game-world," *Neuron*, vol. 110, no. 23, pp. 3952-3969.e8, Dec. 2022, doi: 10.1016/j.neuron.2022.09.001.
- [60] M. Khajehnejad, F. Habibollahi, A. Loeffler, A. Paul, A. Razi, and B. J. Kagan, "Dynamic network plasticity and sample efficiency in biological neural cultures: A comparative study with deep reinforcement learning," *Cyborg Bionic Syst.*, art. cbsystems.0336, Aug. 2025, doi: 10.34133/cbsystems.0336.
- [61] L. Smirnova, B. S. Caffo, D. H. Gracias, et al., "Organoid intelligence (OI): The new frontier in biocomputing and intelligence-in-a-dish," *Front. Sci.*, vol. 1, art. 1017235, Feb. 2023, doi: 10.3389/fsci.2023.1017235.
- [62] F. D. Jordan, M. Kutter, J.-M. Comby, F. Brozzi, and E. Kurtys, "Open and remotely accessible Neuroplatform for research in wetware computing," *Front. Artif. Intell.*, vol. 7, art. 1376042, May 2024, doi: 10.3389/frai.2024.1376042.
- [63] I. Staley, "Substrate-flexible informational quantum mechanics: QBism, RQM, and the question of non-human epistemic systems," *Preprints*, May 2026. [Preprint].

- [64] I. Staley, "Pruning as rendering: A decoherent-histories reformulation of the observer-rendering criterion in simulation-theoretic quantum mechanics," Preprints, May 2026. [Preprint].
- [65] I. Staley, "Informational pruning as a third path: Final-state constraints between consistent quantum causes and quantum causal models," Preprints, May 2026. [Preprint].
- [66] I. Staley, "Pruned histories and the computational multiverse: What does quantum speedup ontologically require?," Preprints, Apr. 2026. [Preprint].
- [67] T. M. Cover and J. A. Thomas, Elements of Information Theory, 2nd ed. Hoboken, NJ, USA: Wiley, 2006.
- [68] M. P. Müller, "Law without law: From observer states to physics via algorithmic information theory," Quantum, vol. 4, art. 301, Jul. 2020, doi: 10.22331/q-2020-07-20-301.
- [69] R. B. Griffiths, "Consistent quantum causes," arXiv:2303.13617 [quant-ph], Dec. 2024. [Preprint].
- [70] T. De Saegher, "Decoherent histories contextuality," seminar, Perimeter Institute, PIRSA:25100163, Oct. 2025.