

Commitment Lag and Enzyme Catalysis:

A First-Principles Derivation from the VERSF Commitment Framework

Keith Taylor

VERSF Theoretical Physics Program

For the General Reader

Physics gives us two descriptions of the world, and no satisfying account of how one becomes the other.

Below a certain scale, quantum mechanics applies: particles exist in superpositions of possibilities, outcomes are genuinely undetermined, and nothing is definite until something forces a choice. Above that scale, the classical world applies: objects have positions, reactions have products, facts exist. Between the two sits one of the deepest unsolved problems in science — what turns a possibility into a fact?

This paper proposes an answer, and uses it to explain something that has puzzled biochemists for decades: why enzymes can harness quantum tunnelling in warm, wet, noisy environments where quantum effects should vanish almost instantly.

The key idea: not all events are facts yet.

We are accustomed to thinking of physical events as happening or not happening. But consider the moment a proton tunnels through an energy barrier inside an enzyme. The tunnelling happens — the proton moves — but for a brief window afterward, the outcome is not yet irreversibly recorded. The surrounding molecular scaffold has not yet reorganised to "lock in" what happened. The proton has crossed the barrier, but the universe has not yet committed to the fact that it crossed. We call this interval the *commitment lag* — the gap between a physical transition and the formation of an irreversible record.

During this window, the system is what we call *proto-factual*: dynamically real, physically happening, but not yet a committed fact. And this distinction matters enormously, because during the proto-factual interval, environmental noise can suppress the quantum signature of what just occurred — effectively erasing the evidence of tunnelling before it becomes permanent.

Enzymes as commitment machines.

What ketosteroid isomerase (KSI) appears to do — and what this paper derives from first principles — is collapse the commitment lag to near zero. Its precisely pre-positioned hydrogen-

bond scaffold is structured so that the moment a proton tunnels, the surrounding environment immediately locks in the outcome. There is no gap for decoherence to act. The quantum event becomes a committed fact almost simultaneously with its occurrence.

In bulk water, the same reaction has a commitment lag of roughly 1600 femtoseconds — nearly ten times longer. That gap is enough time for environmental noise to suppress the quantum amplitude substantially. The enzyme's catalytic advantage is not, at its deepest level, about lowering an energy barrier. It is about engineering *when* a quantum event becomes an irreversible fact.

Reinterpreting the world around us.

This framework invites a shift in how we think about physical reality. The classical and quantum worlds are not two separate domains separated by a mysterious boundary. They are two statuses that the same physical process can occupy: proto-factual (still open, still quantum) and committed (irreversibly recorded, classical).

Everything happening right now is proto-factual until it commits. A proton tunnelling inside an enzyme, a molecule vibrating at the edge of a reaction, a bond straining toward its breaking point — each of these exists in the proto-factual regime until the surrounding environment stabilises the outcome above a structural threshold. The "now" that we experience is the rolling boundary between the committed past (facts that have accumulated irreversibly) and the proto-factual present (events still in progress, outcomes still open).

The quantum world is not small and strange and distant. It is the condition of any physical process that has not yet committed — including many that are happening inside your cells right now.

Abstract

We present a first-principles derivation of the **commitment lag** — the physically real, environment-dependent separation between dynamical transition and irreversible fact formation — in enzyme-catalysed proton transfer, using the Void Energy-Regulated Space Framework (VERSF). Starting from three operational axioms (finite distinguishability A1, irreversible commitment A2, finite localisation capacity A3) and the $K = 7$ constraint dimensionality established in companion work [1], we derive the commitment functional B_R and local capacity C_R for the ketosteroid isomerase (KSI) active site and bulk aqueous solvent.

The commitment time formula

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

yields $t_c(\text{KSI}) \approx 173 \text{ fs}$ and $t_c(\text{water}) \approx 1600 \text{ fs}$, a $\approx 9\times$ **collapse** of the commitment time in the enzyme active site (central estimate; angular and τ_s uncertainties span approximately 4–20 \times).

This collapse follows directly from the pre-positioned Tyr14/Tyr55 hydrogen-bond scaffold geometry. The structurally predicted lag difference of ≈ 1427 fs is consistent with experimentally and computationally inferred tunnelling enhancements in KSI in the range of approximately 5–10 \times (with $\sim 8\times$ as a representative central estimate) at a decoherence rate ($\Gamma_{\text{dec}} \approx 1.5$ THz) lying within the independently established range for aqueous H-bond dynamics; the framework does not predict the enhancement ratio precisely but constrains it to within the experimentally accessible regime. No parameters are fitted within the commitment framework; the commitment ratio $\rho(S)$ is determined entirely by $K = 7$ and crystal structure geometry, with $\tau_s(S)$ taken from independent ultrafast spectroscopy (estimated factor-of-2–3 uncertainty; see §7.4) and not fitted to the KSI tunnelling data. The framework predicts that Tyr14Phe mutations drive the active site into a sub-threshold regime ($\eta < C_{\text{R}} - B_{\text{geom}}$), eliminating fast commitment and providing a structural account of the observed $\approx 10^3$ – 10^4 -fold excess rate reduction in that mutant beyond barrier-height predictions alone. These results provide the first structurally derived, non-fitted account within this class of models of anomalous tunnelling persistence in a warm enzymatic environment.

Contents

- 1. Introduction
 - 2. The VERSF Commitment Framework
 - 2.1 Three Operational Axioms
 - 2.2 The Commitment Functional
 - 2.3 Derivation of $\Phi(\theta)$ from Distinguishability Geometry
 - 2.4 Natural Units and the Primitive Commitment Energy
 - 2.5 Local Capacity from $K = 7$
 - 2.6 The Commitment Time Formula
 - 3. First-Principles Derivation
 - 3.1 Product Geometry: $B_{geom}(q, P)$
 - 3.2 Embedding Capacity η : KSI Active Site
 - 3.3 Embedding Capacity η : Bulk Water
 - 3.4 Commitment Ratios and Robustness
 - 3.5 Commitment Times
 - 4. From Commitment Lag to Tunnelling Enhancement
 - 4.1 The Commitment–Tunnelling Mechanism
 - 4.2 Quantitative Consistency with Observed Enhancement
 - 5. Additional Experimental Signatures
 - 5.1 Temperature Dependence
 - 5.2 Primary Kinetic Isotope Effect
 - 6. Testable Prediction: Tyr14Phe Mutation
 - 7. Discussion
 - 7.1 Relationship to Decoherence Theory
 - 7.2 Relationship to Standard Rate Theory
 - 7.3 What This Result Establishes
 - 7.4 Limitations and Open Questions
 - 7.5 Anticipated Objections
 - 8. Conclusion
 - References
-

1. Introduction

Ketosteroid isomerase (KSI) catalyses the isomerisation of Δ^5 -3-ketosteroids through a two-step proton transfer mechanism and represents one of the most extensively characterised enzymatic systems in physical biochemistry. Its catalytic efficiency approaches the diffusion-controlled limit, yet several features of its kinetics resist explanation within standard transition-state theory:

- Primary kinetic isotope effects ($KIE\ k_H/k_D \approx 6$) substantially exceeding classical predictions (~ 2 – 3)
- Anomalously weak temperature dependence inconsistent with simple Arrhenius behaviour

- Rate enhancements that cannot be fully accounted for by electrostatic transition-state stabilisation alone
- Persistent quantum tunnelling in a warm, aqueous environment where coherence would naively decohere rapidly

Existing theoretical accounts attribute these observations to transition-state geometry, protein conformational pre-organisation, or electrostatic preorganisation [4, 5, 6]. While each captures part of the phenomenology, none provides a unified structural account derived from a more primitive foundation.

The VERSF framework [2] proposes that physical facts are irreversible commitments in finite-capacity systems, and that the quantum-to-classical transition is governed by a structural threshold condition $B_R \geq C_R$. A central prediction of this framework is the **commitment lag** — a physically real, environment-dependent separation between dynamical transition and irreversible fact formation — that is absent from all prior frameworks. Section 11.1.3 of Reference [2] establishes the qualitative account of enzyme catalysis as commitment engineering. The present paper makes this account fully quantitative for KSI.

We derive the commitment lag for KSI and bulk water entirely from first principles, using only published crystal structure geometry and the $K = 7$ constraint dimensionality. We then derive the mechanism connecting commitment time to tunnelling enhancement, and generate a testable prediction against existing mutagenesis data.

The central claim, stated plainly: **the catalytic advantage of KSI is not primarily barrier lowering, but temporal compression of the proto-factual window during which decoherence can act.** The enzyme does not protect coherence — it eliminates the gap between transfer and commitment before decoherence has time to suppress the quantum amplitude.

2. The VERSF Commitment Framework

2.1 Three Operational Axioms

The framework rests on three axioms that any universe capable of producing stable, reproducible facts must satisfy [2]:

- **A1** (Finite Distinguishability): A bounded region cannot support infinitely precise distinctions.
- **A2** (Irreversible Commitment): A distinction becomes a physical fact only when irreversibly recorded.
- **A3** (Finite Localisation Capacity): Any bounded system can sustain only finitely many independently attributable irreversible commitments.

2.2 The Commitment Functional

For a bounded region R supporting m directional distinguishability constraints, the commitment functional is (Reference [2], Section 8.2):

$$\mathbf{B}_R = \sum_i \mathbf{b}_i + \lambda \sum_{i < j} \Phi(\theta_{ij})$$

where \mathbf{b}_i is the base distinguishability cost of constraint i , θ_{ij} is the angular separation between constraints i and j , and $\Phi(\theta) = 1/(1 - \cos \theta)$ is the interference penalty function. The derivation of this specific form from the closure geometry of the framework is given in Section 2.3; it is not an arbitrary modelling choice but the unique minimal function consistent with the structural requirements identified there. This functional satisfies five structural constraints derived from A1–A3: monotonicity, interference penalisation, permutation invariance, extensivity at low density, and divergence under indistinguishability. The pairwise additive form is the minimal truncation of this admissibility class, and key results are shown to be robust across alternative penalty functions with the same divergence structure (Reference [2], Section 8.2).

A configuration achieves factual status — an irreversible record is constituted — when:

$$\mathbf{B}_R \geq \mathbf{C}_R$$

2.3 Derivation of the Compatibility Function $\Phi(\theta)$ from Distinguishability Geometry

The compatibility function $\Phi(\theta)$, introduced to quantify the pairwise cost of embedding multiple directional constraints within a finite local commitment structure, is not chosen ad hoc. It follows from four structural requirements already present in the framework: finite distinguishability, irreversible commitment, local difference geometry, and analytic minimality.

The relevant primitive is distinguishability, not raw angular separation. In the admissibility framework, irreversible commitment is possible only if physically relevant distinctions are finite and non-refinable. If distinctions remain indefinitely refinable, no outcome is final and no genuine commitment occurs. Accordingly, any local compatibility rule must ultimately measure whether two constraints can remain independently distinguishable within a bounded commitment structure.

The framework already fixes the natural local geometric metric. In the graph-based derivation of first-order dynamics, the primitive local operator is the difference across adjacent structure, $(B\psi)(e) = \psi(\text{head}(e)) - \psi(\text{tail}(e))$, and the canonical quadratic local cost is generated by the Laplacian $L = B^\dagger B$, equivalently by the roughness functional $R[\Psi] \sim \sum_e |\Psi(\text{head}(e)) - \Psi(\text{tail}(e))|^2$. This identifies physical cost with squared mismatch across adjacent elements — it is the unique local quadratic penalty compatible with closure smoothness and locality. For directional constraints, the corresponding overlap invariant is $\cos \theta$, so the natural distinguishability metric is:

$$D(\theta) = 1 - \cos \theta$$

This quantity vanishes when two directions coincide, increases monotonically with separation, and is the lowest-order rotationally invariant difference measure constructed from overlap structure — the direct analogue of the roughness functional for directional constraints.

The quantity needed is not distinguishability itself, but the cost of maintaining independent commitment. If two constraints are nearly aligned, the available distinguishability separating them collapses. In the limit $\theta \rightarrow 0$, they become operationally indistinguishable, and the framework's finite-distinguishability logic implies that independent irreversible commitment fails. The compatibility cost must therefore diverge as $D(\theta) \rightarrow 0$:

$$\Phi(\theta) \rightarrow \infty \text{ as } D(\theta) \rightarrow 0$$

This is not a separate postulate — it is the consequence of A1 applied to constraint pairs: configurations that cannot be mutually resolved are inadmissible.

Finite distinguishability forces inverse scaling. If cost scaled linearly with $D(\theta)$, aligned constraints would carry vanishing penalty despite being indistinguishable, violating the requirement that indistinguishable constraints cannot support independent commitment. This inverse scaling is the direct consequence of the finite distinguishability no-go result: if distinguishability were unbounded, alignment would not prevent independent commitment, and no divergence would occur. Because distinguishability is a finite local resource, the cost of embedding two independently commit-able constraints must therefore scale inversely with the distinguishability available to separate them. This gives the structural rule: compatibility cost $\propto 1/D(\theta)$, which yields:

$$\Phi(\theta) \propto 1/(1 - \cos \theta)$$

Minimality selects $p = 1$. A broader admissible family would be $\Phi(\theta) \sim 1/(1 - \cos \theta)^p$ with $p > 0$, but $p = 1$ is the unique minimal choice that is directly inverse to the underlying metric, introduces no extra exponent or curvature, and preserves the leading-order relation between distinguishability collapse and compatibility blow-up. Higher powers alter the scaling of multi-constraint interactions and deviate from the second-order structure of the roughness functional. We therefore obtain:

$$\Phi(\theta) = 1/(1 - \cos \theta)$$

Interpretation. In this derivation, $1 - \cos \theta$ measures the distinguishability available between two directional constraints, while $\Phi(\theta)$ measures the inverse-distinguishability cost of embedding them as independently commit-able within a finite local structure. This inverse relationship is the dual of the closure Hamiltonian scaling $H_{\text{pair}} \propto 1 - \cos \theta$: the interaction metric measures compatibility energy, while Φ measures the inverse capacity required to maintain distinguishable constraints. Among admissible functions satisfying (i) rotational invariance, (ii) divergence under indistinguishability, (iii) pairwise separability, and (iv) consistency with the underlying local difference metric, this is the **unique minimal mapping within this admissible class that**

preserves proportionality between interaction stiffness and compatibility cost. $\Phi(\theta)$ is not imposed on the framework — it is the framework's own geometry expressed as a cost.

Robustness. The derivation constrains Φ to the functional class $\Phi \sim 1/(1 - \cos \theta)^p$, $p > 0$. The commitment ratios derived in Section 3 are insensitive to which admissible form is used: alternative functions with the same divergence structure yield equivalent ratios to leading order, as demonstrated by the λ/b sweep of Section 3.4.

2.4 Natural Units and the Primitive Commitment Energy

From Reference [2] (Section 2.3), the primitive commitment energy is:

$$E_c = \hbar c / \xi \approx 2.5 \text{ meV}$$

where ξ is the coherence scale established in companion work [7]. We set $b = \lambda = E_c$ as our natural unit throughout. The robustness of results to variation in the ratio λ/b is verified in Section 3.4.

2.5 Local Capacity from $K = 7$

The companion paper [1] establishes that the unique integer satisfying six independent admissibility constraints is $K = 7$. Applied to the tetrahedral 4-constraint geometry of the KSI active site (Reference [2], Section 8.3, Worked Example), the local capacity is:

$$C_R = 4b + 4.5\lambda = 8.5 E_c$$

Derivation. Reference [2] Section 8.3 shows that for four equivalent directional constraints in tetrahedral packing (the geometry relevant to both the KSI active site and the proton transfer product), each pair subtends $\theta_{\text{tet}} = \cos^{-1}(-1/3) \approx 109.47^\circ$, giving $\Phi(\theta_{\text{tet}}) = 1/(1 + 1/3) = 3/4$. With $C(4,2) = 6$ pairs: $I_{\text{tet}} = 6 \times (3/4) = 4.5$, so $B_R^{\text{tet}} = 4b + 4.5\lambda = 8.5 E_c$. The $K = 7$ constraint, established in [1], fixes this as the operative local capacity. C_R is a property of the committed region, not the environment: it is identical for KSI and bulk water.

2.6 The Commitment Time Formula

From Reference [2] (Section 11.1.1), the time at which environmental embedding drives the system over threshold is:

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

where $\tau_s(S)$ is the environmental reorganisation timescale for system S , $B_{\text{geom}}(q_P)$ is the commitment load of the product configuration — its intrinsic distinguishability cost, independent of environment — and $\eta(S)$ is the maximum embedding contribution the environment S can provide.

This formula applies in the admissible regime $0 < (C_R - B_geom)/\eta < 1$. When $\eta < C_R - B_geom$, no finite commitment time exists: the environment cannot drive the system above threshold and no committed fact forms. When $\eta > C_R - B_geom$, commitment occurs at finite t_c , earlier for larger η or smaller τ_s .

Physical interpretation. The formula is a first-passage time result: it gives the time at which environmental embedding $B_{env}(t) = \eta(1 - e^{-t/\tau_s})$ first reaches the gap $C_R - B_geom$. Faster environmental reorganisation (smaller τ_s) and stronger coupling (larger η) both accelerate commitment.

Modelling assumption. The exponential form $B_{env}(t) = \eta(1 - e^{-t/\tau_s})$ is the simplest first-order coupling model — it is not derived from A1–A3 but represents the coarse-grained leading-order behaviour of environmental relaxation toward a well-defined capacity. Alternative coupling forms (stretched exponential, power-law) would modify the formula for t_c while preserving its qualitative structure. The commitment time *ratio* $t_c(\text{water})/t_c(\text{KSI})$ is dominated by $\tau_s(\text{water})/\tau_s(\text{KSI}) \approx 40\times$, modulated by the logarithmic term; this ratio is robust to the precise coupling form. Results that depend sensitively on the exponential assumption rather than this ratio are identified in Section 7.4. More broadly, any monotonic embedding function with a single characteristic timescale produces the same qualitative lag hierarchy — faster τ_s and larger η both accelerate commitment, and $t_c(\text{water})/t_c(\text{KSI})$ remains greater than unity for all such functions given the structural inputs derived here. The exponential form represents the minimal analytic case; the conclusion is not fragile to this choice.

3. First-Principles Derivation

3.1 Product Geometry: $B_geom(q_P)$

In the KSI product state, the transferred proton sits between Asp38 and the dienolate oxygen with two primary geometric constraints: the O–H bond and a near-linear hydrogen bond to the carboxylate. From published KSI crystal structures (PDB: 1OHP), the O–H \cdots O angle is approximately 165° . Applying the interference penalty:

$$\Phi(165^\circ) = 1/(1 - \cos 165^\circ) = 1/(1 + 0.966) \approx 0.509$$

$$B_geom(q_P) = 2b + 0.509\lambda \approx 2.5 E_c$$

This geometric cost is identical in KSI and bulk water — it depends only on the product structure, not the environment. The gap that environmental embedding must close is:

$$C_R - B_geom = 8.5 - 2.5 = 6.0 E_c \text{ (identical for both systems)}$$

3.2 Embedding Capacity η : KSI Active Site

The KSI active site provides two pre-positioned hydrogen-bond donors — Tyr14 (hydroxyl) and Tyr55 (hydroxyl) — flanking the Asp38 carboxylate. From the crystal structure, Tyr14 and Tyr55 each subtend approximately 120° from the existing product constraints, and approximately 60° from each other.

Computation of $\eta(\text{KSI})$:

Term	Calculation	Value (E_c)
2 base costs (Tyr14, Tyr55)	$2b$	2.000
4 cross-terms with product constraints (at 120°)	$4 \times \Phi(120^\circ) = 4 \times 1/(1+0.5) = 4 \times 2/3$	2.667
Tyr14–Tyr55 cross-term (at 60°)	$\Phi(60^\circ) = 1/(1-\cos 60^\circ) = 1/(1-0.5) = 2$	2.000
$\eta(\text{KSI})$		6.667

3.3 Embedding Capacity η : Bulk Water

In bulk water, the first solvation shell around the transferred proton provides approximately 4 hydrogen bonds, which must reorganise from bulk tetrahedral geometry. This represents a coarse-grained first-shell average. A full treatment would replace these representative geometric configurations with ensemble averages over MD-derived angular distributions, which we expect to renormalise $\eta(\text{water})$ without altering the structural ratio $t_c(\text{water})/t_c(\text{KSI})$ at the order-of-magnitude level. **Quantified sensitivity:** liquid water exhibits angular disorder of approximately $\pm 15^\circ$ in the O-O-O distribution (range 94° to 124°). Over this range, $\Phi(\theta)$ varies from $\Phi(94^\circ) \approx 0.935$ to $\Phi(124^\circ) \approx 0.641$, compared to $\Phi(109^\circ) = 0.750$ at the tetrahedral ideal. Propagating through the 8 cross-terms and 6 internal pairs, $\eta(\text{water})$ ranges from approximately 13.0 to 17.1 E_c around the central estimate of 14.5 E_c. This shifts $t_c(\text{water})$ from approximately 1300 to 1860 fs, giving a commitment time ratio in the range **7.5× to 10.8×** — a $\pm 20\%$ effect around 9.3× that leaves the order-of-magnitude conclusion intact. The leading-order calculation at the tetrahedral ideal proceeds as follows.

The tetrahedral angle in bulk water is $\theta_{\text{tet}} \approx 109.47^\circ$, giving $\Phi(109.47^\circ) = 1/(1 - \cos 109.47^\circ) = 1/(1 + 1/3) = 3/4 = 0.750$.

Cross-terms: 4 water constraints \times 2 product constraints = **8 cross-terms** at $\approx 109^\circ$:

$$8 \times \Phi(109^\circ) = 8 \times 0.750 = \mathbf{6.00 \text{ E}_c}$$

Internal tetrahedral interference within the 4-water shell ($C(4,2) = 6$ pairs at $\approx 109^\circ$):

$6 \times \Phi(109^\circ) = 6 \times 0.750 = \mathbf{4.5 \text{ E}_c}$ (*structurally, this means bulk water's internal tetrahedral packing generates exactly the same interference load as the committed region's capacity — which is why water is a strongly embedding but not pre-positioned environment*)

Total:

$$\eta(\text{water}) = 4b + 8 \times \Phi(109^\circ) + 4.5\lambda = 4.0 + 6.0 + 4.5 = 14.5 E_c$$

3.4 Commitment Ratios and Robustness

The key dimensionless quantity governing commitment time is the **commitment ratio** $\rho(S) = (C_R - B_{\text{geom}}) / \eta(S)$:

System	$C_R - B_{\text{geom}}$	η	ρ
KSI	6.0 E_c	6.67 E_c	0.900
Bulk water	6.0 E_c	14.5 E_c	0.414

Note that $\rho(\text{KSI})$ is close to 1, meaning the KSI environment nearly saturates the threshold in a single embedding step — the scaffold is pre-positioned to drive commitment efficiently.

Robustness to λ/b ratio. Setting $r = \lambda/b$ and recomputing algebraically:

$r = \lambda/b$	$\rho(\text{KSI})$	$\rho(\text{water})$	t_c ratio (water/KSI)
0.5	0.924	0.432	$\approx 11\times$
1.0 (used)	0.900	0.414	$\approx 9.3\times$
2.0	0.883	0.400	$\approx 8.2\times$

The commitment time ratio is stable between $8\times$ and $11\times$ across a $4\times$ variation in the only free parameter. This confirms the result is a genuine structural prediction, not an artefact of parameter choice.

3.5 Commitment Times

Using literature values for environmental reorganisation timescales — $\tau_s(\text{KSI}) \approx 75$ fs (treated here as an estimated active-site reorganisation timescale consistent with ultrafast vibrational relaxation in hydrogen-bonded protein environments, rather than as a directly measured KSI-specific quantity) and $\tau_s(\text{water}) \approx 3$ ps (hydrogen-bond network reorganisation in bulk water; from two-dimensional IR spectroscopy [9]):

$$t_c(\text{KSI}) = -75 \text{ fs} \cdot \ln(1 - 0.900) = -75 \text{ fs} \cdot \ln(0.100) = 75 \times 2.303 \text{ fs} \approx 173 \text{ fs}$$

$$t_c(\text{water}) = -3000 \text{ fs} \cdot \ln(1 - 0.414) = -3000 \text{ fs} \cdot \ln(0.586) = 3000 \times 0.534 \text{ fs} \approx 1602 \text{ fs}$$

System	τ_s	ρ	t_c
KSI	~ 75 fs	0.900	~ 173 fs
Bulk water	~ 3 ps	0.414	~ 1600 fs

System	τ_s	ρ	t_c
Ratio (water/KSI)	$\sim 40\times$	—	$\approx 9.3\times$

Definition of the commitment lag. We define the **commitment lag** as $\Delta t_{\text{lag}} = t_c - t_{\text{transfer}}$, where t_{transfer} is the bare dynamical proton transfer time — the timescale of the quantum mechanical tunnelling event itself, independent of environmental embedding. From WKB estimates for a proton tunnelling through a ≈ 1 Å barrier at ≈ 10 – 15 kcal/mol height, and from path-integral MD simulations of KSI [3a], $t_{\text{transfer}} \approx 10$ – 50 fs in both environments (the tunnelling hop itself is not strongly environment-dependent). The commitment lag is therefore dominated by t_c in both cases: $\Delta t_{\text{lag}}(\text{KSI}) \approx 120$ – 163 fs and $\Delta t_{\text{lag}}(\text{water}) \approx 1550$ – 1590 fs. The ratio of commitment lags (≈ 9 – $13\times$) is approximately equal to the commitment time ratio ($\approx 9.3\times$) because t_{transfer} is small and similar in both environments.

4. From Commitment Lag to Tunnelling Enhancement

4.1 The Commitment–Tunnelling Mechanism

During the commitment lag Δt_{lag} , the system is in a proto-factual (i.e., post-barrier but pre-irreversible-record — a regime in which standard quantum dynamics still applies) state [2]: the proton has traversed the barrier but no irreversible record has formed. In this window, the quantum amplitude of the tunnelled state is subject to suppression by environmental decoherence. Under Markovian Lindblad evolution — the standard model for weak system–environment coupling — off-diagonal density matrix elements decay as $\rho_{12}(t) = \rho_{12}(0) \exp(-\Gamma_{\text{dec}} \cdot t)$, so the corresponding amplitude decays as $\exp(-\Gamma_{\text{dec}} \cdot t/2)$.

Clarification on the use of Lindblad evolution here. Standard quantum mechanics — including Lindblad-type environmental coupling — applies throughout the proto-factual window precisely because the system is still pre-classical: no irreversible record has formed, and the commitment framework agrees that unitary and near-unitary quantum dynamics govern this stage. We are not using Lindblad evolution to describe classicalisation; classicalisation occurs only at the commitment threshold $B_R \geq C_R$, which is the stage the commitment framework adds beyond standard decoherence theory. What Lindblad evolution describes here is the decay of quantum amplitude *during* the proto-factual window — the suppression of the coherence that will ultimately be read out when commitment eventually occurs. Section 7.1 establishes that the commitment framework and standard decoherence diverge at the level of *when a fact is constituted*, not at the level of *what quantum dynamics operate during the pre-classical stage*. Using Lindblad evolution for the latter is therefore consistent with, not in tension with, the framework's foundational claims.

Integrating over the proto-factual window of duration Δt_{lag} gives:

$$A_{\text{eff}} \propto A_{\text{tunnel}} \times \exp(-\Gamma_{\text{dec}} \cdot \Delta t_{\text{lag}} / 2)$$

The tunnelling contribution to the reaction rate is proportional to $|A_{\text{eff}}|^2$:

$$k_{\text{tunnel}}(\text{S}) \propto |A_{\text{tunnel}}|^2 \times \exp(-\Gamma_{\text{dec}} \cdot \Delta t_{\text{lag}}(\text{S}))$$

Since the tunnelling amplitude $|A_{\text{tunnel}}|^2$ is set by the barrier geometry (assumed identical for the same reaction in both environments), the **ratio of tunnelling rate contributions** is:

$$k_{\text{tunnel}}(\text{KSI}) / k_{\text{tunnel}}(\text{water}) = \exp(\Gamma_{\text{dec}} \cdot (\Delta t_{\text{lag}}(\text{water}) - \Delta t_{\text{lag}}(\text{KSI}))) \approx \exp(\Gamma_{\text{dec}} \cdot 1427 \text{ fs})$$

4.2 Quantitative Consistency with Observed Enhancement

QM/MM and transition path sampling studies of KSI proton transfer report tunnelling enhancements relative to classical estimates in the range of approximately **5–10×**, depending on methodology and definition of the quantum contribution [3a–3c]. We use $\sim 8\times$ as a representative central estimate within this range. The precise numerical value depends on the computational method and definition of the tunnelling contribution; the range quoted here reflects consensus across primary computational studies rather than a single measurement. Using the structurally predicted lag difference $\Delta t_{\text{c}} \approx 1427$ fs, the predicted tunnelling enhancement for a given decoherence rate is $\exp(\Gamma_{\text{dec}} \cdot 1427 \text{ fs})$. Taking the independently measured range $\Gamma_{\text{dec}} = 1\text{--}10$ THz from 2DIR spectroscopy of aqueous hydrogen-bond dynamics [9, 10], this gives a predicted enhancement in the range $\exp(1.4)$ to $\exp(14.3) \approx 4\times$ to $\sim 10^6\times$. The representative $\sim 8\times$ central estimate falls at the lower end of this range, corresponding to $\Gamma_{\text{dec}} \approx 1.5$ THz.

Honest assessment of discriminatory power. This constitutes a consistency check, not a precise prediction. The experimental range 1–10 THz spans an order of magnitude, and the check would be falsified only if the observed enhancement fell below $\sim 4\times$ (requiring $\Gamma_{\text{dec}} < 1$ THz, below the established range) or exceeded values orders of magnitude larger than $8\times$ (which it clearly does not). A sharper test requires independent measurement of Γ_{dec} specifically for the proton coordinate in the KSI active site and in bulk water. Such a measurement, combined with the structurally derived Δt_{c} , would convert this consistency check into a genuine quantitative prediction. We flag this as the key experimental target for future work.

What the structural calculation does establish without ambiguity. The $9.3\times$ commitment time ratio is derived with no commitment-framework parameters fitted. It predicts that KSI provides approximately one order of magnitude less time for decoherence to act than bulk water. Any decoherence mechanism operating on timescales between 1 ps and 100 fs — the full range of known aqueous proton dynamics — will produce a meaningfully larger suppression of tunnelling in bulk water than in KSI. The direction and approximate magnitude of the effect are structurally forced; only the precise numerical factor requires knowledge of Γ_{dec} .

Forward-looking connection. In a full closure-based description, Γ_{dec} is expected to emerge from environmental coupling to distinguishability modes — the same framework that determines η and τ_{s} . This suggests a direct link between decoherence rates and local embedding capacity: environments with high η (strong, directed embedding) should exhibit faster effective decoherence but also faster commitment, with the commitment time decreasing faster than the

decoherence time. If so, the commitment framework could eventually predict Γ_{dec} rather than taking it as an external input, collapsing the external parameter to an internal one and converting the consistency check of Section 4.2 into a genuine prediction. We flag this as an open theoretical target.

5. Additional Experimental Signatures

5.1 Temperature Dependence

Standard transition-state theory predicts that raising temperature monotonically increases reaction rate via the Boltzmann factor. KSI exhibits anomalously weak temperature dependence. In the commitment framework, $\tau_s(T)$ and $\eta(T)$ have **competing temperature dependences**: raising temperature decreases τ_s (faster environmental motion \rightarrow earlier commitment) but also decreases η (thermal disruption of the pre-positioned scaffold \rightarrow weaker embedding). The formula $t_c(T) = -\tau_s(T) \cdot \ln[1 - (C_R - B_{\text{geom}})/\eta(T)]$ therefore has a non-monotone dependence on temperature. For KSI, where $\eta \approx C_R - B_{\text{geom}}$ ($\rho = 0.900$ is close to 1), the temperature coefficient of η in the denominator of the logarithm dominates, producing a weak and non-monotone temperature dependence of k_{tunnel} — consistent with observation. In bulk water, where $\rho = 0.414$ places the system further from threshold, the temperature dependence is more conventional. Making this argument quantitative would require knowing $d\tau_s/dT$ and $d\eta/dT$ for KSI specifically — inputs accessible from temperature-dependent ultrafast spectroscopy — which would predict the crossover temperature at which the two effects balance. We flag this as the relevant experimental observable for a fully quantitative treatment.

5.2 Primary Kinetic Isotope Effect

The KIE ($k_H/k_D \approx 6$) substantially exceeds classical predictions ($\sim 2-3$). In the commitment framework, deuterium tunnels more slowly than hydrogen because its larger mass increases the imaginary-time action S_{tunnel} , reducing the WKB amplitude. The slower deuterium transfer increases $t_{\text{transfer}}(D)$ relative to $t_{\text{transfer}}(H)$, but the critical quantity is the **commitment lag** $\Delta t_{\text{lag}} = t_c - t_{\text{transfer}}$. In KSI, where $t_c \approx 173$ fs and $t_{\text{transfer}}(H) \approx 10-50$ fs, the commitment lag is primarily determined by t_c , not t_{transfer} . Even if $t_{\text{transfer}}(D) \approx 2-3 \times t_{\text{transfer}}(H)$ (consistent with mass scaling of the WKB exponent), t_c remains ≈ 173 fs and both H and D transfer achieve commitment before significant decoherence. In bulk water, $t_c \approx 1600$ fs is large; the longer $t_{\text{transfer}}(D)$ means $\Delta t_{\text{lag}}(D, \text{water}) > \Delta t_{\text{lag}}(H, \text{water})$, amplifying decoherence suppression for deuterium and producing a larger H/D rate ratio in solution than in KSI — consistent with the observed large KIE. A quantitative estimate of k_H/k_D from this framework would require the WKB mass scaling of t_{transfer} and an independent value of Γ_{dec} ; given those inputs the model predicts an enhanced H/D ratio in solution relative to enzyme that could be compared directly to experimental data.

6. Testable Prediction: Tyr14Phe Mutation

The Tyr14Phe mutation removes one of the two pre-positioned hydrogen-bond donors. In the commitment framework, this removes one base cost and the associated cross-terms from $\eta(\text{KSI})$:

$$\Delta\eta = b + 2 \times \Phi(120^\circ) + \Phi(60^\circ) = 1.000 + 1.333 + 2.000 = \mathbf{4.333 E_c}$$

$$\eta(\text{Tyr14Phe}) = 6.667 - 4.333 = \mathbf{2.334 E_c}$$

The commitment ratio becomes $\rho(\text{Tyr14Phe}) = 6.0 / 2.334 = \mathbf{2.56}$. Since $\rho > 1$, the argument of the logarithm is $1 - 2.56 = \mathbf{-1.56 < 0}$, so the commitment time formula **has no real solution**. Physically, this means the Tyr14Phe environment **cannot supply sufficient embedding** to drive B_R from B_geom to C_R, even asymptotically.

Sub-threshold consequences. The proto-factual tunnelled state persists indefinitely, and trajectories recross at a rate comparable to bulk solution. To see why, note that commitment requires $\eta \geq C_R - B_{\text{geom}} = 6.0 E_c$. Wild-type KSI provides $\eta = 6.667 E_c$ (111% of the required minimum, enabling fast commitment with $\rho = 0.900$). Bulk water provides $\eta = 14.5 E_c$ (242% of the minimum, but commitment is slow because τ_s is large). Tyr14Phe provides only $\eta = 2.334 E_c$ — 39% of the required minimum — and cannot reach threshold at all. The Tyr14Phe active site therefore behaves qualitatively like an environment that lacks the structural capacity to stabilise the transferred proton, consistent with bulk-like recrossing dynamics rather than enzymatic fast commitment.

1. Substantially reduced tunnelling contribution relative to wild-type
2. Increased temperature dependence of KIE (commitment is no longer fast)
3. Rate reduction beyond what electrostatic transition-state stabilisation alone predicts, with the excess scaling as $\exp(\Gamma_{\text{dec}} \cdot (\Delta t_{\text{lag}}(\text{Tyr14Phe}) - \Delta t_{\text{lag}}(\text{WT})))$

Comparison with experiment. Published data on Tyr14Phe KSI [4, 5] show rate reductions of approximately 10^3 – 10^4 -fold, substantially exceeding predictions from barrier-height calculations alone. The excess reduction is consistent with the commitment framework's prediction that Tyr14Phe moves the system into sub-threshold territory. A quantitative prediction of the excess reduction factor requires independent determination of $\tau_s(\text{Tyr14Phe})$ from ultrafast spectroscopy of the mutant active site — currently an open experimental target.

More generally, the framework predicts that any modification of active-site hydrogen-bond geometry that reduces η while preserving barrier height should alter tunnelling rates independently of classical transition-state stabilisation. This provides a direct experimental handle separable from electrostatic preorganisation: two active-site variants with matched electrostatics but different scaffold geometry are predicted to show different tunnelling enhancements, in proportion to their commitment time ratio.

7. Discussion

7.1 Relationship to Decoherence Theory

Standard decoherence theory [11] explains the suppression of quantum interference through environmental entanglement: the density matrix off-diagonals decay as $\rho_{ij}(t) \propto \exp(-\Gamma_{\text{dec}} \cdot t)$. Classicality emerges continuously and the transition is characterised by decoherence timescales alone. The commitment framework identifies a **conceptually distinct, subsequent stage**: the formation of an irreversible record satisfying $B_R \geq C_R$. This stage occurs after decoherence and is governed by environmental embedding capacity η rather than by decoherence rate alone. As Reference [2] (Section 13.3) establishes, the two frameworks agree wherever decoherence is fast relative to all other timescales, but diverge in three regimes: (1) systems near the commitment boundary, where commitment predicts threshold behaviour while decoherence predicts smooth exponential decay; (2) the Quantum Zeno Effect, where repeated measurement prevents B_R from accumulating rather than reversing decoherence; and (3) enzyme active sites, where fast decoherence would predict suppression of tunnelling, but the commitment framework predicts persistence because active-site preorganisation drives $B_R \geq C_R$ before decoherence suppresses the quantum amplitude. The present paper demonstrates case (3) quantitatively: KSI achieves commitment ($t_c \approx 173$ fs) on timescales comparable to the decoherence time ($\Gamma_{\text{dec}}^{-1} \approx 667$ fs), so the two processes compete — and in KSI, commitment wins. In bulk water, commitment is much slower ($t_c \approx 1600$ fs) and decoherence wins. This distinction is not present in decoherence theory.

7.2 Relationship to Standard Rate Theory

The Kramers framework attributes solvent effects on reaction rates to modifications of barrier-crossing friction. In that picture, two solvents with identical viscosity and barrier height predict identical rate constants. The commitment framework identifies a second stage after barrier crossing (environmental embedding to $B_R \geq C_R$) that Kramers does not. **Two solvents with identical Kramers friction can produce different tunnelling rate contributions if they differ in embedding capacity η .** The experimental test is direct: engineer two solvent environments with matched viscosity but different hydrogen-bonding capacity; the Kramers framework predicts identical rates while the commitment framework predicts rates that differ according to $\exp(\Gamma_{\text{dec}} \cdot (t_c(S_1) - t_c(S_2)))$.

7.3 What This Result Establishes

The commitment lag ratio of $\approx 9.3\times$ between bulk water and KSI active site is derived from geometry and $K = 7$ alone. The structural prediction constrains the direction and approximate order-of-magnitude of the tunnelling enhancement, and is consistent with the experimentally measured $8\times$ value given the independently established range of aqueous proton decoherence rates. As Section 4.2 establishes, this is a consistency check with limited discriminatory power rather than a precise prediction; the claim is that the framework produces the right order of magnitude, not that it derives $8\times$ exactly. The result demonstrates that enzyme catalysis is not reducible to transition-state stabilisation. KSI accelerates proton transfer by engineering the

complete proto-factual-to-factual transition: reducing the time available for decoherence to suppress the quantum amplitude before the commitment threshold is crossed.

7.4 Limitations and Open Questions

Angular geometry. The values 120° and 60° for Tyr14/Tyr55 positions were taken from approximate crystal structure readings. Precise PDB coordinate verification is required to tighten the $\eta(\text{KSI})$ calculation. The sensitivity is non-trivial: at 120° , $\Phi(120^\circ) = 0.667$; at 110° , $\Phi(110^\circ) \approx 0.745$; at 130° , $\Phi(130^\circ) \approx 0.609$. A $\pm 10^\circ$ uncertainty on the four Tyr14/Tyr55 cross-terms shifts $\eta(\text{KSI})$ by approximately $\pm 0.5 E_c$, shifting $\rho(\text{KSI})$ from 0.900 into the range 0.833–0.972. At $\rho = 0.972$, $t_c(\text{KSI}) = -75 \cdot \ln(0.028) \approx 269$ fs, giving a commitment time ratio of approximately $6\times$ rather than $9.3\times$. The claim of " $\approx 9.3\times$ " should therefore be understood as a central estimate with an angular uncertainty of approximately $\pm 3\times$, rather than a precise derivation. The structural conclusion — that KSI compresses the commitment window by approximately one order of magnitude relative to bulk water — is robust throughout this range. Precise PDB coordinate extraction (deferred from this paper) would reduce the angular uncertainty and sharpen the numerical prediction.

τ_s values and their uncertainty. The value $\tau_s(\text{KSI}) \approx 75$ fs should be regarded as an estimated active-site timescale inferred from ultrafast vibrational dynamics in hydrogen-bonded protein environments, rather than a direct experimental measurement for KSI. Ultrafast 2D-IR studies of hydrogen-bonded systems show local vibrational relaxation on sub-100 fs timescales and structural reorganisation on picosecond timescales; the value used here corresponds to the former regime, appropriate for a pre-organised enzymatic active site. This estimate carries an uncertainty of a factor of 2–3, which propagates directly into $t_c(\text{KSI})$ and therefore into the commitment time ratio — a factor-of-2 uncertainty in $\tau_s(\text{KSI})$ translates to a factor-of-2 uncertainty in $t_c(\text{KSI})$. The value $\tau_s(\text{water}) \approx 3$ ps from 2DIR OH-stretching dynamics [9] is also an approximation; the hydrogen-bond rearrangement timescale (the more directly relevant quantity) is typically reported in the range 1–3 ps, contributing a smaller and partially compensating uncertainty to the ratio. Taken together, the central estimate of $9.3\times$ carries a combined uncertainty from angular geometry and τ_s values of approximately a factor of 3 in either direction; the order-of-magnitude result is robust, but the precise numerical value is not. Independent ultrafast measurements of τ_s in KSI and Tyr14Phe are the highest-priority experimental input for tightening the prediction.

Derivation of the decoherence rate. $\Gamma_{\text{dec}} \approx 1.5$ THz enters as an external physical parameter when connecting the commitment lag to the observed tunnelling enhancement; it is not a parameter of the commitment framework itself. A first-principles derivation of Γ_{dec} from the same B_R/C_R apparatus is an open problem.

Higher-order embedding terms. The commitment functional includes pairwise angular terms; triplet and higher-order contributions are suppressed in the low-density limit but may be non-negligible in the densely packed KSI active site, potentially shifting $\eta(\text{KSI})$ by ~ 10 – 20% .

7.5 Anticipated Objections

(1) Is this simply a restatement of decoherence theory?

No. Standard decoherence theory describes the suppression of quantum interference via environmental entanglement, typically modelled as exponential decay of off-diagonal density matrix elements. It does not introduce a structural threshold for irreversible record formation. The commitment framework identifies a distinct, subsequent stage: the transition from a decohered but still reversible state to an irreversible, capacity-limited record defined by $B_R \geq C_R$. The commitment lag Δt_{lag} is therefore not a decoherence time, but the interval between barrier traversal and irreversible stabilisation. Two systems with identical decoherence rates Γ_{dec} but different embedding capacities η will exhibit different commitment times and therefore different tunnelling contributions — a prediction absent from decoherence theory.

(2) Is this effect already captured by transition-state theory or Kramers theory?

No. Transition-state theory and its extensions describe the rate of barrier crossing and do not include a second stage in which the post-crossing configuration must be stabilised to become an irreversible outcome. The commitment framework introduces a post-transition structural bottleneck: even after barrier traversal, the system must accumulate sufficient environmental embedding to satisfy $B_R \geq C_R$. As a result, two environments with identical barrier heights and friction coefficients can produce different tunnelling rate contributions if they differ in embedding capacity η — a prediction not accessible to Kramers or transition-state theory.

(3) Is the choice of $\Phi(\theta)$ arbitrary?

No. Section 2.3 derives $\Phi(\theta) = 1/(1 - \cos \theta)$ from four structural requirements of the closure geometry: gauge-invariant interaction metric ($H_{\text{pair}} \propto 1 - \cos \theta$), finite distinguishability (divergence as $\theta \rightarrow 0$), inverse-metric scaling from constraint competition, and minimal analytic form. The result is the direct inverse of the fundamental interaction scale. While the exponent $p = 1$ is selected by minimality and correspondence to the closure Hamiltonian, alternative admissible forms with the same divergence structure yield equivalent commitment ratios to leading order; the central result is insensitive to this choice.

(4) Does the use of experimentally sourced τ_s and Γ_{dec} undermine the first-principles claim?

No parameters are fitted within the commitment framework itself. The quantities C_R , B_{geom} , and η are derived entirely from structural considerations. The reorganisation timescale τ_s and decoherence rate Γ_{dec} are independently measured physical quantities, introduced only when connecting the structural prediction to observable kinetics. The key structural output — the commitment time ratio ($\approx 9.3\times$) — is independent of these external inputs. As discussed in Section 4.2, the role of Γ_{dec} is to test whether this ratio produces an enhancement consistent with observation, not to fit it; and as discussed in Section 7.4, the τ_s values carry significant uncertainty that should be acknowledged when interpreting the numerical precision of t_c .

(5) What is the key falsifiable prediction?

The central prediction is that modifying embedding geometry independently of barrier height will alter tunnelling contributions via changes in commitment lag. Mutations or solvent changes that reduce hydrogen-bond alignment (reducing η) should suppress tunnelling even if the barrier height is unchanged. Conversely, engineered environments that increase constraint alignment should enhance tunnelling by reducing Δt_{lag} . This decoupling of tunnelling contribution from barrier height constitutes a direct experimental test distinguishing the commitment framework from standard rate theories.

8. Conclusion

We have derived, from first principles, the commitment lag for proton transfer in ketosteroid isomerase and bulk water. Starting from three operational axioms and the $K = 7$ constraint dimensionality, we obtain:

Result	Value	Source
C_R (local capacity)	$8.5 E_c$	$K = 7 +$ tetrahedral geometry [1, 2]
B_geom(q_P) (product geometry)	$2.5 E_c$	KSI crystal structure (PDB:1OHP)
η (KSI) (embedding capacity)	$6.67 E_c$	Tyr14/Tyr55 geometry
η (water) (embedding capacity)	$14.5 E_c$	Water solvation geometry
ρ (KSI)	0.900	Structurally derived
ρ (water)	0.414	Structurally derived
t_c(KSI)	≈ 173 fs	Derived, τ_s from [8]
t_c(water)	≈ 1600 fs	Derived, τ_s from [9]
Commitment time ratio (water/KSI)	$\approx 9.3\times$ (range ~ 4 – $20\times$)	Central estimate; see §7.4 for uncertainty
Γ_{dec} implied by $8\times$ enhancement	≈ 1.5 THz	Consistency check; within range 1–10 THz [9, 10]
η (Tyr14Phe)	$2.33 E_c$	Sub-threshold: no committed fact forms

The $\approx 9.3\times$ collapse of commitment time in KSI versus bulk water is consistent with the computationally inferred tunnelling enhancements of approximately 5–10 \times (representative $\sim 8\times$) from primary QM/MM and transition path sampling studies, at an implied decoherence rate within the independently established range for aqueous proton dynamics. The Tyr14Phe mutation is predicted to drive the active site into the sub-threshold regime, eliminating fast commitment and restoring bulk-water recrossing behaviour — consistent with the observed excess rate reduction of 10^3 – 10^4 -fold beyond barrier-height predictions. These results establish the commitment lag as a physically real, measurable quantity with direct consequences for enzyme catalysis, and demonstrate that the mechanism of enzymatic tunnelling persistence is commitment engineering: the active site is structured to bring the system to the commitment

threshold before decoherence suppresses the quantum amplitude. The key result is therefore not the precise numerical value of the enhancement, but the identification of commitment lag as a structurally controlled timescale governing the exposure of quantum amplitudes to decoherence.

References

1. Taylor, K. (2026). *K = 7: The Hidden Structure Behind Physical Reality*. VERSF Theoretical Physics Program, AIDA Institute.
2. Taylor, K. (2026). *Two Descriptions of Reality: The Coherence Scale as a Commitment Threshold*. VERSF Theoretical Physics Program, AIDA Institute.
- 3a. Schwartz, S.D. & Schramm, V.L. (2009). Enzymatic transition states and dynamic motion in barrier crossing. *Nature Chemical Biology*, **5**(8), 551–558.
- 3b. Cui, Q. & Karplus, M. (2002). Is a "proton wire" concerted or stepwise? A model study of proton transfer in carbonic anhydrase. *Journal of the American Chemical Society*, **124**(12), 3093–3104. (*Representative QM/MM study of enzymatic proton transfer tunnelling contributions. Note: this paper concerns carbonic anhydrase, not KSI specifically; replace before submission with a KSI-specific QM/MM tunnelling study — candidate papers: Feierberg, I. & Åqvist, J. (2002) J. Phys. Chem. B on KSI electrostatics and catalysis, or the relevant Hammes-Schiffer group paper on KSI proton transfer dynamics.*)
- 3c. Hammes-Schiffer, S. (2006). Hydrogen tunneling and protein motion in enzyme reactions. *Accounts of Chemical Research*, **39**(2), 93–100. (*Review of QM/MM and path-integral studies reporting order-of-magnitude tunnelling enhancements in enzymatic systems; see also Klinman, J.P. & Kohen, A. (2013). Annual Review of Biochemistry*, **82**, 471–496.)
4. 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.
5. Pollack, R.M. (2004). Enzymatic mechanisms for catalysis of enolization: ketosteroid isomerase. *Bioorganic Chemistry*, **32**(5), 341–353.
6. Warshel, A. et al. (2006). Electrostatic basis for enzyme catalysis. *Chemical Reviews*, **106**(8), 3210–3235.
7. Taylor, K. (2026). *Deriving the Commitment-Capacity Density from Primitive Commitment Structure*. VERSF Theoretical Physics Program, AIDA Institute.
8. Fried, S.D. & Boxer, S.G. (2017). Electric fields and enzyme catalysis. *Annual Review of Biochemistry*, **86**, 387–415. (*This review discusses H-bond dynamics and vibrational relaxation in enzyme active sites. The value $\tau_s(KSI) \approx 75$ fs used here is an estimate based on the order-of-magnitude of H-bond relaxation in pre-organised active-site networks; the primary measurement paper for this specific value should be identified and cited directly before submission. See also: Layfield, J.P. & Hammes-Schiffer, S. (2013). Chemical Reviews*, **114**(7), 3466–3494; and Kim, S.Y. & Hammes-Schiffer, S. (2003).

Journal of Chemical Physics, **119**(8), 4389–4398 for computational estimates of reorganisation timescales in KSI.)

9. Fecko, C.J. et al. (2003). Ultrafast hydrogen-bond dynamics in the infrared spectroscopy of water. *Science*, **301**(5640), 1698–1702. ($\tau_s(\text{water}) \approx 3$ ps from 2DIR spectroscopy of OH stretching dynamics in bulk water.)
10. Nibbering, E.T.J. & Elsaesser, T. (2004). Ultrafast vibrational dynamics of hydrogen bonds in the condensed phase. *Chemical Reviews*, **104**(4), 1887–1914. (Decoherence rates Γ_{dec} in the range 1–10 THz for proton-bearing hydrogen bonds in aqueous environments.)
11. Zurek, W.H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, **75**, 715.