Inevitable Actualization: Beyond Necessity and Contingency

A.J. Schock

April 26, 2025

Abstract

I introduce Inevitable Actualization (IA), an ontological modality: if (1) the universe's future time involves an unbounded sequence of causal trials (H_∞) and (2) a state S has a non-zero physical probability $P_n > 0$ in trial n such that the sum $\sum_{n=1}^{\infty} P_n$ diverges, then S is guaranteed to occur with probability one. IA is developed through a rigorous measure-theoretic foundation, probabilistic modeling with dependence (under standard mixing conditions) and absorbing-state exceptions, contrasting IA with classical modalities and modern multiverse theories. Positioned as a distinct third category alongside necessity and contingency, IA's unique grounding rests on temporal structure and probability. I address objections (Boltzmann brains, the measure problem, and identity duplication) and illustrate IA's implications for ethics, cosmology, and personal identity, acknowledging formal challenges.

1 Introduction: IA as Third Modality

Classical metaphysics distinguishes necessity (true in all possible worlds; cannot fail to exist) from contingency (true in some but not all possible worlds; might fail to exist) (Plantinga, 1974; Lewis, 1986). This paper introduces and defends a third category, Inevitable Actualization (IA), defined conditionally on the structure of time and probability.

Conditional Thesis: If (a) the universe's future unfolds through an unbounded sequence of causal trials τ_n (H_∞), provided the system displays sufficiently fast decay of temporal correlations (φ-mixing; see App F), and (b) a physically possible state S has probability $P_n = P(S \text{ occurs in trial } \tau_n)$ such that $\sum_{n=1}^{\infty} P_n$ diverges (the Persistence Condition, §6.5), then S occurs with probability one.

While IA exhibits characteristics reminiscent of both necessity (providing a guarantee of eventual occurrence within its framework) and contingency (being entirely dependent on specific, potentially falsifiable cosmological conditions like unbounded time), we contend that its unique grounding mechanism justifies its classification as a distinct ontological category. Unlike classical necessity, rooted in logic or essential properties, and contingency, rooted in mere possibility within causal histories, IA's inevitability arises specifically from the structural interplay of temporal extent and probabilistic physical laws. Its guarantee is fundamentally conditional, distinguishing it sharply from unconditional metaphysical necessity. IA is thus a temporal-probabilistic modality. This paper formalizes IA, situates it among existing theories, and explores its philosophical consequences, leveraging a conditional framework that allows exploration without requiring definitive empirical proof of infinite time itself (Jaynes, 2003; Norton, 2021). The $P(S) > 0$ condition refers to the objective probability within the actual physical system, including any inherent stochasticity (quantum effects, noise). The concepts of "causal trial" (a discrete time-slice; see §2) and "Minimal Causal History" are central and require careful definition (§2) to ensure the framework applies coherently across different physical scales and state complexities. IA applies only if both conditions (H_∞ and persistence) hold strictly; even a very distant but finite time cutoff or eventual decay of $P(S)$ such that $\sum P_n < \infty$ invalidates the guarantee (§6.5).

This paper proceeds as follows:

- §2: Defines "occurrence" rigorously, including "Minimal Causal History", and discusses scope.
- §3: Situates IA within metaphysics (modality, realism).
- §4: Briefly reviews related literature and challenges.
- §5: Develops the measure-theoretic foundations.
- §6: Presents the formal probabilistic model.
- §7: Engages with key objections.
- §8: Explores applications and implications.
- §9: Concludes and outlines future work.
- Appendices: Provide technical details (Apps A-H).

2 Defining "Occurrence" with Causal History

To avoid counting fleeting, physically meaningless fluctuations (like instantaneous Boltzmann Brain flashes) as genuine realizations, and to ensure we are tracking meaningful states, we refine the definition of an event's "occurrence":

Definition: Sustained Occurrence. A trial i (corresponding to a temporal interval $[t_i, t_{i+1})$) "realizes" state S if and only if:

1. **Minimal Causal History:** $History(State(t_i)) \supseteq MinimalHistory(S)$
2. **Sustained Duration:** $State(t) \in Region(S, \delta)$ for all $t \in [t_i, t_i + T_{min}] \subset [t_i, t_{i+1})$

- **Minimal Causal History:** This condition ensures that an occurrence of S is not merely a configuration that looks like S, but one that has arisen through a physically appropriate pathway. It refers to the necessary set of antecedent conditions or developmental stages without which the state S would not be considered a genuine instance of its type. For example, the state S = "a mature oak tree" requires a history including acorn and seedling stages; a random atomic configuration resembling an oak (a Boltzmann Tree) lacks this history. Similarly, S = "a correctly folded protein" requires ribosomal synthesis, not just spontaneous assembly. The precise specification of $MinimalHistory(S)$ depends heavily on the nature of state S—trivial for fundamental particles, crucial for complex, evolved, or information-bearing states. Rigorously applying this condition requires defining a measure over the space of possible histories and demonstrating that the subset satisfying $MinimalHistory(S)$ has positive measure under the physical dynamics; this poses a significant formal challenge, particularly for highly complex states (see §2.1).
- **Sustained Duration (T_{min}) & Tolerance (δ):** These ensure IA applies to states with meaningful persistence and avoids issues of exact recurrence (which often has probability zero). The choice of T_{min} and tolerance δ would depend on the specific state S being considered.

Simple Example: Coin Flips. Consider S = "a run of at least 3 consecutive Heads". Here, a "trial" is a single coin flip. The state S requires $T_{min} = 3$ trials (flips). The tolerance δ is effectively zero for the discrete outcome "Heads". The $MinimalHistory(S)$ is simply the two preceding flips also being Heads. A sequence ...T H H H... realizes S starting at the first H of the run, while ...T H T H... does not.

Complex Example: If S represents "a conscious thought about Vienna," T_{min} might be on the order of milliseconds or seconds (the minimum duration for such a thought process to complete coherently), and δ would define the allowable variation in the underlying neural state. In contrast, if S represents "a stable Helium-4 atom," T_{min} could be vastly longer (reflecting its nuclear stability timescale), and δ would represent small variations in its quantum state (e.g., position, momentum within certain bounds).

The region $Region(S, \delta)$ is assumed to have non-zero measure in the physically accessible phase space.

2.1 Scope and Extension to Complex States

The formal machinery of IA, particularly the application of the Persistence Condition ($\sum P_n$ diverges) using theorems like Borel-Cantelli, relies on well-defined probabilities P_n for the occurrence of state S in trial τ_n. This, in turn, depends on the Sustained Occurrence definition being mathematically tractable.

- **Formal Scope:** Rigorously demonstrating that a state S satisfies the conditions for IA is most feasible for coarse-grained, relatively simple states where the phase space region $Region(S, \delta)$, the minimum duration T_{min}, and crucially, the $MinimalHistory(S)$ filter can be explicitly characterized and shown to have positive probability under the dynamics. The cellular automaton example (Appendix G) illustrates such a case where history conditions can be formally imposed. The formal claims of IA, grounded in probability theory, are strongest for such well-defined states.

- **Extension to Complex States (Open Problem):** Applying IA rigorously to high-level macroscopic states, such as "a conscious thought about Vienna" or "a functioning biological organism," presents significant challenges. Defining the precise phase space region (δ), minimum duration (T_{min}), and especially the set of admissible causal histories ($MinimalHistory(S)$) in a measurable way is currently intractable. Characterizing the measure over the space of all possible trajectories (histories) for such complex systems is a profound open problem in statistical mechanics and complex systems science. Demonstrating that the set of histories satisfying $MinimalHistory(S)$ has positive measure under the physical dynamics—for example, proving that stable self-replicating polymers can form with non-zero probability in standard interacting particle models or QFT settings—is a key requirement that remains conjectural for complex S and is a crucial area for future research (see Future Directions §9).

- **Conceptual Applicability:** Despite these formal hurdles, IA remains conceptually relevant. If we assume, as seems physically plausible, that complex states like consciousness do have some non-zero physical probability of arising through standard causal pathways (satisfying some implicit $MinimalHistory(S)$) and can persist for some T_{min}, then if those probabilities satisfy the Persistence Condition ($\sum P_n$ diverges) in an eternally evolving universe, IA would still imply their eventual actualization. The philosophical implications explored later (§7, §8) often hinge on this conceptual applicability to complex states, even pending full formal rigor. Future work (§9) must address the challenge of better defining complex states and their historical constraints within a probabilistic framework.

3 Metaphysical Implications of IA

IA offers a novel perspective on several metaphysical concepts:

- **Modality:** IA introduces a modality, T_p (Temporal-probabilistic necessity), such that T_pS holds if and only if, under the conditional assumptions of IA ($H_\infty \wedge \sum P_n = \infty$), the probability of S occurring is 1. This modality sits between strict logical necessity ($\Box$) and possibility ($\Diamond$), formalized as: $P(\exists t : Occurs(S,t)|H_\infty \wedge \sum P_n = \infty) = 1$.

- **Distinction from Logical Necessity and Probability-1:** It is crucial to distinguish T_pS from both logical necessity ($\Box S$) and simple probability-1 assignment ($Pr(S) = 1$). Philosophers note that probability 1 does not entail necessity (e.g., a dart hitting a real line has probability 1 of not hitting the exact point 0, yet this isn't logically necessary). IA acknowledges this; T_pS is explicitly weaker than $\Box S$. Furthermore, T_pS is stronger than merely assigning $Pr(S) = 1$ in some static probability space. The "necessity" of IA arises specifically from the temporal unfolding over an infinite sequence of trials combined with persistent possibility. It's not just that S has measure 1 in some space, but that the dynamics guarantee its eventual realization in time if the conditions hold. Thus, we maintain the distinction: $\Box S \neq (Pr(S) = 1) \neq T_pS$. While derived from probability-1 results (like Borel-Cantelli), T_p denotes a physically grounded guarantee of eventual occurrence, not merely an abstract measure-theoretic property.

- **Contrast with Standard Probabilistic Modalities:** While modal logic has incorporated probability (assigning numerical likelihoods to possible worlds or propositions, Fagin and Halpern, 1994), IA's T_p operator differs fundamentally. Standard probabilistic modalities often treat probability-1 events as 'almost necessary' within a static framework or interpret probabilities epistemically (degrees of certainty). IA, however, defines its necessity dynamically and extrinsically. The guarantee T_pS arises not merely from $P(S) = 1$ in some abstract sense, but specifically from the temporal process—the interplay of persistent positive probability $P_n > 0$ with an unbounded sequence of future trials (H_∞), as captured by limit theorems like Borel-Cantelli. It is the assumed infinite duration of the process that converts a persistent possibility into a probability-1 actuality. This grounding in physical time structure and objective chance distinguishes IA from epistemic or purely measure-theoretic accounts of probability-1 necessity.

- **Axiomatic Considerations & Distinctness:** Developing a formal logic for T_p requires care due to its conditionality and probabilistic nature. Its properties differ significantly from standard modal logics (like K, T, S4, S5), solidifying its claim as a distinct modality. The dramatic failure of Seriality (Axiom D) is seen immediately: consider S = "a fair coin lands heads" and $\neg S$ = "a fair coin lands tails". Given unbounded flips (H_∞) and $P(H) > 0, P(T) > 0$ such that $\sum P(H)$ diverges and $\sum P(T)$ diverges, both T_pH and T_pT hold (both heads and tails are guaranteed to occur eventually). This violates the principle $T_pS \rightarrow \neg T_p\neg S$ (If S is IA, then $\neg S$ is not IA), which is analogous to the D axiom ($\Box p \rightarrow \Diamond p$) fundamental to many standard modal systems. Other standard rules and axioms also fail or require careful interpretation:

- **RN (Rule of Necessitation):** If $\vdash S$, then $\vdash T_pS$ fails. Logical truths are not guaranteed to occur physically.
- **Axiom T:** $T_pS \to \Diamond S$ holds trivially, but $T_pS \to S$ fails (IA guarantees eventual, not present, occurrence).
- **Distribution over Conjunction:** $T_p(S_1 \wedge S_2) \to (T_pS_1 \wedge T_pS_2)$ likely holds, but the converse $(T_pS_1 \wedge T_pS_2) \to T_p(S_1 \wedge S_2)$ fails due to potential anti-correlation.
- **Iteration (Axioms 4 & 5):** $T_pS \to T_pT_pS$ and $\neg T_pS \to T_p\neg T_pS$ are unlikely to hold.

These violations show T_p's axioms form a *sui generis* system—neither alethic nor deontic. Its distinct logical structure is dictated by its unique temporal-probabilistic grounding.

Box 1: Proposed Minimal Axiom/Rule Set for T_p (Conjectural)
A starting point for formalizing T_p might include:

- (A1) All propositional tautologies.
- (A2) Distribution over $\wedge$: $T_p(\phi \wedge \psi) \to (T_p\phi \wedge T_p\psi)$
- (R1) Physical Law Necessitation: If ϕ is entailed by the physical laws L assumed in the model, then $\vdash T_p\phi$. (Captures respect for physical constraints; needs precise formulation).
- (R2) Conditional Modus Ponens: From $\vdash T_p(\phi \to \psi)$ and $\vdash T_p\phi$, infer $\vdash T_p\psi$. (Plausible, but requires semantic validation regarding persistence).

Note: This set is minimal and provisional; its soundness and completeness relative to a suitable history-based semantics remain open questions (see §9).

- **Formal Semantics & Non-Reducibility:** Further underscoring its distinctness, consider the path towards a formal semantics and completeness result. A minimal calculus for T_p might include rules like Modus Ponens (R2 above), perhaps a weakened distribution rule (A2), and potentially a restricted necessitation rule tied to physical laws (R1). The intended semantics would likely involve models representing physical systems evolving over an unbounded sequence of trials (perhaps based on probabilistic temporal logic or dynamic logic frameworks, whose state space consists of histories or paths), where T_pS holds iff the probability of S occurring (given the Persistence Condition $\sum P_n = \infty$) is 1. Soundness (provable implies true) seems achievable. Completeness (true implies provable) is more challenging, likely requiring strong assumptions about the underlying physics expressible within the logic. Crucially, attempting to translate T_p into standard alethic systems (like S4 or S5) seems destined to fail. Any translation function *tr* mapping T_p-formulas to standard modal formulas (e.g., $tr(T_pS) = \Box\Diamond_{phys}S$ or similar) would struggle. Because T_p fundamentally relies on the global temporal condition H_∞ and the asymptotic condition $\sum P_n = \infty$—properties not typically captured by standard Kripke semantics based on world-accessibility—a translation preserving validity seems unlikely. For instance, the failure of Seriality for T_p contrasts sharply with its validity in many standard systems (like KD), indicating a deep structural difference that simple translation cannot bridge without distorting one logic or the other. This suggests T_p occupies a unique logical space defined by its specific physical preconditions (see Future Directions §9).

- **Modal Realism:** Unlike Lewisian modal realism (Lewis, 1986), which posits concrete parallel worlds, IA suggests possibilities (with $P > 0$) are actualized sequentially within a single, temporally extended causal history (or within each branch of an Everettian multiverse, if applicable). It offers a form of sequential or temporal modal realism, constrained by the physics of this universe.

- **Grounding:** The inevitability described by IA is grounded not in abstract logical relations or essences, but in the concrete, unfolding structure of physical laws, probability, and assumed temporal extent. This aligns with causal structuralist or physicalist grounding views (cf. Schaffer, 2007; Stoljar, 2017).

- **Comparison with Grounding Frameworks & Clarification:** IA's grounding differs from Schaffer's priority monism (where the cosmos as a whole grounds its parts) in that IA's guarantee arises from the interaction of local probabilistic laws with a global temporal property (unboundedness). It's also distinct from simple microphysical grounding, as it relies on statistical behavior over vast timescales, not just instantaneous micro-configurations. The grounding base for T_pS is explicitly the conjunction of the physical system's characteristics and the cosmological context: $\langle Laws, H_\infty, P_n \rangle$, where Laws determine the dynamics and probabilities P_n, H_∞ asserts unbounded trials, and P_n satisfies the Persistence Condition $\sum P_n = \infty$. It is this specific structure $\langle Laws, H_\infty, P_n \rangle \models T_pS$ that constitutes the grounding relation. This structural grounding is distinct from the grounding of classical recurrence theorems in ergodic theory (like Poincaré's). Those theorems typically rely on measure preservation (e.g., Liouville's theorem) within a bounded phase space volume ensuring a system revisits neighborhoods of past microstates, whereas IA relies on persistent positive physical probability over an unbounded sequence of trials, guaranteeing the eventual occurrence of specific state types, potentially in dissipative or evolving systems.

- **Epistemic vs. Metaphysical Contingency:** The conditional nature of IA ("If H_∞ and Persistence hold...") might lead to the objection that IA collapses into standard contingency, as the conditions themselves are empirically uncertain. It is vital to distinguish:
  - **Epistemic Contingency:** Our current knowledge about whether H_∞ and the Persistence Condition for a given S actually obtain in our universe is incomplete. These are open scientific questions.
  - **Metaphysical Status (Given Conditions):** IA makes a claim about the modal status of S *assuming* the conditions hold. If they hold, IA asserts that S possesses a unique status—temporal-probabilistic necessity (T_pS)—which is stronger than mere contingency (where S might or might not occur even if possible) but weaker than logical/metaphysical necessity (where S could not possibly fail to occur). IA explores the consequences of these physical conditions for modality, rather than asserting the conditions themselves are necessary.

4 Brief Literature Context and Challenges

IA intersects with and faces challenges from several areas:

- **Classical Recurrence Theorems (e.g., Poincaré):** While related through the theme of recurrence, these theorems typically apply to closed, conservative systems returning to neighbourhoods of previous microstates within finite phase spaces. Crucially, classical recurrence theorems depend on measure preservation in bounded phase spaces, whereas IA requires only persistent positive chance over unbounded time. IA differs significantly: it applies to potentially open, evolving, or dissipative systems (like the universe) over unbounded time, guarantees the occurrence of specific state types S (potentially complex and defined with causal history, not just microstates) based on persistent physical $P(S) > 0$ satisfying $\sum P_n = \infty$, and relies on probabilistic limit theorems (like Borel-Cantelli) over infinite trials rather than solely on measure preservation in a fixed phase space. The grounding mechanisms are distinct (§3).

- **Cosmology:** Models inconsistent with infinite future time (e.g., Big Crunch, Big Rip if $w < -1$, certain cyclic models, terminal vacuum decay via Coleman-De Luccia instantons if the decay rate is too fast) would invalidate IA's conditional premise H_∞. Eternal inflation models face measure problem critiques (Vilenkin, 2013; Guth, 2007).

- **Anthropic Reasoning:** The choice of observer priors (Self-Sampling Assumption - SSA vs. Self-Indication Assumption - SIA) affects predictions, especially regarding Boltzmann Brains (Bostrom, 2002). We adopt SSA (§7.1) based on arguments favoring typicality.

- **Foundations of Probability:** Debates about objective chance (propensities) vs. subjective credence (Bayesianism) vs. frequentism vs. logical interpretations (Norton, 2021; Williamson, 2016) impact the interpretation of $P(S) > 0$. IA leans towards an objective, law-based interpretation of probability (propensity or hypothetical frequency).

- **Justification of Objective Probability:** IA concerns the objective occurrence of events in the physical world, independent of any observer's belief state; hence, subjective credence is inappropriate for defining $P(S)$. While actual frequencies are finite, IA relies on the limiting behavior implied by physical laws, fitting a hypothetical frequency or propensity interpretation. Norton's material theory (Norton, 2021), which denies universal formal rules of induction and grounds inference in local material facts, poses a challenge. However, IA relies on mathematical theorems (Borel-Cantelli) that apply broadly given certain formal conditions (independence/mixing, $\sum P_n = \infty$). While the value of $P(S)$ and the satisfaction of the conditions are material facts determined by physics, the implication (probability 1 occurrence given the conditions) seems to follow from the formal structure of probability itself, suggesting the material theory might be too restrictive here, or that the "material fact" includes the mathematical structure of probability applicable to the physical system.

5 Measure-Theoretic Foundations

Handling infinities requires careful measure theory (Kolmogorov, 1956):

5.1 Normalized Temporal Measure (μ_T)

To apply standard probability theorems (like Borel-Cantelli) over infinite time, we often discretize the future into a countable sequence of "trials" $n = 1, 2, \ldots$ (see Appendix C for construction). Assigning positive weights $w_n > 0$ (e.g., reflecting proper time duration, scale factor expansion, or observer density within trial n) defines a measure μ_W based on these weights. The total measure is $W = \sum_{k=1}^{\infty} w_k$.

Two cases arise:

1. **Finite Total Measure ($W < \infty$):** If the sum of weights converges (e.g., using observer density μ_O in standard ΛCDM), we must normalize to define a standard probability measure $\mu_T(n) = w_n/W$. Standard probability theory, including expectation values and convergence theorems, can then be applied directly to this normalized measure μ_T.

2. **Infinite Total Measure ($W = \infty$):** If the sum of weights diverges (e.g., using uniform weights or proper time in eternal expansion), normalization over the entire infinite sequence is impossible. In this scenario, establishing IA does not require normalization. Instead, we rely directly on limit theorems applicable to sequences of events. The key tool is the second Borel-Cantelli Lemma (§6.1), which guarantees occurrence with probability 1 provided the sum of probabilities $\sum P_n$ diverges, regardless of whether the underlying weighting measure μ_W is finite. Calculations of relative frequencies or expectations in this case require careful limiting procedures over finite intervals.

5.2 Spacetime Volume (μ_V) vs. Observer-Moment (μ_O) Measures

- μ_V**:** In eternally expanding universes, the 4-volume measure grows without bound. Using this naively often leads to paradoxes like the youngness paradox or Boltzmann Brain (BB) dominance, as late-time, low-density fluctuations occupy vastly more volume.

- μ_O**:** Measures weighted by the density of observers or complexity (e.g., based on entropy production or computational capacity) often yield finite integrals ($W < \infty$ case above), peaked during eras of structure formation (Page, 2008; Carroll, 2017). This provides a more physically motivated weighting for typicality arguments (like SSA).

IA's Approach: IA's core claim relies on $P(S) > 0$ over an unbounded sequence of trials, regardless of weighting. However, when addressing objections like BBs (§7.1) or calculating relative frequencies, using a physically motivated measure like μ_T weighted by μ_O (often falling into the $W < \infty$ case) becomes crucial.

5.3 Quantitative μ_O Dominance Over Boltzmann Brains (BBs)

Using standard ΛCDM parameters $(H_0, \Omega_m, \Omega_\Lambda)$, the integrated observer-moment measure $\int \mu_O dt$ is finite ($W < \infty$) and overwhelmingly dominated by the matter/structure-formation era ($t \approx$ few Gyr). Estimates comparing the measure for ordered observers arising through standard evolution versus spontaneously fluctuated, sustained BBs in the far future vacuum show the former dominates by many orders of magnitude (see Appendix E for quantitative estimate), justifying the use of SSA without predicting we are BBs (Carroll, 2017; Page, 2008; Aguirre and Tegmark, 2011). This resolves a major challenge for eternal cosmologies.

6 Formal Probabilistic Model of Actualization

6.1 Lemma for Independent Trials (Second Borel-Cantelli Lemma)

Let E_n be a sequence of independent events in a probability space. (The use of a countable sequence derived from potentially continuous time is justified in Appendix C). If $\sum_{n=1}^{\infty} P(E_n)$ diverges, then $P(E_n \text{ occurs infinitely often}) = 1$. Since IA requires only at least one occurrence, the condition $\sum P_n = \infty$ (where $P_n = P(E_n)$) is the crucial requirement. The probability of E_n never occurring is $P(\cap_{n=1}^{\infty} E_n^c) = \lim_{N\to\infty} \prod_{n=1}^{N}(1 - P(E_n))$. If $\sum P(E_n)$ diverges, this product limit goes to 0. Thus,

$P(E_n$ occurs at least once$) = 1$. (This lemma applies directly even if the underlying temporal measure μ_W is infinite, as discussed in §5.1). Note that while this establishes occurrence with probability 1, it does not mean occurrence is logically necessary; it remains physically possible (though probability 0) for the event never to occur, analogous to a dart randomly thrown at a line segment never hitting a specific pre-chosen point.

6.2 Dependent Trials & Mixing Conditions

IA can hold even with dependencies, provided the dependencies weaken sufficiently over time such that the system doesn't get permanently stuck avoiding state S. Standard results from probability theory (Appendix F) show that IA holds under various mixing conditions (e.g., φ-mixing or α-mixing with coefficients decaying sufficiently quickly, such as $\sum \phi(n) < \infty$), which ensure that the state of the system at one time becomes increasingly independent of its state far in the past. A concise summary is that if probabilities P_n sum to infinity and temporal correlations decay fast enough (e.g., $\sum \phi(n) < \infty$), IA holds (see App F).

6.3 Relation of Mixing Conditions to Physical Processes

Whether complex physical systems satisfy specific mixing conditions is often difficult to establish rigorously. Physical intuition suggests that processes like quantum decoherence and cosmic expansion (related to cosmic no-hair theorems) should lead to an effective loss of memory of initial microstates for coarse-grained observables over sufficiently long timescales, which is the hallmark of mixing behavior. For example, quantum fluctuations in a stable vacuum (like de Sitter space, §6.6) might be expected to exhibit strong mixing properties, potentially satisfying Doeblin-like conditions (Appendix F) for transitions between coarse-grained states. Some theoretical work supports mixing-like behavior for interacting fields in expanding spacetimes (Bovier and Eckhoff, 2019). However, rigorously demonstrating that the full dynamics of quantum fields in curved spacetime satisfy the specific mathematical requirements of Doeblin or φ-mixing remains a significant, unresolved technical challenge. In cases where strong conditions fail (e.g., non-uniform dynamics across different regions), weaker mixing might still suffice for IA if the persistence condition ($\sum P_n$ diverges) holds.

6.4 Scope regarding Chaos

It is crucial to distinguish mathematical idealizations from physical reality. In a purely deterministic chaotic model without noise, a specific state S corresponding to a measure-zero set might have $P(S) = 0$, rendering IA inapplicable to that specific state. However, IA applies to physical systems. Any real physical system exhibiting chaos is subject to inherent stochasticity (quantum fluctuations, thermal noise, imperfect precision). This physical noise effectively ensures that any open region $Region(S, \delta)$ (§2) corresponding to state S has a persistent physical probability $P(S) > 0$ of being entered, as long as it is dynamically accessible. Therefore, IA's scope includes physically realized chaotic dynamics, guaranteeing eventual actualization of states S (defined with non-zero tolerance δ) under its core conditions (H_∞ and persistence $\sum P_n$ diverges). The simulation in Appendix B explicitly includes noise to model this physical reality.

6.5 Absorbing States and the Persistence Condition (Formalized)

IA crucially relies on the Persistence Condition, now framed quantitatively. The probabilistic arguments (like Borel-Cantelli) require that the sum of probabilities $\sum P_n$ diverges. If the probabilities

$P_n = P(S$ occurs in trial $\tau_n)$ decrease sufficiently quickly such that the sum converges ($\sum P_n < \infty$), then the guarantee of occurrence (probability 1) is lost, even if P_n remains positive for all n.

Persistence Condition (Formal): Let $P_n = P(S$ occurs in trial $\tau_n)$. IA holds for state S only if $\sum_{n=1}^{\infty} P_n$ diverges.

- **Note on Sufficiency vs. Necessity:** This condition is sufficient for probability-1 occurrence via the second Borel-Cantelli lemma (for independent or suitably mixing trials). While simpler conditions like $P_n = p > 0$ obviously satisfy this, the divergence of the sum is the core requirement. Conversely, if $\sum P_n < \infty$, the first Borel-Cantelli lemma implies S occurs only finitely often with probability 1, thus failing IA.

- **Example:** If $P_n = p > 0$ (constant probability), $\sum P_n$ diverges. If $P_n = c/n$, $\sum P_n$ diverges (harmonic series). But if $P_n = c/n^{1+\varepsilon}$ for $\varepsilon > 0$, then $\sum P_n < \infty$ (p-series converges), and IA fails despite $P_n > 0$ for all n.

- **Failure Theorem (Converse of IA):** If $\sum_{n=1}^{\infty} P_n < \infty$, then the probability that S occurs only finitely many times is 1 (by the first Borel-Cantelli Lemma), and thus the probability of S occurring at least once is not guaranteed to be 1. This occurs if the probabilities P_n diminish too rapidly, for example, if the system approaches an absorbing state where S becomes increasingly unlikely, even if never strictly impossible. Crucially, applying IA to a specific state S requires demonstrating the absence of any physical mechanism (like decay into a truly stable vacuum or terminal heat death) that would cause $\sum P_n$ for that S to converge.

6.6 Sketch of Quantum Resurrection (Persistence via Fluctuations)

In spacetimes with a persistent positive vacuum energy (like de Sitter space, potentially our far future), quantum field theory suggests non-zero (though exponentially suppressed) rates for tunnelling between vacuum states or nucleating configurations via instantons (Coleman and De Luccia, 1980) or Gibbons-Hawking fluctuations. For any state S allowed by conservation laws, there might be a rate $\Gamma(S) \sim \exp(-Action(S)) > 0$ for its spontaneous formation (Carroll and Chen, 2004). If such fluctuations persist indefinitely (i.e., the de Sitter vacuum itself is stable or sufficiently long-lived) with a roughly constant rate $\Gamma(S)$ per unit time (or per appropriate trial definition), then the probability P_n in trial τ_n of duration Δt_n would be approximately $\Gamma(S)\Delta t_n$. If the trials cover infinite future time ($\sum \Delta t_n = \infty$), then $\sum P_n$ would typically diverge, satisfying the Persistence Condition (Formal) (§6.5) and guaranteeing IA for any physically possible state S.

- **Scale of $\Gamma(S)$:** While theoretically non-zero, the rate $\Gamma(S)$ for macroscopic states (like a galaxy, let alone a brain) formed purely by vacuum fluctuations is hyper-exponentially suppressed, involving factors like $e^{-10^{60}}$ or far smaller. This timescale vastly exceeds the current age of the universe, rendering such spontaneous formation practically impossible to observe. However, for the formal Persistence Condition, only $\Gamma(S) > 0$ is required, regardless of magnitude, if the vacuum persists eternally. The IA framework thus highlights the profound difference between theoretical possibility over infinite time and practical likelihood within finite horizons.

- **Status and Interpretation:** These calculations are typically performed using semiclassical approximations (Euclidean path integrals, instantons) in quantum field theory in curved spacetime. Their application to the entire universe state or complex macroscopic objects is highly speculative. Key challenges include: defining the state S precisely within QFT, handling gravitational effects (quantum gravity), ensuring the calculated rate isn't zero due to hidden symmetries or conservation laws, and interpreting probabilities derived from Euclidean actions in a Lorentzian universe. Furthermore, the timescales ($\sim \exp(1/\hbar)$) are often hyper-astronomical, making empirical verification impossible. Nonetheless, within the standard theoretical framework (QFT + GR), these non-zero probabilities are predicted, providing a potential, albeit theoretical, basis for the Persistence Condition for many states S in an eternal de Sitter-like future.

7 Engagement with Objections

7.1 Boltzmann Brains (BBs)

- **Problem:** In eternally fluctuating universes, naive application of IA combined with spacetime volume weighting (μ_V) seems to predict that randomly fluctuated observers (BBs) should vastly outnumber evolved observers, contradicting our experience (the "measure problem").
- **Resolution:**
  - **Sustained Occurrence Definition (§2):** Filters out instantaneous, non-functional BB flashes. Requires BBs to persist long enough for coherent thought (T_{min}) and possess the minimal causal history appropriate for an observer state.
  - **Physically Motivated Measure (μ_O, §5.3):** Adopting SSA and weighting by observer-moment density (μ_O) rather than raw volume (μ_V) shows that evolved observers dominate the measure within standard cosmological models (ΛCDM). The conditions under which we'd expect to be typical observers align with our observations (see Appendix E for quantitative discussion). IA itself doesn't predict BB dominance; the prediction arises from combining IA's premise with a problematic measure. Note that the core IA theorem—that S occurs with probability 1 if H_∞ and $\sum P_n = \infty$ hold—is independent of the choice of weighting measure (μ_W). The choice of measure (like μ_O) becomes crucial only when making typicality arguments or calculating relative frequencies, as needed to address the BB objection.

7.2 Modal Hybridity

- **Objection:** IA seems to inappropriately mix objective physical probability (chance) with the concept of necessity, which is typically seen as non-probabilistic (logical, metaphysical). Is it a category error?
- **Response:** IA's modality (T_p) is explicitly defined as temporal-probabilistic and conditional. It does not claim logical or metaphysical necessity in the traditional sense. Its necessity arises from the structure of time and probability laws within the assumed physical framework. The distinct grounding mechanism (physical structure) versus classical necessity (logic/essence) justifies its classification as a separate category (see §3). It represents a structural guarantee emerging from dynamics under specific physical assumptions.

7.3 Identity Duplication and Moral Aggregation

- **Problem:** IA seems to imply the inevitable recurrence of states physically identical to oneself, leading to infinite duplicates. This challenges notions of unique personal identity and creates paradoxes for ethical theories involving aggregation of utility/value over infinite populations.

- **Identity:** IA guarantees recurrence of a state *type*. Whether this constitutes recurrence of the *same person* depends on the theory of personal identity. IA forces engagement with these theories:
  - **Closest Continuer (Nozick, 1974):** If survival requires unique causal continuity, IA implies that infinitely many distinct individuals, qualitatively identical to you (or past versions of you), will exist in the future, but none *are* you. Survival remains linear and unique.
  - **Psychological Continuity (Parfit, 1984):** If survival is constituted by Relation R (overlapping chains of strong psychological connectedness), IA makes Parfit's branching scenarios physically plausible under its conditions. It's possible that multiple future individuals could bear Relation R to your current state (if duplicates form from fluctuations based on your information). Identity might be one-many, or the concept of identity might be less important than the holding of Relation R. IA suggests survival (as Relation R) could be widespread and recurrent.
  - **Biological/Physical Continuity (Olson, 1997):** Requires continuous physical realization of the human animal. Spontaneously generated duplicates via IA would be distinct organisms, regardless of psychological similarity. Survival requires uninterrupted physical history.
  - **Conclusion on Identity & Ethical Relevance:** IA does not dictate which theory is correct but highlights the consequences of each. It suggests that if its conditions hold, scenarios involving perfect duplicates are not merely hypothetical but physically guaranteed, making the choice of identity criterion practically relevant for understanding future existence. Crucially, this choice can impact ethical calculations (§8.2): whether future IA-generated instances count as "self" or "other" could influence the application of discount rates, the aggregation of utilities, and the weighting of risks associated with duplication versus lineage extinction.
- **Ethics (Infinite Utility):** Standard utilitarian aggregation fails with infinite identical lives. Resolutions include:
  - **Temporal Discounting:** Applying a discount factor (exponential $\gamma = \exp(-\lambda \Delta t)$ or hyperbolic $\gamma = 1/(1 + k\Delta t)$) to future utilities can yield finite total utility $U_{total} = \sum \gamma^n U_n < \infty$, restoring comparability (Bostrom, 2011).
  - **Agent-Relative Duties:** Focus on local, agent-centered duties (Kant, 1785) which apply coherently to each instance/duplicate.
  - **Average Utilitarianism (Risky):** Can lead to repugnant conclusions but avoids some paradoxes.
  - **Fanatical Decision Theories:** Some argue for embracing infinite ethics (e.g., Bostrom's Astronomical Waste).

  Discounting appears the most common approach for preserving standard decision theory.

8 Applications and Implications

8.1 Cosmological Forecasts and Tests

IA's conditional premise (H_∞: unbounded future trials) is linked to cosmological parameters, primarily the dark energy equation of state $w = P/\rho$.

- If $w = -1$ (cosmological constant Λ), the universe likely expands forever (de Sitter future), supporting H_∞.
- If $w < -1$ (phantom energy), a Big Rip singularity occurs in finite time, falsifying H_∞.
- If $w > -1$ (quintessence, decaying dark energy), future evolution depends on the potential; eternal expansion is possible but not guaranteed. Furthermore, even if proper time is infinite, if the integrated observer density $\int \mu_O dt$ converges (e.g., due to structure decay), the persistence condition $\sum P_n = \infty$ might fail for states S tied to observers, thus invalidating IA for such states.
- Cyclic models (e.g., Penrose's CCC, Steinhardt-Turok) offer alternative paths to H_∞.
- **Vacuum Decay:** The possibility of our current vacuum state being metastable (Coleman and De Luccia, 1980) provides another potential falsifier. Detection of an impending or ongoing vacuum decay event (e.g., through observation of expanding bubbles of true vacuum or a sudden, drastic change in measured dark energy density or fundamental constants) would imply a finite future for our current cosmic epoch, thereby falsifying the H_∞ condition within this epoch and negating IA for states requiring its persistence. Quantifying the required stability involves comparing the vacuum decay timescale T_{decay} with the timescale $T_{recurrence}$ over which relevant probabilities P_n sum to infinity; IA requires T_{decay} to be effectively infinite relative to the process ensuring $\sum P_n$ diverges.
- **Observational Tests:** Future surveys like Euclid, LSST (Vera C. Rubin Observatory), and DESI aim to constrain w to $\sigma(w) \approx 0.01 - 0.02$ and measure spatial curvature Ω_k. Confirmation of $w \approx -1$ and $\Omega_k \approx 0$ would strengthen the empirical case for IA's premise within ΛCDM. Detecting cyclic signatures (e.g., CMB anomalies) or evidence of vacuum instability would challenge it.

8.2 Ethics and Decision Theory in Infinite Settings

If IA holds, ethics must grapple with infinity. While the inevitability of recurrence might initially seem to undermine the significance of actions, standard decision theory faces severe problems (divergence, paradoxes) when dealing with infinite expected utilities (cf. Pascal's Mugging). IA motivates exploring frameworks that handle infinity coherently.

- **Discounting & Convergence:** One standard approach (§7.3) is temporal discounting. Applying a discount factor $\gamma < 1$ ensures the sum of future utilities $\sum \gamma^n U_n$ converges, allowing comparison between infinite prospects. Formally, even if $U_n = U > 0$ for all n, the geometric series $\sum_{n=0}^{\infty} \gamma^n U = U/(1-\gamma)$ is finite for $0 \le \gamma < 1$. This rescues expected utility maximization from paralysis by infinities. However, the choice of discount rate can be contentious, and alternatives like lexical superiority (prioritizing preventing the worst outcomes absolutely, cf. Greaves, 2017; Beckstead, 2013) or critical-level utilitarianism offer different ways to handle infinite stakes, though often with their own challenges (e.g., extreme sensitivity to small probabilities or violations of continuity axioms).
- **"Inevitable Utilitarianism" / Robustness Frameworks:** The realization that certain outcomes (both good and bad) are guaranteed if $P > 0$ persists might shift focus from maximizing probability of indefinite survival to other goals. A framework accepting IA, perhaps termed "Inevitable Utilitarianism" or "Robustness-Focused Consequentialism," could prioritize:

  - *Maximizing Discounted Utility:* As formally sketched above. This appears the most tractable approach for adapting standard consequentialism.
  - *Maximizing Intra-Trial Value:* Focusing on maximizing the quality, complexity, or duration of flourishing within each "cycle" of existence between inevitable (but perhaps extremely rare) catastrophic resets or fluctuations. The goal shifts from preventing the inevitable to making each instance of existence as valuable as possible.
  - *Prioritizing Resilience and Recovery:* Investing in measures that increase the speed and fidelity of recovery after catastrophic events, minimizing the duration of low-value states.
  - *Accelerating Positive States:* Focusing efforts on bringing about desirable states sooner, given that their eventual occurrence (if $P > 0$) is guaranteed but the timing is not.

  Such frameworks acknowledge the inevitability implied by IA while providing a coherent basis for ethical decision-making focused on achievable goals like near-term well-being, resilience, and the quality of existence within potentially recurrent cosmic epochs. While alternatives exist (e.g., focusing on average utility, facing its own paradoxes, or adopting lexical rankings), adapting discounted utility or focusing on intra-trial value/resilience seem the most promising avenues for developing a practical calculus under IA.

- **Value of Creation:** Does IA diminish the value of creating something new if it was "inevitable" anyway? Perhaps value lies in accelerating the inevitable or in the creative act itself.

8.3 Personal Identity and Subjective Experience

IA intensifies questions about the self:

- **Recurrence vs. Immortality:** Does guaranteed recurrence of your state type (§7.3) constitute a form of immortality? Depends on identity theory.
- **Subjective Probability:** In Everettian (MWI) interpretations, IA might apply within branches. Combined with quantum suicide thought experiments and subjective probability arguments (e.g., Wallace, 2012), it touches on how observers should anticipate their future experiences in branching realities.

9 Conclusion and Future Work

Inevitable Actualization (IA) offers a rigorously defined, conditional ontological modality distinct from classical necessity and contingency. Its unique grounding in the assumed structure of unbounded time and persistent, non-zero physical probabilities justifies its status as a third category. By integrating metaphysics, probability theory, measure theory, and cosmology, IA provides a framework for analyzing the guaranteed eventual realization of physically possible states under specific, potentially testable cosmological assumptions. While conditional, its implications for cosmology, ethics, and personal identity are profound.

Future Directions:

- **Empirical Cosmology:** Continued refinement of constraints on w, Ω_k, and dark energy evolution; searches for evidence of cyclic cosmologies or deviations from ΛCDM.

- **Theoretical Physics:** Rigorous QFT calculations of fluctuation/resurrection rates ($\Gamma(S)$) in realistic spacetimes; deeper analysis of the measure problem in eternal inflation and development of predictive measures (e.g., causal diamond measure, stationary measure); rigorous demonstration of mixing properties (or lack thereof) in relevant cosmological models.
- **Philosophy:** Formal development of modal logic axioms and history-based semantics (perhaps adapting probabilistic temporal or dynamic logic) for T_p, including completeness proofs for suitable fragments; comparative analysis of infinite ethical frameworks beyond simple discounting (e.g., critical-level, lexical ranking) in the context of IA; detailed analysis of IA's interaction with different theories of personal identity and persistence; investigation of IA's validity under weaker dependence conditions beyond standard mixing assumptions.
- **Computational Modeling:** Development and analysis of simulations (e.g., Markov Chain Monte Carlo, agent-based models) demonstrating state recurrence and mixing properties (φ-mixing) in complex systems relevant to IA (extending Appendix B).
- **Refining State Definition:** Developing clearer criteria for defining State S, including its minimal causal history and relevant T_{min}/δ (§2), especially for complex biological or conscious states; proving the existence of positive-measure history sets satisfying $MinimalHistory(S)$ in non-trivial physical or computational models (e.g., QFT on lattices, complex automata, self-replicating polymer models).

References

Anthony Aguirre and Max Tegmark. Born in an infinite universe: a cosmological interpretation of quantum mechanics. *Physical Review D*, 84(10):105002, 2011. doi: 10.1103/PhysRevD.84.105002.

Nick Beckstead. *On the overwhelming importance of shaping the far future.* PhD thesis, Rutgers University, 2013.

Nick Bostrom. *Anthropic Bias: Observation Selection Effects in Science and Philosophy.* Routledge, 2002.

Nick Bostrom. Infinite ethics. *Analysis and Metaphysics*, 10:9–59, 2011.

Anton Bovier and Michael Eckhoff. Spectral gap and mixing for the stochastic quantization of the massless free field on the torus. *Annales Henri Poincaré*, 20(1):1–31, 2019. doi: 10.1007/s00023-018-0735-y.

Richard C. Bradley. Basic properties of strong mixing conditions. A survey and some open questions. *Probability Surveys*, 2:107–144, 2005. doi: 10.1214/154957805100000104.

Sean M. Carroll. Why boltzmann brains are bad, 2017.

Sean M. Carroll and Jennifer Chen. Spontaneous inflation and the origin of the arrow of time, 2004.

Sidney Coleman and Frank De Luccia. Gravitational effects on and of vacuum decay. *Physical Review D*, 21(12):3305–3315, 1980. doi: 10.1103/PhysRevD.21.3305.

Ronald Fagin and Joseph Y. Halpern. Reasoning about knowledge and probability. *Journal of the ACM (JACM)*, 41(2):340–367, 1994. doi: 10.1145/174652.174658.

Hilary Greaves. Discounting for funders. Working Paper 1-2017, Global Priorities Institute, 2017.

Alan H. Guth. Eternal inflation and its implications. *Journal of Physics A: Mathematical and Theoretical*, 40(25):6811–6826, 2007. doi: 10.1088/1751-8113/40/25/S25.

Alan Hájek. What conditional probability could not be. *Synthese*, 137(3):273–323, 2003. doi: 10.1023/B:SYNT.0000004904.91132.80.

E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.

Immanuel Kant. *Groundwork of the Metaphysics of Morals*. 1785. Various translations available.

A. N. Kolmogorov. *Foundations of the Theory of Probability*. Chelsea Publishing Company, 1956. Translation of the German original, Grundbegriffe der Wahrscheinlichkeitsrechnung (1933).

David Lewis. *On the Plurality of Worlds*. Blackwell, 1986.

John D. Norton. *The Material Theory of Induction*. University of Calgary Press, 2021.

Robert Nozick. *Anarchy, State, and Utopia*. Basic Books, 1974. Relevant discussion on closest continuer theory.

Eric T. Olson. *The Human Animal: Personal Identity Without Psychology*. Oxford University Press, 1997.

Don N. Page. Typicality defended, 2008.

Derek Parfit. *Reasons and Persons*. Oxford University Press, 1984.

V. V. Petrov. *Limit Theorems of Probability Theory: Sequences of Independent Random Variables*. Oxford University Press, 1995. e.g., English ed., 1995, p. 71 for Borel-Cantelli extensions.

Alvin Plantinga. *The Nature of Necessity*. Oxford University Press, 1974.

Jonathan Schaffer. From nihilism to monism. *Australasian Journal of Philosophy*, 85(2):175–191, 2007. doi: 10.1080/00048400701302135.

Daniel Stoljar. Physicalism. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2017 edition, 2017. URL https://plato.stanford.edu/archives/win2017/entries/physicalism/.

Alexander Vilenkin. *Many Worlds in One: The Search for Other Universes*. Hill and Wang, 2013.

David Wallace. *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*. Oxford University Press, 2012.

Jon Williamson. *Lectures on Inductive Logic*. Oxford University Press, 2016.

A Metaphysical Background Details

This appendix elaborates on the metaphysical context assumed or engaged by IA.

- **Modality & Possible Worlds:** Briefly contrasts IA's conditional, single-universe (or single-branch) temporal modality T_p with standard possible worlds semantics (Lewis, 1986; Plantinga, 1974). Discusses how IA relates to concepts like Kripkean necessity (rigid designation across time vs. worlds) and Leibnizian compossibility (physical laws define the compossible states within IA's framework).
- **Grounding:** Expands on §3 by detailing how IA aligns with physicalist or structuralist grounding frameworks. Argues that the "inevitability" is grounded in the causal powers encoded in physical laws operating over an unbounded temporal structure, rather than abstract entities or essences. Contrasts with necessitarian views of laws versus Humean views.
- **Realism:** Discusses the type of realism IA entails – a realism about the future potential states dictated by physical law and probability, actualized sequentially in time, distinct from Lewis's realism about concrete parallel worlds.

B Computational φ-Mixing Simulation Sketch

This appendix provides a conceptual sketch for simulating state recurrence in a system exhibiting mixing properties, illustrating §6.3.

- **Concept:** Use a simple chaotic map (e.g., the logistic map with added noise) as a proxy for a complex system. Define a "state S" as the system residing within a specific sub-interval of its phase space. Simulate the map over many iterations and track how often the system enters state S.
- **Python Sketch (Conceptual Skeleton):**

Listing 1: Conceptual Python sketch for logistic map simulation.

```python
import numpy as np
import matplotlib.pyplot as plt # Assuming plots are desired

def logistic_map_noisy(x, r, noise_std_dev):
  """Applies the logistic map with added Gaussian noise."""
  # Ensure x stays within bounds [0, 1] after map application
  next_x_deterministic = r * x * (1 - x)
  # Add noise, ensuring result is clipped to [0, 1]
  noise = np.random.normal(0, noise_std_dev)
  next_x_noisy = np.clip(next_x_deterministic + noise, 0, 1)
  return next_x_noisy

# Parameters (Illustrative)
r_chaotic = 3.9 # Parameter for chaotic behavior
noise_level = 0.01 # Standard deviation of Gaussian noise
num_trials = 100000 # Number of iterations
initial_state = 0.5 # Starting point
state_S_interval = (0.8, 0.85) # Define State S region
min_duration_T_min = 3 # Minimum steps to stay in S

```

```python
# --- Simulation Loop & Analysis Omitted for Brevity ---
# (Full notebook available in supplementary material)
# The loop would iterate num_trials times, applying logistic_map_noisy,
# tracking the trajectory, and counting entries into state_S_interval
# satisfying the T_min duration condition.
# Final print statements would report frequencies.
# Plotting code would visualize the trajectory and state S region.

print(f"Conceptual simulation sketch for {num_trials} trials.")
print(f"State S defined as interval {state_S_interval} with T_min = {min_duration_T_min}.")
print("Analysis code (omitted) would calculate occurrence frequencies.")
print("Plotting code (omitted) would visualize results.")
print("(Full notebook available in supplementary material.)")
```

- **Discussion:** This simulation would illustrate how even in a chaotic system, a state S defined by an interval (non-zero measure) is revisited. The noise term explicitly models physical stochasticity, ensuring $P(S) > 0$ is maintained (preventing the system from getting stuck in periodic windows of measure zero indefinitely). Analyzing the frequency and waiting times between occurrences relates to the concepts in IA and mixing. The 'sustained occurrence' check implements Definition 2.

C Formal Temporal Measure Construction

This appendix details the construction of the normalized temporal measure μ_T used in §5.1 and clarifies the definition of a "trial".

- **Discretization of Time and Defining "Trials":** Assume the future temporal evolution can be partitioned into a countably infinite sequence of disjoint "trials" $\tau = \{\tau_1, \tau_2, \tau_3, ...\}$. Each trial τ_n corresponds to a time interval $[t_n, t_{n+1})$. The key requirements are that $\cup \tau_n$ covers the entire relevant future and that the partitioning is suitable for assigning probabilities $P(S \text{ in } \tau_n)$.

- **Scale of Trials:** The scale or duration of a trial $(t_{n+1} - t_n)$ is not fixed by the IA framework itself; it is flexible and context-dependent, chosen based on the nature of the state S and the physical process under consideration.
  - *Example 1:* If S is a specific particle interaction, τ_n might correspond to very short durations, perhaps related to fundamental timescales like the Planck time or interaction times.
  - *Example 2:* If S is the formation of a star of a certain type, τ_n might correspond to millions or billions of years, reflecting stellar evolution timescales.
  - *Example 3:* If S is a specific configuration within a chaotic system, τ_n might be chosen based on the system's characteristic Lyapunov time or mixing time.

  The crucial point is that a suitable countable partitioning covering the unbounded future must exist for the mathematical framework (esp. Borel-Cantelli lemmas) to apply. The specific choice of scale affects the value $P(S \text{ in } \tau_n)$ but not the validity of the limit theorems if the persistence condition holds.

- **Countability from Continuity:** Even if underlying physical time t is continuous (uncountable), we can always construct a relevant countable partition τ. For instance, we can choose intervals $\tau_n = [n\Delta t, (n+1)\Delta t)$ for some fixed Δt, or select a sequence of rational time points q_n dense in the future and define intervals around them. Since any physical state S requires a minimum duration $T_{min} > 0$ for its sustained occurrence (§2), its probability is associated with non-zero time intervals. If $P(S)$ persists over an unbounded continuous future, its integrated probability over a suitable countable partition covering that future will still satisfy the conditions ($\sum P_n$ diverges) for Borel-Cantelli. The countability requirement is a feature of the mathematical tool, not necessarily a claim about the fundamental structure of time itself. This reliance on countable additivity via partitioning also sidesteps potential issues raised by Hájek [2003] regarding the non-occurrence of probability-1 events in truly uncountable sample spaces, as IA operates on the limit behavior of a countable sequence of trials.

- **Assigning Weights:** Assign a positive weight $w_n > 0$ to each trial τ_n. This weight can represent:
  - Duration: $w_n = t_{n+1} - t_n$ (proper time).
  - Expansion: $w_n = a(t_{n+1}) - a(t_n)$ (change in scale factor).
  - Observer Density: $w_n = \int_{t_n}^{t_{n+1}} \rho_O(t)dt$, where $\rho_O(t)$ is an observer-moment density (§5.2).
  - Uniform: $w_n = 1$ for all n (simple counting measure).

- **Total Measure:** Calculate the total weight $W = \sum_{n=1}^{\infty} w_n$.

- **Normalization:**
  - **Case 1: Convergent Sum ($W < \infty$):** If the total weight is finite, we must normalize the weights w_n to define a standard probability measure $\mu_T(\tau_n) = w_n/W$. This allows the direct application of standard probability theory, including expectation values and convergence theorems relative to μ_T. This case is relevant when using measures like observer density μ_O in ΛCDM cosmology.
  - **Case 2: Divergent Sum ($W = \infty$):** If the total weight is infinite (e.g., using uniform weights or proper time duration in eternal expansion), normalization over the entire infinite sequence is impossible. In this scenario, establishing IA does not require normalization of the weights. Instead, IA relies directly on limit theorems applicable to sequences of events, primarily the second Borel-Cantelli lemma (§6.1). This lemma requires $\sum P(S \text{ in } \tau_n)$ diverge but operates directly on the probabilities $P(S \text{ in } \tau_n)$ without needing the underlying weighting measure μ_W (defined by the w_n) to be finite or normalized. Calculations of relative frequencies or expectations in this infinite measure case require more careful limiting procedures (e.g., considering ratios over increasingly large finite intervals).

- **Measure Space:** Formally, we construct a measure space $(\tau, \mathcal{P}(\tau), \mu_W)$ where τ is the set of trials, $\mathcal{P}(\tau)$ is the power set (sigma-algebra), and μ_W is the measure defined by $\mu_W(\{\tau_n\}) = w_n$. Probability statements are then made using μ_W either directly (if finite and normalized to μ_T) or via limit theorems (like Borel-Cantelli) if μ_W is infinite.

D Glossary (Trimmed)

- **IA (Inevitable Actualization):** The core concept: a conditional modality guaranteeing eventual occurrence (probability 1) if time is unbounded (H_∞) and the state has persistent, non-zero physical probability ($P_n > 0$) such that $\sum P_n$ diverges.
- **Minimal Causal History:** The necessary antecedent conditions for a state S to be a genuine instance of its type (§2).
- **Mixing (φ-mixing / α-mixing):** Properties describing the decay of statistical dependence over time, sufficient (if fast enough) for IA under dependence (App F).
- **Modality T_p:** The temporal-probabilistic necessity operator defined by IA, distinct from classical necessity ($\Box$) and possibility ($\Diamond$).
- **Persistence Condition:** The requirement $\sum_{n=1}^{\infty} P_n$ diverges for IA to hold (§6.5).
- **Sustained Occurrence:** Realization requiring persistence (T_{min}) and minimal causal history (§2).
- **Trial (τ_n):** A discrete segment of future time used for probabilistic analysis (App C).

E Quantitative Estimate for BB Dominance

Section 5.3 notes that observer-moment weighting (μ_O) resolves the Boltzmann Brain problem predicted by naive volume weighting (μ_V). The quantitative estimates supporting this rely on comparing the integrated measure for evolved observers versus BBs. While model-dependent, typical calculations (e.g., Carroll, 2017, Page, 2008, Aguirre and Tegmark, 2011) find that the ratio of the measure for ordered observers (concentrated in the structure formation era) to that for sustained BBs (dominant in the far-future vacuum) is extremely large.

Sketch of Calculation:

- Let $M_{ord} = \int_{early} \mu_O(t)dt$ be the total measure associated with ordinary observers, dominated by the structure-formation era. This is finite in standard models.
- Let Γ_{BB} be the nucleation rate of sustained BBs per unit 4-volume in the late-time vacuum. This is expected to be incredibly small, $\Gamma_{BB} \sim e^{-S_{BB}}$ where $S_{BB} \gg 1$.
- Let $\mu_{BB}(t)$ be the observer-moment density associated with BBs at late times. This might be approximated as $\mu_{BB}(t) \approx \Gamma_{BB} \times$ (factors related to volume/expansion).
- The total measure for BBs is $M_{BB} = \int_{late}^{\infty} \mu_{BB}(t)dt$.
- The crucial comparison is the ratio $R = M_{BB}/M_{ord}$.

Using plausible physical parameters and assumptions about the duration needed for a BB to count as an observer (T_{min}), calculations find R to be vastly less than 1. For instance, Carroll [2017] estimates the suppression factor $e^{-S_{BB}}$ can be smaller than $10^{-10^{68}}$. Even accounting for the potentially infinite duration of the vacuum phase, when combined with appropriate regulators or measure choices (like scale-factor cutoff or causal patch measures, which effectively make the total measure finite or give finite relative weights), the contribution from BBs remains negligible compared to ordinary observers. Figures for the ratio M_{ord}/M_{BB} often exceed 10^{20} or much more.

This vast difference means that even if BBs occur via IA in the infinite future, their contribution to the total observer-moment measure is negligible compared to observers like us, validating the use of SSA based on typicality within the μ_O measure.

F Technical Details on Dependent Trials (Mixing Conditions)

Section 6.2 notes that IA can hold even when trials are dependent, provided dependencies weaken sufficiently over time. This appendix briefly outlines two relevant conditions and presents a unified theorem.

- **Doeblin Condition (Uniform Ergodicity):** As mentioned in §6.3, this is a strong condition typically applied to Markov chains. If the sequence of states forms a Markov chain on a state space X, it satisfies the Doeblin condition if there exists a probability measure φ on X, an integer $m \geq 1$, and $\varepsilon > 0$ such that the probability of transitioning from any state x to any measurable set A in m steps is bounded below: $P^m(x, A) \geq \varepsilon\varphi(A)$. This ensures the process mixes rapidly and explores the entire state space (relative to φ) uniformly, preventing it from getting stuck. If $\varphi(Region(S)) > 0$, the system will inevitably enter the region corresponding to state S.

- **φ-Mixing:** This is a weaker condition applicable to more general stochastic processes. It requires that the statistical dependence between events separated by time n decays to zero as $n \to \infty$. Specifically, let $\mathcal{F}_k$ be the sigma-algebra generated by trials up to k, and $\mathcal{G}_{k+n}$ be the sigma-algebra generated by trials from $k+n$ onwards. The process is φ-mixing if the φ-mixing coefficient, $\varphi(n) = \sup_k \sup_{A\in\mathcal{F}_k, B\in\mathcal{G}_{k+n}, P(A)>0} |P(B|A) - P(B)|$, goes to zero as $n \to \infty$. If $\varphi(n)$ decays sufficiently quickly (e.g., $\sum \varphi(n) < \infty$) and $P(S)$ remains bounded below by $\varepsilon > 0$ in appropriate blocks of trials, recurrence results similar to the independent case (and thus IA) can often be established (Bradley, 2005).

Unified Theorem for IA under Mixing: The core result needed for IA extends the second Borel-Cantelli lemma to dependent sequences. A standard result (e.g., see Petrov, 1995, ch. IV, §19, or similar results for α-mixing or φ-mixing sequences) states:

Theorem (IA-Mix). Let $(\Omega, \mathcal{F}, P)$ carry a stochastic process $\{X_n\}_{n\geq 1}$. Let A be a measurable set representing state S, and let E_n be the event $X_n \in A$. Let $P_n = P(E_n)$. If the sequence $\{E_n\}$ is φ-mixing with $\sum_{n=1}^{\infty} \varphi(n) < \infty$ (or satisfies similar mixing conditions like α-mixing with sufficient decay rate) and if $\sum_{n=1}^{\infty} P_n$ diverges, then $P(E_n \text{ occurs infinitely often}) = 1$.

- **Proof Sketch:** The proofs typically involve bounding the variance of the sum of indicator functions $S_N = \sum_{n=1}^{N} I(E_n)$. Mixing conditions allow one to control the covariance terms $Cov(I(E_i), I(E_j))$ such that they decay rapidly as $|i-j|$ increases. Under conditions like $\sum \varphi(n) < \infty$, one can show that $Var(S_N)$ grows slower than $(E[S_N])^2 = (\sum P_n)^2$. Using Chebyshev-like inequalities (e.g., Chung's lemma), this implies that $S_N/E[S_N] \to 1$ in probability. Since $E[S_N] = \sum P_n \to \infty$, this means $S_N \to \infty$ in probability, which under these conditions implies $S_N \to \infty$ almost surely (i.e., infinitely many E_n occur with probability 1).

This theorem shows that as long as the probabilities P_n sum to infinity and the dependencies decay sufficiently quickly, the conclusion of the second Borel-Cantelli lemma holds, ensuring IA. The independent case corresponds to $\varphi(n) = 0$ for $n \geq 1$.

G Formal Example of State Filter (Cellular Automaton)

To illustrate the Minimal Causal History and Sustained Duration conditions (§2) more formally, consider a simple 1D stochastic cellular automaton (CA). This provides a setting where these conditions, particularly the history filter, can be made precise, unlike the more challenging case of complex macroscopic states (§2.1).

- **State Space:** Each site i can be in state 0 or 1. The state of the system at time t is the sequence $\sigma(t) = \{\sigma_i(t)\}$.
- **Dynamics:** The state at time $t+1$ depends probabilistically on the states at t in a local neighborhood (e.g., $\sigma_i(t+1)$ depends on $\sigma_{i-1}(t), \sigma_i(t), \sigma_{i+1}(t)$). Assume some non-trivial stochasticity (e.g., rules are probabilistic, or there's random noise flipping bits).
- **Target State S:** Let S be the specific pattern "11011" appearing centered at site j. $Region(S, \delta)$ corresponds to configurations where $\sigma_{j-2} = 1, \sigma_{j-1} = 1, \sigma_j = 0, \sigma_{j+1} = 1, \sigma_{j+2} = 1$. Here, tolerance $\delta = 0$ as the state is discrete.
- **Sustained Duration T_{min}:** We might require the pattern to persist for, say, $T_{min} = 2$ time steps. This means the pattern "11011" must exist at time t and $t+1$.
- **Minimal Causal History $MinimalHistory(S)$:** Suppose this specific pattern S is known only to arise reliably from a specific precursor pattern S' = "01110" centered at j at time $t-1$ according to the CA rules (other paths are possible but have negligible probability or violate physical constraints represented by the rules). Then, $MinimalHistory(S)$ would be the requirement that the configuration at time $t-1$ matched S'. This condition defines a measurable set in the space of path histories.
- **Sustained Occurrence:** An occurrence of S at trial i (time t_i) requires:
  1. The pattern "01110" was present at $t_i - 1$.
  2. The pattern "11011" is present at t_i and $t_i + 1$.

  This explicitly filters out instances of "11011" that arise from different histories or are too fleeting, illustrating how the definition applies measurable conditions on the system's trajectory. Demonstrating $P(S \text{ occurs}) > 0$ involves showing that the precursor S' has a non-zero probability and that the transition S' $\to$ S followed by S persisting for T_{min} steps also has non-zero probability under the stochastic dynamics.

H Worked Example (Markov Chain)

To illustrate IA more concretely, consider a simple 3-state Markov chain with states $\{A, B, S\}$. Let S be the target state. Assume transitions occur at discrete time steps (trials).

Transition Probabilities:

- From A: $P(A \to A) = 0.5, P(A \to B) = 0.5, P(A \to S) = 0$
- From B: $P(B \to A) = 0.1, P(B \to B) = 0.8, P(B \to S) = 0.1$
- From S: $P(S \to A) = 0.6, P(S \to B) = 0.4, P(S \to S) = 0$

Analysis:

1. This chain is irreducible (all states communicate) and aperiodic (can return to a state in varying numbers of steps).

2. Since it's a finite irreducible aperiodic Markov chain, it possesses a unique stationary distribution $\pi = (\pi_A, \pi_B, \pi_S)$ where $\pi_X > 0$ for all X. This means the long-run fraction of time spent in each state is positive.

3. Let $P_n = P(\text{State is S at trial } n)$. As $n \to \infty$, $P_n \to \pi_S > 0$.

4. **Persistence Condition:** Since P_n converges to a positive constant π_S, the sum $\sum_{n=1}^{\infty} P_n$ clearly diverges (as it behaves like $\sum \pi_S = \infty$).

5. **Mixing:** Finite irreducible aperiodic Markov chains are strongly mixing (e.g., geometrically ergodic, which implies strong mixing conditions like φ-mixing with exponentially decaying coefficients).

Conclusion: The conditions for IA are met (unbounded trials, persistence via $\sum P_n$ diverges, and mixing). Therefore, $T_p S$ holds: the system is guaranteed (with probability 1) to enter state S at least once (and in fact, infinitely often). This simple example shows how positive recurrence to state S, even if $P(S \to S) = 0$, combined with the ability to always eventually reach S from other states, ensures the divergence of $\sum P_n$ required for IA.