

Two Descriptions of Reality:

The Coherence Scale as a Commitment Threshold

A Structural Derivation of Quantum Theory, Thermodynamics, and Time from Irreversible Commitment

VERSF Theoretical Physics Program

For the General Reader

Why does time have a direction? Why can't you un-burn a piece of paper? Why does quantum mechanics produce definite outcomes at all, rather than leaving everything in permanent superposition? And why do these three puzzles — the arrow of time, irreversibility, and quantum measurement — feel like they ought to be connected, even though physics treats them as separate problems?

This paper argues that they are connected, and that the connection runs deep. All three are consequences of a single principle: **the universe can only hold a finite number of distinguishable facts in any given region, and once a fact is created, it cannot be erased.**

Think of it this way. A fact is something that has happened and left a trace — a record that persists and can be used to influence what happens next. Before a fact is created, a physical system exists in a state of *open possibility*: it may be heading toward many different outcomes, and none of them has been committed to yet. We call this a *proto-factual* state. The moment an irreversible record is produced — a particle hits a detector, a molecule bonds, a neuron fires — a fact comes into existence. That transition, from open possibility to committed record, is what this paper is about.

Physics has two descriptions of reality and no principled account of where one ends and the other begins. Below some threshold, quantum mechanics applies: systems exist in superpositions, outcomes are probabilistic, and nothing is definite until measured. Above some threshold, the classical world applies: objects have positions, reactions have products, decisions get made. Between the two sits the measurement problem — the embarrassing absence of any first-principles rule for where to draw the line. For nearly a century, physicists have drawn it wherever it was convenient: at the detector, at consciousness, at decoherence, at some vaguely specified macroscopic scale. Decoherence theory provides a scale-dependent account that is physically well-grounded and empirically successful — but it describes the suppression of quantum interference rather than the formation of definite facts, and it does not identify a threshold at which outcomes are constituted rather than merely becoming effectively classical. The cut, on the deepest question of what makes something a fact rather than a possibility, is still drawn by hand.

This paper argues that the cut does not need to be placed at all — because quantum and classical are not two kinds of reality. They are the same degrees of freedom in two different states of commitment. A physical system is in the quantum regime when the cost of maintaining all its open possibilities within a bounded region is still affordable. It crosses into the classical regime when that cost exceeds the region's capacity — at which point an irreversible record is created, a definite outcome exists, and a fact comes into being. The coherence scale is not a boundary between two worlds; it is a threshold within one.

What does this mean concretely? The degrees of freedom of a particle — its position, momentum, spin — do not change character as the system grows larger or more entangled with its environment. What changes is their *status*: below the threshold, those degrees of freedom are in open superposition, evolving reversibly, with no outcome yet committed. Above the threshold, the same degrees of freedom have produced an irreversible record; the outcome is fixed and cannot be undone. The quantum-to-classical transition is not a change in the physics; it is a change in the commitment status of the same physical variables. This is why quantum effects — tunnelling, interference, entanglement — can persist in warm biological systems like enzyme active sites: those systems are operating below the commitment threshold despite being macroscopic. And it is why a detector produces a definite click rather than a superposition: the detector's degrees of freedom are immediately driven above threshold by their coupling to a vast thermal environment. There is no cut. There is only the threshold, and which side of it you are on.

From just three requirements for any universe capable of producing stable facts, we derive:

- **Entropy** — not as a postulate, but as the only measure of how many possible facts are consistent with what has already happened
- **The second law of thermodynamics** — not as a statistical tendency, but as a structural necessity: facts accumulate and cannot be cancelled
- **Quantum mechanics** — the mathematical framework of superposition and probability emerges as the unique description of the pre-commitment layer
- **The Born rule** — the probability formula at the heart of quantum theory follows from consistency conditions once quantum structure is established
- **The arrow of time** — past and future are not symmetric because facts are irreversible
- **The appearance of classicality** — why the world looks definite at large scales is not a mystery requiring a separate postulate; it is a consequence of the commitment threshold being generically exceeded by macroscopic systems coupled to large environments

These are not restatements of known physics in new language. They are derivations from a more primitive foundation — one that asks not "what are the laws?" but "what must be true for any universe that has laws at all?"

The paper is written to be self-contained. Technical sections develop the mathematics with care. The discussion explicitly states what is proven, what is established by existing theorems, and what remains open.

What the coherence scale actually is.

Every bounded physical system at the capacity threshold — the smallest system capable of sustaining an independently attributable irreversible commitment — can sustain exactly one: one bit. This is not an assumption. A companion paper proves it formally (the No Multi-Primitive Occupancy Theorem, Taylor, in preparation). The proof proceeds by deriving the complete observable algebra of the threshold cell and showing that every logically possible configuration of two primitive events falls into exactly one of three cases — fully separable (Class D), fully identified (Class I), or partially resolved (Class P) — and each case leads to contradiction. Class D: a fully separable pair of events implies resolvable structure within the minimal cell, which contradicts the definition of the cell as the minimum commitment-supporting scale. Class I: fully identical events share one admissibility class and cannot constitute two independently attributable facts — claiming they are two facts without any observable distinguishing them is operationally empty. Class P: partial resolution requires an entropy increment strictly between $\ln 2$ and $2 \ln 2$, but the admissible entropy spectrum contains only integer multiples of $\ln 2$; no such intermediate value is physically realisable. All three cases are exhaustive and each closes. The result is one independently attributable primitive commitment per minimal causal support. *A reader cannot evaluate this claim without access to the companion paper; this foundational result is conditional on that unpublished proof, and the main paper's use of it should be understood accordingly. The result should therefore be read as conditional on the validity of that proof.*

The coherence scale is not a spatial size. It is a threshold in distinguishability capacity — the point at which a bounded system's degrees of freedom can no longer sustain open alternatives and must produce an irreversible record. While that threshold is not yet crossed, the system is quantum: it exists in superposition, multiple outcomes remain genuinely open, and no fact yet exists. The moment the threshold is crossed, one possibility becomes a permanent irreversible record. That moment is the coherence scale.

What determines when the commitment happens? The rate at which the system's interaction with its environment uses up that single bit. An atom in a vacuum interacts with almost nothing, so its bit stays uncommitted — it remains in superposition indefinitely. A molecule in air is jostled constantly; its bit is committed almost instantly. A proton in an enzyme active site is in an environment engineered to delay commitment for precisely long enough that the quantum event — tunnelling — can complete. The enzyme does not give the proton extra time in any conventional sense. It controls when the bit commits.

What we call coherence time in conventional physics is simply how long it takes for a system's minimal region to have its one bit committed by its environment. Time itself, in this framework, is what the sequence of committed bits creates — it does not exist prior to them. The coherence scale comes first. Apparent duration is a consequence.

The framework also predicts the characteristic scale at which stable structures can exist in our universe. A system with too few degrees of freedom is destroyed by quantum fluctuations before its commitment can stabilise. A system with too many is torn apart by cosmic expansion before it can maintain coherence. Between these failure modes lies a capacity sweet spot — which, in

physical systems with approximately uniform degrees-of-freedom density (typical physical matter), maps onto a characteristic length of roughly 30–100 micrometres: approximately the width of a human hair. This is not a fundamental spatial granularity but the capacity threshold expressed in familiar units. In strongly correlated quantum systems or engineered low-dimensional environments, the same capacity condition may correspond to very different physical sizes. The numerical value reflects typical degrees-of-freedom densities, not a fixed spatial boundary. That the same scale corresponds to typical cell sizes is noted as a consistency check, not claimed as biological evidence for the framework.

The framework derives both the threshold condition and what must mathematically live at it: the structure that describes a system whose single bit is still open is necessarily quantum mechanics — complex numbers, superposition, interference, and the Born rule all follow from the geometry of an uncommitted binary choice. And it predicts something no prior theory predicts: a measurable gap between when a physical transition happens and when the local bit commits, detectable with the fastest optical experiments currently available, that collapses in enzyme active sites and opens up in weakly coupled solvents. That gap is the paper's central experimental prediction.

The coherence scale is not a mystery. It is a threshold. And it has been derived.

Technical Abstract

Modern physics provides precise laws for how systems evolve, but lacks a first-principles account of when and how physical outcomes are constituted as irreversible records. Quantum mechanics describes superpositions of possibilities, thermodynamics describes irreversible records, and the arrow of time distinguishes past from future — yet these frameworks are introduced independently, without a unifying principle governing the transition from possibility to fact.

We propose that this transition is governed by a single structural condition: irreversible commitment in a finite-capacity system. A physical configuration remains proto-factual while the cost of maintaining its open alternatives is below a local capacity threshold. When this cost exceeds the threshold, an irreversible record is formed and a fact is constituted.

From three minimal operational constraints — finite distinguishability (A1), irreversible commitment (A2), and finite localisation capacity (A3) — we derive a unified structure in which entropy emerges as the unique measure of closure multiplicity, the second law follows from the non-cancellability of commitment, and the arrow of time arises from the irreversible ordering of commitment events. The reversible layer admits a complex Hilbert space representation via a Galois-invariance principle (proven in a companion paper: \mathbb{C} is the unique amplitude field compatible with interference, isotropy, and automorphism invariance), with the Born rule derived by four converging routes detailed in Section 9; the primary derivation (Double Square Rule) works prior to Hilbert space and in all dimensions, resolving the $\dim=2$ gap; Gleason's theorem and the Tick-Bit first-passage mechanism provide independent consistency checks.

The central empirical claim of the framework is that dynamical evolution and irreversible record formation are distinct physical processes. A system may undergo a transition without producing a fact until environmental embedding drives it across a commitment threshold $B_R \geq C_R$. This predicts the existence of a measurable commitment lag between dynamical change and outcome formation — a quantity absent from all prior frameworks and independently accessible in multiple experimental domains.

We demonstrate this structure across domains. In ultrafast proton transfer, solvent-dependent separation between transfer and product stabilisation timescales reveals a commitment lag tuneable by solvent embedding capacity. In enzyme active sites, this lag collapses due to environmental preorganisation, providing a structural account of anomalous tunnelling persistence and weak temperature dependence in enzymes such as ketosteroid isomerase. In quantum systems, the Quantum Zeno Effect arises as suppression of commitment through repeated interruption of environmental embedding — a distinct mechanism from standard decoherence that predicts partial suppression when measurement coupling is weak, independently of measurement frequency.

These systems are structurally identical: in each case, a proto-factual process evolves toward a threshold $B_R \rightarrow C_R$, and the presence, absence, or timing of commitment determines whether a physical outcome is constituted. The framework therefore makes a unified prediction: systems with identical dynamical evolution may differ in physical outcome depending solely on whether the commitment threshold is reached.

The result is a unified account in which thermodynamics, quantum theory, and time are not independent foundations, but consequences of a single principle:

Physical facts are irreversible commitments in a finite-capacity system.

Table of Contents

1. Introduction
2. Minimal Conditions for Physical Reality
 - 2.1 The Three Axioms
 - 2.2 Proto-Factual Configurations and Facts
 - 2.3 The Capacity Constraint
3. Commitment, Closure, and Entropy
 - 3.1 Commitment Events
 - 3.2 Closure Structure
 - 3.3 Entropy as Closure Multiplicity
 - 3.4 Non-Cancellability of Commitment
4. Thermodynamics from Irreversible Commitment
 - 4.1 The Second Law
 - 4.2 Equilibrium Distribution
 - 4.3 Temperature and the First Law
5. Necessity of a Commitment Boundary
6. Minimal Structure of the Fold Interface
 - 6.1 Binary Commitment Structure
 - 6.2 Orientation Structure (with Taylor Limit assumption)
 - 6.3 Why Complex Structure is the Unique Choice (via Galois theory, CHD)
7. Hilbert Space from Reversible Distinguishability
8. Measurement as Commitment
 - 8.1 The Three-Stage Structure of Fact Formation
 - 8.2 The Commitment Functional
 - 8.3 Worked Example: Tetrahedral vs Square-Planar Packing
 - 8.4 Measurement Defined
 - 8.5 The Quantum Zeno Effect as Suppression of Commitment
9. The Born Rule
10. Emergence of Time
11. Extension to Chemical and Biological Systems
 - 11.1 Chemical Reactions
 - 11.1.1 Worked Example: Threshold Crossing Along a Reaction Path
 - 11.1.2 Concrete System: Solvent-Dependent Proton Transfer
 - 11.1.3 Proton Transfer in Enzymes versus Solution: Collapse of the Commitment Lag

- 11.1.4 Concrete System: Ketosteroid Isomerase
- 11.2 Mechanism Transitions and OER
- 11.3 Biological Threshold Systems
- Gene regulatory switches
- Neuronal action potentials
- Protein folding as distinguishability collapse
- 11.4 Physical Threshold Systems
- 11.5 Environmental Modulation of the Commitment Threshold
- 11.6 Scope of the Extension

12. Unified Structure

13. Discussion

- Theorem (Threshold Necessity) and Corollaries
- 13.0 Core Empirical Claim of the Framework
- 13.1 What Is Established
- 13.2 What Remains Conditional
- 13.3 Relationship to Prior Work
- 13.4 Anticipated Objections

14. Conclusion

15. References

1. Introduction

Modern physics rests on several foundational frameworks: thermodynamics describes macroscopic irreversibility, quantum mechanics describes probabilistic dynamics, and spacetime provides the arena for physical events. Despite their empirical success, these frameworks are not derived from a single underlying principle. Entropy is introduced axiomatically, measurement is treated as a postulate, and the arrow of time is imposed rather than explained.

This paper takes a different starting point. Rather than assuming physical structures, we ask:

What minimal conditions must hold for a universe to support stable, reproducible facts?

The Coherence Scale as a Commitment Threshold — What This Paper Establishes

Existence. Any physical theory describing a universe with both reversible dynamics and irreversible records must contain a nontrivial condition separating the two regimes. This is not a modelling choice but a structural necessity, proven in Section 5 and formalised in the Threshold Necessity Theorem of the Discussion. The coherence scale — the scale at which reversible

quantum dynamics gives way to irreversible classical records — is not a hand-placed cut but a derived threshold.

Form. The threshold is characterised by the commitment condition $B_R = C_R$, where B_R is the accumulated distinguishability cost of maintaining open alternatives in a bounded system and C_R is the finite local capacity guaranteed by A3. Below threshold ($B_R < C_R$): the system is proto-factual, evolving reversibly, with no outcome yet committed. At threshold: an irreversible record is constituted, a fact enters the causal ledger, and the arrow of time advances by one commitment event. Quantum and classical are not two ontologies; they are the same degrees of freedom in two commitment regimes.

The threshold is a capacity condition, not a spatial boundary. The coherence scale is not the size of a region but the point at which a system's degrees of freedom can no longer sustain open alternatives. In systems with approximately uniform degrees-of-freedom density, this maps onto a characteristic length scale — but the underlying quantity is distinguishability capacity. The same threshold operates in a quantum dot (few degrees of freedom, small physical size), a molecule (more degrees of freedom, intermediate size), an enzyme active site (structured capacity), and a macroscopic detector (immediately exceeded) — not because these are the same size, but because they each reach the same capacity condition $B_R = C_R$ under their respective environmental couplings.

Structure. The minimal structure of the threshold interface is necessarily complex Hilbert space. Complex amplitudes are the unique viable representation — \mathbb{R} excluded by the absence of continuous phase, \mathbb{H} by the conflict between its automorphism group and configuration-space isotropy — as proven in companion paper CHD via the Galois Invariance Principle. The Born rule $P(i) = |\langle i|\psi\rangle|^2$ is the unique probability law consistent with this structure, derived by four converging routes with distinct roles in Section 9.

Empirical signature. The threshold predicts a quantity absent from all prior frameworks: the commitment lag $t_c - t_{\text{transfer}}$, a physically real, environment-dependent separation between the dynamical transition of a system and the formation of a committed fact. Decoherence theory separates dynamical timescales from classicality emergence, but the commitment lag identifies a conceptually distinct stage: environmental stabilisation to $B_R \geq C_R$, which occurs after dynamical transition and after decoherence, and which can be manipulated independently of both. In chemical systems this lag is tuneable by solvent identity and collapses in enzymatic active sites. In quantum systems, suppression of fact formation by repeated interruption of environmental embedding is the structural account of the Quantum Zeno effect. Systems with identical dynamical evolution and identical decoherence profiles may differ in physical outcome depending solely on whether environmental stabilisation reaches the commitment threshold — a distinction not present in decoherence theory.

Physical constants. When the VERSF structural conditions are applied at the gauge-invariant interface at which pre-factual commitments become physical facts, they constrain the electromagnetic coupling strength without continuous free parameters. A companion paper establishes a leading-order structural value matching the observed inverse fine structure constant to within 0.08%. This result belongs to that companion paper; it is noted here as a downstream

consequence of the present framework's structural conditions, not as a result established in this paper.

The coherence scale is not a mystery to be interpreted. It is a threshold to be derived — and it has been derived.

We show that once these conditions are precisely stated, the core structures of physics follow necessarily. The approach is operational: we make no assumptions about fields, particles, or spacetime geometry. We require only that a system be capable of producing records that are stable, distinguishable, and irreversible. From these requirements alone, the full machinery of thermodynamics and quantum mechanics is recoverable.

The central reframing. The conventional picture treats quantum and classical as two distinct regimes of nature, separated by a boundary whose location must be specified by hand — the Heisenberg cut, placed wherever it is convenient. This paper reframes the question entirely. Quantum and classical are not two kinds of physical reality. They are the same degrees of freedom operating in two different commitment regimes. Below the commitment threshold $B_R < C_R$, the dynamics of any physical system are reversible, possibilities remain open, and superposition is the natural state — this is the quantum regime, not because quantum mechanics is a special low-energy theory, but because no irreversible record has yet been produced. Above the threshold, irreversible records accumulate, outcomes become definite, and thermodynamics governs — this is the classical regime, not a different ontology but a different operational status of the same underlying degrees of freedom. The coherence scale is not a mystery to be interpreted; it is a threshold to be derived: the point at which the accumulated distinguishability cost of an evolving configuration exceeds the capacity of its local region to sustain open alternatives.

Three clarifications about scope are important at the outset.

First, not all results here are novel. The logarithmic entropy form, Hilbert space reconstruction, and the Born rule are known results. What is new is the *unified derivation pathway* — the demonstration that a single set of three primitive operational constraints suffices to ground all of them simultaneously, without independent postulates for each.

Second, we distinguish what is *proven* within this framework from what is *recovered* by showing that our axioms entail the hypotheses of existing theorems. Both are legitimate contributions; we mark them clearly throughout.

Third, the framework extends beyond fundamental physics to chemistry and biology via the Bit Conservation and Balance (BCB) formalism (Taylor, in preparation), which derives chemical structure rules — valence limits, molecular geometry, aromaticity, reaction admissibility — as consequences of the same capacity constraints. Section 11 summarises this connection.

The paper proceeds as follows. Section 2 states the three operational axioms, introduces the proto-factual / factual distinction, and specifies the capacity constraint. Section 3 derives entropy as the unique measure of closure multiplicity. Section 4 develops thermodynamics. Section 5 establishes the necessity of a commitment boundary. Section 6 derives its minimal structure. Section 7 connects the reversible layer to Hilbert space. Section 8 identifies measurement with commitment. Section 9 derives the Born rule. Section 10 addresses the emergence of time. Section 11 extends the framework to chemical and biological systems. Sections 12–14 present the unified structure, discussion, and conclusions.

2. Minimal Conditions for Physical Reality

2.1 The Three Axioms

We adopt three operational constraints that any universe capable of supporting stable, reproducible facts must satisfy. These are not assumptions about specific physical laws; they are requirements for the coherent existence of physical records.

A1. Finite Distinguishability. A physical system cannot support infinitely precise distinctions. Any observable configuration admits only finitely many resolvable states within a bounded context.

Motivation. An infinite-resolution record cannot be produced, stored, or transmitted in finite time with finite resources. Any workable notion of observable fact requires a finite state count at any given scale. This follows from purely operational constraints: distinguishing infinitely many configurations would require infinite resources — infinite memory to store the record, infinite time to communicate it, infinite precision to act on it. A fact that cannot be distinguished from other facts by any finite process is not a fact in any operationally meaningful sense. The finite distinguishability condition is therefore a precondition for any theory of facts, prior to and independent of quantum mechanics. That quantum mechanics happens to instantiate this constraint (through the action quantum and Bekenstein-type bounds) is a consistency check on the framework, not its motivation.

A2. Irreversible Commitment. A distinction becomes a physical fact only when it is irreversibly recorded. Reversible distinctions do not constitute physical reality.

Motivation. A distinction that can be completely undone leaves no trace and cannot ground a reproducible observation. Physical facts require records that persist. Reversible dynamics — whether classical or quantum — can evolve without producing any fact whatsoever.

A3. Finite Localisation Capacity. Any bounded physical system can sustain only a finite number of independently attributable irreversible commitments.

Motivation. This is the operational analogue of the Bekenstein entropy bound. It ensures that commitment structure is locally finite. Crucially, A3 guarantees the *existence* of a finite local capacity C_R for every bounded system, but does not determine its value. The magnitude of C_R is domain-specific and requires calibration from physical data.

Remark on capacity vs. geometry. C_R is a capacity in degrees of freedom, not a geometric measure of spatial size. The coherence scale at which $B_R = C_R$ is first reachable is not a fixed spatial length but a threshold in distinguishability capacity. In systems with approximately uniform degrees-of-freedom density ρ_{dof} , this capacity maps onto a characteristic length $\xi \sim (C_R/\rho_{\text{dof}})^{1/3}$. In systems with higher or lower effective degrees-of-freedom density — strongly correlated quantum systems, engineered low-dimensional structures, or highly constrained biological environments — the same capacity condition may correspond to very different physical sizes. The coherence scale is a threshold in distinguishability capacity, not a fixed spatial boundary.

Remark on the upper bound on distinguishability. A3 is consistent with known entropy bounds (Bekenstein, holographic limits), which place finite limits on the information content of bounded systems. Detailed discussion of these bounds, their hierarchy (holographic ceiling, Lloyd instantiation limit, Landauer access floor), and their relationship to A3's local capacity C_R is deferred to Section 13.3.

2.2 Proto-Factual Configurations and Facts

The axioms A1–A3 motivate a distinction that will be used throughout the paper.

Definition (Proto-factual configuration). A proto-factual configuration is a physically realised state that evolves under the governing dynamics of the system, may carry energy, structure, and correlations, but does not yet constitute an independently attributable outcome. Such configurations are dynamically real but outcome-indeterminate: their evolution is compatible with multiple future commitment histories, and no local observable certifies a committed result.

Examples. A quantum system in superposition before measurement; a reacting molecule traversing a transition state; an excited biological signalling molecule before irreversible amplification.

Definition (Fact). A fact is a configuration satisfying three criteria:

1. *Independent attribution:* the state can be identified without reference to its future evolution.
2. *Closure minimality:* the outcome is not decomposable into sub-events that independently constitute committed results.

3. *Environmental embedding*: the configuration is stabilised through interaction with its environment such that its record is persistent and not erasable by local operations.

Facts are not merely instantaneous states. They are stabilised outcomes embedded in the causal structure of the universe.

Definition (Primitive and composite facts). A *primitive fact* is an irreducible binary commitment, associated with a $\ln 2$ increment in closure structure. A *composite fact* is a higher-level structure — a molecular bond, a detector click, a biological decision — formed from many primitive commitments in sequence or in parallel.

2.3 The Capacity Constraint

For a bounded region R , define the committed bit count B_R as the operational count of irreversible distinguishability commitments maintained over a specified timescale τ . A commitment counts toward B_R only if it remains distinguishably instantiated over τ with failure probability $\leq \varepsilon_{rec}$.

The BCB Admissibility Inequality. Physical structures are admissible if and only if:

$$B_R \leq C_R$$

where C_R is the finite local capacity guaranteed by A3. This single inequality governs what can exist in any bounded region.

Timescale dependence. All admissibility statements are indexed by a stability timescale τ . A commitment counts toward B_R only on timescales over which it remains distinguishably sustained. This is not a weakness but a feature: it explains why "stable" means different things in different physical contexts.

Coherence scale and primitive commitment energy. The companion paper *Deriving the Commitment Barrier* (Taylor, in preparation) establishes that the primitive commitment energy — the minimum energy cost of one irreversible commitment event at the coherence scale ξ — takes the form:

$$E_c = r \cdot \hbar c / \xi$$

where r is a dimensionless closure-stiffness coefficient. Under the physical entropy identification (identifying VERSF closure entropy with thermodynamic entropy, conditional on this interpretation), $r = 1$ and the primitive consistency theorem yields:

$$E_c \approx 2.5 \text{ meV}$$

This energy scale falls in the millielectronvolt range, accessible to low-temperature quantum experiments including Josephson junction physics and precision atom interferometry. It constitutes a testable prediction of the framework: if the commitment threshold is a physically

real condition rather than a formal construct, deviations from standard quantum statistics should appear in engineered systems at or below this energy scale. *Note: this prediction is conditional on the physical entropy identification and on the coherence scale formula $\xi = (\hbar c/\rho)^{1/3}$. The cubic formula is established by the No Multi-Primitive Occupancy companion paper (Theorem 2.3a and Corollary 2.3b): the coherence cell has area ξ^2 (not volume ξ^3) because the primitive spatial substrate is two-dimensional; with $E_c = \hbar c/\xi$ and cell area ξ^2 , the primitive density is $\rho_c = E_c/\xi^2 = \hbar c/\xi^3$, giving $\xi = (\hbar c/\rho)^{1/3}$. An earlier companion paper (the Commitment Barrier paper) used a quartic formula $\xi = (\hbar c/\rho)^{1/4}$, derived without the 2D cell geometry result. These two formulas are inconsistent and yield different numerical predictions for E_c . The cubic formula is the operative one in this paper because it rests on the formally proved 2D dimensionality result (Theorem 2.3a); the quartic formula is superseded by that result. The Commitment Barrier paper requires revision to adopt the cubic formula consistently. Until that revision is complete, the energy prediction $E_c \approx 2.5$ meV should be treated as conditional on the cubic formula being correct.*

Domain	Typical τ	Example
Quantum measurement	10^{-15} – 10^{-9} s	Decoherence timescale
Molecular geometry	10^{-13} – 10^{-12} s	Vibrational timescale
Chemical transition state	10^{-13} – 10^{-6} s	Barrier crossing
Catalytic turnover	10^{-3} –1 s	Electrode kinetics
Biological decision	10^{-3} – 10^3 s	Signal amplification

3. Commitment, Closure, and Entropy

3.1 Commitment Events

A **commitment event** is an irreversible process that produces a new distinguishable record. By A2, such an event cannot be undone without producing compensating structure elsewhere. Section 3.4 shows that any apparent cancellation is itself a new commitment, making the process strictly monotone.

3.2 Closure Structure

Let Γ denote the set of **closure-equivalence classes** of distinguishable configurations. Two configurations c_1, c_2 are closure-equivalent if no admissible finite sequence of commitment events can produce an observable distinction between them. By A1 and A3, $|\Gamma|$ is finite within any bounded context.

We write $\Gamma_{\text{compatible}}(c)$ for the set of closure classes consistent with the recorded history up to configuration c .

3.3 Entropy as Closure Multiplicity

Define the entropy of a configuration c as:

$$S(c) = k_B \ln W, \quad W = |\Gamma_{\text{compatible}}(c)|$$

This form is not postulated. It is the unique form consistent with the structural constraints on any measure of closure multiplicity.

Proposition 1 (Uniqueness of the Logarithmic Form).

Any entropy functional $S: \Gamma \rightarrow \mathbb{R}$ satisfying:

1. *Monotonicity under refinement:* if Γ' refines Γ , then $S(\Gamma') \geq S(\Gamma)$,
2. *Additivity for independent subsystems:* $S(A \times B) = S(A) + S(B)$,
3. *Normalisation:* S vanishes for a single distinguishable state,

must take the form $S = k_B \ln W$.

Proof. Additivity requires:

$$f(W_1 W_2) = f(W_1) + f(W_2)$$

on \mathbb{N}^+ . This is the multiplicative Cauchy equation. Its unique measurable solution is $f(W) = C \ln W$ for constant C . Monotonicity excludes $C \leq 0$. Normalisation requires $f(1) = 0$, satisfied identically for all C . The constant $C = k_B > 0$ is fixed by thermodynamic calibration.

Remark. The novelty is not the logarithm — this is standard — but the grounding of the entropy measure in closure-record structure rather than observer uncertainty, phase-space volume, or Shannon information. The same functional form is derived from a more primitive foundation.

3.4 Non-Cancellability of Commitment

Lemma 1 (Commitment Cannot Be Cancelled).

Let e be a commitment event producing record r . Any process e' that erases r must itself be a commitment event, producing at least one new record r' . Therefore no finite sequence of commitment events can restore the pre- e closure structure.

Proof. Suppose e' completely restores the closure structure to its pre- e state without producing any new record. Then e' is reversible by definition — it leaves no trace of the erasure. But a reversible process is not a commitment event by A2. Contradiction. Hence e' must produce at

least one new record r' . The closure structure after e, e' contains r' and is therefore strictly larger than the pre- e structure. Iterating: no finite sequence of commitment events can reduce $|\Gamma_{\text{compatible}}|$.

Remark. This lemma is the structural heart of the paper. It is closely related to Landauer's principle — erasing information requires dissipation — but is derived here from A2 alone, without invoking thermodynamic quantities.

4. Thermodynamics from Irreversible Commitment

4.1 The Second Law

Proposition 2 (Arrow of Entropy).

For any sequence of commitment events, $\Delta S \geq 0$. Moreover, $\Delta S > 0$ for each commitment event — entropy strictly increases.

Proof. By Lemma 1, every commitment event e produces at least one new record r that cannot be cancelled without producing a further record r' . Therefore every commitment event strictly increases $|\Gamma_{\text{compatible}}|$: there is no commitment event that leaves the closure structure unchanged. The case $\Delta S = 0$ corresponds to the absence of any commitment event — no new fact is created and the system remains in proto-factual evolution. The logarithm is monotone, so each commitment event produces a strict positive increment in S . \square

Remark. The previous version of this proof stated "each commitment event either increases $|\Gamma_{\text{compatible}}|$ or leaves it unchanged." This was imprecise. Lemma 1 establishes that commitment events always strictly increase $|\Gamma_{\text{compatible}}|$ — the "leaves unchanged" case corresponds to no commitment occurring, not to a commitment event with zero effect. The distinction matters: the second law as stated here is strictly stronger than $\Delta S \geq 0$ for commitment-containing sequences; it is $\Delta S > 0$ for any interval containing at least one commitment event.

Interpretation. The second law is not statistical. It is structural: entropy increases because irreversible facts are created and cannot be undone. There is no need to invoke large numbers, ergodicity, or coarse-graining. The asymmetry holds for a single commitment event.

4.2 Equilibrium Distribution

Given a system subject to the constraint $\langle E \rangle = U$ (fixed mean energy), the maximum-entropy distribution over states i is found by the method of Lagrange multipliers. We extremise:

$$\mathcal{L} = -k_B \sum_i p_i \ln p_i - \lambda_0 (\sum_i p_i - 1) - \lambda_1 (\sum_i p_i E_i - U)$$

Setting $\partial \mathcal{L} / \partial p_i = 0$ gives:

$$-k_B (\ln p_i + 1) - \lambda_0 - \lambda_1 E_i = 0 \quad p_i \propto e^{(-\beta E_i)}$$

where $\beta = \lambda_1 / k_B$. Normalising:

$$p_i = e^{(-\beta E_i)} / Z, \quad Z = \sum_i e^{(-\beta E_i)}$$

This is the canonical (Boltzmann) distribution, derived as the unique maximum-entropy distribution consistent with a mean energy constraint. It is not postulated.

4.3 Temperature and the First Law

Temperature emerges as the conjugate variable to entropy under the constraint:

$$\beta = 1/(k_B T) = \partial(\ln Z) / \partial(-U) |_{\{X_a\}}$$

The first law follows from the structure of constrained entropy maximisation:

$$dU = T dS + \sum_a Y_a dX_a$$

where Y_a, X_a are generalised forces and displacements. Thermodynamics, on this account, is the macroscopic description of closure dynamics under admissible coarse-graining. Temperature, pressure, and chemical potential are not primitive quantities; they are conjugate variables that emerge from the structure of constrained entropy maximisation.

5. Necessity of a Commitment Boundary

The commitment framework is not introduced as an interpretive choice. It is forced by a minimal set of empirical constraints that any physical theory describing our universe must satisfy. This section establishes that any framework consistent with these constraints must include a threshold separating reversible dynamics from irreversible record formation — and that any framework lacking such a threshold must abandon at least one established empirical fact.

Three empirical constraints. We take as given:

1. *Finite distinguishability* (from A1): Physical systems admit only finite distinguishability at any operational scale. No measurement can distinguish infinitely many configurations. This is established by the finite entropy bounds of Bekenstein and Hawking, by the quantisation of physical observables, and by the operational precision limit of any physical apparatus.
2. *Existence of irreversible records* (empirical fact): Persistent, stable, reproducible measurement outcomes exist. Detectors click. Photographic plates develop. Chemical reactions produce stable products. Neurons fire. Whatever their underlying dynamics,

these outcomes constitute records that cannot be erased by local operations on the system alone.

3. *Finite localisation capacity* (from A3): Any bounded region of space supports only finite information content. Quantum information theory, thermodynamics, and holographic entropy bounds all confirm this. No bounded region can maintain an unbounded number of independent distinguishability commitments.

These are not theoretical assumptions. They are experimentally established properties of physical systems. A theory that denies any of them fails to describe the observed universe before any further structure is introduced.

Proposition 3 (Existence of the Fold Interface).

If (i) reversible distinctions exist, and (ii) physical facts must exist, then there is a boundary at which reversible distinctions become irreversible records. Any framework consistent with constraints 1–3 above must include this boundary.

Proof. Suppose no such boundary exists. One of two cases must hold:

Case A: All physical configurations remain reversible. Then no irreversible records can form. But irreversible records exist (constraint 2). Contradiction. A theory in which all dynamics are reversible cannot describe measurement outcomes, thermodynamic irreversibility, or any stable fact about the world. It fails to describe the observed universe.

Case B: All physical configurations are irreversibly committed. Then no reversible dynamics and no quantum superposition exist. But interference experiments demonstrate reversible superposition at the quantum scale (double-slit, Stern-Gerlach, Bell-inequality violations). A theory in which all dynamics are irreversible cannot describe quantum coherence. It contradicts the empirical success of quantum mechanics.

The only remaining possibility: physical systems admit *both* reversible and irreversible regimes, and there exists a boundary at which reversible configurations become irreversibly recorded.

That boundary is the **fold interface** \mathcal{F} . \square

What alternatives must give up. The argument does not merely show that a threshold exists in *this* framework. It shows that any framework lacking a threshold must pay one of two costs:

Framework type	What it abandons	Empirical failure
Purely reversible (unitary only)	Irreversible records	Cannot account for measurement outcomes or thermodynamic irreversibility
Purely irreversible (collapse everywhere)	Quantum superposition	Cannot account for interference, Bell violations, quantum coherence in engineered systems
Threshold-denying (facts as external)	Unity of physics	Facts are introduced by hand, not derived from physical dynamics

The last row is the standard Copenhagen position: the boundary between quantum and classical is imposed by fiat (the Heisenberg cut), located wherever convenient, and not derived from the theory. This is not a refutation of Copenhagen — it remains an internally consistent position — but it requires treating facts as external to the physical description rather than emerging from it.

A commitment threshold is therefore not an additional postulate of this framework. It is a structural necessity for any physical theory that aims to derive fact formation from physical dynamics without treating the quantum-classical boundary as an unexplained primitive.

Remark on the strength of the argument. The proof is by excluded middle, and a referee may object that such arguments are weak — that they establish existence without structural information. This is correct and is acknowledged: Proposition 3 proves *that* a boundary exists, not *where* it is, *what structure it has*, or *how it operates*. Section 6 derives the minimal structure of the fold interface; Section 8 identifies it with the commitment functional $B_R = C_R$; Sections 11 and 13 develop its experimental consequences. The argument here is foundational, not complete. Its value is that it forces any alternative framework to be explicit about which empirical constraint it abandons.

Definition. The boundary established by Proposition 3 is the **fold interface** \mathcal{F} — the structural boundary in the space of physical processes at which reversible possibility becomes irreversible record.

6. Minimal Structure of the Fold Interface

The fold interface must support two operations: (i) reversible orientation dynamics in the pre-commitment layer, and (ii) binary commitment at the interface itself. This section establishes why complex structure is the *unique* viable home for these operations, using the Galois-theoretic argument developed in a companion paper (Taylor, *Complex Hilbert Space from Distinguishability Principles*, in preparation; hereafter CHD). Section 6 previously described this as a motivational argument; with the companion result available, it is now a theorem citation.

6.1 Binary Commitment Structure

By A1, the minimal commitment distinguishes exactly two alternatives. We assign a binary label $\sigma \in \{0, 1\}$ to the commitment outcome.

6.2 Orientation Structure (with Continuity Assumption)

Pre-commitment dynamics support reversible transformations between states. The requirement that these transformations be *continuous* — the **Taylor Limit** of CHD — is an additional structural assumption not derived from A1–A3 alone. It requires that probability functionals be analytic and continuous, which excludes pathological discontinuous automorphisms with no physical content. We adopt it explicitly; without it, the Galois-theoretic field selection argument cannot proceed.

6.3 Why Complex Structure is the Unique Choice

The Galois Invariance Principle. Any amplitude field supporting the pre-commitment dynamics must satisfy: *observable predictions are invariant under all automorphisms of the amplitude field that fix the real numbers* (CHD, §4.2). This is the requirement that physics not depend on arbitrary internal conventions of the number system — exact analogue of coordinate invariance in spacetime physics.

Candidate fields. Frobenius's theorem (1878) classifies the finite-dimensional *associative* division algebras over \mathbb{R} : these are exactly \mathbb{R} , \mathbb{C} , and \mathbb{H} . Octonions are a division algebra but non-associative, so they fall outside Frobenius's classification. They are excluded separately: reversible composition of physical processes requires associativity — if applying transformation A then B then C must equal applying A then (B then C) — and octonions violate this, making them inadmissible as an amplitude field regardless of their other properties. Finite fields and p-adic numbers are excluded by the absence of a norm compatible with probability. This reduces the menu to exactly three associative candidates: \mathbb{R} , \mathbb{C} , and \mathbb{H} (CHD, §3.2).

Elimination of \mathbb{R} (CHD, §4.4). The continuous automorphism group $\text{Aut}(\mathbb{R}/\mathbb{R})$ is trivial — only the identity. A trivial automorphism group means no non-trivial phase transformations exist. But interference — the superposition of commitment alternatives with relative phase — requires continuous phase. Real amplitudes have only discrete phase (± 1), which cannot support the continuous interference patterns that the fold interface must sustain. \mathbb{R} is excluded.

Elimination of \mathbb{H} (CHD, §4.5, Appendix A.3). The continuous automorphism group $\text{Aut}(\mathbb{H}/\mathbb{R}) \cong \text{SO}(3)$ — the full group of three-dimensional rotations on the quaternionic imaginary subspace. This group is too large: it conflicts with the permutation symmetry of distinguishable configurations (isotropy). For three symmetrically-distinguishable configurations with amplitudes (i, j, k) , the joint requirement of isotropy under S_3 and Galois invariance under $\text{SO}(3)$ forces any probability functional to depend only on $|\psi_1|^2 + |\psi_2|^2 + |\psi_3|^2$, destroying phase sensitivity and eliminating interference. The proof is by explicit construction: the permutation $\sigma = (12)$ and the rotation $R: i \rightarrow j, j \rightarrow -i, k \rightarrow k$ together force $P(j, i, k) = P(j, -i, k)$, which, combined with the full $S_3 \times \text{SO}(3)$ constraint, reduces P to a function of total norm only. Quaternionic amplitudes cannot simultaneously satisfy Galois invariance, isotropy, and phase-sensitive interference. \mathbb{H} is excluded.

Selection of \mathbb{C} (CHD, Theorem 1). The continuous automorphism group $\text{Aut}(\mathbb{C}/\mathbb{R}) = \{\text{identity, conjugation}\}$ — a two-element abelian group. This is the Goldilocks structure: rich enough to support continuous phase (the full $U(1)$ circle), constrained enough that conjugation commutes with all configuration permutations (isotropy is preserved), and compatible with the Taylor Limit (both automorphisms are continuous). No clash between Galois invariance and isotropy arises.

Theorem (CHD, Theorem 1). *Let (A, d) be a finite metric space with $|A| \geq 3$. Under the requirements of interference, Galois invariance, isotropy for symmetric distinguishability, and the Taylor Limit, the complex numbers \mathbb{C} are the unique viable amplitude field.*

Theorem (CHD, Theorem 2). *Under the same conditions, with the isometry group G acting transitively on Λ , the unique G -invariant inner product on $\mathbb{C}^{|\Lambda|}$ is (up to scaling) $\langle \psi | \varphi \rangle = \sum_{\lambda} \psi(\lambda) \varphi(\lambda)$, making the state space a complex Hilbert space.**

The fold interface therefore carries the structure of a qubit — $\mathcal{H}_{\text{fold}} \cong \mathbb{C}^2$ — with $SU(2)$ as its reversible symmetry group and the Bloch sphere as its state space. This is no longer a motivated claim; it follows from Theorems 1 and 2 of CHD, given that A1–A3 plus the Taylor Limit entail all six assumptions of those theorems. The companion paper CHD constitutes the completion of the proof that Section 6 previously left open. *This result is conditional on the Taylor Limit (Section 6.2). The necessity of this condition is established at the level of impossibility — finite holonomy is inadmissible under TPB and bit conservation (PN Appendix C, Layer 1, complete) — while the upgrade to continuous $U(1)$ structure depends on four technical items currently completed in outline form in PN Appendix C, Layer 2. All Hilbert-space results in this paper should be read as conditional on the Taylor Limit unless otherwise stated.*

7. Hilbert Space from Reversible Distinguishability

Proposition 5 (Hilbert Representation).

The reversible pre-commitment layer admits a representation as a complex Hilbert space \mathcal{H} .

Derivation pathway. The result is established through three complementary routes, each illuminating a different aspect:

Route 1: Galois-theoretic derivation (CHD). The companion paper *Complex Hilbert Space from Distinguishability Principles* (Taylor, in preparation) derives Hilbert structure directly from a finite metric space of distinguishable configurations and the Galois Invariance Principle. The algebraic foundation for why \mathbb{R} , \mathbb{C} , and \mathbb{H} are the only candidates is provided by a complementary route through the amplitude algebra. The companion paper *Closing the Structural Gaps* (CSG) first derives (A5) locality of composition from admissibility via the operational congruence principle, and derives (S2) universal additive inverses under IAC, establishing abelian group structure on the amplitude algebra. The companion papers *Internal Admissible Closure from Pre-Factual Reversibility* (IAC) and *On the Structural Status of Algebraic Reversibility and Compositional Completeness* (PAR/CC) jointly prove that every non-null pre-factual contribution must admit a cancellation partner (IAC), from two principles: Pre-Factual Algebraic Reversibility (PAR) and Compositional Completeness (CC). PAR is the requirement that no internal transition within the pre-factual sector is one-way — because any such transition would constitute irreversibility before any fact is grounded, violating the Separation of Levels. The PAR/CC paper derives PAR through three independent routes: from admissibility plus the Pre-Factual Symmetry Condition (no ungrounded directional asymmetry), from admissibility plus information preservation (no pre-factual distinguishability can be destroyed), and from the Landauer-type argument that erasure requires a physically grounded cost-bearing outlet which does not exist before facts form. All three routes converge on the same conclusion. IAC, combined with Frobenius's theorem, restricts the amplitude field to $\{\mathbb{R}, \mathbb{C}, \mathbb{H}\}$,

which the Galois Invariance Principle then reduces to \mathbb{C} . As demonstrated in Section 6, A1 (finite distinguishability) provides the configuration space (Λ, d) with $|\Lambda| \geq 3$; A2 (irreversible commitment) establishes the isometry group G of reversible dynamics; A3 (finite localisation) bounds $|\Lambda|$; and the Taylor Limit provides analyticity. Under these conditions, CHD Theorem 1 proves \mathbb{C} is the unique amplitude field and CHD Theorem 2 proves the G -invariant inner product gives a complex Hilbert space. This route is complete given the assumptions stated in Section 6.2.

Route 2: Hardy (2001) correspondence. Hardy derives quantum theory from five operational axioms. A1–A3 directly entail three of them:

- *Simplicity* (degrees of freedom = K for a K -state system): from A1 and A3.
- *Subspace axiom* (subsystems are valid systems): from the closure structure of Section 3.2.
- *Probability axiom* (normalised probability distributions): from A1 and binary commitment. Two axioms — *continuity* and *composite systems* — require the Taylor Limit and non-overlapping A3 regions respectively. These are entailed given the stated assumptions.

Route 3: Chiribella–D'Ariano–Perinotti (2011) correspondence. CDP's critical principles — *local tomography* (from A3 applied to composite regions) and *purification* (the fold interface \mathcal{F} provides the purification boundary) — are entailed by A1–A3. The remaining CDP principles (atomicity, causality) require additional argument.

What this establishes. Route 1 (CHD) provides a complete derivation of \mathbb{C} and Hilbert structure from the conditions that A1–A3 plus the Taylor Limit entail, closing the gap that was flagged as an open problem in previous versions of this paper. Routes 2 and 3 provide independent consistency checks through different axiomatic frameworks. Together they establish Proposition 5 with considerably more than the "strongly motivated" status it previously held.

Under the combined result, the reversible layer carries the structure of a complex Hilbert space \mathcal{H} , with pure states as unit rays, mixed states as density operators $\rho \geq 0$, $\text{Tr } \rho = 1$, and reversible dynamics as unitary evolution.

Dynamics. Continuous reversible evolution in \mathcal{H} is governed by Stone's theorem: any strongly continuous one-parameter unitary group $U(s)$ takes the form $U(s) = e^{(-iHs/\hbar)}$ for self-adjoint generator H . This gives:

$$i\hbar \cdot d|\psi\rangle/ds = H|\psi\rangle$$

as the unique structure for reversible pre-commitment evolution. The Schrödinger equation is not postulated; it is the unique continuous reversible dynamics on the Hilbert space established by A1–A3.

8. Measurement as Commitment

8.1 The Three-Stage Structure of Fact Formation

Physical processes that produce facts exhibit a universal three-stage structure that is domain-independent:

Stage	Description	Commitment status
Proto-factual evolution	System evolves under reversible dynamics	$B_R < C_R$
Threshold approach	System enters outcome-indeterminate region near \mathcal{F}	$B_R \rightarrow C_R$
Stabilisation	Outcome becomes embedded in environment	$B_R \geq C_R$ (fact)

This structure is not specific to quantum measurement. It governs chemical reactions (Section 11.1), biological decisions (Section 11.3), and any physical process that produces a stable record.

8.2 The Commitment Functional

Definition. For a bounded region R supporting m directional distinguishability constraints, define the *commitment functional*:

$$B_R = \sum_{i=1}^m b_i + \lambda \sum_{i < j} \Phi(\theta_{ij})$$

where:

- b_i is the base distinguishability cost of constraint i (classified by type: bond, lone pair, directional interaction)
- θ_{ij} is the angular separation between constraints i and j
- $\Phi(\theta)$ is an interference penalty function satisfying: (i) symmetry $\Phi(\theta) = \Phi(-\theta)$, (ii) divergence as $\theta \rightarrow 0$ (co-directional constraints cannot be simultaneously distinguished), (iii) finite minimum at $\theta = \pi$
- λ controls interference strength

A natural representative satisfying these conditions is $\Phi(\theta) = 1/(1 - \cos\theta)$, which gives $\Phi(\theta) \sim 2/\theta^2$ as $\theta \rightarrow 0$ and $\Phi(\pi) = 1/2$. Key results are robust across alternative penalty functions with the same divergence structure (see BCB companion paper).

Commitment condition. A configuration achieves factual status when:

$$B_R \geq C_R$$

where C_R is the finite local capacity guaranteed by A3. The value of C_R scales with the effective distinguishability bandwidth of the region and requires domain-specific calibration. For atomic valence regions, $C_R \approx \alpha \cdot \Omega_R$ where Ω_R is an effective solid-angle bandwidth and α is a calibratable scaling constant.

Admissibility Class of Commitment Functionals. The specific additive-pairwise form of B_R should not be interpreted as uniquely derived from A1–A3. What *is* derived is the existence of a commitment functional satisfying the following structural constraints:

1. *Monotonicity*: B_R must increase under the addition of independent distinguishability constraints.
2. *Interference penalisation*: co-directional or overlapping constraints must incur super-additive cost.
3. *Permutation invariance*: B_R must be invariant under relabelling of constraints.
4. *Extensivity at low density*: for weakly interacting constraints, $B_R \sim \sum_i b_i$.
5. *Divergence under indistinguishability*: configurations that cannot be mutually resolved must be inadmissible.

These five conditions define an *admissibility class* of functionals, of which the pairwise form $\sum_i b_i + \lambda \sum_{i < j} \Phi(\theta_{ij})$ is the minimal truncation. Higher-order terms — triplet interactions and beyond — are not excluded by A1–A3, but are suppressed in the low-density limit and do not alter leading-order admissibility conditions. The specific interference penalty $\Phi(\theta) = 1/(1-\cos\theta)$ is a minimal representative of the class satisfying conditions (2) and (5); results are robust across alternative penalties with the same divergence structure (see BCB companion paper, Appendix B).

The present functional should therefore be understood as a lowest-order effective theory, analogous to a Landau expansion, rather than a uniquely determined expression. The structural constraints above are what A1–A3 actually fix; the specific form is the simplest functional consistent with them.

Remark on C_R . The existence of C_R is guaranteed by A3. Its magnitude is not derived from first principles within this framework; it is an empirical parameter analogous to the order-parameter coefficients in Landau–Ginzburg theory. Just as Landau–Ginzburg theory captures universal phase behaviour from symmetry and order-parameter constraints without specifying microscopic pairing mechanisms, the commitment functional captures universal fact-formation behaviour without specifying the microscopic degrees of freedom that enforce the capacity limit. Detailed calibration of C_R for chemical systems is developed in the BCB companion paper.

8.3 Worked Example: Why Four Equivalent Constraints Prefer Tetrahedral Packing

To show explicitly how the commitment functional operates, consider the simplest nontrivial case: a bounded region R supporting four equivalent directional constraints. Three constraints can always be arranged coplanarly without ambiguity; four require a genuine spatial packing solution and are therefore the minimal case in which geometry is non-trivial.

We take $m = 4$, equal base costs $b_i = b$, and the representative interference penalty $\Phi(\theta) = 1/(1 - \cos\theta)$. This Φ satisfies the structural conditions of Section 8.2: symmetric, divergent as $\theta \rightarrow 0$, finite and minimal at $\theta = \pi$. The example therefore rests not on a special choice but on the generic requirement that near-collinear distinguishability constraints become prohibitively expensive.

Two candidate packings are compared. The base term $\sum b_i = 4b$ is identical in both; the distinction comes entirely from the interference term.

Tetrahedral packing. In a regular tetrahedron every pair of directions subtends $\theta_{\text{tet}} = \cos^{-1}(-1/3) \approx 109.47^\circ$, giving:

$$\Phi(\theta_{\text{tet}}) = 1/(1 + 1/3) = (3/4)$$

With $4\{2\} = 6$ pairs:

$$I_{\text{tet}} = 6 \cdot (3/4) = 4.5 \Rightarrow B_{\text{R}^{\text{tet}}} = 4b + 4.5\lambda$$

Square-planar packing. Four adjacent pairs at 90° and two opposite pairs at 180° :

$$\Phi(\pi/2) = 1/(1 - \cos(\pi/2)) = 1/(1 - 0) = 1, \Phi(\pi) = 1/(1 - \cos(\pi)) = 1/(1 - (-1)) = 1/2$$

$$I_{\text{sq}} = 4(1) + 2((1/2)) = 5 \Rightarrow B_{\text{R}^{\text{sq}}} = 4b + 5\lambda$$

Result.

$$B_{\text{R}^{\text{sq}}} - B_{\text{R}^{\text{tet}}} = 0.5\lambda > 0 \quad (\lambda > 0 \text{ by construction})$$

Therefore $B_{\text{R}^{\text{tet}}} < B_{\text{R}^{\text{sq}}}$ for all positive λ , independent of the absolute scale of b .

For four equivalent directional constraints, tetrahedral packing has strictly lower commitment cost than square-planar packing.

Admissibility regimes. The capacity condition $B_{\text{R}} \leq C_{\text{R}}$ then produces three distinct regimes:

Condition	Consequence
$C_{\text{R}} > B_{\text{R}^{\text{sq}}}$	Both packings admissible; tetrahedral preferred as lower cost
$B_{\text{R}^{\text{tet}}} \leq C_{\text{R}} < B_{\text{R}^{\text{sq}}}$	Tetrahedral structurally selected; square-planar inadmissible
$C_{\text{R}} < B_{\text{R}^{\text{tet}}}$	Neither packing admissible; four constraints cannot be sustained

The middle regime is the critical one. Here geometry is not merely preferred — it is the only admissible structure. The threshold behaviour is not inserted by hand; it follows from the competition between distinguishability cost and finite local capacity.

Scope of the example. This should be read as a worked illustration of the framework's logic, not a full microscopic derivation of molecular geometry. The claim is narrower and more precise: once a finite local capacity and a positive penalty for angular crowding are assumed, the tetrahedral arrangement is selected over the square-planar one for four equivalent constraints. The framework generates an ordering between competing configurations from its formal structure rather than redescribing an observed geometry after the fact.

8.4 Measurement Defined

Definition. A **measurement** is a commitment event at the fold interface \mathcal{F} : a process by which a proto-factual state $|\psi\rangle \in \mathcal{H}$ undergoes irreversible commitment to an outcome $i \in \mathcal{P}$ (a pointer-basis record) when $B_R \geq C_R$:

$$\Phi: \mathcal{H} \rightarrow \mathcal{P}, |\psi\rangle \rightarrow i \text{ when } B_R \geq C_R$$

Interpretation. Measurement is not wavefunction collapse imposed from outside the formalism. It is irreversible commitment intrinsic to the dynamics at \mathcal{F} . Before threshold: the system is proto-factual, evolving unitarily. After threshold: an irreversible fact exists. The apparent discontinuity of quantum measurement is the structural signature of crossing \mathcal{F} .

Remark on macroscopic measurement. For a macroscopic apparatus, $B_R \geq C_R$ is reached rapidly and robustly: the interference sum $\sum_{i < j} \Phi(\theta_{ij})$ grows with the number of entangled degrees of freedom, and environmental coupling drives B_R above C_R on decoherence timescales. This accounts for the absence of observable superposition at macroscopic scales without invoking a separate collapse postulate.

Physical mechanism: the Tick-Bit race. The definition above is structural — it identifies when commitment occurs ($B_R \geq C_R$) but not the microscopic mechanism by which one outcome wins. The companion paper *Quantum Measurement as a Tick Race* (Taylor, in preparation; hereafter Tick-Bit) provides this mechanism.

After decoherence, each branch A of the post-measurement state $|\Psi\rangle = \sum_A \psi_A |A\rangle|E_A\rangle$ (with $\langle E_A|E_B\rangle \approx 0$) generates microscopic fluctuation events — **ticks** — at a rate λ_A . A tick is a micro-event that accumulates commitment load toward threshold: it is proto-factual. When a tick crosses the metastability threshold of the apparatus, it triggers an irreversible macroscopic signal — a **bit** — which is the committed fact. In the framework's language:

Tick-Bit term	Commitment framework term	Status
Tick	Proto-factual micro-event	$B_R < C_R$, approaching threshold
Bit	Committed fact	$B_R \geq C_R$, threshold crossed

Tick-Bit term	Commitment framework term	Status
Race	Competition among branches to reach $B_R = C_R$	First-passage process

The tick rate for each branch is forced by Fermi's golden rule and U(1) symmetry to scale as $\lambda_A = \kappa |\psi_A|^2$. The apparatus is branch-blind (its Hamiltonian cannot condition its response on which branch it is coupled to), so all branches share the same hazard shape $h_0(t)$, differing only in scale. The outcome is determined by the first branch to produce a threshold-crossing tick — physically forced by the thermodynamic structure of metastable amplifying detectors, where a single supercritical fluctuation triggers deterministic irreversible relaxation (Tick-Bit, §2.3).

This mechanism is not posited independently of the commitment framework — it is the commitment framework at the micro-level. The first threshold-crossing tick is precisely the event that drives B_R from below C_R to above it. The irreversibility of the bit is the irreversibility of commitment. The race among branches is the competition among environmental embedding trajectories to produce a committed record.

Remark on macroscopic measurement. For a macroscopic apparatus, $B_R \geq C_R$ is reached rapidly and robustly, making the race essentially instantaneous. The commitment lag $t_c - t_{\text{transfer}}$ is negligible, consistent with the general principle that preorganised environments collapse the commitment lag (Section 11.1.3).

Remark on the Zeno effect. Repeated interventions that maintain $B_R < C_R$ — by interrupting the environmental coupling before stabilisation — prevent commitment and hold the system in its proto-factual state. In Tick-Bit's language: repeated measurement resets the tick race, returning each branch to its initial state before any branch accumulates enough ticks to cross the metastability threshold. This is the structural account of the quantum Zeno effect: forced sub-threshold interruption of the tick race.

8.5 The Quantum Zeno Effect as Suppression of Commitment

The quantum Zeno effect provides a direct experimental test of the commitment framework within quantum mechanics itself, and its structural account within this framework unifies it with chemical, biological, and condensed-matter threshold phenomena.

Standard account. In the quantum Zeno effect, a system subjected to measurements at interval Δt fails to evolve away from its initial state as $\Delta t \rightarrow 0$. Experimentally demonstrated in trapped ions, superconducting qubits, and atomic systems, the effect is standardly interpreted as a consequence of the measurement postulate: each measurement collapses the wavefunction back toward the initial state, interrupting unitary evolution. The effect is real and well-confirmed; the interpretation is postulate-dependent.

Commitment account. In the present framework, the Zeno effect requires no separate measurement postulate. Evolution toward a new outcome corresponds to the accumulation of commitment load B_R through environmental coupling. In the absence of intervention, the system evolves from a proto-factual state toward commitment as $B_R \rightarrow C_R$. The commitment time t_c is the time at which this threshold is first crossed.

Frequent measurement at interval Δt interrupts the environmental embedding before B_R reaches C_R . Each intervention resets the environmental coupling — coupling to the measurement apparatus replaces and erases the accumulated embedding toward the prospective committed outcome, returning the system to a proto-factual configuration. The result:

$$B_R(t) < C_R \text{ for all } t$$

No commitment event occurs. The system remains dynamically active — unitary evolution is continuously attempted — but no irreversible record forms. The Zeno effect is not freezing by observation; it is suppression of commitment through repeated interruption of embedding.

Two regimes. The framework predicts a natural division:

Regime	Condition	Consequence
Zeno	$\Delta t \ll t_c$	Embedding interrupted before threshold; commitment suppressed; system remains proto-factual
Free evolution	$\Delta t \gg t_c$	Embedding completes between interventions; commitment occurs; fact formed
Anti-Zeno	$\Delta t \sim t_c$	Measurement timing can assist embedding by forcing partial environmental coupling at the optimal moment; commitment may be accelerated

The critical quantity governing Zeno behaviour is not measurement frequency alone but the ratio $\Delta t/t_c$ — the measurement interval relative to the intrinsic commitment time. This reframes the Zeno effect as a relationship between two timescales rather than a limiting consequence of observation frequency.

Prediction distinguishing accounts. In the standard account, Zeno suppression depends on the measurement basis and apparatus only through the overlap of the post-measurement state with the evolved state. In the commitment account, suppression depends additionally on whether the measurement interaction is sufficient to reset the environmental embedding — to drive B_R back below C_R — before the next evolution interval. These predictions diverge in systems where the measurement apparatus has weak environmental coupling: a measurement that produces a record in the quantum-information sense but does not fully reset environmental embedding should show partial Zeno suppression, with a Zeno plateau at $B_R < C_R$ rather than complete suppression. This is experimentally distinguishable from the standard account using systems where measurement coupling strength can be tuned independently of measurement frequency.

Unification. In standard quantum mechanics, the Zeno effect is a consequence of the measurement postulate. In the present framework, it follows from the same structural condition that governs chemical reactions, neuronal threshold crossing, and phase transitions: the system fails to reach the commitment threshold $B_R = C_R$. The measurement postulate is not needed as a separate axiom; the Zeno effect is the quantum instance of a universal pattern of threshold-controlled fact formation.

9. The Born Rule

Proposition 6 (Probability Law).

Given a system in state $|\psi\rangle \in \mathcal{H}$ and a measurement in basis $|i\rangle$, the probability of outcome i is:

$$P(i) = |\langle i|\psi\rangle|^2$$

Primary derivation: the Double Square Rule. The companion paper *The Double Square Rule* (Taylor, in preparation; hereafter DSR) derives this result from a foundation prior to Hilbert space, working directly from the distinguishability geometry established by A1–A3. The derivation proceeds in three steps, each of which is a theorem in DSR.

Step 1 (DSR §3–4): Bilinear structure is forced. Irreversible selection cannot act on individual paths without violating either gauge invariance (Axiom A8: physics cannot depend on arbitrary phase conventions) or interference (empirical: destructive interference is observed). The only selection mechanism compatible with gauge invariance, factorization for independent systems, and interference is one that acts on correlation structures between *pairs* of paths. The probability functional must therefore take the bilinear form:

$$P(A) = \sum_{\{P, P' \in R_A\}} W(P, P')$$

where $W(P, P')$ is a kernel encoding the geometric relationship between path pairs.

Step 2 (DSR §5.2–5.3): Factorization forces rank-one structure. The kernel W must be positive semidefinite (positivity of probability), symmetric (relabelling invariance), and factorising (A9: independent systems have independent probabilities). By the Mercer decomposition theorem, a positive semidefinite kernel factorising as $W(P \otimes Q, P' \otimes Q') = W(P, P') \cdot W(Q, Q')$ must be rank-one: $W(P, P') = \varphi(P) \cdot \varphi(P')^*$ for some function φ . This is proven in DSR Theorem 5.2 by showing that a higher-rank Mercer decomposition produces cross-terms under the factorisation condition that are mutually inconsistent. No tensor product structure is assumed — it emerges from the factorisation constraint on the kernel.

Step 3 (DSR §5.4–5.5): Geometric dependence selects $\varphi(P) = e^{i\theta(P)}$. The kernel must depend only on the geometric phase difference $\theta(P) - \theta(P')$ between paths (Axiom A6: physics determined by distinguishability geometry alone). Combined with the rank-one structure and multiplicativity under path composition, this forces $\varphi(P) = e^{in\theta(P)}$ for integer n . The

fundamental representation $n = 1$ is selected by the requirement that all physically distinguishable phases remain distinguishable in the kernel — higher n destroys information about paths whose phase difference is less than $2\pi/n$ (DSR Theorem 5.3; see also Physical Necessity companion, Lemma 6.1). The resulting probability is:

$$P(A) = \sum_{\{P, P' \in R_A\}} e^{i(\theta(P) - \theta(P'))} = |\sum_{\{P \in R_A\}} e^{i\theta(P)}|^2 = |\psi_A|^2$$

This derivation works in all dimensions — including $\dim = 2$ — without invoking Gleason's theorem, and does not require non-contextuality as an additional input. The bilinear structure is physically forced (not assumed), and the squaring emerges from the pairwise summation rather than from a measurement postulate.

Why higher-order alternatives fail. The Physical Necessity companion (Taylor, in preparation; hereafter PN) shows that trilinear and higher-order ($k > 2$) kernels are inadmissible: any k -linear kernel with $k > 2$ either reduces to a product of bilinear terms or involves higher harmonic orders $|n| > 1$, in which case the normalization $\sum_A P(A)$ is not preserved under reversible (unitary) evolution (PN Theorems 4.1–4.2). This is a direct consequence of the representation theory of $U(d)$: the n -th symmetric power representation (for $|n| > 1$) is inequivalent to the fundamental, so mixed-harmonic probability functionals are not unitary invariants. The Born rule's quadratic form is the unique normalisation-preserving choice (PN Theorem 8.3).

Consistency check via Gleason. Once Hilbert structure is established (Section 7), the Born rule also follows as a consistency check from Gleason's theorem (1957) for $\dim \geq 3$: the only frame-function probability measure on closed subspaces of a Hilbert space is the Born measure $P(i) = \text{Tr}(\Pi_i \rho)$. For $\dim = 2$ (the qubit), Busch (2003) extends this to POVMs without requiring non-contextuality. The Double Square Rule derivation and Gleason/Busch are independent routes to the same conclusion; their agreement provides mutual confirmation.

Fourth route: first-passage statistics (Tick-Bit). The companion paper *Quantum Measurement as a Tick Race* (Taylor, in preparation) provides a fourth derivation, independent of both the geometric and Hilbert-space routes, grounded directly in the microphysics of measurement. After decoherence, each branch A generates ticks at rate $\lambda_A = \kappa |\psi_A|^2$ (forced by Fermi's golden rule and $U(1)$ symmetry). Under the proportional hazards model forced by detector branch-blindness, and with first-tick selection forced by the thermodynamic structure of metastable amplifying detectors, the first-passage probability for branch A is:

$$P(A) = \lambda_A / \sum_B \lambda_B = |\psi_A|^2 / \sum_B |\psi_B|^2$$

This is an exact theorem of first-passage statistics (Tick-Bit, Theorem 2.1), not an axiom. The key constraints (A1')–(A3') — proportional hazards, amplitude-squared rates, first-tick selection — are each independently forced by known physics rather than posited. The tick rate scaling $\lambda_A \propto |\psi_A|^2$ does not assume the Born rule; it is a dynamical statement about transition rates derived from perturbation theory, conceptually prior to any probabilistic interpretation of $|\psi|^2$.

The four routes together constitute a robust convergence, each with a distinct role:

- **(1) DSR** — primary derivation from distinguishability geometry: derives $P = |\psi|^2$ prior to Hilbert space from pairwise kernel uniqueness forced by gauge invariance, factorisation, and positivity
- **(2) Gleason/Busch** — Hilbert-space consistency check: confirms the Born rule is the unique frame-function probability measure on closed subspaces of the established complex Hilbert space
- **(3) PN** — physical admissibility justification: justifies each structural assumption of DSR by admissibility elimination, establishing that no alternative probability assignment survives
- **(4) Tick-Bit** — dynamical realisation: shows the Born probabilities emerge from first-passage statistics of the tick-race mechanism, grounding the rule in the microphysics of measurement

Each addresses a different aspect of the result — derivation, consistency, admissibility, and dynamical mechanism — and their agreement provides cross-validation rather than strict independence. $P(i) = |\langle i|\psi\rangle|^2$ is not an arbitrary postulate but the unique probability law consistent with the full structure of quantum dynamics, distinguishability geometry, and irreversible measurement.

Composition. The tensor product structure for composite systems follows from the rank-one structure of the kernel (DSR §5.3): for independent systems, the factorisation condition $W(P\otimes Q, P'\otimes Q') = W(P, P')\cdot W(Q, Q')$ is precisely the condition that commitment events in non-overlapping bounded regions (A3) compose independently. This resolves the gap identified in earlier versions of this section, where tensor products were invoked without derivation.

Remark on the nature of quantum randomness. The Tick-Bit mechanism suggests that quantum probability is epistemic rather than ontic: the outcome of a measurement is determined in principle by the microstate configuration of the environment (which branch's tick configuration produces the earliest threshold-crossing time), but this configuration is epistemically inaccessible due to the exponential complexity of environmental degrees of freedom after decoherence and the chaotic sensitivity of the subsequent dynamics. This parallels classical statistical mechanics, where Boltzmann's distribution emerges from deterministic Newtonian dynamics plus epistemic uncertainty about microstates. The commitment framework takes no mandatory position on this metaphysical question — the structural results ($B_R \geq C_R$ as the condition for fact formation) hold independently of whether the underlying dynamics are ultimately deterministic or stochastic — but the Tick-Bit mechanism provides a concrete physical picture in which randomness is a consequence of necessary ignorance rather than a brute fact about nature.

Interpretation. The Born rule is not a postulate. It is the unique probability assignment consistent with: (i) the bilinear structure forced by gauge invariance and interference; (ii) the rank-one structure forced by factorisation and positivity; and (iii) the geometric phase structure established in Section 6. The appearance of $|\cdot|^2$ is the mathematical signature of pairwise selection — probability is counting pairs of paths rather than individual paths, because irreversible selection respects the relational structure of distinguishability geometry.

10. Emergence of Time

Definition. The **physical time-order** of a system is the partial order on commitment events induced by causal closure: event $e_1 < e_2$ if the record produced by e_1 is accessible to the commitment context of e_2 .

Proposition 7 (Irreversibility of Time-Order).

The time-order $<$ is strict: if $e_1 < e_2$, then $e_2 \not< e_1$.

Proof. Suppose $e_2 < e_1$. Then the record of e_2 is accessible to e_1 , and e_1 could in principle cancel e_2 . But Lemma 1 shows that cancellation of e_2 requires at least one new commitment e' , producing a closure structure strictly larger than the pre- e_2 state. The resulting configuration is therefore distinguishable from the pre- e_2 state, and the supposed restoration of temporal order fails. Hence $<$ is strict.

Interpretation. Time is not a background parameter over which physical events are distributed. It is the irreversible ordering of commitment events. The felt asymmetry of past and future — the sense that the past is fixed and the future open — is the subjective expression of the structural asymmetry proven in Proposition 2: entropy increases because new facts accumulate and cannot be erased.

Remark on temporal experience. The distinction between proto-factual and factual states invites a connection to the phenomenology of time — the sense that the present is the boundary between fixed past and open future. This connection is suggestive: what physics calls the commitment boundary \mathcal{F} , and what phenomenology calls "now," may be structurally related. However, the paper does not develop this connection beyond the structural level. Moving from the ordering of commitment events to the experienced asymmetry of time involves philosophical commitments — about the relationship between physical structure and conscious experience — that go well beyond what the formal framework establishes. The relevant literature (Price 1996 on the asymmetry of time; Barbour 1999 on the timelessness of the block universe; Callender 2017 on the philosophy of time) is engaged elsewhere. The structural claim of this section is limited to what is proven: that the irreversible ordering of commitment events is asymmetric, and that this asymmetry is what physics calls the arrow of time.

11. Extension to Chemical and Biological Systems

Scope note. Everything in this section is Tier 2: it requires the BCB functional form for B_R (the additive-pairwise commitment functional, Taylor, in preparation) plus empirical calibration of C_R from chemical data. None of it follows from A1–A3 alone. A1–A3 establish the existence of a commitment functional satisfying structural constraints; the BCB form is the minimal representative of that admissibility class, not a derived consequence. The connection between the foundational physics of Sections 2–10 and the chemistry of this section is that BCB is a concrete

instantiation of the same threshold condition $B_R = C_R$ applied to molecular degrees of freedom — the LG-analogy relationship described in Section 13.1.

The commitment framework is not restricted to quantum measurement. Fact formation — the transition from proto-factual evolution to irreversible record — is a universal structural process that appears wherever physical dynamics must produce stable, causally effective outcomes.

11.1 Chemical Reactions

A chemical reaction decomposes naturally into the three-stage structure of Section 8.1:

Stage	Chemical description	Commitment status
Encounter	Molecular collision and deformation	$B_R < C_R$
Transition state	Saddle-point traversal	$B_R \rightarrow C_R$
Product formation	Relaxed, stabilised molecule	$B_R \geq C_R$ (composite fact)

The **transition state** — a first-order saddle point on the potential energy surface, with lifetimes of order 10^{-13} s — is a proto-factual configuration: dynamically real, carrying energy and structure, but outcome-indeterminate (trajectories may proceed to products or return to reactants) and lacking the environmental embedding required for independent attribution.

The **transmission coefficient** κ of transition-state theory — which corrects the classical rate for trajectories that recross the dividing surface — acquires a direct interpretation: it is the fraction of proto-factual trajectories that achieve $B_R \geq C_R$ stabilisation and become composite facts.

Femtochemistry and the direct observation of proto-factual configurations. The development of femtosecond spectroscopy — laser pulses shorter than 10^{-13} s — made it possible to probe molecular systems on timescales shorter than a single bond vibration, directly accessing transition-state configurations. In Zewail's classic experiments on alkali halide dissociation and cycloaddition reactions, wavepackets were tracked as they evolved through the energy barrier in real time. In this framework, these experiments are observations of proto-factual configurations: the transition-state geometries accessed by femtosecond pump-probe spectroscopy are sub-threshold structures whose environmental embedding is incomplete. What femtochemistry observes is not the fact but the proto-factual process from which the fact emerges. The stabilised products — detected after vibrational relaxation and solvent reorganisation — are the committed records. The experimental impossibility of observing the transition state as a persistent object is not a technical limitation; it is a structural consequence of the transition state failing the attribution and coherence conditions for facthood.

The Arrhenius equation under the commitment framework. Each term of the Arrhenius rate equation $k = A e^{(-E_a/k_B T)}$ acquires a precise interpretation. The activation energy E_a is the energy cost of accessing the proto-factual zone — the minimum energy required to drive a molecular encounter to the dividing surface. It is not the cost of forming a fact; it is the cost of reaching the configuration from which fact formation can, with some probability, proceed. The pre-exponential factor A encodes the frequency and geometric probability of encounters that

enter the proto-factual zone productively — reactions with strict orientation requirements (concerted multi-bond rearrangements) have small A because the proto-factual zone is geometrically narrow; reactions with loose transition states have large A because it is wide. The rate constant k itself is the probability per unit time that a proto-factual configuration achieves the full environmental stabilisation required for commitment — it is the fact-formation rate, not merely the barrier-crossing rate.

Chemical structure as admissibility. The BCB framework (Taylor, in preparation) shows that chemical structure rules emerge as consequences of the capacity constraint $B_R \leq C_R$:

- *Molecular geometry*: The minimum-interference packing of directional constraints under $B_R \leq C_R$ reproduces VSEPR geometries. For four equivalent σ -constraints, tetrahedral packing gives total interference $I_{\text{tet}} = 4.5$ versus $I_{\text{sq}} = 5.0$ for square-planar packing, making tetrahedral geometry strictly less capacity-costly.
- *Valence limits*: Capacity $C_R \approx \alpha \cdot \Omega_R$ scales with atomic size. Second-row atoms (C, N, O) have small Ω_R , making $m = 6$ coordination inadmissible ($B_R > C_R$). Third-row atoms (S, P) have larger Ω_R , making hypervalency admissible. This derives the octet rule as a capacity bound, not a fiat.
- *Aromaticity*: Cyclic delocalized systems with $4n+2$ electrons achieve $g = 1$ (no frontier degeneracy to resolve), giving zero closure penalty. Systems with $4n$ electrons have $g > 1$, requiring additional commitment to resolve degeneracy, with closure penalty $B_{\text{cl}} = \beta \ln g > 0$.

Reactive intermediates as a commitment spectrum. Not all species that appear between reactant and product in a reaction mechanism achieve the same factual status. A long-lived, isolable intermediate — a tropylium cation characterised by NMR, a persistent radical crystallographically confirmed — satisfies all three criteria for facthood: independent attribution, closure minimality, and environmental embedding. A short-lived species whose "existence" is inferred only from kinetic data and whose lifetime is shorter than solvent reorganisation timescales (~ 1 ps in solution) does not. The spectrum from proto-factual to factual is continuous across reactive intermediates, and the experimentally meaningful question — "is this intermediate a real chemical species?" — is precisely the question of whether it crosses the commitment threshold. Trapping experiments, which intercept an intermediate by a fast scavenging reaction, are experiments that probe whether the intermediate has achieved sufficient factual status to sustain an independent commitment with the trapping agent.

Concerted versus stepwise mechanisms. The distinction between concerted mechanisms (a single transition state, bond breaking and forming simultaneously — $\text{S}_{\text{N}}2$ substitution, Diels-Alder cycloaddition) and stepwise mechanisms (one or more intermediates — $\text{S}_{\text{N}}1$ ionisation, addition-elimination) has a precise commitment-framework interpretation. A concerted mechanism involves a single proto-factual zone traversal separating the reactant fact from the product fact: no intermediate entry is created in the causal record between them. A stepwise mechanism involves multiple traversals, with a committed intermediate fact produced between them — provided the intermediate achieves sufficient stability and environmental embedding. When the intermediate does not achieve facthood (lifetime too short for solvent reorganisation), the stepwise and concerted mechanisms are structurally indistinguishable at the level of fact

formation, regardless of what the potential energy surface looks like. Woodward-Hoffmann orbital symmetry rules determine which concerted pathways have low electronic barriers; the commitment framework determines which of those pathways are structurally admissible given finite local capacity — the two constraints are independent, and both must be satisfied.

Enzyme active sites as commitment scaffolds. An enzyme active site performs three distinct commitment-relevant functions, each corresponding to a different component of the fact-formation process. First, it reduces the configuration-space volume of proto-factual trajectories by constraining substrate orientation: random solution encounters generate geometrically diverse proto-factual configurations of which most are non-productive; the active site selects only those encounter geometries that are compatible with subsequent commitment, increasing A in the Arrhenius equation. Second, it lowers E_a through electrostatic complementarity to the transition-state charge distribution, increasing the probability that a proto-factual encounter has sufficient energy to reach the dividing surface. Third — and least discussed in standard accounts — it provides a commitment scaffold: an ordered network of hydrogen bonds, metal cofactors, and oriented water molecules that couple the emerging product configuration to environmental degrees of freedom more rapidly than in bulk solution, driving $B_R \geq C_R$ immediately after the barrier crossing rather than after a diffusion-limited solvation delay. This third function directly increases κ , the transmission coefficient. Enzyme catalysis is therefore not reducible to transition-state stabilisation; it is engineering of the complete proto-factual-to-factual transition, with the active site accelerating each of the three stages that determine whether and how rapidly a committed product fact forms.

Quantum tunnelling. During barrier traversal, no classical trajectory exists; intermediate positions are not independently attributable; and measurement within the barrier destroys the process. In this framework, tunnelling is proto-factual throughout. Commitment occurs only upon emergence and stabilisation on the product side. Kinetic isotope effects — the slower tunnelling rate of deuterium relative to hydrogen — are differences in fact-formation probability arising from the mass-dependence of the action budget component of B_R .

Interpretive scope. This classification does not modify the quantitative predictions of quantum mechanics for tunnelling rates; it provides a structural account of tunnelling trajectories within the proto-factual regime. The claim is ontological, not computational.

Testable implication. Within this framework, commitment occurs when environmental coupling drives $B_R \geq C_R$. For tunnelling processes this predicts a distinction absent from standard decoherence accounts: increasing environmental coupling should not merely decohere the tunnelling system but should shift the effective tunnelling-to-classical crossover point, altering the relative contribution of tunnelling versus over-barrier activation in a way that depends on the capacity C_R of the local environment, not only on coherence timescales. Specifically, systems engineered to remain below threshold — by suppressing environmental coupling and thereby keeping $B_R < C_R$ — should exhibit enhanced tunnelling persistence *beyond* standard decoherence expectations, because commitment is suppressed rather than merely coherence degraded. This suggests an experimental distinction between decoherence-limited behaviour (standard quantum mechanics) and capacity-limited commitment behaviour (this framework): the two accounts predict the same tunnelling rate under identical environmental conditions but

diverge in their predictions for how that rate shifts as the environment is systematically varied. Enzyme active sites, where tunnelling persists in warm aqueous environments that would naively decohere quantum coherence rapidly, may provide a natural test case.

11.1.1 Worked Example: Threshold Crossing Along a Reaction Path

The geometry example of Section 8.3 shows static admissibility — which configurations can be sustained. The following example shows dynamic commitment — how a system crosses from proto-factual evolution into stabilised record along a time-evolving reaction coordinate. This is the central claim of the framework made concrete.

Let q denote a reaction coordinate with three characteristic regions: reactant basin $q = q_R$, barrier / transition-state region $q = q^\ddagger$, and product-side stabilisation region $q = q_P$. We model the commitment load as:

$$B_R(q, t) = B_{\text{geom}}(q) + B_{\text{env}}(q, t)$$

where $B_{\text{geom}}(q)$ is the intrinsic distinguishability cost of the instantaneous structural configuration, and $B_{\text{env}}(q, t)$ is the contribution from environmental embedding — solvent reorganisation, vibrational relaxation, coupling to surrounding degrees of freedom. Commitment requires the total load to reach threshold: $B_R(q, t) \geq C_R$.

Reactant basin. The reactant is already a composite fact — a previously committed stable structure. But the reactive transformation has not yet occurred. The local evolution toward reaction therefore begins in a proto-factual regime relative to the prospective outcome: $B_R(q_R, t) < C_R$ for the incipient product-forming process. This is an important distinction: the reactant is factual as a stabilised structure, while the reactive event remains proto-factual until a new committed outcome forms.

Barrier region. As the system approaches the transition state, bond lengths distort, angular constraints tighten, and distinguishability costs rise:

$$B_{\text{geom}}(q_R) < B_{\text{geom}}(q) < B_{\text{geom}}(q^\ddagger)$$

However, at the transition state the environment has not yet stabilised the emerging outcome. Solvent shells have not reorganised, excess vibrational energy has not been dissipated, and no persistent record of "product formed" yet exists. Despite high geometric strain, $B_{\text{env}}(q^\ddagger, t)$ remains insufficient for commitment:

$$B_R(q^\ddagger, t) = B_{\text{geom}}(q^\ddagger) + B_{\text{env}}(q^\ddagger, t) < C_R$$

The transition state is therefore not the location of fact formation. It is the point of maximum sub-threshold structural strain. The outcome remains open: trajectories may proceed to products or return to reactants.

Product-side descent. Once past the barrier, geometric strain relaxes: $B_{\text{geom}}(q_P) < B_{\text{geom}}(q^\ddagger)$. But environmental coupling now rises sharply as the system begins embedding the new outcome into its surroundings — vibrational energy redistributes, solvent molecules reorganise, a persistent local record begins to form:

$$\partial B_{\text{env}} / \partial t > 0$$

The total commitment load therefore crosses threshold after the barrier. There exists a commitment time t_c such that $B_R(q_P, t_c) = C_R$, and for $t > t_c$, $B_R > C_R$: an irreversible fact has formed.

Minimal schematic model. The logic can be made explicit with a simple illustrative parametrisation. Let the structural term peak at the barrier:

$$B_{\text{geom}}(q) = B_0 + A e^{-(q - q^\ddagger)^2/\sigma^2}$$

and let the environmental term grow only on the product side:

$$B_{\text{env}}(q, t) = H(q - q^\ddagger) \eta(1 - e^{-t/\tau_s})$$

where H is the Heaviside function, $\eta > 0$ is the maximum embedding contribution, and τ_s is a stabilisation timescale. Before the barrier, $B_{\text{env}} = 0$; at the barrier, B_{geom} is maximal but commitment may still fail; after the barrier, B_{env} increases with time and can drive the total above threshold. The commitment condition becomes:

$$B_0 + A e^{-(q - q^\ddagger)^2/\sigma^2} + H(q - q^\ddagger) \eta(1 - e^{-t/\tau_s}) \geq C_R$$

This form captures the central structural claim: fact formation need not occur at the barrier top. It occurs when post-barrier stabilisation drives the total commitment load across threshold.

Three regimes. The example reveals three physically distinct cases:

Regime	Condition	Consequence
Sub-threshold recrossing	$B_R(q_P, t) < C_R$ for all t	Trajectory returns to reactant; no fact formed
Threshold-crossing	$\exists t_c : B_R(q_P, t_c) = C_R$	Fact forms at t_c ; committed product
Rapid commitment	$\tau_s \ll$ recrossing timescale	Commitment immediately after barrier crossing

Interpretation of κ . Not every trajectory that crosses q^\ddagger achieves sufficient stabilisation to satisfy $B_R(q_P, t) \geq C_R$. Some recross before environmental embedding completes. In this framework, the transmission coefficient κ is the fraction of barrier-crossing trajectories that proceed far enough into the product-side stabilisation regime to reach commitment threshold.

This is stronger than saying κ corrects for recrossing: it identifies what successful trajectories have in common — they are the ones for which environmental embedding outpaces reversal.

Environmental modulation of the commitment time. The commitment time t_c is not fixed solely by the intrinsic reaction coordinate. It depends on the rate and strength of product-side environmental embedding. Making this explicit, write the environmental term with explicit dependence on temperature T and solvent or environmental conditions S :

$$B_R(q, t; T, S) = B_{\text{geom}}(q) + H(q - q^\ddagger) \eta(T, S) (1 - e^{-t/\tau_s(T, S)})$$

where $\eta(T, S)$ is the maximum embedding strength available from that environment and $\tau_s(T, S)$ is the stabilisation timescale. Commitment occurs when $B_R(q_P, t_c; T, S) = C_R$, giving:

$$t_c(T, S) = -\tau_s(T, S) \cdot \ln[1 - (C_R - B_{\text{geom}}(q_P))/\eta(T, S)]$$

provided $0 < (C_R - B_{\text{geom}}(q_P))/\eta(T, S) < 1$.

Three physically distinct cases follow directly.

Faster environmental relaxation (smaller τ_s): the environmental term rises to threshold sooner, reducing t_c . A solvent that reorganises rapidly around the emerging product geometry commits the product fact earlier, even if barrier crossing occurs at the same point on the potential surface.

Stronger embedding (larger η): the environmental term reaches a higher asymptote and crosses C_R sooner, again reducing t_c . Polar or strongly reorganising environments should generally move commitment earlier than weakly coupled ones.

Temperature: the effect is non-monotone, and that is a strength rather than a weakness. Raising T can decrease τ_s by speeding environmental motion (earlier commitment) but can also decrease η by degrading stabilisation persistence (later or no commitment). The framework therefore predicts that temperature does not act uniformly on commitment: it shifts t_c in either direction depending on whether environmental reorganisation or environmental destabilisation dominates. This is a falsifiable, domain-specific prediction distinguishable from purely kinetic accounts in which higher temperature monotonically accelerates reaction.

No commitment. The condition $\eta(T, S) < C_R - B_{\text{geom}}(q_P)$ defines a regime in which the environment is too weak ever to drive the system over threshold, even after barrier crossing. The product-side configuration then remains proto-factual indefinitely — susceptible to recrossing and decay — and no committed fact forms. This corresponds to the sub-threshold recrossing regime of the three-regime table above, but now with an explicit environmental criterion. A reaction that is energetically downhill and passes its transition state may still fail to produce a committed product if the environment cannot supply sufficient embedding.

This expression makes explicit that external control parameters shift the commitment moment rather than merely altering the barrier-crossing rate. The transition state stays where the potential-energy surface puts it; the fact-formation time t_c is separately determined by

environmental capacity. The lag $t_c - t_{\text{barrier}}$ between barrier crossing and fact formation is an experimentally tunable quantity — accessible in principle through ultrafast spectroscopy, solvent isotope effects, and pressure-dependent kinetics — that provides a direct handle on commitment dynamics independently of transition-state structure.

Distinction from standard kinetic accounts. Standard reaction theory treats environmental effects as modifications to barrier-crossing rates, typically through changes in activation energy or frictional damping — the Kramers framework. In that picture, the decisive event is passage over the barrier, and environmental parameters affect how often that passage occurs. Faster solvent relaxation lowers friction and increases the crossing rate; stronger coupling raises recrossing corrections. The barrier is the event; everything else is a correction to it.

In the present framework, the decisive event is not barrier crossing but commitment: the moment at which $B_R \geq C_R$. Environmental parameters therefore act not only on the rate of barrier crossing but on the existence and timing of a distinct commitment event occurring after the barrier. This produces a structural distinction with testable consequences:

	Standard kinetic account	Present framework
Decisive event	Barrier crossing	Commitment ($B_R \geq C_R$)
Role of environment	Modifies crossing rate	Determines whether and when commitment occurs
Barrier and reaction event	Coincide	Conceptually distinct; lag $t_c - t_{\text{barrier}}$ is physical
Two environments with equal crossing rates	Predict identical reaction	May differ in commitment timing or occurrence

The last row is the sharpest experimental handle. Two solvents engineered to produce identical barrier-crossing rates — identical Kramers friction coefficients, identical transmission coefficients — but differing in their embedding capacity $\eta(T,S)$ should produce measurably different product formation times and, in the sub-threshold regime, different product yields. Standard theory has no mechanism to predict this divergence; the present framework predicts it directly from $t_c(T,S) = -\tau_s \cdot \ln[1 - (C_R - B_{\text{geom}})/\eta]$. This is not solvent effects in new language. It is a new quantity — the commitment lag — with a distinct experimental signature.

Scope. This is a worked illustration of threshold crossing, not a full derivation of rate theory. The parametric model is schematic; calibration of B_0 , A , η , and τ_s from molecular data is developed in the BCB companion paper. The structural point stands independently: the transition state and the commitment event are conceptually distinct, even when they are experimentally difficult to resolve. This distinction becomes practically consequential in systems where environmental coupling is highly structured — enzyme active sites, confined catalytic pores, low-temperature tunnelling systems — where the barrier may be crossed while the outcome remains proto-factual for an extended interval.

11.1.2 Concrete System: Solvent-Dependent Proton Transfer

This section identifies a specific experimental system where the commitment lag is measurable and solvent-tuneable. Section 11.1.3 then shows what happens when this lag collapses — the enzyme case — and Section 11.1.4 grounds that prediction in a specific enzyme with existing data.

Excited-state proton transfer from photoacid donors — 7-hydroxyquinoline and related compounds — in solvents of varying hydrogen-bond structure provides a natural experimental testbed for the commitment lag. This class of reaction has been characterised by ultrafast spectroscopy across water, alcohols, DMSO, and aprotic proton-accepting media, and exhibits exactly the environmental sensitivity the framework predicts: initial proton transfer can occur on timescales of ~ 100 fs in water, while other solvents show markedly different dynamics associated with altered solvent reorganisation and proton-accepting structure.

The system is valuable precisely because it separates three stages that are conflated in standard kinetic accounts:

Stage	Description	Commitment status
Photoexcitation	Donor enters proto-factual excited state; proton transfer becomes dynamically allowed	$B_R < C_R$
Transfer and early solvent response	Proton relocates through local hydrogen-bond structure; solvent not yet equilibrated	$B_R < C_R$ (proto-factual throughout)
Product-side embedding	Solvent reorganises around transferred proton and conjugate base; stable local record forms	$B_R \geq C_R$ (fact)

The decisive point is that the transfer event and the commitment event need not coincide. A proton may leave the donor on an ultrafast timescale while the solvent environment is still far from equilibrium. In that window, the system is proto-factual after transfer but before embedding: the new charge distribution exists dynamically, but the product state is not yet a stabilised, irreversible record. This interpretation is consistent with studies emphasising the role of nonequilibrated solvent environments in ultrafast proton transfer — the "spectroscopic product" appears rapidly, but its factual status requires the subsequent solvent reorganisation.

The $t_c(S)$ expression derived in Section 11.1.1 applies directly. Writing the environmental term with explicit solvent dependence:

$$B_R(q, t; S) = B_geom(q) + H(q - q^\ddagger) \eta(S)(1 - e^{-(t/\tau_s(S))})$$

commitment occurs when $B_R(q_P, t_c; S) = C_R$, giving:

$$t_c(S) = -\tau_s(S) \cdot \ln[1 - (C_R - B_geom(q_P))/\eta(S)]$$

A solvent that reorganises rapidly and stabilises the transferred proton strongly (large η , small τ_s) produces a small t_c : commitment follows closely after transfer. A solvent that accepts the proton but reorganises weakly or more slowly produces a larger t_c — even if the initial transfer step is equally fast. In the limiting case $\eta(S) < C_R - B_geom(q_P)$, the transferred configuration

never accumulates sufficient environmental embedding to cross threshold: no committed product form forms, and the proton may return without net reaction.

The framework therefore predicts that systematic variation of solvent across a matched series — holding initial transfer rates fixed while varying η and τ_s through solvent choice — should produce measurable variation in the lag $t_c - t_{\text{transfer}}$ between the spectroscopic appearance of the transferred proton and the formation of the fully committed product state. This lag is not predicted by conventional accounts, in which solvent effects enter primarily through barrier modification and the transfer event is identified with the reaction itself.

Experimental signature of the commitment lag. Ultrafast spectroscopy typically identifies the proton-transfer event through changes in emission or absorption associated with proton relocation — a fast signal, often sub-picosecond in water. In the present framework, this spectroscopic signature marks the transfer event t_{transfer} , not the commitment event t_c . These need not coincide. The framework therefore predicts the existence of two potentially distinct timescales in time-resolved spectra:

Signal	Identity	Framework interpretation
Fast emission/absorption change	Proton relocation	Transfer event t_{transfer} (proto-factual)
Slower spectral equilibration	Solvent-reorganised product	Commitment event t_c (fact formation)

Systematic variation of solvent should modify the relationship between these timescales in a specific pattern:

- In strongly embedding solvents (large η , small τ_s): $t_c \approx t_{\text{transfer}}$ — commitment follows transfer immediately; the two signatures are unresolved
- In weakly embedding or slow-reorganising solvents: $t_c > t_{\text{transfer}}$ — a measurable gap opens between the initial transfer signal and the fully stabilised product feature
- In the sub-threshold regime $\eta(S) < C_R - B_{\text{geom}}(q_P)$: t_c does not exist — the transferred configuration decays without producing a stable product signal, and the fast spectroscopic change is followed by recovery rather than product formation

This pattern is falsifiable and solvent-series specific. If time-resolved spectra across a matched solvent series show no systematic separation between transfer and product-equilibration timescales — if every solvent shows the same lag structure regardless of embedding capacity — the commitment framework's account of this system is wrong. If the separation tracks η and τ_s as predicted, it provides a direct experimental probe of the commitment condition $B_R \geq C_R$, independent of barrier-crossing kinetics and inaccessible to standard Kramers-type accounts.

Scope. The claim here is that this class of system constitutes a natural experimental testbed for the commitment lag, not that existing published data have been fitted to the model. Calibration of $\eta(S)$ and $\tau_s(S)$ from published ultrafast transient spectra — resolving the lag between spectroscopic transfer and solvent-equilibrated product — is the immediate experimental task. The framework makes the prediction; the measurement determines whether it is correct.

11.1.3 Proton Transfer in Enzymes versus Solution: Collapse of the Commitment Lag

Enzyme-catalysed proton transfer presents one of the most persistent anomalies in physical chemistry. Reactions in enzymatic active sites often exhibit kinetic isotope effects, temperature dependences, and tunnelling contributions inconsistent with classical transition-state theory — even in warm, strongly coupled aqueous environments where quantum coherence would naively be expected to decohere rapidly. The commitment framework provides a unified structural account of these observations, and generates a direct experimental prediction distinguishing it from standard explanations.

Bulk solution. In bulk solvent, proton transfer proceeds through solvent-mediated pathways in which the environment must reorganise around the transferred proton to stabilise the new charge distribution. In the commitment framework, transfer and commitment are distinct events. The proton may relocate rapidly — on femtosecond to picosecond timescales — while environmental embedding lags behind. The commitment lag $t_c - t_{\text{transfer}}$ is finite and governed by $\eta(S)$ and $\tau_s(S)$: in weakly embedding or slowly reorganising solvents, this lag is large enough that a significant fraction of transfer trajectories recross before $B_R \geq C_R$ is achieved. Tunnelling contributions are suppressed in this regime because the proto-factual tunnelling trajectory must compete with recrossing on timescales comparable to or shorter than the environmental embedding time.

Enzymatic active sites. An enzyme active site provides a precisely preorganised environment: proton acceptors positioned by evolution, hydrogen-bond networks locked into specific geometries, electrostatic fields pre-aligned to the transition-state charge distribution. In the commitment framework, this preorganisation has two direct effects on the commitment parameters:

- *Increased η :* the active site provides strong, directed stabilisation of the transferred proton, coupling the emerging product configuration to a dense, ordered scaffold of environmental degrees of freedom
- *Decreased τ_s :* the scaffold is already positioned — no diffusive reorganisation is required. Environmental embedding begins immediately upon transfer, not after a solvation delay

The combined effect drives $t_c \rightarrow t_{\text{transfer}}$: the commitment lag collapses. Transfer and commitment become effectively simultaneous. The proton does not enter a proto-factual transferred-but-uncommitted state; it crosses the commitment threshold as it crosses the spatial gap.

Unified interpretation of enzyme tunnelling anomalies. This collapse of the commitment lag provides a structural account of three otherwise disparate observations:

Persistent tunnelling in warm environments. Tunnelling is a proto-factual process: the barrier is traversed before commitment. If $t_c \rightarrow t_{\text{transfer}}$, the commitment event follows the tunnelling event before environmental fluctuations can suppress the quantum amplitude. The enzyme does not protect coherence in the conventional sense; it eliminates the window during which

decoherence could compete with commitment. Tunnelling persists not because the enzyme is quantum-coherent but because the commitment is fast.

Anomalous temperature dependence. In bulk solution, the temperature dependence of the rate reflects both barrier crossing (Boltzmann factor) and embedding kinetics ($\tau_s(T)$). In enzyme active sites where the commitment lag is collapsed, the rate-determining step shifts: barrier crossing may become subordinate to the embedding-controlled commitment process, producing temperature dependences that do not follow simple Arrhenius behaviour and that differ markedly from the same reaction in solution.

Catalytic efficiency beyond barrier lowering. Standard accounts attribute enzyme catalysis primarily to transition-state stabilisation — lowering E_a . The commitment framework adds a second mechanism: reduction of the commitment lag. An enzyme that eliminates recrossing by forcing $t_c \rightarrow t_{\text{transfer}}$ achieves a higher effective κ than the same reaction in solution, even at identical barrier heights. Catalytic power is not solely an energetic quantity; it includes the probability that a barrier-crossing trajectory achieves commitment before reversal.

Prediction. The framework predicts that enzyme active sites should systematically eliminate the measurable commitment lag relative to the same reaction in bulk solution. In terms of observable timescales:

System	Transfer timescale t_{transfer}	Commitment timescale t_c	Lag
Bulk weakly-embedding solvent	Fast	Slow	Large; recrossing probable
Bulk strongly-embedding solvent	Fast	Moderate	Reduced; fewer recrossings
Enzyme active site	Fast	$\approx t_{\text{transfer}}$	Near-zero; commitment simultaneous

Ultrafast spectroscopic comparison of proton transfer in enzyme active sites, enzyme-mimetic environments (cyclodextrin hosts, synthetic hydrogen-bond cages), and bulk solvents should reveal this convergence: the separation between transfer and product-stabilisation signatures should decrease systematically as environmental preorganisation increases, approaching zero in the enzyme case. Failure of this convergence — enzyme and solution showing the same lag structure — would falsify the commitment account of enzyme tunnelling.

Enzymes are not catalysts in the traditional sense. They are devices that force proto-factual dynamics to cross the commitment threshold before reversal becomes possible.

This reframing has consequences beyond proton transfer: it implies that enzyme specificity is partly a commitment phenomenon — the active site is shaped not only to bind the transition state but to ensure that correct-product trajectories achieve commitment while incorrect-product trajectories do not. Selectivity is commitment engineering.

11.1.4 Concrete System: Ketosteroid Isomerase

Ketosteroid isomerase (KSI) is one of the most extensively characterised enzymes in physical biochemistry and provides the sharpest available test of the commitment-lag framework for enzymatic proton transfer. KSI catalyses isomerisation of steroid substrates through a concerted proton-transfer mechanism, and has been studied with kinetic isotope effects, temperature-dependent rate measurements, and ultrafast spectroscopy. The kinetic data present a standing interpretive challenge for classical transition-state theory: KSI exhibits unusually large primary kinetic isotope effects (k_H/k_D substantially exceeding classical predictions), comparatively weak temperature dependence, and rate enhancements that cannot be fully accounted for by electrostatic transition-state stabilisation alone.

In the present framework, these observations receive a unified structural account through the collapse of the commitment lag.

Bulk-solution reference. In bulk aqueous solution, analogous proton transfer from carbon acids proceeds through a solvent-mediated pathway in which water molecules must reorganise around the developing charge distribution. The commitment lag $t_c - t_{\text{transfer}}$ is determined by $\eta(\text{water})$ and $\tau_s(\text{water})$: hydrogen-bond network reorganisation around a newly transferred proton requires multiple solvent shells to equilibrate, typically on picosecond timescales. A significant fraction of transfer trajectories recross during this lag, reducing the effective κ and suppressing tunnelling contributions through environmental decoherence.

KSI active site. The KSI active site contains two tyrosine residues (Tyr14 and Tyr55) and an aspartate (Asp38) arranged to provide precisely positioned hydrogen-bond donors and acceptors along the proton-transfer coordinate. This preorganisation has direct consequences for the commitment parameters:

- $\eta(\text{KSI}) \gg \eta(\text{water})$: the tyrosine hydroxyl groups and active-site electrostatics provide immediate, directed stabilisation of the transferred proton, with no diffusive reorganisation required
- $\tau_s(\text{KSI}) \ll \tau_s(\text{water})$: the scaffold is structurally committed before the substrate arrives; coupling of the emerging product configuration to the environmental scaffold is near-instantaneous

The predicted consequence is $t_c \rightarrow t_{\text{transfer}}$: the commitment lag collapses to near zero. The proton-transfer event and the commitment event become experimentally indistinguishable, occurring on the same ultrafast timescale.

Interpretation of KSI anomalies.

Observation	Standard account	Commitment framework
Large kinetic isotope effects	Extended tunnelling through low barrier	Tunnelling completes before decoherence — commitment lag collapsed

Observation	Standard account	Commitment framework
Weak temperature dependence	Protein conformational pre-organisation	Embedding-controlled commitment already maximised — temperature modulates η weakly
Rate enhancement beyond electrostatics	Unknown additional contributions	Collapse of κ -reducing commitment lag — fewer recrossings, more committed trajectories

The standard account attributes the large KIE to a narrow, symmetric barrier admitting significant tunnelling. The commitment framework offers a complementary account: the KIE is large not only because tunnelling is geometrically favoured but because the commitment lag is short enough that tunnelling trajectories achieve commitment before environmental fluctuations can suppress the quantum amplitude. In solution, tunnelling is geometrically similar but the commitment lag allows decoherence to compete; in KSI, the lag is collapsed and tunnelling trajectories commit before decoherence acts.

Prediction. The framework makes a specific prediction distinguishable from standard accounts: the rate-limiting step in KSI should shift between bulk solution and active site not primarily in barrier height but in the embedding-controlled commitment process. This predicts:

1. Mutants that disrupt active-site preorganisation (e.g., Tyr14Phe, which removes a hydrogen-bond donor) should increase the commitment lag — increasing recrossing and suppressing tunnelling — even if the energetic barrier is only modestly affected
2. The temperature dependence of the KIE (the ratio A_H/A_D in the Arrhenius pre-exponential factors) should correlate with active-site preorganisation metrics rather than with barrier geometry alone
3. Ultrafast measurements in KSI versus water-exposed mutants should reveal a separation between transfer and product-stabilisation timescales that is absent or minimal in the wild-type active site

Prediction (1) is testable against existing mutagenesis data; predictions (2) and (3) are accessible to current ultrafast spectroscopy. KSI is a sufficiently well-characterised system that if the commitment-lag account is correct, the signatures should be resolvable with existing techniques. If the mutant KIE data and ultrafast measurements show no systematic relationship to preorganisation metrics — if disrupting the active-site scaffold does not increase the commitment lag as measured by spectroscopic separation of transfer and product timescales — the commitment account of KSI is falsified.

11.2 Mechanism Transitions and OER

The oxygen evolution reaction (OER) in electrochemical water splitting involves two competing mechanisms: the adsorbate evolution mechanism (AEM), in which O–O bond formation occurs via surface-bound intermediates, and the lattice oxygen mechanism (LOM), in which lattice oxygen atoms participate directly.

In BCB terms, as operating current density and surface coverage increase, the cumulative ledger cost B_R of maintaining all required adsorbate commitments on the catalyst surface approaches C_R . When $B_R^{(AEM)} > C_R$, the surface-only pathway becomes inadmissible, and the system must redistribute commitments into the lattice — activating LOM as a capacity-relief mechanism.

This reframes the AEM→LOM transition as an *admissibility transition* rather than a purely energetic one. BCB predicts: (i) a sharp, thresholded onset of lattice oxygen participation correlated with surface congestion proxies rather than adsorption energies; (ii) hysteresis if the transition involves structural reorganisation; (iii) delayed LOM onset for catalysts engineered to increase C_R via surface morphology or lattice compliance. These predictions are experimentally falsifiable and distinct from purely energetic accounts.

11.3 Biological Threshold Systems

The following systems share a common structural pattern with the commitment framework: a continuous input accumulates until it crosses a threshold, at which point an irreversible — or effectively irreversible — outcome is produced. We do not claim these systems *prove* the framework. We claim they exhibit the same structural pattern the framework predicts, and that the framework provides a unifying description across domains where none previously existed.

Gene regulatory switches. Cells make fate decisions using thresholded gene expression networks. Below a signal threshold, transcription factors produce transient mRNA that is degraded before a stable protein record accumulates. Above threshold, positive feedback loops sustain transcription, protein product accumulates, and the system locks into a stable ON state. The classic lac operon exhibits exactly this bistability: the same concentration of inducer produces different outcomes depending on prior state, because the commitment threshold has or has not been crossed.

In the framework's terms: molecular fluctuations and transient transcriptional events are proto-factual — dynamically real, carrying information, but not yet constituting a stable committed record. The commitment event is not a single molecular binding but the accumulation of sufficiently stabilised transcriptional records:

$B_R \sim$ (number of stabilised transcriptional commitments)

Below threshold, $B_R < C_R$: transient mRNA is produced and degraded, leaving no persistent record. Above threshold, environmental embedding through protein feedback drives $B_R \geq C_R$: a stable cellular state is constituted. The biological irreversibility of cell-fate decisions — their resistance to perturbation once made — is not kinetic coincidence but structural consequence: once $B_R \geq C_R$ at sufficient causal closure, the commitment is not erasable by local operations.

Neuronal action potentials. A neuron integrates incoming signals as graded membrane potential changes. Below the firing threshold, perturbations decay and no propagating signal is produced. Above threshold, voltage-gated sodium channels open, producing a self-amplifying depolarisation — the action potential — that propagates along the axon and is transmitted to

downstream neurons. The outcome is all-or-none and effectively irreversible on neurobiological timescales.

The mapping is precise:

Stage	Neuronal description	Commitment status
Sub-threshold integration	Graded potential, decaying	$B_R < C_R$
Threshold crossing	Channel opening, runaway depolarisation	$B_R = C_R$
Action potential	Self-propagating, irreversible signal	$B_R > C_R$ (fact)

A neuron fires not when a single input arrives, but when accumulated distinguishability — the integrated history of synaptic inputs stabilised against membrane noise — exceeds the local capacity threshold, producing a self-propagating irreversible record. Neural coding, on this account, is commitment engineering: the nervous system controls which proto-factual accumulations reach threshold and which do not.

Protein folding as distinguishability collapse. Levinthal's paradox is one of the most cited puzzles in molecular biology: a protein of 100 residues has more possible conformations than there are atoms in the observable universe, yet real proteins fold to unique three-dimensional structures in milliseconds to seconds. Random search is computationally impossible. The standard resolution appeals to energy funnels — a biased landscape that guides the chain toward its native state. This is correct as far as it goes, but it does not explain the *speed, reliability, and irreversibility* of folding without additional assumptions.

The commitment framework provides a structural resolution. Folding is not a search; it is a sequence of irreversible distinguishability commitments. Each time a structural contact locks in — two residues come together and are stabilised by their environment — it partitions the previously admissible configuration space, eliminating a large class of indistinguishable alternatives. The admissible set \mathcal{A}_k after k committed contacts satisfies:

$$\mu(\mathcal{A}_k) \leq \mu(\Omega) \cdot c^k, \quad c < 1$$

where $\mu(\Omega)$ is the total configuration measure. This exponential collapse — proven in a simplified polymer model under a mild average-case survival condition — rules out Levinthal-style combinatorial search: the effective configuration space vanishes exponentially with the number of committed structural facts, not because high-energy states are penalised but because they become structurally inadmissible.

In this framework, B_R for the protein folding context is given concretely by the **Late-Collapse Penalty (LCP)**:

$$\text{LCP} = \sum_{k=1}^M w_k \cdot (1 + D_{k-1})$$

where w_k is the distinguishability weight of the k -th committed contact (proportional to its sequence separation and topological impact), and $D_{k-1} = \sum_{m \leq k-1} w_m$ is the cumulative

distinguishability already committed. The LCP weights each structural commitment by how constrained the system already is when that commitment must be made: large commitments late in folding — when the chain is already confined — contribute super-linearly to kinetic cost. This is a concrete, computable instantiation of B_R within a specific physical domain.

Classical contact-order scaling as derived result. The empirical observation that folding rate correlates with relative contact order ($RCO = \langle \Delta s \rangle / N$, where Δs is sequence separation of native contacts) has been treated as a phenomenological input to folding models since Plaxco et al. (1998). Within the commitment framework, it emerges as a theorem: under a confinement hypothesis — that after partial collapse, loop-closure probability is exponentially suppressed with sequence separation — the total folding time satisfies:

$$\ln T \sim \gamma \cdot N \cdot RCO + \text{LCP corrections}$$

Contact order appears not as an energetic postulate but as the leading cost of acquiring long-range distinguishability commitments in a progressively constrained system.

Empirical consistency check. The commitment ordering predicted by this framework — contacts with lower sequence separation and lower topological impact should commit first — generates qualitative predictions about transition-state structure. These were tested against published Φ -value data for chymotrypsin inhibitor 2 (CI2). The commitment ordering score was computed from the native contact graph using purely topological parameters (sequence separation and bridge status) with no fitting to CI2 data; the weights were calibrated from separate chemical benchmark data (H_2O geometry and the CF_6/SF_6 existence boundary). Spearman rank correlation between predicted commitment order and experimental Φ -values gives:

$$\rho = 0.52, p = 0.008 \text{ (n = 24 residues, permutation test)}$$

This is a consistency check, not a prediction test: the result confirms that commitment ordering is not orthogonal to known transition-state structure, but a moderate correlation ($\approx 27\%$ of rank variance explained) at $n = 24$ cannot distinguish the commitment ordering from other topology-based orderings. The key clarification regarding Plaxco et al. (1998): the contact-order correlation with folding rate established by Plaxco and colleagues is an empirical observation at the level of full-protein rates. The present framework derives contact-order scaling as a conditional theorem (under the confinement hypothesis) from the commitment dynamics — this is a re-derivation from a different starting point, not an independent prediction. The novelty of the CI2 result is not that it confirms something Plaxco et al. did not already know, but that commitment ordering predicts *which residues* form the transition state, not merely the overall rate. Whether this residue-level prediction outperforms other topology-based transition-state predictors (Go-model Φ -value predictions, for instance) requires a larger benchmark dataset and is a target for future work.

Chaperones as commitment-sequence engineers. Molecular chaperones accelerate protein folding and prevent misfolding. In standard accounts they are described as lowering energy barriers or preventing aggregation. In the commitment framework, a more precise description is

available: chaperones act by reordering and timing the commitment sequence, increasing the probability that high-impact constraints (w_k large) form early — when D_{k-1} is small — rather than late, thereby reducing LCP. The ribosome, during co-translational folding, functions as the first chaperone of this kind: by restricting the available configuration space as the nascent chain emerges, it enforces an ordering of commitments that reduces late large-constraint formation. This predicts that chaperone effectiveness is not primarily energetic but sequential: the key quantity is not how much they lower a barrier but how they shift the ordering of structural commitments.

A full treatment of protein folding kinetics as distinguishability collapse — including proofs, LCP calibration, MAP projection for structure prediction, and comparison with AlphaFold — is developed in a companion paper (Taylor, in preparation, *Protein Folding as Distinguishability Collapse*).

11.4 Physical Threshold Systems

The same pattern appears in condensed matter physics and photonics, where it has been studied quantitatively and is well-understood at the level of mechanism. The commitment framework does not add to this mechanistic understanding; it provides a common structural description.

Superconducting phase transition. Above the critical temperature T_c , a metal supports only normal electron transport: thermal fluctuations disrupt any incipient Cooper pairing, and no macroscopic ordered state forms. Below T_c , Cooper pairs condense into a global coherent ground state — a macroscopic quantum fact — that is stable against local perturbations and exhibits emergent phenomena (zero resistance, Meissner effect) absent in the normal phase.

In the framework's terms: above T_c , microscopic fluctuations are proto-factual — real, correlated, but not constituting a stable committed record at the macroscopic scale. Below T_c , the global condensate crosses the commitment threshold: $B_R \geq C_R$ for the collective order parameter, and a macroscopic fact is constituted. Phase transitions, generally, are collective commitment events in which a system crosses from reversible fluctuation to globally stabilised structure.

This framing connects naturally to the renormalisation group: the flow toward a fixed point is the approach to $B_R = C_R$ in the space of collective degrees of freedom, and the ordered phase is the committed state that results.

Laser threshold. Below the lasing threshold, a pumped optical medium produces only spontaneous emission: photons are emitted independently, with random phases, producing no coherent field. Above threshold, stimulated emission dominates, a coherent mode is selected and amplified, and a self-sustaining coherent field is established.

The threshold is sharp: pump power versus output intensity shows a pronounced kink at threshold, with near-linear scaling above and negligible output below. Below threshold: proto-factual photons, each an independent emission event, no global record. Above threshold: coherent field, a collective commitment in which a single mode is selected and stabilised against

fluctuations. The mode selection — which of many possible lasing modes is committed to — is determined by initial fluctuations amplified past the threshold, exactly as the proto-factual evolution of a reacting molecule determines which product fact is eventually committed.

Percolation. In a random network where edges are added with probability p , below a critical threshold p_c only finite disconnected clusters exist. Above p_c , a system-spanning connected component appears discontinuously. The transition is sharp, universal in its critical exponents, and structurally irreversible in the sense that the spanning cluster, once formed, cannot be destroyed by local edge removal.

This is perhaps the purest mathematical realisation of the commitment pattern: local links are proto-factual, each adding incrementally to B_R ; the spanning cluster is the committed fact, constituted when global connectivity — causal closure across the system — is achieved. The universality of percolation critical exponents across physical systems suggests that the threshold structure is not domain-specific but a general feature of fact formation in networks of interacting degrees of freedom.

11.5 Environmental Modulation of the Commitment Threshold

The examples in Sections 11.1–11.4 treat C_R as a fixed property of a given region and B_R as something a system accumulates through its dynamics. But physical systems reveal a further structural feature: external parameters — temperature, pressure, phase state, solvent environment — can shift the relationship between B_R and C_R directly. In the framework's terms, these parameters do not merely accelerate or decelerate processes; they move the threshold at which fact formation occurs.

Two distinct mechanisms are operative.

Capacity modulation. Parameters that reduce the accessible configuration space or suppress environmental degrees of freedom lower C_R . Freezing is the clearest example. As a liquid transitions to a solid, reconfiguration pathways collapse and molecular motion is suppressed. Structures that were dynamically reversible — sub-threshold, proto-factual — in the liquid phase become effectively irreversible in the solid. Freezing does not add energy to the system; it reduces the system's capacity to sustain open alternatives, thereby forcing previously sub-threshold configurations into the factual regime. Pressure acts similarly: increasing pressure constrains distinguishability bandwidth and reduces the number of admissible configurations, effectively lowering C_R and inducing structural commitments at lower intrinsic complexity than would occur at ambient conditions.

Load modulation. Parameters that alter the rate of distinguishability constraint generation or the persistence of environmental embedding shift B_R . Temperature is the canonical example, but its effect is subtle: high temperature simultaneously raises B_R transiently (increased fluctuations generate more constraint attempts) and suppresses stabilisation (thermal noise prevents environmental embedding from persisting). The net effect on threshold crossing therefore depends on which mechanism dominates — a prediction that differs from purely kinetic accounts, which treat temperature as uniformly accelerating. Low temperature enhances

stabilisation persistence, allowing environmental embedding to accumulate and drive $B_R \geq C_R$ for structures that would fluctuate back at higher temperatures. Solvent polarity, ionic strength, and confinement geometry modulate B_R through their effect on environmental coupling strength.

This gives a unified account of several otherwise disparate phenomena:

Protein folding and denaturation. A folded protein is a composite fact: a stabilised structure embedded in and supported by its aqueous environment. Thermal denaturation drives B_R below C_R by disrupting the environmental embedding — breaking hydrogen bonds, exposing hydrophobic residues, collapsing the organised solvent shell — without necessarily changing the intrinsic covalent structure. Denaturation is therefore not the destruction of a structure but the dissolution of its factual status: the protein returns to a proto-factual ensemble of fluctuating conformations. Refolding is the re-establishment of $B_R \geq C_R$ through environmental re-embedding.

Pressure-induced transitions. High-pressure phases of materials — including high-pressure polymorphs, pressure-induced superconductors, and pressure-denatured proteins — represent cases where lowered C_R forces commitment to structural configurations that are sub-threshold at ambient conditions. The same intrinsic degrees of freedom cross threshold at different external conditions because the capacity of the local region has been modified.

Catalyst action as threshold engineering. Catalysts lower activation barriers in the conventional account. In this framework, a more complete description is available: catalysts modify both B_R (by organising transition-state geometry, reducing constraint interference through active-site complementarity) and C_R (by providing a structured local environment with high embedding capacity). Enzymatic rate enhancement is therefore threshold engineering — the active site is precisely shaped to bring B_R to C_R at the lowest possible structural complexity, maximising the probability of commitment for specific substrates while remaining sub-threshold for others.

The general principle is:

Fact formation is not determined solely by intrinsic system dynamics. External control parameters act by shifting the balance between commitment load B_R and local capacity C_R , moving the threshold at which irreversible records are constituted.

This is experimentally consequential: it predicts that the same intrinsic system can be driven between proto-factual and factual regimes by environmental manipulation alone, independently of any change to the system's intrinsic Hamiltonian. The clearest tests are systems near the commitment boundary — where small changes in temperature, pressure, or solvent produce sharp, discontinuous changes in outcome — as opposed to systems far from threshold, where the same manipulations produce smooth, continuous responses.

11.6 Scope of the Extension

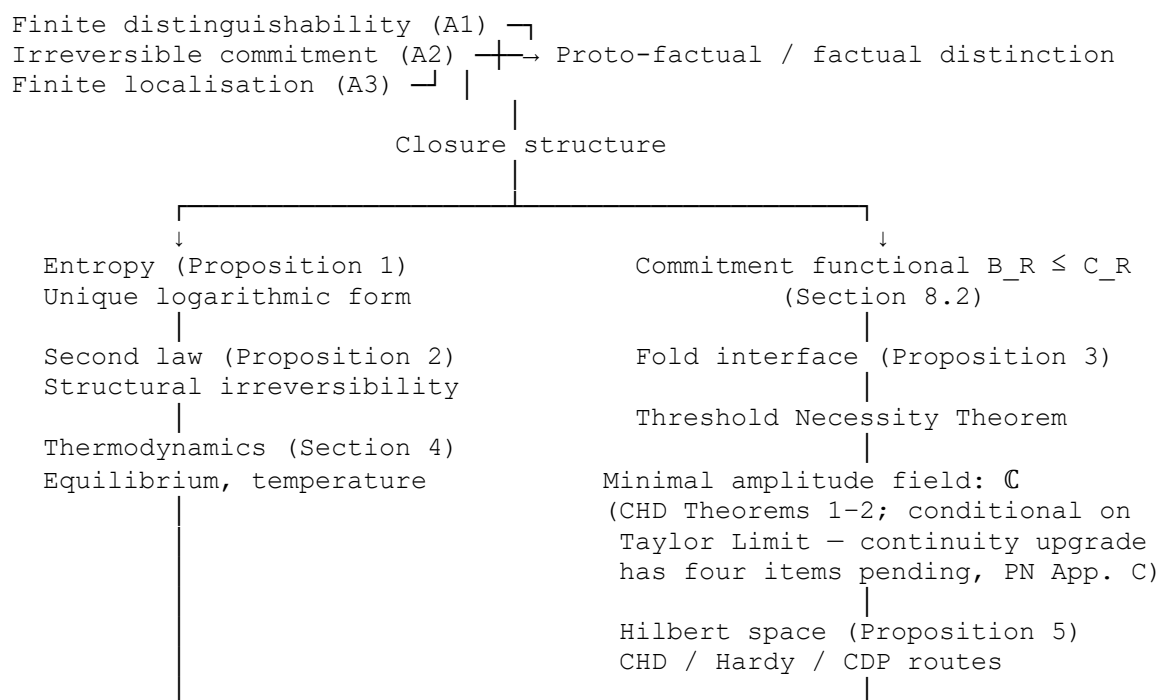
This extension does not replace quantum mechanics, chemical kinetics, or thermodynamics. It provides an ontological interpretation of when and how physical processes produce stable, causally effective outcomes. The commitment functional B_R is in each domain inferred from measurable quantities — decoherence rates, reaction lifetimes, stabilisation timescales, order parameters — rather than directly observed.

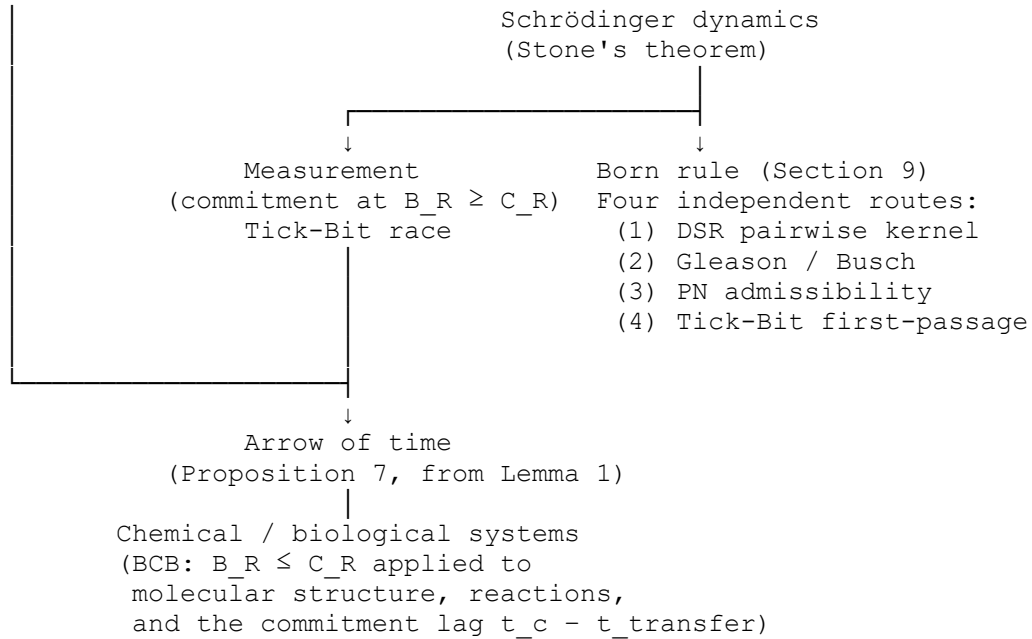
The systems discussed in Sections 11.3–11.4 are not presented as evidence for the framework in the sense of independent experimental confirmation. They are presented as instances of the same structural pattern — continuous accumulation leading to sharp, irreversible threshold crossing — that the framework describes and unifies. What the framework adds is not a new account of each individual system but a common language in which the gene switch, the action potential, the phase transition, and the laser threshold are all recognised as the same kind of event: a proto-factual process crossing the commitment threshold $B_R = C_R$.

A detailed treatment of the BCB framework, including calibration of C_R from molecular data and connection to Marcus theory and Kramers escape rates, is developed in a companion paper (Taylor, in preparation). The systems discussed in Sections 11.3–11.5 are presented not as independent experimental confirmation but as instances of the same structural pattern — continuous accumulation leading to sharp, irreversible threshold crossing — that the framework describes.

12. Unified Structure

All major physical structures derived in this paper follow from a single chain:





13. Discussion

Theorem (Threshold Necessity)

Before developing the framework's specific empirical claims, we state a structural result about the space of possible theories. This theorem does not assert the correctness of the present framework; it identifies a constraint that any adequate theory of a universe like ours must satisfy.

Theorem (Threshold Necessity). *Let a physical theory describe a universe satisfying:*

1. *Finite distinguishability: within any bounded physical context, only finitely many distinctions are operationally resolvable.*
2. *Existence of irreversible records: some physical processes produce records that persist and can serve as inputs to later processes.*
3. *Existence of reversible dynamics: some physical processes evolve without immediately producing irreversible records.*

Then the theory must contain a nontrivial condition separating reversible, outcome-indeterminate evolution from irreversible record formation. Any universe satisfying (1)–(3) must admit a boundary — structural, dynamical, or operational — at which configurations that are not yet facts become facts.

Proof. By condition (3), reversible processes exist. By condition (2), irreversible records exist. Suppose no separating condition exists between these two regimes. Then either (Case A) all physical evolution is reversible — contradicting (2), since no irreversible records could form — or (Case B) all physical configurations are already irreversibly committed at all times,

eliminating any regime of reversible superposition or interference — contradicting (3) and the observed interference phenomena that quantum mechanics correctly describes. Both exhaustive alternatives lead to contradiction. Therefore a nontrivial separating condition must exist. \square

The significance of this result is not that it proves the present framework correct. It proves something narrower and stronger: any theory adequate to a universe like ours must contain a condition separating reversible dynamics from irreversible record formation. The dispute is therefore not over whether such a condition exists, but over how it is to be characterised. The present framework characterises it as the commitment threshold $B_R = C_R$. Alternative frameworks must supply a functionally equivalent boundary under whatever terminology they prefer.

Corollary (Trilemma for Threshold-Denying Theories). *Any theory that denies the existence of a condition separating reversible evolution from irreversible record formation must reject at least one of the following: (i) the existence of reversible physical evolution; (ii) the existence of irreversible records; (iii) the applicability of the theory to a universe with both.*

This trilemma identifies the precise cost of denying a threshold. A threshold-denying theory must choose: deny reversible evolution (incompatible with interference phenomena and quantum coherence), deny irreversible records (incompatible with measurement outcomes and thermodynamic irreversibility), or treat one regime as external to the theory — leaving the transition from possibility to fact as an unexplained primitive. Standard Copenhagen occupies the third position: the Heisenberg cut is placed by fiat, not derived from physical dynamics. This is internally consistent but purchases consistency at the cost of leaving fact formation outside the physical description.

Corollary (No-Cut Alternatives Still Require a Functional Threshold). *A theory may reject the language of a "cut" or "threshold," yet if it distinguishes in practice between reversible evolution and constituted records, it contains a condition functionally equivalent to a commitment threshold.*

Proof sketch. Any such theory must specify, explicitly or implicitly, when a process counts as producing a record rather than remaining in reversible evolution. That criterion may be expressed in terms of decoherence, amplification, consistency of histories, redundancy, environmental embedding, coarse-graining, or observer-relative interaction. But if it plays the role of separating the two regimes, then it is functionally a threshold condition, regardless of the terminology used. \square

This corollary closes the terminological escape route. Einselection identifies pointer states through decoherence but must specify when decoherence is sufficient to constitute a record — that specification is a threshold condition. Consistent histories must specify which sets of histories are consistent — that specification functions as a threshold. Relational quantum mechanics must specify which interactions constitute facts for a given observer — that specification is a threshold. The present framework differs from these not in requiring a threshold (all adequate theories do) but in making the threshold a derived quantity — $B_R = C_R$ — with calculable experimental consequences.

A theory without a fact threshold must either forbid reversible dynamics, forbid irreversible records, or leave their coexistence unexplained.

13.0 Core Empirical Claim of the Framework

The examples developed in Sections 8.5 and 11.1.2–11.1.4 are not analogies. They are demonstrations of a single structural claim that constitutes the primary empirical content of the framework:

The formation of a physical outcome is not identical to the dynamical transition that precedes it.

In chemical systems this appears as a separation between barrier crossing and product stabilisation. The commitment lag $t_c - t_{\text{transfer}}$ is a physically real, environment-dependent quantity: it is large in weakly embedding solvents, reduced in strongly embedding ones, and collapsed to near zero in preorganised enzyme active sites. Two systems with identical barrier-crossing dynamics can produce different outcomes — or the same outcome at different times — depending solely on whether the commitment threshold $B_R = C_R$ is reached.

In quantum systems the same structure appears in the Quantum Zeno effect. Repeated measurement prevents the system from accumulating sufficient environmental embedding to cross the commitment threshold. The system continues to evolve dynamically — unitary evolution is not suspended — but no irreversible record forms. The effect is not a consequence of the measurement postulate; it is the quantum instance of commitment suppression, structurally identical to a chemical reaction that recrosses before environmental embedding completes.

These are not analogous systems that happen to share a mathematical form. They are physically distinct realisations of the same underlying threshold condition:

$B_R(t) < C_R \rightarrow$ proto-factual: dynamical evolution without fact formation $B_R(t) \geq C_R \rightarrow$ commitment: an irreversible record is constituted

The framework's unified prediction follows directly:

Systems with identical dynamical evolution may differ in physical outcome depending solely on whether the commitment threshold is reached.

This prediction is experimentally testable in two independent domains simultaneously. In ultrafast chemistry: matched solvent series with identical barrier heights but differing embedding capacity should show solvent-dependent separation between transfer and product-stabilisation timescales, with the separation tracking $\eta(S)$ and $\tau_s(S)$ as the framework predicts. In quantum systems: measurement apparatus with tunable coupling strength should produce partial rather than complete Zeno suppression when coupling is insufficient to reset environmental embedding, with suppression depth tracking coupling strength rather than measurement frequency.

If both predictions are confirmed, the same threshold condition is governing fact formation across quantum and chemical systems — which is the central claim. If either is violated in a way that cannot be reconciled with the commitment framework's parameters, the framework requires revision. This dual falsifiability is what distinguishes the framework from an interpretive gloss on existing physics.

13.1 What Is Established

Important structural note on claim tiers. The paper contains two distinct types of result that must not be conflated. Tier 1 results follow from A1–A3 alone. Tier 2 results follow from A1–A3 plus the specific BCB instantiation of the commitment functional $B_R = \sum b_i + \lambda \sum \Phi(\theta_{ij})$. The tetrahedral geometry result, chemical structure predictions, and chemical applications are Tier 2 — they do not follow from A1–A3 alone but from A1–A3 plus the specific additive-pairwise functional form plus empirical calibration of the parameters. The conclusion that "the irreversible commitment of distinguishability in a finite-capacity system" derives chemistry is accurate only with this qualification understood: A1–A3 establish the existence of a commitment functional satisfying structural constraints; the BCB form is the minimal representative of the resulting admissibility class, not a derived consequence of A1–A3 alone.

Tier 1 — Proven from A1–A3 within this framework:

- *Entropy* takes the unique logarithmic form given by closure-multiplicity counting (Proposition 1, via Cauchy functional equation). Independently derived as Theorem A in the Commitment Barrier companion paper.
- *The second law* holds structurally: commitment events strictly increase entropy (Lemma 1, Proposition 2, corrected proof)
- *Thermodynamics* (equilibrium distribution, temperature, first law) follows from entropy maximisation under constraints (Section 4)
- *A commitment boundary* must exist given the simultaneous existence of reversible dynamics and physical facts (Proposition 3)
- *The arrow of time* follows from the irreversibility of commitment ordering (Proposition 7, from Lemma 1)
- *The existence* of a commitment functional satisfying monotonicity, interference penalisation, permutation invariance, extensivity, and divergence under indistinguishability (Section 8.2 — structural constraints only)
- *One bit per minimal causal region*: the No Multi-Primitive Occupancy Theorem (companion paper) proves by exhaustion over the observable algebra that a coherence cell at threshold supports exactly one independently attributable primitive commitment event — no more can coexist without violating primitive minimality, operational resolvability, or the discrete entropy spectrum
- *Primitive entropy quantum* $\Theta_0 = k_B \ln 2$: binary minimality (the smallest irreversible refinement is $1 \rightarrow 2$) combined with the forced logarithmic form gives the primitive entropy cost of one commitment event (Commitment Barrier companion, Lemma B and Corollary C; conditional on physical entropy identification)

Tier 2 — Established via companion papers plus A1–A3 plus Taylor Limit:

- *Algebraic reversibility of the pre-factual sector* (IAC: every non-null pre-factual contribution admits a cancellation partner): proven in the IAC and PAR/CC companion papers. PAR + CC \rightarrow IAC is proven in both companions (IAC Theorem 7.1; PAR/CC Theorem 3.1). The PAR/CC paper additionally provides three independent routes to deriving PAR from deeper conditions: (1) Admissibility + the Pre-Factual Symmetry Condition (no directional asymmetry before any fact is grounded), Theorem A.2; (2) Admissibility + Observational Faithfulness (no pre-factual information loss), Theorem A.4; (3) The Landauer-type Pre-Factual No-Erasure Condition: genuine irreversibility requires a cost-bearing outlet; before facts form no such outlet exists; therefore no pre-factual irreversibility, Theorem B.2. All three routes converge on PAR from distinct physical directions — directional asymmetry, distinguishability, thermodynamic cost — providing strong independent grounds that PAR is a structural necessity. IAC + Frobenius restricts the amplitude field to $\{\mathbb{R}, \mathbb{C}, \mathbb{H}\}$. A2 alone does not derive PAR (explicit counterexample: PAR/CC Proposition 4.1).
- *Complex amplitude field \mathbb{C}* is the unique viable amplitude field: proven in CHD Theorem 1 (Galois invariance eliminates \mathbb{R} and \mathbb{H} from the Frobenius candidates) and independently via the Galois-theoretic argument of Section 6
- *Hilbert space structure* is the unique G-invariant inner product space: proven in CHD Theorem 2
- *The Schrödinger equation* follows from Stone's theorem applied to the resulting Hilbert space
- *The Born rule* $P(i) = |\langle i|\psi\rangle|^2$ is established by two independent routes: (a) the Double Square Rule (DSR Theorem 5.9) derives it prior to Hilbert space from pairwise kernel uniqueness — this route works in all dimensions without invoking Gleason and resolves the dim=2 issue; (b) as a consistency check, Gleason's theorem (dim ≥ 3) and Busch's POVM extension (all dimensions) confirm the result given Hilbert structure. The DSR derivation is the primary route; Gleason/Busch provide independent confirmation.

Status note. The Taylor Limit (analyticity and continuity of probability functionals) remains an explicit additional assumption beyond A1–A3. Given the Taylor Limit, the derivation from A1–A3 to complex Hilbert space and Born rule is complete. The Physical Necessity companion (PN) justifies each structural ingredient of the DSR derivation by showing alternatives are physically inadmissible.

Tier 2 — BCB instantiation (A1–A3 plus specific functional form plus calibration):

- *Tetrahedral geometry selection* over square-planar for four equivalent constraints (Section 8.3)
- *Chemical structure rules* — valence limits, aromaticity, hypervalency — as capacity constraints
- *Mechanism transitions and OER* predictions
- *Protein folding kinetics* via LCP

Introduced and applied:

- *The proto-factual / factual distinction* (Section 2.2): a domain-independent criterion for when physical dynamics produces stable outcomes
- *The commitment functional B_R* (Section 8.2): with the explicit acknowledgment that A1–A3 constrain its admissibility class but do not uniquely determine its form
- *Extension to chemical, biological, and quantum threshold systems* (Sections 8.5, 11)

13.2 What Remains Conditional

- **Derivation of C_R from first principles.** A3 guarantees the existence of a finite local capacity; its value requires calibration. A microscopic derivation — from Landauer-type arguments about distinguishability lifetimes under decoherence — would complete the foundational picture but is not attempted here.
- **Uniqueness of the commitment functional form.** The additive structure $B_R = \sum b_i + \lambda \sum \Phi(\theta_{ij})$ and the choice of interference penalty Φ are motivated but not uniquely proven. The BCB companion paper demonstrates robustness across alternative penalty functions; a uniqueness theorem analogous to Proposition 1 remains an open problem.
- **The closure cost scale β .** For chemical systems, the aromaticity closure penalty $B_{cl} = \beta \ln g$ is constrained ($\beta \approx 1.5\text{--}3$) but not precisely calibrated, pending a first-principles derivation of the delocalization benefit.
- **Relativistic extension.** The fold interface has been described non-relativistically. A covariant formulation is required to connect with quantum field theory and cosmological contexts.
- **Calibration of the thermodynamic scale.** The identification of closure entropy with thermodynamic entropy at the scale of k_B involves a calibration step presupposing contact with a thermal bath. This is standard but not derived from first principles here.
- **The threshold condition $B_R = C_R$ is structurally characterized but not yet fully quantified in physical units.** The structural form is established: B_R is the accumulated distinguishability cost of maintaining open alternatives in a bounded region; C_R is the finite local capacity guaranteed by A3; and the threshold is crossed when B_R reaches C_R . The dimensional form is $\chi(\xi) = \rho L^3 / \hbar c = 1$, giving the coherence scale $\xi = (\hbar c / \rho)^{1/3}$. However, the absolute value of ξ in SI units — whether the coherence scale corresponds to the Planck length, the electroweak scale, or something else — is not pinned within this paper.

Two independent calibration routes exist in companion papers. First, the chemical route: for a second-row atomic valence region (carbon), the BCB companion paper calibrates $C_C \approx 16$ effective commitment units, fixed by the hypervalency boundary between CF_6 (non-admissible) and SF_6 (admissible). This connects C_R to measurable atomic properties but does not directly give ξ in SI units. Second, the electromagnetic route: the Interface companion paper (*Taylor, in preparation*) connects ξ to the Z-boson mass scale through the bare coupling $g_0^{-2} = 2^7 = 128$, suggesting $\xi \approx \hbar c / m_Z c^2 \approx 2 \times 10^{-18}$ m at leading order. This identification is preliminary — formally establishing the relationship between ξ and the electroweak scale is designated as open target P13 in that paper.

Until both calibration routes converge on a single consistent value of ξ , the threshold is quantified domain-by-domain (chemical valence: $C_C \approx 16$; electromagnetic: ξ at

electroweak scale) rather than from first principles. This is an open problem and is not concealed as one.

Consequences and reframings. Several results that appear elsewhere as separate empirical facts or interpretive puzzles follow as consequences of the framework and are worth stating explicitly:

The transition state is not observable because it is not a fact. The impossibility of directly observing a transition state — as opposed to observing its temporal neighbourhood via femtosecond spectroscopy — is not a technical limitation of current instrumentation. It is a structural consequence of the transition state failing the attribution and coherence conditions for facthood. Transition-state theory's dividing surface is real and predictively powerful, but it describes the geometry of the proto-factual zone, not the location of a committed event. Femtochemical observations are observations of proto-factual evolution, not of the facts that emerge from it.

Quantum coherence in chemistry is proto-factual coherence. Electronic and vibrational coherences observed in ultrafast spectroscopy of photosynthetic systems, molecular electronics, and potentially olfaction reflect coherent evolution within the proto-factual regime — quantum mechanical superpositions that have not yet achieved environmental decoupling and committed outcome selection. Their decoherence marks the transition from proto-factual to factual status. These coherences are physically significant precisely because they operate below the commitment threshold, allowing the system to explore configuration space before committing. This reframes the quantum biology debate: the question is not whether quantum coherence exists in biological systems (it does, in the proto-factual regime) but whether it persists long enough to influence which committed outcome is produced.

Reaction mechanisms describe proto-factual evolution, not sequences of facts. A reaction mechanism — arrow-pushing in organic chemistry, a potential energy surface in computational chemistry, a free-energy diagram in physical chemistry — is a description of how proto-factual configurations evolve toward committed outcomes. The mechanism is not a sequence of facts; it is a map of the proto-factual landscape that determines which committed outcomes are produced, and with what probabilities.

13.3 Relationship to Prior Work

This framework does not claim to derive results unknown to physics. It provides a unified derivation pathway and positions that pathway honestly relative to prior work.

Entropy. The logarithmic form is standard (Boltzmann, Khinchin). The novelty is its grounding in closure-record structure rather than observer uncertainty or phase-space measure. The Cauchy functional equation argument follows Khinchin's uniqueness theorem.

Second law. The structural (non-statistical) derivation connects to the resource-theoretic treatments of irreversibility (Brandão et al., 2015), which show that the second law follows from the structure of quantum operations without invoking statistical assumptions. Lemma 1 provides the more primitive version of this argument.

Hilbert space. The reconstruction programme of Hardy (2001) and Chiribella–D'Ariano–Perinotti (2011) is the state of the art. Our contribution is showing that three operational axioms — finite distinguishability, irreversible commitment, finite capacity — are sufficient to entail their hypotheses, reducing the axiomatic base.

Born rule. Gleason's theorem (1957) is the standard result. Our contribution is its placement within a commitment-based measurement theory that identifies measurement with threshold-crossing rather than collapse.

Measurement. The identification of measurement with decoherence-driven environmental embedding connects to Zurek's einselection programme (2003). The commitment functional B_R gives this a concrete multi-component form, distinguishing entropy export from the additional requirements of causal closure and action budget.

Comparison with standard decoherence. Standard decoherence is characterised by suppression of off-diagonal density matrix elements: $\rho_{ij}(t) = \rho_{ij}(0) \cdot e^{(-\Gamma_{ij} t)}$, where Γ_{ij} is the decoherence rate from environmental coupling. Classicality emerges continuously as off-diagonals decay. The commitment framework adds three structural requirements beyond entropy export: (i) causal closure — the outcome must be independently attributable; (ii) action budget — sufficient phase accumulated; (iii) a threshold — B_R must reach C_R , not merely approach it.

The two frameworks agree wherever decoherence is fast relative to all other timescales — in macroscopic room-temperature systems, the threshold is exceeded almost instantaneously. They diverge in three regimes: (a) systems near the commitment boundary, where the commitment framework predicts discrete threshold behaviour while decoherence predicts smooth exponential decay; (b) the Quantum Zeno effect, where interruption prevents B_R from accumulating rather than reversing already-occurring decoherence; (c) enzyme active sites, where fast decoherence would predict suppression of tunnelling but the commitment framework predicts its persistence because active-site preorganisation drives $B_R \geq C_R$ before decoherence suppresses the quantum amplitude. These three cases define the distinguishing experimental programme.

Consistent histories and relational quantum mechanics. The consistent histories programme (Griffiths 2002; Omnès 1994) addresses when a quantum history constitutes a definite outcome through the consistency condition on history sets. The commitment framework differs in being operational: instead of asking which sets of histories are logically consistent, it asks which processes produce irreversible records. The two approaches are likely compatible but have not been formally connected. Rovelli's relational quantum mechanics (Rovelli 1996) locates facts at interactions between systems and observers, with facts being observer-relative. The commitment framework shares the spirit of this approach but provides a physical threshold criterion ($B_R = C_R$) for when an interaction constitutes a fact-forming event. Connecting the commitment threshold to both programmes is a direction for future work.

Chemical rules. The BCB framework (Taylor, in preparation) connects to Landau–Ginzburg theory as an effective admissibility theory: BCB captures universal structural admissibility from distinguishability constraints without specifying the microscopic capacity mechanism. The

relation to VSEPR geometry is explicit: VSEPR is recovered as the phenomenology of minimum-interference packings; BCB specifies the capacity constraint that makes those packings survive.

Holographic and distinguishability bounds. The finite localisation capacity (A3) is consistent with and bounded above by established entropy bounds. The holographic entropy bound of the causal patch, $\mathcal{L} = \pi(R_U/\ell_P)^2 \approx 2.3 \times 10^{123}$, established by 't Hooft (1993) and Susskind (1995), is the geometric ceiling on distinguishable states within any causal horizon. A companion paper (Taylor, *The Operational Structure of Physical Distinguishability*, in preparation) reframes this as a bound on operationally distinguishable states — states separable by any physical process requiring finite energy, time, and entropy production — and identifies a three-layer hierarchy: the holographic ceiling ($\sim 10^{123}$) as geometric maximum state-space capacity; the Lloyd instantiation limit ($\sim 10^{120}$) as the count of configurations physically realised over cosmic history; and the Landauer access floor ($\sim 10^{90}$ at current CMB temperature) as the operationally accessible count under current thermodynamic conditions. These three layers are not competing approximations but answers to different questions: capacity, history, and current access respectively. The commitment threshold $B_R = C_R$ of the present paper operates locally, at the scale of individual bounded systems, far below the causal-patch ceiling; A3 is consistent with all three layers.

Physical coupling constants. A companion paper (*Interface Realization and Physical Constants*, Taylor, in preparation) extends the framework to the gauge-invariant interface at which pre-factual commitments become physical facts. The structural conditions of the present paper — binary irreducibility, minimal commitment cost, algebraic reversibility, and compositional completeness — when applied at this interface, determine the electromagnetic coupling strength without continuous free parameters. The leading-order structural value $g_{\text{eff}}^{-2} = 2^7 \cdot (14+1)/14 \approx 137.14$ matches the observed inverse fine structure constant to within 0.08%, with the residual identified as a first-order approximation gap. Coupling constants that the standard model treats as unexplained empirical inputs emerge as structural outputs of the interface architecture. Extensions to non-Abelian gauge groups (weak and strong couplings) are designated as open targets in that paper.

The principal novel contribution is the *unified derivation chain* together with the *commitment lag* as an experimentally accessible quantity: the demonstration that three operational constraints suffice to ground entropy, thermodynamics, quantum theory, time, and the domain-general structure of fact formation simultaneously, and that the lag between dynamical transition and fact formation is a physically real, independently measurable quantity absent from all prior frameworks.

13.4 Anticipated Objections

The following objections are those a referee at a foundations journal is most likely to raise. Each is addressed in turn, with a statement of what is true in the objection and what it does not establish.

O1. "This is decoherence in new language."

What is true. Decoherence theory and the commitment framework share the core insight that environmental coupling drives the transition from quantum to effectively classical behaviour. Both are environment-driven, both scale with environmental degrees of freedom, and both produce the appearance of definite outcomes at macroscopic scales.

What the objection does not establish. The frameworks make different predictions in three regimes. First, decoherence predicts continuous exponential suppression of off-diagonal density matrix elements; the commitment framework predicts a threshold at $B_R = C_R$ where fact formation occurs discontinuously. In systems near the commitment boundary, these predictions diverge and are in principle distinguishable by precision measurement of outcome statistics near threshold. Second, in the Quantum Zeno effect, decoherence accounts predict that repeated measurement reverses accumulated decoherence; the commitment framework predicts that repeated measurement prevents B_R from accumulating toward C_R — a different mechanism that predicts partial suppression when measurement coupling is weak (Section 8.5). Third, in enzyme active sites, decoherence theory predicts rapid quantum-to-classical transition in warm wet environments; the commitment framework predicts persistence of tunnelling contributions because active-site preorganisation drives $B_R \geq C_R$ before decoherence suppresses the quantum amplitude. At least one of these three divergences is testable with current technology.

O2. " B_R is a free parameter. The framework can accommodate any observation by adjusting it."

What is true. The specific additive-pairwise form of B_R is not uniquely derived from A1–A3; it is the minimal representative of the admissibility class. The parameters (b_i, λ, Φ) require calibration from physical data and are not predicted from first principles.

What the objection does not establish. The framework is not unfalsifiable. A1–A3 constrain B_R to a specific admissibility class — it must satisfy monotonicity, interference penalisation, permutation invariance, extensivity at low density, and divergence under indistinguishability. Not every functional satisfies these constraints. Furthermore, C_R is not freely adjustable post hoc: for chemical systems it is anchored to two independent empirical constraints (the CF_6/SF_6 existence boundary and the H_2O bond angle) from which it is uniquely bracketed. Once calibrated in one domain, the parameters make predictions in others. The BCB companion paper (Taylor, in preparation) demonstrates predictions on 11 held-out chemical species with 100% classification accuracy at the calibrated parameter values — this result appears in the BCB companion, not the present paper, and a referee should consult that paper to assess it directly. The question is not whether B_R has free parameters — all effective theories do — but whether those parameters are independently constrained and whether the resulting predictions are falsifiable. They are.

O3. "The Hilbert space derivation in Sections 6–7 is incomplete. Section 6 doesn't prove what it claims and Section 7 doesn't establish that A1–A3 fully entail Hardy's axioms."

What is true in the original version of this objection. Previous versions of Sections 6–7 had this gap: Section 6 provided motivation without proof, and Section 7 acknowledged that two of Hardy's five axioms were not fully entailed by A1–A3.

Current status. The Hilbert space gap is closed by CHD (Sections 6–7). The Born rule gap — specifically the $\dim=2$ issue and the tensor product derivation — is closed by the Double Square Rule companion (DSR). DSR derives $P = |\psi|^2$ prior to Hilbert space, from pairwise kernel uniqueness forced by gauge invariance, factorisation, and positivity. This derivation works in all dimensions without invoking Gleason and produces the tensor product structure from the factorisation constraint rather than assuming it. The two companion papers together (CHD for Hilbert space, DSR for Born rule) close all the gaps that the referee correctly identified in previous versions.

What the objection establishes going forward. The Taylor Limit remains an explicit additional assumption beyond A1–A3. The Physical Necessity companion (PN) attempts to justify it through the TPB \rightarrow continuous holonomy argument (PN Appendix C). That argument has two layers. Layer 1 — proving that purely finite holonomy is incompatible with TPB, bit conservation, and temporal extensibility — is a complete proof (PN Appendix C, Section C.4). Layer 2 — upgrading "infinite" to "continuous" holonomy, and identifying the minimal admissible case as $U(1)$ — has four specific technical points that remain to be completed to journal-standard rigor. These are stated explicitly in PN Section C.8:

- (i) Formal definition of the operational topology d_{op} as a metric, and proof that A7 (total boundedness) follows from BCB finite distinguishability
- (ii) Rigorous proof that compositions of near-inverse loops produce holonomies within $\varepsilon \cdot \delta_{step}$ of identity, and that \bar{H} is a locally compact metrizable topological group
- (iii) Formal proof that independent holonomy dimensions correspond to independent commitment channels, establishing $d_H = 1$ from TPB irreducibility
- (iv) Rigorous proof that BCB finite distinguishability forces the holonomy connected component to be compact

A referee should assess these for themselves. They are technical completions of a structurally outlined argument, not conceptual gaps — Layer 1 establishes that holonomy must be infinite, and the logical architecture for why it must be continuous is in place. But they are real outstanding items, and a referee is right to treat them as such. The Taylor Limit therefore rests on: (a) a complete impossibility theorem (Layer 1) showing finite holonomy is physically inadmissible, and (b) a partially completed continuity upgrade (Layer 2) identifying $U(1)$ as the minimal continuous case. Until Layer 2 is complete, the Taylor Limit's justification is partial — it is well-motivated and the core necessity argument is proven, but the specific identification with $U(1)$ carries four unclosed technical obligations, which a dedicated paper (*Why Finite Distinguishability Forces Continuous Phase*, Taylor, in preparation) will complete.

O4. "The commitment lag is not distinguishable from standard Kramers theory. You are just relabelling frictional effects."

What is true. In the limit where environmental reorganisation is fast and strong (τ_s small, η large), the commitment lag is negligible and the commitment framework and Kramers theory make identical predictions for chemical reaction rates.

What the objection does not establish. The frameworks diverge in the opposite limit — weak or slow embedding — and their divergence is structurally specific, not a matter of numerical adjustment. Kramers theory predicts that two solvents with identical friction coefficients produce identical rate constants. The commitment framework predicts that two solvents with identical Kramers friction — identical barrier-crossing rates — can produce different product yields if they differ in embedding capacity η . This is because the commitment framework identifies a second stage after barrier crossing (environmental stabilisation to $B_R \geq C_R$) that Kramers does not. The experimental test is direct: engineer two solvent environments with matched viscosity but different hydrogen-bonding capacity and measure product yield. Kramers predicts identical yields; the commitment framework predicts different yields correlated with η . No parameter adjustment can make Kramers reproduce this prediction because the prediction concerns a stage Kramers does not model.

O5. "The axioms A1–A3 are not derived — they are assumed. Why should we accept them?"

What is true. A1–A3 are axioms, not theorems. They are not derived from more primitive principles within this paper.

What the objection does not establish. The axioms are not assumptions about specific physical laws — they are preconditions for any universe capable of supporting stable, reproducible facts. A1 (finite distinguishability) is not a contingent physical fact; it is a requirement for the concept of "observable state" to be coherent. A2 (irreversible commitment) is not an assumption about quantum mechanics; it is the definition of what makes a difference a fact rather than a fluctuation. A3 (finite localisation capacity) is the operational content of the Bekenstein bound, which is independently supported by black hole thermodynamics and holographic arguments. The correct question to ask is not "are these assumptions?" — all frameworks have axioms — but "are there plausible universes that violate these conditions and still support stable physics?" If there are, this framework applies only to our kind of universe. If there are not, the axioms are necessary preconditions rather than contingent choices. We believe the latter, but acknowledge that establishing this claim rigorously is a further research task.

O6. "The framework makes no new predictions in quantum mechanics — the Zeno effect is already predicted by standard QM."

What is true. The existence of the Quantum Zeno effect is predicted by standard quantum mechanics via the measurement postulate. The Double Square Rule derivation of $P = |\psi|^2$ recovers standard quantum predictions in all currently accessible regimes.

What the objection does not establish. The framework makes several predictions that standard QM does not:

First, the degree of Zeno suppression should depend on measurement coupling strength independently of measurement frequency. Standard QM predicts complete projection for any measurement; the commitment framework predicts partial suppression when coupling is insufficient to reset B_R .

Second, the Double Square Rule companion (Section 9.9) predicts distinguishability-gradient decoherence: two systems with identical masses, coupling strengths, and environmental temperatures but different state-space geometries (e.g., coherent vs. squeezed motional states in optomechanics) should decohere at different rates, with the ratio $\Gamma_1/\Gamma_2 = |\nabla d|_1^2 / |\nabla d|_2^2$. Standard QM, GRW, CSL, and Diósi-Penrose all predict equal rates for these systems.

Third, the Physical Necessity companion (Section 9.1, Theorem 8.3) shows that any non-quadratic probability rule $P \propto |\psi|^p$ with $p \neq 2$ fails to preserve normalisation under reversible evolution. This provides a direct experimental constraint: if third-order interference is ever observed (Sorkin's hierarchy), the framework requires revision at that order. Null results on third-order interference (Sinha et al. 2010) are consistent predictions, not post-hoc accommodations.

The commitment framework gives a unified account of *why* the Zeno effect exists — sub-threshold interruption of commitment — that requires no measurement postulate. This is a novel theoretical contribution even if the phenomenology is already known.

O7. "The connection between the foundational physics (Sections 1–10) and the chemistry/biology (Section 11) is not derivational — it is analogical. The BCB formalism is a separate theory."

What is true. The BCB formalism involves calibrated parameters not derived from first principles, and constitutes an independent effective theory. The connection to A1–A3 is that BCB is a concrete instantiation of the admissibility class established by A1–A3, not a theorem of it.

What the objection does not establish. The paper does not claim that chemistry is derived from A1–A3. It claims that BCB is the minimal representative of the functional class that A1–A3 admit, and that applying this representative to chemical systems generates falsifiable predictions.

This is the same relationship Landau-Ginzburg theory bears to statistical mechanics: LG is not derived from the partition function, but it is the minimal effective theory consistent with the symmetry constraints, and its predictions are falsifiable. The analogy is explicit and the paper does not overclaim beyond it. The novel contribution of the chemistry sections is the *commitment lag* — a quantity not present in standard chemical kinetics, generated by the same threshold structure as the quantum applications, and testable in both domains simultaneously. This cross-domain falsifiability is what makes the connection more than analogical.

O8. "The framework is not Lorentz-covariant and cannot account for relativistic quantum field theory."

What is true. The framework as presented is non-relativistic. The fold interface is defined without reference to spacetime structure, and the commitment time t_c is a non-relativistic concept.

What the objection does not establish. Non-relativistic foundations are a standard starting point for reconstructing quantum theory (Hardy 2001, CDP 2011 both begin non-relativistically). The paper does not claim to derive QFT. The Lorentz-covariant extension — identifying what the commitment boundary becomes in a relativistic setting, and how t_c transforms — is explicitly listed as an open problem. A framework that successfully grounds thermodynamics and non-relativistic quantum mechanics from three axioms is a contribution even if QFT extension remains to be done.

14. Conclusion

We have shown that a universe capable of producing stable facts necessarily exhibits thermodynamics, quantum structure, and an arrow of time. These are not independent features of physics requiring independent postulates. They are consequences of a single underlying requirement:

The irreversible commitment of distinguishability in a finite-capacity system.

The three axioms A1–A3 are not laws of nature. They are structural preconditions for the existence of nature as a system of facts. From them, the major frameworks of physics — entropy, thermodynamics, quantum mechanics, measurement, time — follow as necessary structure.

The proto-factual / factual distinction, and the commitment functional $B_R \leq C_R$ that governs it, extend this structure beyond fundamental physics to chemistry, catalysis, and biology. Chemical structure rules, mechanism transitions, and biological irreversibility all instantiate the same threshold: the point at which reversible possibility becomes irreversible record.

This suggests that the apparent plurality of physical foundations — the separate postulation of entropy, quantum structure, and time — is an artefact of their historical discovery in isolation. There is a single structural principle underneath, and it is simpler than any of the frameworks it generates.

A companion paper (*Interface Realization and Physical Constants*, Taylor, in preparation) applies the same structural conditions at the gauge-invariant interface level and derives a leading-order expression for the electromagnetic coupling strength without continuous free parameters, matching the observed inverse fine structure constant to within 0.08%. This result belongs to that companion paper and is noted here only as evidence that the structural conditions of the present paper have consequences beyond quantum mechanics — the scope of those consequences is an open question.

15. References

Thermodynamics and Entropy

- Boltzmann, L. (1877). Über die Beziehung zwischen dem zweiten Hauptsatze der mechanischen Wärmetheorie. *Wien. Ber.*, 76, 373–435.
- Khinchin, A.I. (1957). *Mathematical Foundations of Information Theory*. Dover.
- Bekenstein, J.D. (1973). Black holes and entropy. *Physical Review D*, 7(8), 2333.
- Brandão, F.G.S.L., et al. (2015). The second laws of quantum thermodynamics. *PNAS*, 112(11), 3275–3279.

Philosophy of Time

- Price, H. (1996). *Time's Arrow and Archimedes' Point*. Oxford University Press.
- Barbour, J. (1999). *The End of Time*. Oxford University Press.
- Callender, C. (2017). *What Makes Time Special?* Oxford University Press.
- Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5, 183–191.

Quantum Reconstruction

- Hardy, L. (2001). Quantum theory from five reasonable axioms. arXiv:quant-ph/0101012.
- Chiribella, G., D'Ariano, G.M., Perinotti, P. (2011). Informational derivation of quantum theory. *Physical Review A*, 84, 012311.
- Gleason, A.M. (1957). Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics*, 6(6), 885–893.
- Busch, P. (2003). Quantum states and generalized observables: a simple proof of Gleason's theorem. *Physical Review Letters*, 91, 120403.
- Stone, M.H. (1932). On one-parameter unitary groups in Hilbert space. *Annals of Mathematics*, 33(3), 643–648.

- Frobenius, F.G. (1878). Über lineare Substitutionen und bilineare Formen. *Journal für die reine und angewandte Mathematik*, 84, 1–63.
- Solèr, M.P. (1995). Characterization of Hilbert spaces by orthomodular spaces. *Communications in Algebra*, 23(1), 219–243.
- Dakić, B.; Brukner, Č. (2011). Quantum theory and beyond: Is entanglement special? In H. Halvorson (Ed.), *Deep Beauty*. Cambridge University Press.
- Sorkin, R.D. (1994). Quantum mechanics as quantum measure theory. *Modern Physics Letters A*, 9, 3119–3127.
- Sinha, U., et al. (2010). Ruling out multi-order interference in quantum mechanics. *Science*, 329, 418–421.

Measurement and Decoherence

- Zurek, W.H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75, 715.
- Cox, D.R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society B*, 34(2), 187–220.
- Griffiths, R.B. (2002). *Consistent Quantum Theory*. Cambridge University Press.
- Omnès, R. (1994). *The Interpretation of Quantum Mechanics*. Princeton University Press.
- Rovelli, C. (1996). Relational quantum mechanics. *International Journal of Theoretical Physics*, 35, 1637–1678.
- Cramer, J.G. (1986). The transactional interpretation of quantum mechanics. *Reviews of Modern Physics*, 58, 647.

Chemical Structure

- Gillespie, R.J.; Nyholm, R.S. (1957). Inorganic stereochemistry. *Quarterly Reviews, Chemical Society*, 11, 339–380.
- Hückel, E. (1931). Quantentheoretische Beiträge zum Benzolproblem. *Zeitschrift für Physik*, 70, 204–286.
- Woodward, R.B.; Hoffmann, R. (1969). The conservation of orbital symmetry. *Angewandte Chemie International Edition*, 8, 781–853.
- Reed, A.E.; Schleyer, P.v.R. (1990). Chemical bonding in hypervalent molecules. *JACS*, 112, 1434–1445.

OER and Catalysis

- Grimaud, A., et al. (2017). Activating lattice oxygen redox reactions in metal oxides. *Nature Chemistry*, 9, 457–465.
- Zagalskaya, A.; Alexandrov, V. (2020). Role of defects in the interplay between adsorbate evolving and lattice oxygen mechanisms. *ACS Catalysis*, 10, 3650–3657.
- Dionigi, F., et al. (2020). In-situ structure and catalytic mechanism of NiFe layered double hydroxides. *Nature Communications*, 11, 2522.

Quantum Biology

- Lambert, N., et al. (2013). Quantum biology. *Nature Physics*, 9, 10–18.

Companion Paper

- Taylor, K. (in preparation). Predicting chemical rules from Bit Conservation and Balance (BCB). VERSF Theoretical Physics Programme, AIDA Institute.
- Taylor, K. (in preparation). Protein folding as distinguishability collapse: a Bit Conservation framework for folding kinetics. VERSF Theoretical Physics Programme, AIDA Institute.
- Taylor, K. (in preparation). Proto-factual structure and fact formation in chemical reactions. VERSF Theoretical Physics Programme, AIDA Institute.
- Taylor, K. (in preparation). Complex Hilbert space from distinguishability principles. VERSF Theoretical Physics Programme, AIDA Institute. [CHD]
- Taylor, K. (in preparation). The Double Square Rule: a derivation of quantum probability from discrete informational geometry. VERSF Theoretical Physics Programme, AIDA Institute. [DSR]
- Taylor, K. (in preparation). Physical necessity of quantum probability structure. VERSF Theoretical Physics Programme, AIDA Institute. [PN]
- Taylor, K. (in preparation). Quantum measurement as a tick race: deterministic outcome selection via first-passage dynamics. VERSF Theoretical Physics Programme, AIDA Institute. [Tick-Bit]
- Taylor, K. (in preparation). Deriving the commitment-capacity density from primitive commitment structure. VERSF Theoretical Physics Programme, AIDA Institute. [No Multi-Primitive Occupancy]
- Taylor, K. (in preparation). Interface realization and physical constants: structural consequences of the VERSF framework. VERSF Theoretical Physics Programme, AIDA Institute. [Interface/Coupling]
- Taylor, K. (in preparation). Internal admissible closure from pre-factual reversibility in fact-producing universes. VERSF Theoretical Physics Programme, AIDA Institute. [IAC]
- Taylor, K. (in preparation). Closing the structural gaps: marginal compositional consistency and algebraic closure in fact-producing universes. VERSF Theoretical Physics Programme, AIDA Institute. [CSG]
- Taylor, K. (in preparation). Multiple structural derivations of the mesoscopic coherence scale ξ_{meso} : overdetermination of the two-Planck window from independent closure principles. VERSF Theoretical Physics Programme, AIDA Institute. [Mesoscopic]
- Taylor, K. (in preparation). Taylor's Number: the computational boundary of physical reality. VERSF Theoretical Physics Programme, AIDA Institute. [Computational Boundary]
- Taylor, K. (in preparation). The operational structure of physical distinguishability: reframing, convergence, and hierarchy at the causal-patch bound. VERSF Theoretical Physics Programme, AIDA Institute. [Distinguishability Hierarchy]
- Taylor, K. (in preparation). On the structural status of algebraic reversibility and compositional completeness in fact-producing universes. VERSF Theoretical Physics Programme, AIDA Institute. [PAR/CC]

- Taylor, K. (in preparation). Deriving the commitment barrier from closure entropy and the coherence scale. VERSF Theoretical Physics Programme, AIDA Institute. [Commitment Barrier]

Protein Folding

- Plaxco, K.W., Simons, K.T., Baker, D. (1998). Contact order, transition state placement and the refolding rates of single domain proteins. *Journal of Molecular Biology*, 277, 985–994.
- Itzhaki, L.S., Otzen, D.E., Fersht, A.R. (1995). The structure of the transition state for folding of chymotrypsin inhibitor 2 analysed by protein engineering methods. *Journal of Molecular Biology*, 254, 260–288.
- Levinthal, C. (1969). How to fold graciously. *Mössbauer Spectroscopy in Biological Systems*, 67–69.
- Bryngelson, J.D., et al. (1995). Funnels, pathways, and the energy landscape of protein folding. *Proteins*, 21, 167–195.

Femtochemistry and Reaction Dynamics

- Zewail, A.H. (2000). Femtochemistry: Atomic-scale dynamics of the chemical bond. *Journal of Physical Chemistry A*, 104(24), 5660–5694.
- Polanyi, J.C.; Zewail, A.H. (1995). Direct observation of the transition state. *Accounts of Chemical Research*, 28(3), 119–132.
- Tolbert, L.M.; Solntsev, K.M. (2002). Excited-state proton transfer: From constrained systems to "super" photoacids to superfast proton transfer. *Accounts of Chemical Research*, 35(1), 19–27.
- Pérez-Lustres, J.L., et al. (2007). Ultrafast proton transfer to solvent: Molecularity and intermediates from solvation- and diffusion-controlled regimes. *Journal of the American Chemical Society*, 129(19), 5408–5409.
- Klinman, J.P.; Kohen, A. (2013). Hydrogen tunneling links protein dynamics to enzyme catalysis. *Annual Review of Biochemistry*, 82, 471–496.
- Hammes-Schiffer, S. (2006). Hydrogen tunneling and protein motion in enzyme reactions. *Accounts of Chemical Research*, 39(2), 93–100.
- Scrutton, N.S.; Hay, S. (2007). Probing coupled motions in enzymatic hydrogen tunnelling reactions. *Biochemical Society Transactions*, 35, 1350–1355.
- Hawkinson, D.C.; Eames, T.C.M.; Pollack, R.M. (1991). Energetics of 3-oxo- Δ^5 -steroid isomerase: source of the catalytic power of the enzyme. *Biochemistry*, 30(43), 10498–10506.
- Fried, S.D.; Boxer, S.G. (2017). Electric fields and enzyme catalysis. *Annual Review of Biochemistry*, 86, 387–415.
- Layfield, J.P.; Hammes-Schiffer, S. (2013). Hydrogen tunneling in enzymes and biomimetic models. *Chemical Reviews*, 114(7), 3466–3494.