Redefining Entropy from First Principles: A BCB–TPB Formulation

Keith Taylor *VERSF Theoretical Physics Program*

For the General Reader

Entropy is one of physics' most important — and most misunderstood — concepts. It is often described as *disorder*, or as a measure of how much we don't know about a system. But these descriptions have always felt unsatisfying, because disorder is not a physical thing. You cannot weigh it, build with it, or point to it in an experiment.

This paper proposes a different answer: **entropy is the cost of keeping a record**.

Every time nature makes a permanent, irreversible distinction — a particle detected, a measurement made, a bit of information locked into the world — that act has a physical cost. Entropy is simply the running total of those costs. It grows because once a record is written, it cannot be un-written without writing something else. The universe's ledger only ever gets longer.

This has some striking consequences. The reason what we call "time" only runs in one direction is not that the universe began in a statistically rare low-entropy state and then drifted toward disorder. In this framework, the direction of time's arrow is the direction in which committed distinctions accumulate. A committed record constrains what can happen next, but it cannot be undone without further commitment elsewhere. The ledger grows, and that growth defines "after."

Time is therefore not treated as an independent background parameter in this framework. Instead, the ordering of events is primary, and the familiar notion of clock time emerges from that ordering through the TPB conversion: a system that commits distinctions more densely generates a denser sequence of clock-ticks, while a system near equilibrium generates few new committed distinctions and exhibits minimal temporal progression. In the emergent regime, this reproduces standard relativistic time dilation, but explains it as a consequence of record-formation constraints rather than as a primitive axiom. Coordinate time and proper time remain valid effective parameters in this emergent layer and are recovered exactly when the TPB mapping is applied.

The paper develops this idea rigorously, shows that it reproduces all the standard results of thermodynamics exactly, and identifies new predictions — particularly near black holes — where this deeper picture departs from conventional physics in testable ways.

No specialist background is required to follow the main argument. Mathematical details are confined to Sections 7–9 and the appendix.

Abstract

We propose a rigorous reformulation of entropy grounded in the Bit-Conservation & Balance (BCB) and Ticks-Per-Bit (TPB) frameworks. Rather than defining entropy as disorder or missing information, we redefine it as the **physical ledger cost of irreversible commitment** required to maintain distinguishability under law-invariant dynamics. This formulation replaces abstract microstate counting with physically committed distinctions, connects entropy directly to irreversible record formation, and recovers all standard thermodynamic results — including Boltzmann entropy, Shannon entropy, Landauer's principle, and the second law — as limiting cases. Physical time is not assumed as a primitive background parameter. Instead, the ordering of committed distinctions is primary, and emergent clock time arises from that ordering through the TPB conversion: the familiar coordinate time and proper time of standard physics are recovered as effective descriptions in the emergent layer. The framework introduces a physically motivated correction to the microstate count, Ω_FD, arising from finite distinguishability constraints, yielding measurable deviations from standard statistical mechanics in high-commitment-density regimes. The maximum committed-bit density per emergent second is shown to saturate the Bekenstein information bound, bridging thermodynamics, relativity, and information theory under a single causal-record accounting principle.

Table of Contents

---

1. Foundational Premise

In the BCB framework, physical reality is constituted not by abstract mathematical states but by **committed records** — irreversible distinctions written into the world's causal ledger. Two principles anchor the framework:

- **Finite Distinguishability (FD):** The number of physically resolvable distinctions within any finite causal region is bounded. Not all mathematically defined microstates correspond to physically distinguishable configurations.
- **Irreversible Commitment (IC):** Once a distinction is committed to the causal ledger, erasing it requires writing at least one new committed distinction elsewhere. Commitment therefore accumulates monotonically.

The physical state of a system at ledger time τ is defined entirely by the set of committed distinctions it has stabilized — not by a wavefunction or phase-space point, but by a finite, causally grounded record.

Note on ledger time. Ledger time τ is the fundamental ordering parameter — the index of sequential committed distinctions. It is not a background stage on which events occur; it is the accumulation of events itself. TPB is the conversion factor between ledger ordering and emergent clock time: the emergent duration associated with a given increment in ledger index is proportional to TPB, with the proportionality fixed by calibration in the emergent layer. Coordinate time and proper time are recovered as effective parameters in this emergent description, and standard relativistic results hold exactly in the regimes where that description applies.

🔍 **In plain language —** Physics usually describes the world using abstract mathematical states: wavefunctions, probability distributions, points in phase space. This framework takes a different starting point. The only things that are *physically real* are events that have been permanently recorded — distinctions that have been made and cannot be undone. Think of a footprint in concrete: the moment the concrete sets, a distinction is committed. The framework is built entirely on these committed distinctions, and everything else — entropy, what we experience as time, even geometry — is derived from them. At the fundamental level, time is not a container in which events happen; it is the sequence of committed events. At the everyday level of physics, familiar coordinate time and proper time are recovered exactly as emergent descriptions of that sequence.

---

2. The Committed Bit: Formal Definition

Before defining entropy, we must define its fundamental unit precisely. A **committed bit** is a physical event satisfying all four of the following conditions:

Distinguishability. Two outcomes O_1 and O_2 remain distinguishable above a threshold ε under admissible readout operations: the best-achievable statistical distance (total variation distance for classical records; trace distance for quantum states) exceeds ε over the readout window.

Remark: In many detector-limited settings, ε is governed by an energy–bandwidth–integration-time constraint. This can be heuristically expressed as an energy–time tradeoff $\Delta E \cdot \Delta t\_readout \gtrsim \hbar/2$, but the operational statistical-distance criterion above is the general definition and applies equally in classical and quantum regimes.

Stability. The record remains readable above threshold ε across at least two independent admissible readouts separated by a nonzero ledger depth. Equivalently: there exists an integer $N \geq 1$ such that for all $n \in \{0, \ldots, N\}$, the record remains operationally distinguishable (distance $\geq \varepsilon$) under admissible evolution and readout. This defines stability in terms of ledger succession rather than presupposed clock duration.

Irrecoverability. No admissible physical operation — one consistent with the laws of the system — restores the pre-record microstate without incurring at least one further committed bit elsewhere.

Causal usability. The record can influence at least one downstream event in the forward light cone.

This definition anchors the committed bit in operational physics. Abstract mathematical states that fail any of these four criteria do not count as physical distinctions and do not contribute to entropy. The definition also clarifies why quantum superpositions are not committed bits: they fail the stability and irrecoverability conditions until decoherence or measurement collapses them into a stable, causally usable record.

🔍 **In plain language —** Not every event in physics counts as a "committed bit." A quantum particle hovering in superposition — simultaneously in two states — hasn't committed to anything yet. Only when it interacts with the world and produces a stable, irreversible, usable result does a committed bit get written. The four conditions are a checklist for deciding whether something genuinely happened, in the physical sense: Was it distinguishable from noise? Did it last? Can it influence the future? And is it impossible to un-do without cost? If yes to all four, it's a committed bit. If not, it's just a mathematical possibility.

---

3. Commitment Entropy (S_c)

Let C(τ) denote the cumulative number of committed bits up to ledger time τ. We define **Commitment Entropy** as:

$$S_c = k_{IC} \cdot C$$

The **commitment constant** k_IC is fixed physically — not fitted — by Landauer's principle:

$$k_{IC} = k_B \ln 2$$

This is the minimum thermodynamic cost per committed bit. With this choice, S_c carries units of J/K and is directly comparable to thermodynamic entropy without any free parameters.

Entropy is therefore the **accumulated irreversible ledger cost** of a system's causal history.

🔍 **In plain language —** Entropy, in this framework, is simply the total bill. Every time the universe makes an irreversible distinction — every committed bit — it adds one entry to the ledger. Entropy is the running total of that ledger. The constant k_IC = k_B ln 2 is the cost of a single entry, set by Landauer's principle: it's the minimum energy a computer must dissipate as heat when erasing one bit of information. So the framework isn't inventing a new number — it's connecting entropy directly to a measured physical quantity.

---

4. Record-Resolvable Entropy (S_R) and the FD Correction

Standard statistical mechanics defines entropy as $S = k_B \ln \Omega$, where Ω is the number of microstates compatible with macroscopic constraints. Under finite distinguishability, many of these mathematical microstates are physically indistinguishable — they require more committed bits to resolve than are available within the causal region.

Definition of Ω_{FD}. Let C_{min} denote the minimum committed bits required to distinguish adjacent microstates in a given system, and C_{max} denote the total committed capacity available to the causal region.

Operational definition of C_{min}. $C_{min}(\varepsilon)$ is the minimum number of committed bits required to reliably discriminate two neighbouring microstates within the admissible process class and above the distinguishability threshold ε. Formally:

$$C_{min}(\varepsilon) := \min\{ C : \exists\, \Pi \in \mathcal{O}(\mathscr{R}, \varepsilon) \text{ such that } D(\Pi(\rho_1), \Pi(\rho_2)) \geq \varepsilon \}$$

where D is total variation distance (classical) or trace distance (quantum), and ρ_1, ρ_2 are adjacent microstate ensembles under the chosen macroscopic constraints. In practice, C_{min} is the cost of the minimal record carrier that is (i) stable over the required ledger depth N (Section 2), (ii) irrecoverable without additional commitment, and (iii) distinguishable above ε.

Estimation heuristics. In laboratory detectors, C_{min} is set by the minimal decision-relevant record (threshold crossing, latch state, discriminator output), scaling with required false-alarm and miss probabilities. In thermal systems, C_{min} is set by the minimal stable basin separation required for metastability at ambient temperature. In gravitational and horizon settings, C_{min} is naturally tied to the smallest physically separable record unit consistent with the local causal resolution — i.e., the minimal distinguishable partition supported by the region's admissible record algebra.

Definition (FD suppression). Let $\alpha := C_{min} / C_{max}$. Then the physically resolvable microstate count is:

$$\Omega_{FD} = \Omega \cdot \exp(-\alpha \ln \Omega) = \Omega^{(1-\alpha)}$$

This is the capacity-thinning result derived in Appendix B: each nat of distinguishability extracted consumes a fraction α of the available commitment capacity, producing a multiplicative suppression across the $\ln \Omega$ refinement depth. The $\Omega^{(1-\alpha)}$ form is clean and exact — no absorptions or approximations. Note that C_{min} here refers to the committed bits required to resolve a *single adjacent microstate distinction* (not the full identification cost); the full refinement over $\ln \Omega$ distinctions generates the exponent $\alpha \ln \Omega$.

The ratio $\alpha = C_{min} / C_{max}$ represents the fraction of available causal record capacity consumed per unit of distinguishability. Its physical content depends on context:

Physical setting	What determines C_{max}
Isolated thermodynamic system	Bekenstein bound: $C_{max} = 2\pi RE/\hbar c \ln 2$
Black hole horizon	Bekenstein–Hawking entropy: $C_{max} = S_{BH}/(k_B \ln 2) = A/(4\ell_P^2 \ln 2)$
Laboratory detector	Detector capacity: channel bandwidth × integration time
Planck-scale region	Holographic bound: $C_{max} = 1$ (one bit per Planck area)

We define **Record-Resolvable Entropy**:

$$S_R := k_B \ln \Omega_{FD} = k_B (1 - \alpha) \ln \Omega$$

so the FD correction is:

$$\Delta S = S_R - k_B \ln \Omega = -k_B \alpha \ln \Omega$$

Note the convention: S_R uses the natural-log form directly, consistent with the standard Boltzmann formula $S = k_B \ln \Omega$. Commitment entropy $S_c = k_B \ln 2 \cdot C$ uses the per-bit counting convention. The two are connected via the consistency checks in Section 9, where $C = \log_2 \Omega$ recovers $S_c = k_B \ln \Omega$ exactly. For non-uniform ensembles, replace $\log_2 \Omega$ with the Shannon entropy H; see Section 9.4.

The FD correction $\Delta S = -k_B \alpha \ln \Omega$ is negligible under ordinary laboratory conditions ($\alpha \ll 1$) but becomes significant — suppressing the effective phase volume as $\Omega^{(1-\alpha)}$ — near black hole horizons and Planck-scale densities where $\alpha \to 1$.

4.1 The Two Entropies: Physical Interpretation

The two entropy measures are complementary and together exhaust the thermodynamic account of a system:

Quantity	Physical interpretation
S_c	Irreversible causal history — depth of the ledger
S_R	Remaining physical possibility — available phase volume

Capacity accounting (inequality; equality under maximal-mixing closure). Let C_{max} denote the committed capacity of the causal region (e.g., the Bekenstein or holographic bound), and define the capacity entropy budget $S_{cap} := k_B \ln 2 \cdot C_{max}$. The ledger entropy satisfies $S_c \leq S_{cap}$. The record-resolvable entropy satisfies:

$$S_R \leq S_{cap} - S_c$$

and therefore:

$S_c + S\_R \leq S\_cap$

with equality holding under a maximal-mixing closure in which the remaining capacity is fully saturated by the resolvable phase volume at the current macroscopic constraints. In open systems, S_cap itself varies with capacity exchange across the boundary, and only the inequality is guaranteed.

Worked closure example (toy). If a region has capacity C_max and a system has already committed C bits, then at most C_max − C further bits can be stably committed. Under maximal mixing, ln Ω_FD = (C_max − C) ln 2, giving S_R = k_B (C_max − C) ln 2 = S_cap − S_c, so S_c + S_R = S_cap. Without maximal mixing, only the inequality S_c + S_R ≤ S_cap is guaranteed.

This frames thermodynamics as **record accounting**: as causal history accumulates (S_c increases), physical possibility space contracts (S_R decreases). The second law is therefore not a statement about disorder but about the inexorable growth of the causal ledger.

🔍 **In plain language** — The framework has two complementary accounts of a system. S_c is how much has already happened — the depth of the causal record, like a diary that only gets longer. S_R is how many different things could still happen — the remaining open possibilities, like blank pages. As the diary fills up, the blank pages diminish. The two always trade off. This is what the second law really says: not that things become messier, but that the universe's diary inevitably fills in.

The FD correction is a subtle but important addition: it says that near extreme physical limits — like the edge of a black hole — some of those "open possibilities" aren't actually distinguishable from each other, so the real number of blank pages is smaller than the standard formula suggests.

---

5. The Second Law and the Arrow of Time

The IC principle implies:

$dS_c/d\tau \geq 0$

because committed distinctions do not un-write. This is a structural consequence of commitment irreversibility, not a statistical one.

Crucially, this avoids the circularity present in purely statistical derivations. The BCB account of the arrow of time rests on three points:

1. **Direction of τ** is defined operationally as the direction of increasing ledger depth C(τ). Time is the index of causal record accumulation.
2. **Irreversibility** arises because commitment events are causally asymmetric: a committed distinction constrains future states but is not constrained by them. This asymmetry is built into the definition of a committed bit (Section 2).

3. **Boundary condition.** The minimal boundary condition for the framework is a state with no committed distinctions, $C = 0$. This is not a fine-tuned statistical selection from an ensemble; it is the unique state in which no physical distinctions have yet been written, and it serves as a natural lower bound on ledger depth. Entropy grows because records can only accumulate from this zero-commitment state, replacing the fine-tuned past hypothesis of standard thermodynamics with a structurally necessary minimum.

🔍 **In plain language —** Why does what we call "time" only run in one direction? The standard answer is unsatisfying: the universe just happened to start in a very special low-entropy state, and it's been getting messier ever since. That's statistically true but requires assuming that time was already there, flowing forward, when the story began. This framework gives a different answer. At the fundamental level, there is no pre-existing time — only the sequence of committed bits. Each committed bit defines "after" with respect to everything that preceded it in the ledger. Forward is not a direction in time; it is the direction of the growing ledger. There is no "backwards" in which records un-accumulate, because un-writing costs at least as much as writing. At the emergent level, all of this reproduces the familiar arrow of time exactly — but now it is derived from the structure of irreversible record-keeping rather than assumed as a primitive feature of the universe.

---

6. Ticks-Per-Bit (TPB), Entropy Rate, and the Bekenstein Bound

6.1 TPB and Temporal Emergence

Define the Ticks-Per-Bit ratio:

$$\mathbf{TPB = d\tau / db}$$

where db is an increment of committed bits and $d\tau$ the corresponding increment of ledger index. Here τ is treated as the primary ordering parameter. The emergent clock time t_{em} is linked to the ledger index by the TPB calibration:

$$\mathbf{dt_{em} = TPB \cdot dC}$$

or equivalently:

$$\mathbf{dC/dt_{em} = 1/TPB}$$

This is the defining equation of temporal emergence: each committed bit generates TPB units of emergent clock time. When we later introduce emergent clock time t_{em} (defined via the TPB calibration in Section 6.2), rates transform via $d/dt_{em} = (d\tau/dt_{em}) \cdot d/d\tau$, with $d\tau/dt_{em}$ determined by the TPB mapping. The entropy accumulated per ledger step follows:

$$dS_c/d\tau = k\_IC \cdot db/d\tau = k\_IC / TPB$$

This equation encodes the emergence of temporal experience from commitment density. **Time is not an independent background parameter in this framework: the accumulation of committed distinctions generates the ordering we experience as "time," and TPB provides the conversion between ledger progression and clock time.** A system committing many bits per ledger step (low TPB) generates a dense sequence of emergent clock intervals — it advances quickly through what observers experience as time. A system near thermal equilibrium, committing few new bits per ledger step (high TPB), generates sparse clock intervals. In the limit $db/d\tau \to 0$, no new ledger entries are written and no temporal progression occurs — consistent with the VERSF picture of time as entirely emergent from the bit-accumulation sequence.

Denser entropy accumulation per ledger step produces denser emergent time; sparser accumulation produces sparser emergent time. The TPB ratio is the conversion factor between the fundamental ledger ordering and the emergent clock — both layers are legitimate descriptions, at their respective levels.

Relativistic scaling (conditional on Lorentzian emergent layer). In the companion record-theoretic spacetime derivation, the emergent causal cone and Lorentz transformations fix the relation between proper time and coordinate time. Under the TPB mapping, the conversion factor must satisfy $TPB \propto dt/d\tau\_proper$. Therefore $TPB(v) = TPB_0 \cdot \gamma$ in special relativity, and $TPB \propto 1/\sqrt{g_{00}}$ in static gravitational fields. This paper uses that result as an interface condition: it does not re-derive Lorentz invariance here, but shows how entropy-as-commitment makes time dilation appear as a record-cost scaling consistent with standard relativistic results. If the companion paper's Lorentzian emergence is not assumed, the weaker statement is that the TPB framework is consistent with and naturally reinterprets standard relativistic time dilation.

6.2 The Bekenstein Bound on Committed Bits

Let t_em denote emergent physical time defined by the TPB-calibrated clock rate — that is, coordinate time as recovered in the emergent layer, not a primitive background. Within this emergent description, the Bremermann–Bekenstein bound constrains the maximum committed bits per unit emergent physical time for a system of mass m:

$$db/dt\_em \leq 2\pi mc^2 / (\hbar \ln 2)$$

where t_em denotes emergent physical time (not a background parameter). Substituting into the entropy accumulation formula:

$$dS_c/dt\_em \leq 2\pi k\_B\, mc^2 / \hbar$$

This is the **Bekenstein information bound** expressed as a maximum entropy accumulation per unit emergent time. Importantly, since t_em itself emerges from the bit sequence via the TPB mapping, this bound is ultimately a constraint on the density of committed bits per ledger step relative to the system's mass-energy — both relativistic (mc^2) and quantum ($\hbar$) constraints limiting how many bits can be distinguished and committed per cycle. The BCB framework does not merely recover this bound — it regrounds it as a consequence of the finite distinguishability capacity of any bounded physical system.

🔍 **In plain language** — These two ideas belong together. First: what we experience as "time passing" emerges from the bit sequence being written. A bonfire is committing bits densely per ledger step; a crystal at equilibrium is committing almost none. The bonfire generates many physical clock intervals per ledger cycle — it "advances through time" more quickly, but only because time itself is being generated by those commitments. The crystal isn't stuck in time — it's barely producing time at all. Second: there's a ceiling on how many bits per emergent second any physical system can commit, set jointly by its mass (relativity) and the uncertainty principle (quantum mechanics). This is the Bekenstein bound. The framework doesn't borrow this result — it derives it as a consequence of what a committed bit physically is.

---

7. Maxwell's Demon and Committed Memory

The BCB framework yields a structural derivation of the Landauer–Bennett resolution of Maxwell's Demon. The demon operates a trap door based on molecular velocity measurements. Each measurement commits one bit to the demon's internal memory register. At the end of the cycle, to return to its initial state, it must erase that memory — and by IC, erasing a committed bit requires writing at least one new committed bit elsewhere, dissipating at least $k_{IC} = k_B \ln 2$ of free energy per bit erased.

This is the Landauer–Bennett resolution, but the BCB framing adds structural clarity: **memory erasure is itself a commitment event**. The ledger cost of un-writing is at least as large as the ledger cost of the original distinction. There is no thermodynamic free lunch in a universe governed by IC. This result is not imported as a separate postulate — it follows directly from the definition of a committed bit in Section 2 and the definition of k_{IC} in Section 3.

🔍 **In plain language** — Maxwell's Demon is a famous thought experiment: a tiny creature sorts fast and slow molecules, seemingly creating order for free and defeating the second law. The resolution has been known since the 1980s — the demon has to erase its memory to complete the cycle, and that erasure costs energy. But *why* does erasure cost energy? In standard physics this is imported as a separate principle. Here it falls out automatically: erasing a committed bit just means writing a new one somewhere else. You can't un-write without writing. The demon's bookkeeping trick has to balance. There is no free lunch in a universe where records can only accumulate.

---

8. Worked Example: Photon Detection

To ground the framework concretely, consider a single photon striking a photodetector.

Physical sequence:

1. Photon (energy hν) arrives at photocathode
2. Photoelectric effect ejects one electron
3. Electron cascade (avalanche multiplication, gain $G \sim 10^6$)
4. Macroscopic current pulse: pointer state registered

BCB account:

Committed bits generated: The energetic distinction between the photon signal and the thermal noise floor corresponds to several committed bits, the precise number depending on detector bandwidth, integration time, and threshold. Formally, the committed bit count is $C \approx \log_2(E\_signal / E\_noise)$, where the noise scale must be evaluated against the specific detector's capacity parameters — it is not simply the ratio hν/k_B T, which does not account for bandwidth or threshold effects. For a well-designed optical detector operating well above threshold, C is typically in the range of a few bits to tens of bits per detection event.

Causal record: The avalanche current pulse satisfies all four committed bit conditions: it is distinguishable above noise (ε ~ thermal fluctuation), stable across the required ledger depth N, irrecoverable without further energy input, and causally usable (it triggers the downstream counting circuit).

FD correction: For a macroscopic detector at room temperature, C_min/C_max ≪ 1, so Ω_FD ≈ Ω and S_R ≈ S_Boltzmann. The FD correction is negligible — consistent with the expectation that deviations appear only at extreme commitment densities.

This example illustrates how the BCB framework translates a concrete physical process into a precise entropy account without invoking abstract microstate counting.

🔍 **In plain language** — This example grounds the abstract ideas in something tangible. A single photon hits a detector. An avalanche of electrons cascades through the device. A current pulse fires. A number ticks up on a counter. At each stage, the BCB framework asks: was a committed bit written? The answer is yes — the current pulse is distinguishable from noise, it persists long enough to be read, it can't be un-done without further energy, and it causes the counter to increment. The photon detection event is, in this framework, an irreversible ledger entry. The entropy produced is the energy cost of making that entry permanent.

---

9. Mathematical Consistency

9.1 Extensivity

For independent subsystems A and B, their commitment histories share no records. Therefore C(A+B) = C(A) + C(B) and:

$$\mathbf{S_c(A+B) = S_c(A) + S_c(B)}$$

For correlated or entangled systems, shared committed records must not be double-counted. The correct generalization is:

$$\mathbf{C(A+B) = C(A) + C(B) - C(A \cap B)}$$

where $C(A \cap B)$ counts distinctions jointly committed to both subsystems' records. This mirrors the subadditivity of von Neumann entropy and reduces to classical extensivity when A and B are uncorrelated.

9.2 Second Law

By IC: $dC/d\tau \geq 0$. Therefore $dS_c/d\tau = k\_IC \cdot dC/d\tau \geq 0$. □

9.3 Recovery of Boltzmann Entropy

If a system has Ω equally distinguishable microstates, specifying one requires $C = \log_2(\Omega)$ committed bits. Therefore:

$$\mathbf{S_c = k\_IC \cdot \log_2(\Omega) = k\_B \ln 2 \cdot \log_2(\Omega) = k\_B \ln(\Omega)}$$

This exactly reproduces Boltzmann entropy. For non-uniform ensembles, replace $\log_2(\Omega)$ with the Shannon entropy $H = -\Sigma_i\, p_i \log_2 p_i$; see Section 9.4. □

9.4 Shannon Entropy Bridge

For a probabilistic ensemble $\{p_i\}$, the expected committed bits to identify an outcome is the Shannon entropy $H = -\Sigma_i\, p_i \log_2 p_i$. Therefore:

$$\mathbf{\mathbb{E}[S_c] = k\_IC \cdot H = k\_B \ln 2 \cdot H = k\_B\ H\_nat}$$

bridging directly to Gibbs entropy. Both thermodynamic and information-theoretic entropy measure committed bit-cost under the same constant k_IC. □

9.5 Landauer Consistency

Landauer's principle: erasing one committed bit dissipates at least $Q\_min = k\_B\ T \ln 2$. In BCB: $\Delta S_c = k\_IC = k\_B \ln 2$ per bit erased. Therefore:

$$\mathbf{Q\_min = T \cdot \Delta S_c = k\_B\ T \ln 2}$$

Landauer's principle is a theorem of IC, not an independent postulate. □

9.6 Finite Distinguishability Correction

In the limit $C_{min}/C_{max} \to 0$ (ordinary conditions): $\Omega_{FD} \to \Omega$ and $S_R \to k_B \ln \Omega$, recovering standard statistical mechanics. Near $C_{min}/C_{max} \to 1$ (Planck-scale or near-extremal black hole): the FD correction becomes $O(k_B)$, predicting a systematic downward deviation of entropy from the Bekenstein–Hawking value. □

🔍 **In plain language** — This section is the technical backbone of the paper, but the takeaway is simple: every standard result in thermodynamics — Boltzmann's entropy formula, Shannon's information measure, Landauer's erasure cost, the second law — drops out of the BCB framework automatically. None of them need to be assumed. They are all consequences of the single idea that entropy counts committed bits at cost $k_B \ln 2$ each. If you accept the definition, you get all the results for free. The consistency checks in this section demonstrate that the framework is not replacing thermodynamics with something incompatible — it is providing a deeper foundation that contains thermodynamics as a special case.

---

10. Comparison With Standard Entropy

Property	Standard statistical mechanics	BCB–TPB framework
Entropy definition	Missing information; log of microstates	Accumulated committed distinctions
Physical basis	Abstract ensemble; phase-space volume	Causal ledger; irreversible records
Arrow of time	Statistical improbability of decrease	Structural irreversibility of commitment
Initial conditions	Fine-tuned low-entropy past hypothesis	Minimal boundary $C = 0$; structurally necessary
Second law	Overwhelming probability argument	Theorem from IC axiom
Landauer principle	External postulate	Derived consequence of k_{IC}
FD correction	Absent	$\Delta S = -k_B \alpha \ln \Omega$ (suppression $\Omega_{FD} = \Omega^{(1-\alpha)}$)
Bekenstein bound	Separate result	Derived as maximum committed bits per emergent second
Time	Background parameter	Emergent: physical clock intervals generated by the bit-accumulation sequence

---

11. The Role of Entropy Across Physics

Entropy is often introduced through thermodynamics, but in modern physics it appears across a remarkable range of domains — from gravitational dynamics to quantum information theory. In the BCB–TPB framework, these appearances share a common origin: all of them involve the formation, transport, or redistribution of **committed distinctions**. Because entropy counts committed bits, any physical process that generates, stores, or processes stable records contributes to entropy. This section surveys those appearances and shows how the BCB framework provides a unifying interpretation.

🔍 **In plain language —** Entropy turns up everywhere in physics — in heat engines, in black holes, in computer memory, in the measurement of quantum systems. This has always seemed suspicious: why should the same quantity govern such different things? The answer, in this framework, is that they are not different things. They are all the same process — the universe writing irreversible records — viewed from different angles. This section takes each domain in turn and shows how the committed-bit picture unifies them.

11.1 Thermodynamics

In classical thermodynamics, entropy governs heat flow, equilibrium, and irreversibility. In the BCB framework, these phenomena arise because macroscopic processes generate committed distinctions — molecular collisions, phase transitions, and radiative emission events all produce irreversible records in the causal ledger. Heat flows and chemical equilibria therefore correspond to the redistribution of commitment activity until spatial gradients in entropy density diminish. The second law $dS_c/d\tau \geq 0$ follows directly from the irreversibility of commitment, not from a statistical argument.

11.2 Information and Computation

In information theory, entropy measures the information content of a message. Landauer's principle connects informational entropy directly to thermodynamic cost: erasing one bit dissipates at least k_B T ln 2 of heat. The BCB framework unifies these perspectives by identifying entropy with the physical cost of record formation itself. Computational processes generate entropy precisely because they commit distinctions into physical memory registers. Information entropy and thermodynamic entropy are therefore two descriptions of the same quantity: committed bit cost, measured in units of k_IC = k_B ln 2 per bit.

11.3 Chemistry and Statistical Physics

Chemical reactions and molecular processes are governed by free energy differences that trace back to entropy changes. In the BCB picture, reactions transform potential distinctions in reactant configurations into committed distinctions in product configurations and environmental records — emitted photons, phonons, or molecular rearrangements. Reaction equilibria correspond to states where further commitment offers no net increase in record-resolvable accessible configurations. The standard Gibbs criterion for equilibrium ($\Delta G = 0$) is a consequence of commitment saturation.

11.4 Quantum Measurement and Decoherence

Quantum systems evolve reversibly until interactions with the environment produce stable records — a process described as decoherence. Decoherence supplies robust, effectively classical records, satisfying the stability and distinguishability conditions of Section 2. A committed bit, as defined here, additionally requires irrecoverability and causal usability; whether decoherence alone suffices for irrecoverability depends on the admissible recovery operations defined for the system — this paper remains neutral on that aspect of the measurement problem. The entropy increase associated with measurement reflects the cost of turning quantum possibilities into classical records. Before the committed bit is written: no ledger entry, no entropy contribution, no ledger advance. After it is written: ledger entry recorded, commitment cost paid, one step of emergent time generated.

11.5 Gravity and Spacetime

Entropy plays a central role in gravitational physics. Black holes possess entropy proportional to their horizon area, and thermodynamic arguments (Jacobson 1995) can reproduce Einstein's field equations from entropy considerations alone. Within the BCB framework, these connections arise because entropy density and commitment density are the same field (Section 12). Spatial gradients in commitment density determine the spacetime volume element and therefore generate curvature. Gravity emerges as the geometric response to variations in causal record formation, not as a separate force imposed on spacetime.

11.6 Cosmology and the Arrow of Time

The arrow of time is closely linked to entropy increase in cosmology. Standard accounts require a fine-tuned low-entropy past. In the BCB framework, the arrow emerges directly from the structure of irreversible record formation: the causal ledger grows monotonically as new distinctions are written into the universe, starting from the minimal boundary condition $C = 0$. Cosmological entropy growth reflects the ongoing accumulation of committed records as the universe evolves — structurally necessary, not statistically accidental.

11.7 Summary: One Quantity, Many Appearances

Across thermodynamics, information theory, chemistry, quantum measurement, and gravity, entropy plays the same fundamental role: it counts the irreversible distinctions that have become part of the physical record.

Domain	Entropy manifestation	BCB interpretation
Thermodynamics	Heat, equilibrium, irreversibility	Redistribution of committed distinctions
Information theory	Message content, erasure cost	Committed bit cost at k_IC per bit
Chemistry	Free energy, reaction equilibrium	Potential → committed distinctions
Quantum physics	Decoherence, measurement	Superposition → committed record
Gravity	Black hole entropy, spacetime curvature	Commitment density → spacetime geometry

Domain	Entropy manifestation	BCB interpretation
Cosmology	Arrow of time, entropy growth	Monotonic ledger accumulation from C = 0

The BCB framework does not merely reinterpret entropy within thermodynamics. It identifies a single physical quantity underlying all of these appearances: the cost of causal record formation.

---

12. Entropy Gradients and Physical Dynamics

The second law establishes that entropy cannot decrease globally:

$$dS_c/d\tau \geq 0$$

But the second law alone does not explain why physical systems *move*. The driver of dynamics is not entropy itself but its gradient — the spatial variation in commitment density that causes systems to evolve toward equilibrium. In the BCB framework, this has a direct and concrete interpretation.

Since commitment entropy is defined as $S_c = k\_B \ln 2 \cdot C$, the spatial entropy gradient is directly proportional to the gradient of commitment density:

$$\nabla S_c = k\_B \ln 2 \cdot \nabla C$$

Regions of space that generate committed records more densely per ledger step possess higher entropy density. When neighbouring regions exhibit different commitment densities, a gradient forms. Physical systems evolve so as to reduce these gradients — producing the familiar irreversible processes of thermodynamics. In near-equilibrium thermodynamics, this reduction of ∇s_c corresponds to the standard Onsager gradient-flow structure: thermodynamic fluxes are proportional to thermodynamic forces, $\mathbf{J = L\ X}$, where the force $X = -\nabla(\partial F/\partial x)$ is the negative gradient of the conjugate potential (so that $J = -L\ \nabla(\partial F/\partial x)$, with the minus sign arising from the potential convention rather than the force convention). Here reinterpreted as redistribution of commitment activity across the causal ledger.

Physical process	What drives it in standard terms	BCB interpretation
Heat conduction	Temperature gradient	Entropy production and record creation redistribute under local coupling to reduce ∇s_c
Diffusion	Concentration gradient	Probability mass flows toward configurations with larger record-resolvable phase volume S_R
Chemical reaction	Free energy gradient	Reactant distinctions (potential) → product distinctions (committed); gradient in S_R drives equilibration

Physical process	What drives it in standard terms	BCB interpretation
Radiative cooling	Entropy production	Commitment occurs at emission and absorption events; spatial redistribution reduces gradients in commitment density
Spacetime curvature	Stress-energy tensor	Commitment density gradient ∇C sources gravitational field via the volume-form constraint

The last row is particularly significant. In the companion VERSF spacetime paper, commitment density ρ_c already appears as the source of spacetime curvature — playing the role of the stress-energy tensor in the emergent Einstein equations. The entropy gradient $\nabla S_c = k_B \ln 2 \cdot \nabla C$ is therefore the same gradient that curves spacetime. This means:

commitment density gradients → entropy gradients → physical forces

This is not an analogy. It is a direct identification: the gradient responsible for thermodynamic transport is the same gradient that generates gravitational dynamics in the BCB framework. Entropy gradients are not abstract statistical tendencies — they are physical gradients in causal record formation, and the dynamics of macroscopic processes can be understood as the universe redistributing commitment activity until those gradients diminish.

Heat source example. A burning fuel source commits bits at high rate; the surrounding air commits bits at lower rate. The commitment density gradient ∇C points outward from the source. Heat flows down this gradient as the high-commitment region drives record formation into the low-commitment surroundings — exactly as thermodynamics requires, but now with a causal-record interpretation of why.

Black hole example. Near a black hole horizon, the commitment density becomes extreme — bounded above by $C_{max} = S_{BH}/(k_B \ln 2) = A/(4\ell_P^2 \ln 2)$, where S_{BH} is the Bekenstein–Hawking entropy and A the horizon area. The gradient ∇C near the horizon is steep, producing both a strong thermodynamic entropy gradient and, via the BCB gravity mechanism, strong spacetime curvature. In the near-extremal limit, $C_{min}/C_{max} \to 1$ and the FD correction to entropy (Section 4) becomes significant precisely where the gradient is steepest — the two effects are not independent but manifestations of the same underlying saturation of commitment capacity.

🔍 **In plain language —** Entropy doesn't just accumulate — it *pulls*. Wherever there are more records being made on one side than another, there's a gradient, and nature tends to even those gradients out. That's why heat flows from hot to cold (the hot region is making records faster), why gases diffuse to fill a room (the dense side has a higher commitment rate), why chemical reactions run toward equilibrium. These aren't separate phenomena with separate explanations. They are all the same thing: the universe smoothing out its record-making rate. And near a black hole, where the commitment density is at its maximum, the gradient is steepest — which is exactly where gravity is strongest. The connection between thermodynamics and gravity is not a coincidence in this framework. It is the same gradient, viewed from two different angles.

13. Entropy Gradients and Emergent Geometry

The previous section established that entropy gradients drive thermodynamic transport. This section takes the next step: spatial variations in entropy density do not merely move heat and particles — they **shape spacetime geometry itself**. This is the bridge between the BCB thermodynamics framework and the record-theoretic derivation of spacetime in the companion VERSF paper.

13.1 Entropy Density and Commitment Density Are the Same Field

Commitment entropy is defined globally as $S_c = k\_B \ln 2 \cdot C$, where C is the total accumulated committed bits. Locally, the entropy density s_c at a point in the emergent spacetime description is directly proportional to the local commitment density ρ_c — the committed records per causal cell:

$$\mathbf{s_c = k\_B \ln 2 \cdot \rho_c}$$

Here $\rho_c(x)$ is the proper commitment density expressed within the emergent spacetime: committed bits per unit emergent proper 4-volume, where the 4-volume itself is derived from the record ensemble via the geometric framework of the companion VERSF paper. The time slicing in this expression is not a pre-given background — it is fixed by the local proper-time slicing that *emerges* from the record definition. This convention locks the field equation to the companion paper's spacetime framework while making explicit that proper time is a derived, not primitive, quantity.

This is not a model assumption. It follows directly from the definitions: entropy counts committed bits, and commitment density counts committed bits per emergent causal 4-cell. The two quantities are the same physical field, related by a fixed constant.

13.2 The Geometric Consequence

In the record-theoretic derivation of spacetime geometry, the spacetime volume element is determined by the Jeffreys distinguishability density J(C) of the local record ensemble. In general:

$$\mathbf{\sqrt{|g|} \propto J(C)}$$

In the Poisson-like near-equilibrium universality class — the natural universality class for dilute, weakly-correlated record ensembles — the Jeffreys density satisfies $J(C) \propto \rho_c$, giving:

$$\mathbf{\sqrt{|g|} \propto \rho_c} \quad \textit{(Poisson-like near-equilibrium regime)}$$

Universality class criteria. The regime $J(C) \propto \rho_c$ holds when: (i) record counts in small causal cells have independent-increment statistics (variance ≈ mean, i.e., Poisson-like); (ii) drift and flux parameters enter through finite-variance transport families; and (iii) correlations between record events are short-ranged compared to the coarse-graining scale. These are the standard conditions under which Fisher scaling yields $J \sim \rho_c$. Departures from these conditions — for example, in strongly correlated quantum systems or near criticality — require the full J(C) expression and may produce non-linear scaling of $\sqrt{|g|}$ with ρ_c.

meaning the local spacetime measure is proportional to the local commitment density in that regime. Substituting the entropy density relation:

$$\sqrt{|g|} \propto s_c$$

The spacetime volume element is proportional to local entropy density. This is a deep structural result: **geometry is the metric expression of how densely the universe is writing records**.

13.3 The Full Chain

The programme from irreversible records to gravity is now complete and explicit:

```
irreversible committed records
          ↓
commitment entropy  Sc = k_B ln 2 · C
          ↓
local entropy density  sc = k_B ln 2 · ρc
          ↓
spacetime volume element  √|g| ∝ sc
          ↓
metric variation → curvature
          ↓
gravity
```

The single compact expression unifying the chain is:

$$\sqrt{|g|} \propto J(C) \sim \rho_c = s_c / (k_B \ln 2)$$

The first proportionality is the record-calibrated volume-form constraint from the spacetime derivation; the middle relation holds in the Poisson-like near-equilibrium universality class; the last equality is the definition of entropy density from committed-bit density. Every symbol appears in either the entropy framework or the spacetime framework — their identification is the statement that thermodynamics and geometry are two descriptions of the same underlying process of causal record formation.

Scope and dependence on the micro-model. The robust statement is the record-calibrated volume-form constraint $\sqrt{|g|} \propto J(C)$, where J(C) is the Jeffreys density determined by the Fisher metric induced by the local record micro-model p(r|C). The simplification $J(C) \propto \rho_c$ — and hence $\sqrt{|g|} \propto s_c$ — holds in the Poisson-like near-equilibrium universality class, where Fisher scaling yields $J \sim \rho_c$. Outside that regime, the entropy–geometry dictionary remains valid in the general form $\sqrt{|g|} \propto J(C)$; what changes is the constitutive relation J(C) implied by p(r|C). This keeps the mapping structurally fixed while making its coefficient and functional form micro-model dependent, as expected for a universality-class claim.

13.4 Physical Interpretation

Regions of higher entropy density possess greater distinguishability capacity — more committed distinctions per emergent causal 4-cell. In the Poisson-like near-equilibrium regime, this raises the local spacetime volume element $\sqrt{|g|}$, which specifies the conformal factor of the metric.

From volume element to curvature (one-step unpacking). Given a fixed causal cone structure, the metric is determined up to a conformal factor: $g^{uv} = \Omega^2(x)\, \tilde{g}^{uv}$, where $\tilde{g}^{uv}$ is any representative of the cone-fixed conformal class. The volume element satisfies $|g| = \Omega^8 |\tilde{g}|$ in four spacetime dimensions (since $\sqrt{|g|} = \Omega^4 \sqrt{|\tilde{g}|}$), so specifying $\sqrt{|g|}$ fixes $\Omega(x)$ locally as $\Omega \propto (\sqrt{|g|})^{(1/4)}$ in a local inertial chart with $|\tilde{g}| = 1$. Spatial variation in $\Omega(x)$ yields nonzero derivatives of the metric — a nontrivial Levi-Civita connection — and therefore nonzero curvature. Thus record-capacity variation (encoded in $\sqrt{|g|} \propto J(C)$) is sufficient to induce curvature once the cone structure is fixed. The cone structure itself is fixed by the invariant influence speed and Lorentz structure in the record-theoretic kinematic derivation of the companion paper (or assumed as standard in the emergent layer if that derivation is not invoked).

The result is that **geometry bends where the universe is writing more records**.

This reframes gravity not as the response of spacetime to energy-momentum, but as the geometric response to spatial gradients in causal record formation. More precisely, curvature arises from spatial variation of the record-capacity field — equivalently, entropy density in the Poisson-like near-equilibrium regime — via the volume-form constraint within the cone-fixed conformal class. This is structurally related to Jacobson's thermodynamic derivation of Einstein's equations from horizon entropy, but grounded in a deeper substrate — commitment density rather than horizon area — and derived from first principles rather than assumed thermodynamic identities.

13.5 Summary of the Entropy–Geometry Dictionary

Thermodynamic concept	Geometric concept	Connecting relation
Committed bit C	Spacetime event	Unit of causal record
Commitment density ρ_c	Spacetime source field	$\rho_c = s_c / k_B \ln 2$
Jeffreys density J(C)	Spacetime volume element	√
Entropy density s_c	Local spacetime measure	√
Entropy gradient ∇s_c	Conformal factor variation	Source of connection and curvature

Thermodynamic concept	Geometric concept	Connecting relation
Entropy accumulation per ledger step $dS_c/d\tau$	Temporal emergence (TPB)	Physical clock interval per step = 1/TPB

This dictionary shows that the entropy paper and the spacetime paper are not parallel developments — they are **the thermodynamic and geometric sectors of a single unified framework**, both grounded in the irreversible accumulation of committed distinctions.

🔍 **In plain language —** This is perhaps the most striking result in the paper. Einstein's theory of gravity says that mass and energy curve spacetime. This framework says something deeper: what really curves spacetime is the *density of irreversible record-making*. Regions where the universe is committing records at high density have a larger local spacetime measure — more room, in a sense, for events to happen. Where that density varies from place to place, the metric bends. That bending is gravity. Mass curves spacetime because mass is associated with high commitment density. Entropy and geometry aren't separate ideas — they are two windows onto the same underlying process of the universe writing its own history. The table above shows exactly how the vocabulary of thermodynamics maps onto the vocabulary of geometry, term by term.

---

14. Experimental Predictions

The BCB–TPB framework makes three classes of predictions distinguishable from standard statistical mechanics:

P1 — Near-extremal black hole entropy suppression. As a black hole approaches extremality, $\alpha = C_{min}/C_{max} \to 1$ and the FD correction becomes significant. The $\Omega^{(1-\alpha)}$ suppression mechanism of Section 4 applied to the black hole microstate count Ω_{BH} gives:

$$\Omega_{FD} = \Omega_{BH}^{(1-\alpha)}$$

Taking the logarithm, the effective thermodynamic entropy is:

$$S_{effective} = k_B \ln \Omega_{FD} = (1-\alpha)\, k_B \ln \Omega_{BH} = (1-\alpha)\, S_{BH}$$ *(indicative; not a precision prediction)*

The predicted suppression is **linear** in $(1-\alpha)$: entropy diminishes proportionally as extremality is approached, vanishing as $\alpha \to 1$. This is qualitatively distinct from both the standard Bekenstein–Hawking result (no suppression) and the logarithmic corrections predicted by loop quantum gravity and string-theoretic approaches — both of which add subleading terms but do not suppress the leading area term. A linear suppression by factor $(1-\alpha)$ provides a clear empirical signature for discrimination between frameworks, pending a full quantum-gravitational treatment of C_{min} and C_{max} in curved spacetime.

P2 — Bekenstein ceiling on committed bits. Systems approaching the maximum committed-bit density per unit emergent mass-energy should exhibit a hard ceiling on entropy accumulation per emergent second: $dS_c/dt\_em \leq 2\pi k\_B\, mc^2/\hbar$. Any apparent violation of this bound would falsify the identification of entropy with committed bits within the emergent physical description.

P3 — Additional decoherence channel from commitment activity (substrate-conditional). If commitment events correspond to objective irreversibility in the record substrate, then in addition to standard environmental decoherence the framework predicts an extra decoherence contribution proportional to local commitment density:

$$\Gamma\_tot = \Gamma\_env + \lambda\, \rho_c$$

where Γ_env is the standard Lindblad/Caldeira–Leggett environmental rate, ρ_c is the local commitment density, and λ is a proportionality constant set by the micro-model. This extra channel produces a measurable deviation from standard decoherence scaling in regimes where environmental coupling is held fixed but commitment activity differs — for example, identical interferometers operated in environments with different record-formation activity. Null results bound λ from above; a positive detection would provide direct evidence for the objective commitment substrate. The prediction is conditional on the micro-model connecting commitment density to the Lindblad environment: if that connection is not independently established, Γ_add represents a conjecture rather than a derivation, but it is in principle falsifiable.

🔍 **In plain language** — A good theory doesn't just explain what we already know — it sticks its neck out and says something new. This section does that in three ways. First: near the edge of a black hole, the entropy should be slightly *lower* than Einstein's formula predicts, because at such extreme densities, fewer microstates are actually distinguishable. The suppression follows a specific mathematical form (exponential, not logarithmic) that would let us tell this prediction apart from those of competing theories. Second: no physical system can accumulate committed bits faster than a ceiling set jointly by its mass and quantum mechanics — and if any system appeared to exceed this, the whole framework would be falsified. Third: a quantum system can't "commit" to a definite state until a minimum number of ledger steps have elapsed — enough to satisfy the stability condition. What we usually call "decoherence time" is, in this framework, a minimum ledger depth: a minimum count of committed bits required before a definite record can exist at all.

---

15. Discussion

The BCB–TPB reformulation of entropy is not interpretational window-dressing. It makes structurally distinct claims:

1. **Entropy is a ledger count.** $S_c = k\_IC \cdot C$ is grounded in physical events satisfying the four committed-bit conditions, not in abstract probability distributions.
2. **The constant k_IC is derived, not fitted.** Setting $k\_IC = k\_B \ln 2$ follows from Landauer's principle and fixes the framework without free parameters.

3. **The FD correction predicts new physics.** The suppression of Ω_FD near extreme commitment densities is in principle measurable, and has a distinctive exponential form.
4. **The arrow of time is structural.** The asymmetry of commitment events — defined by causal usability — explains temporal directionality without a fine-tuned initial condition.
5. **The Bekenstein bound is a consequence.** The maximum entropy rate $dS_c/dt \leq 2\pi k\_B\, mc^2/\hbar$ is not imported but derived from the BCB commitment rate combined with the relativistic energy bound.
6. **Temporal emergence is quantified.** What we call "time" is the ledger being written: physical clock intervals emerge from the bit-accumulation sequence via $dS_c/d\tau = k\_IC/TPB$, making time a derived quantity of the commitment process rather than its container.

---

16. Conclusion

We have reformulated entropy as the accumulated physical cost of irreversible commitment: $\mathbf{S_c = k\_B \ln 2 \cdot C}$. This definition recovers Boltzmann and Gibbs entropy exactly in the classical limit, bridges to Shannon entropy with a fixed physically motivated constant, derives Landauer's principle as a theorem, and provides a structurally grounded account of the second law. Crucially, what we call physical time is not assumed — it emerges from the sequential ordering of committed bits via the TPB mapping, making time a derived consequence of the commitment process rather than its background. The Bekenstein information bound emerges naturally as the maximum committed-bit density per unit emergent mass-energy, and a finite-distinguishability correction to microstate counts predicts measurable deviations at extreme commitment densities.

Entropy is not a measure of disorder. It is the physical cost of causal history. Thermodynamics therefore emerges as the bookkeeping system of irreversible record formation — a ledger that the universe has no choice but to keep, and that grows without bound as long as physical distinctions continue to be made.

Appendix A: Notation Summary

Symbol	Definition
$C(\tau)$	Cumulative committed bits up to ledger time τ
k_IC	Commitment constant = $k\_B \ln 2$
S_c	Commitment Entropy = $k\_IC \cdot C$
S_R	Record-Resolvable Entropy = $k\_B \ln \Omega\_FD = k\_B (1-\alpha) \ln \Omega$
S_cap	Capacity entropy budget = $k\_B \ln 2 \cdot C\_max$
$S_c + S\_R \leq S\_cap$ (equality under maximal-mixing closure only)	
Ω	Standard microstate count

Symbol	Definition
Ω_{FD}	Physically resolvable microstates = $\Omega^{(1-\alpha)}$, $\alpha := C_{min}/C_{max}$
α	FD suppression parameter = C_{min}/C_{max}
C_{min}	Minimum committed bits to resolve one adjacent microstate distinction
C_{max}	Total committed capacity of causal region
TPB	Ticks-Per-Bit = $d\tau/db$
ε	Distinguishability threshold (statistical distance)
ℓ_P	Planck length
$J(C)$	Jeffreys distinguishability density of the record ensemble
ρ_c	Proper commitment density (bits per unit emergent proper 4-volume)
s_c	Local entropy density = $k_B \ln 2 \cdot \rho_c$

Appendix B: Exponential Capacity-Thinning and the FD Suppression

This appendix provides a minimal derivation showing why exponential suppression of the physically resolvable microstate count is the natural form of the finite distinguishability (FD) correction.

B.1 Setup: Distinguishability Constraints Under Finite Capacity

Consider a macroscopic constraint set defining a mathematical microstate count Ω. Standard statistical mechanics assumes all Ω microstates are physically resolvable in principle. Under FD, only a subset is resolvable: discriminating between alternatives requires committed records, and the causal region has finite committed capacity C_{max}.

Let "resolving microstates" be represented as imposing a sequence of distinguishability constraints that progressively refine the state description. The number of refinement steps required to identify a microstate among Ω alternatives is proportional to $\ln \Omega$ (natural-log counting), or equivalently $\log_2 \Omega$ in bit counting.

We model the FD effect as follows:

- Each refinement step requires committing at least C_{min} bits of record capacity to remain operationally resolvable above threshold ε.
- As available capacity is consumed, the probability that the next refinement step remains physically realisable decreases in proportion to the fraction of remaining capacity.

This is the minimal **capacity-thinning** assumption: resolvability is multiplicatively suppressed as capacity saturates.

B.2 Capacity-Thinning Model

Let n index refinement steps and C(n) be the cumulative committed capacity consumed after n steps. Take $C(n) = n \cdot C_{min}$. The survival probability that the $n \to n+1$ refinement remains physically resolvable is:

$$p_{surv}(n) = 1 - C(n)/C_{max} = 1 - n \cdot C_{min}/C_{max}$$

valid in the regime $n \cdot C_{min} \le C_{max}$. The probability that a full refinement of depth N is physically realisable is the product:

$$P_{surv}(N) = \prod_{n=0}^{N-1} (1 - n \cdot C_{min}/C_{max})$$

For $C_{min}/C_{max} \ll 1$ and N not close to saturation, taking the log and expanding $\ln(1-x) \approx -x$ gives:

$$\ln P_{surv}(N) \approx -(C_{min}/C_{max}) \cdot N(N-1)/2$$

Two points are immediate: (i) the suppression is multiplicative and therefore naturally exponential in the continuum limit; (ii) suppression strength is controlled by the capacity fraction C_{min}/C_{max}. The quadratic dependence on N here is an artifact of the linear thinning choice; the following subsection shows how a per-log-refinement assumption produces the single-exponential form used in the main text.

B.3 Single-Exponential Form from Constant Fractional Thinning

Index refinement by a continuous description-depth variable $\xi \in [0, \ln \Omega]$, where ξ increases by one unit for each nat of distinguishability extracted. Assume each unit increase $d\xi$ requires committing a fraction C_{min}/C_{max} of available capacity in expectation. The rate of loss of resolvable alternatives is then proportional to the remaining alternatives:

$$d\Omega_{FD}/d\xi = -\alpha \cdot \Omega_{FD}, \quad \alpha := C_{min}/C_{max}$$

This is the standard multiplicative depletion equation — the same structure that yields exponential decay in reliability models and channel-capacity thinning. Solving with initial condition $\Omega_{FD}(0) = \Omega$:

$$\Omega_{FD}(\xi) = \Omega \cdot e^{-\alpha\xi}$$

Evaluating at full refinement depth $\xi = \ln \Omega$:

$$\Omega_{FD} = \Omega \cdot \exp(-\alpha \ln \Omega) = \Omega^{(1-\alpha)}$$

This is the exact closed-form result derived from the constant fractional thinning model. The power-law form $\Omega^{(1-\alpha)}$ is cleaner and fully equivalent: it states directly that finite capacity reduces the effective phase volume by a power law with exponent $1-\alpha$. Both expressions express the same physical claim — finite record capacity causes multiplicative contraction of the physically resolvable phase volume — and the main text uses the $\Omega^{(1-\alpha)}$ form throughout for consistency.

B.4 Interpretation and Limits

Unsaturated regime. If $C_{min}/C_{max} \ll 1$, the FD correction is negligible and $\Omega_{FD} \to \Omega$. Standard statistical mechanics is recovered exactly.

Saturation regime. As $C_{min}/C_{max} \to 1$, the suppression becomes order-unity and departures from standard statistical mechanics become significant — precisely the regime of near-extremal black holes and Planck-scale densities.

Model dependence. Different micro-models $p(r|C)$ correspond to different thinning functions and therefore potentially different functional forms for Ω_{FD}. The exponential form is the canonical result for independent multiplicative capacity-thinning and is therefore a natural baseline prediction. Deviations from exponential suppression would indicate correlated distinguishability constraints or non-Poisson record statistics — themselves testable differences between universality classes.

B.5 Summary

The FD correction is a capacity effect: the physically resolvable alternative set contracts because stable record formation consumes finite causal capacity. Under the minimal assumption that resolvability depletes multiplicatively with capacity consumption, the physically resolvable microstate count follows a power-law suppression:

$$\Omega_{FD} = \Omega^{(1-\alpha)}, \quad \alpha := C_{min}/C_{max}$$

This compact closed form — used throughout the main text — states directly that finite capacity reduces the effective phase volume by an exponent $1-\alpha$. In the unsaturated regime $\alpha \ll 1$, $\Omega_{FD} \to \Omega$ and standard statistical mechanics is recovered. At saturation $\alpha \to 1$, $\Omega_{FD} \to 1$ — all microstate distinctions are suppressed.