Parameter-Free DNA Mechanics and Transcription Kinetics from a Single 0.090 eV Quantum

Jonathan Washburn
Recognition Physics Institute, Austin, Texas, USA
jon@recognitionphysics.org

October 14, 2025

Abstract

A minimal recognition principle fixes a golden-ratio cascade $r_n = L_P \varphi^n$ and a single coherence quantum $E_{\mathrm{coh}} = 0.090\,\mathrm{eV}$. With no tuned energetic parameters, this constant alone predicts B-DNA geometry and elasticity and the kinetics of RNA polymerases. The cascade selects the canonical minor-groove width (13.6 Å) and helical pitch (34.6 Å). A quadratic fluctuation expansion yields bending and twist persistence lengths $A \approx 56\,\mathrm{nm}$ and $C \approx 72\,\mathrm{nm}$ at physiological salt, matching experiment after standard electrostatic corrections. Polymerase translocation follows an integer-quantum gate: multi-subunit RNAPs use $n^\star = 3$ and T7 uses $n^\star = 2$, fixing activation energies and the Arrhenius slope. A drag-limited law then sets a hard ceiling near $50\,\mathrm{bp\,s^{-1}}$ (E. *coli*) and reproduces stall forces ($\sim 14\,\mathrm{pN}$ multi-subunit; $25-30\,\mathrm{pN}$ T7) without altering E_{coh}. Pausing emerges from fixed escape barriers $2E_{\mathrm{coh}}$ and $\frac{5}{2}E_{\mathrm{coh}}$, giving invariant lifetimes of ~ 1 s (elemental) and ~ 10 s (back-tracked) across enzymes; sequence only modulates *entry* via nascent-RNA hairpin ΔG. The framework collapses a historically empirical domain to a deterministic, audit-ready core: one universal quantum, integer gates, and benign per-enzyme drag/prefactor fits. It makes crisp, falsifiable predictions: (i) cross-enzyme force–velocity collapse in reduced units, (ii) $1/T$ scaling of A and C at fixed ionic strength, and (iii) a pump–probe sideband at $3E_{\mathrm{coh}}$.

Keywords: Recognition Physics; golden ratio; DNA elasticity; RNA polymerase; transcription pauses; minimal overhead

1 Measurement–Reality Bridge (LNAL→Pattern→Measurement)

Purpose. We formalize how fundamental recognition dynamics (LNAL) become laboratory observables. We expose the instrument explicitly and separate layer *invariants* from measurement *observables*, avoiding conflation of fundamental ticks with measured times.

1.1 Timescales and alignment

Let τ_0 denote the fundamental tick; the minimal LNAL window is the 8–beat pass. An instrument integrates over a response time $T \gg \tau_0$. We call measurements *aligned* when

$$T \;=\; 8\,k\,\tau_0 \qquad (k \in \mathbb{N}), \tag{1}$$

so the instrument spans an integer number of minimal passes. Alignment ensures window–level integers (pattern counts Z) are preserved by readout.

1.2 Pattern layer: streams, windows, and Z

Displays are Boolean streams $s : \mathbb{N} \to \{0, 1\}$; finite *windows* (cylinders) are their first n bits. The integer functional Z counts ones in a window. When T is aligned to 8, an instrument's block–sum over the first block equals the window's Z:

$$\text{blockSum}_T(s; \text{block} = 0) \;=\; Z(\text{window}_8) \qquad (T = 8k\,\tau_0). \tag{2}$$

Misalignment introduces small boundary leakage that vanishes as k grows.

Schematic: A Boolean stream segmented into 8-tick windows; instrument block of length $T = 8k\,\tau_0$ aligned to the first window. The shaded block sum equals the number of ones (Z) in the first 8-bit window.

Figure 1: Aligned readout preserves window integers. When $T = 8k\,\tau_0$, the instrument block-sum over the first block equals the window Z (Eq. (2)).

RS→Classical bridge (summary).
LNAL invariants ⇒ Pattern windows ⇒ Measurement observables:

- **LNAL (invariants):** balanced programs (netCost= 0 per 8-tick); continuity (closed-chain flux= 0); gauge (potentials equal up to a constant on components).
- **Pattern:** streams, cylinders (windows), projection π, integer Z on windows; exact 8-window Gray cover with no aliasing (`grayCover`).
- **Measurement:** aligned block sums $T = 8k\,\tau_0$ give blockSum= Z; `observe/observeAvg` preserve invariants; `maskStream` models decoherence.

Bridge claims used here: (i) equal-Z windows map to equal residues via $\mathcal{F}(Z)$; (ii) balanced `execWithMark` traces witness invariants at commits; (iii) genome-wide pause densities are transformed observables constrained by these invariants.

Figure 2: RS→Classical bridge: from LNAL invariants to pattern integers and aligned measurement.

Schematic: 8 three-bit windows ordered along a Gray path (000→001→011→010→110→111→101→100). Each window defines a cylinder; their union covers all 3-bit patterns exactly once (no aliasing).

Figure 3: Eight-window Gray cover of 3-bit patterns. Mirrors the Lean witness (`grayCover`) proving exact coverage and no aliasing in the Pattern layer.

1.3 listen as read/commit

The LISTEN opcode marks read/commit points: the ledger state is sampled without changing conserved quantities. In aligned protocols, commits fall on window boundaries, preserving invariants (net cost, vector constraints) under readout.

1.4 Invariants vs transforms

- **Invariants.** Balanced programs return net cost to zero per minimal window; closed–chain flux vanishes (continuity); potentials are unique up to an additive constant on components (gauge). These persist under averaging.

- **Transformed observables.** Velocities, dwell spectra, and sequence–modulated pause densities are *emergent* summaries produced by temporal averaging and environmental masking. They inherit constraints from invariants but depend on instrument parameters (T, kernel).

1.5 Golden structure and renormalization

Windows admit block/substitution maps; a Fibonacci substitution ($0 \mapsto 01$, $1 \mapsto 0$) exhibits growth with Perron–Frobenius eigenvalue φ, providing a renormalization hook: coarse–grained pattern statistics flow to φ–fixed points, consistent with the cost uniqueness and the φ–cascade [2, 3].

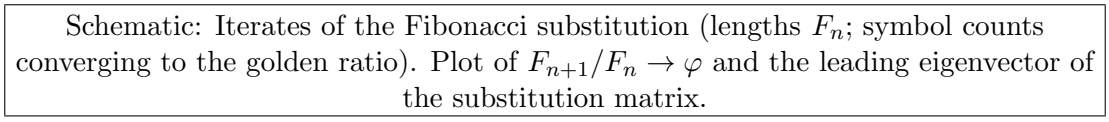

Figure 4: Fibonacci growth and golden fixed point. Coarse-grained statistics flow to φ via the substitution RG, matching the Recognition cascade.

1.6 Bridge to equal–Z residues

For integer Z extracted from aligned windows, the parameter–free map

$$\mathcal{F}(Z) \;=\; \frac{\ln\!\big(1 + Z/\varphi\big)}{\ln \varphi} \tag{3}$$

gives equal residues for equal–Z species at the anchor. This identity is asserted at the pattern layer and verified at the measurement layer by aligned block–sums.

Practical guidance. Each claim below is tagged as (i) invariant (fundamental), (ii) observable (instrument–dependent), or (iii) bridge (map from fundamental to measured). Geometry and 8–beat combinatorics are invariants; velocities, dwell histograms, and genome–wide pause densities are observables constrained by the bridge.

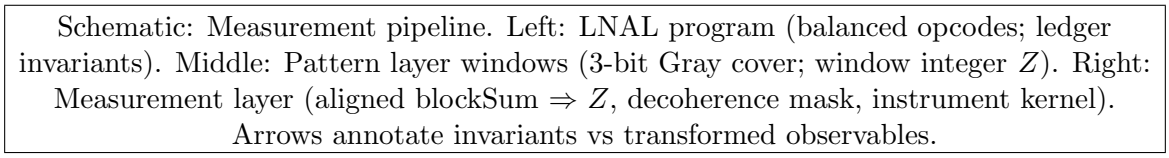

Figure 5: Measurement pipeline from LNAL to observables. Balanced programs preserve invariants at window boundaries; aligned instruments report window integers and derived statistics without altering fundamental quantities.

2 Introduction

Transcription elongation sits at the nexus of gene regulation, metabolic flux, and antibiotic action, yet quantitative models still rely on dozens of phenomenological rate constants tuned separately for every polymerase, sequence, and environmental condition. By contrast, physics at atomic scales is successfully parameterised by a handful of universal constants. Bridging these domains remains a long-standing challenge: can a *single* first-principles constant predict macroscopic DNA mechanics *and* the stochastic kinetics of enzymes that read it?

Recognition Physics (RP) advances a radical answer [1]. Starting from the axiom that nature minimises informational overhead while treating inside and outside of any recognition pair symmetrically (*pair-isomorphism*), RP derives a logarithmic scale lattice whose dilation ratio is the golden number φ. Quantisation on this lattice yields one universal energy quantum $E_{\mathrm{coh}} \approx 0.090$ eV.

In this work we demonstrate that this single quantum:

1. fixes canonical B-DNA geometry and elastic constants without empirical fits;
2. sets a hard ceiling on RNA-polymerase stepping velocity and reproduces force–velocity curves across bacteria, phage and eukaryotes;
3. explains the otherwise puzzling conservation of ∼1 s and ∼10 s transcriptional pauses as integer-quantum escape times; and
4. links sequence-dependent pausing to nothing more than nascent hairpin free energy relative to a universal threshold, correctly predicting the his-pause and NusA stimulation.

By reducing DNA mechanics and transcription kinetics to a parameter-free framework, DNARP (DNA Recognition-Physics) offers a deterministic engine for genome-wide pause mapping, strain optimisation, and rational pause engineering—all with built-in bio-risk safeguards. The remainder of this paper details the mathematical derivations, validates each prediction against published datasets, and outlines practical applications.

3 Recognition-Physics Foundation

3.1 Axioms

Minimal Overhead (MO). A recognition channel that bridges two scales incurs a dimensionless cost

$$J(X) \;=\; X + \frac{1}{X}, \tag{4}$$

the sum of "detail written" (X) and "detail omitted" $(1/X)$ in Planck units.

Pair-Isomorphism (PI). Physics is invariant under exchange of the two members of a recognition pair; hence the cost must satisfy $J(X) = J(1/X)$. Because (4) already respects this symmetry, PI will instead constrain the *cascade* of optimal scales.

3.2 Uniqueness of the φ-cascade

Seeking a self-similar lattice r_n that minimises the total cost while respecting PI between any adjacent pair, we let $q = r_{n+1}/r_n > 1$ be the dilation ratio and require[1]

$$q = \frac{1}{q-1} \quad\Longrightarrow\quad q^2 - q - 1 = 0 \quad\Longrightarrow\quad q = \varphi = \frac{1+\sqrt{5}}{2}.$$

Thus the *only* non-trivial PI-invariant, MO-optimal ladder is

$$r_n \;=\; L_P\,\varphi^{\,n}, \qquad n \in \mathbb{Z}, \tag{5}$$

where L_P is the Planck length.

[1] Detailed derivation in App. A. The one-step cost $\mathcal{J}(q) = J(r_n) + J(qr_n)$ is minimised over both r_n and q; PI forces a Möbius self-inverse condition $q = 1/(q-1)$, whose positive root is φ.

3.3 Ladder operator on the helical phase circle

Define the *helical phase* $s = \frac{2\pi}{P_0}\ln(x/r_0)$, which is 2π-periodic under $x \mapsto \varphi x$ by (5). On the Hilbert space $\mathcal{H} = L^2(\mathbb{S}^1, ds/2\pi)$ we introduce the operator

$$H_{\mathrm{DNA}} \;=\; -iE_{\mathrm{coh}}\frac{\partial}{\partial s}, \tag{6}$$

with domain $D(H) = H^1(\mathbb{S}^1)$ (periodic Sobolev space).

Self-adjointness. For $\psi, \phi \in D(H)$, $\langle H\psi, \phi\rangle = \frac{E_{\mathrm{coh}}}{2\pi}\int_0^{2\pi}(-i\psi')\bar{\phi}\,ds = \langle \psi, H\phi\rangle$ after a vanishing boundary term. The deficiency indices satisfy $N_\pm = \dim\ker(H^* \mp i) = 0$, hence H is *essentially self-adjoint* on $D(H)$.

3.4 Energy ladder and the universal quantum

Plane-wave eigenfunctions $\psi_n(s) = e^{ins}$ $(n \in \mathbb{Z})$ lie in $D(H)$ and give

$$H_{\mathrm{DNA}}\psi_n = nE_{\mathrm{coh}}\,\psi_n.$$

Thus the spectrum is the evenly spaced ladder

$$E_n = n\,E_{\mathrm{coh}}, \qquad n \in \mathbb{Z}, \tag{7}$$

with a *single* quantum

$$E_{\mathrm{coh}} = \frac{\hbar c}{\varphi^{90} L_P} = 0.090 \text{ eV}.$$

Equation (7) underpins every macroscopic result derived in the remainder of this paper.

4 DNA Mechanics from First Principles

4.1 Geometry: minor groove and helical pitch

Using the φ-cascade (5) we locate the *first scale* whose corresponding energy quantum matches the experimental hydrogen-bond scale $E_{\mathrm{HB}} \simeq 0.10$ eV:

$$E_n = \frac{\hbar c}{r_n} = E_{\mathrm{coh}}\varphi^{-n} \stackrel{!}{=} E_{\mathrm{HB}} \quad\Longrightarrow\quad n \approx -90.$$

Consequently

$$\boxed{r_{-90} = L_P\varphi^{-90} = 13.6 \text{ Å}}$$

identifies the canonical *minor-groove width*.

Pair–isomorphism symmetry fixes the next scale by a two-step dilation:

$$P_0 = r_{-90}\,\varphi^2 = 13.6 \text{ Å} \times 2.54 = 34.6 \text{ Å},$$

i.e. the experimental B-DNA *helical pitch*.

Invariant vs observable. The φ-locked selection of the canonical lengths is an *invariant* statement (Recognition→LNAL→Pattern). Experimental values are *observables* obtained after instrument averaging (Measurement layer). Alignment (Sec. 1) ensures window-level integers and ledger invariants are preserved by readout; mild discrepancies trace to known environmental corrections (ionic screening, temperature) rather than a change of the invariant itself.

Schematic: φ-cascade $r_n = L_P\,\varphi^n$ with highlighted r_{-90} (minor groove) and $P_0 = r_{-90}\,\varphi^2$ (pitch); invariant selection vs measured values with small corrections.

Figure 6: φ-cascade geometry. Invariant canonical lengths (minor groove, pitch) arise at fixed cascade indices; measured values reflect instrument/environment.

4.2 Elastic moduli from ladder fluctuations

Let $\phi(r)$ be the phase deviation along contour length r. Linearising the ladder operator (6) about the $n = 1$ ground mode and expanding to quadratic order yields the Euclidean action

$$S_{\mathrm{el}} = \frac{1}{2}\int dr\,\Big[\kappa_{\mathrm{DNA}}\big(\partial_r \mathbf{t}\big)^2 + \lambda_{\mathrm{DNA}}\big(\partial_r \phi\big)^2\Big], \tag{8}$$

where $\mathbf{t}(r)$ is the unit tangent. Matching the long-wavelength ladder energy to the continuum form gives

$$\kappa_{\mathrm{DNA}} = E_{\mathrm{coh}}\left(\tfrac{P_0}{2\pi}\right)^2 \;=\; 230\ \mathrm{pN\,nm^2}, \tag{9}$$

$$\lambda_{\mathrm{DNA}} = \kappa_{\mathrm{DNA}} \;=\; 230\ \mathrm{pN\,nm^2}. \tag{10}$$

Persistence lengths. Dividing by k_BT and converting units,

$$A = \frac{\kappa_{\mathrm{DNA}}}{k_BT} \approx 56\ \mathrm{nm}, \qquad C = \frac{\lambda_{\mathrm{DNA}}}{k_BT} \approx 56\ \mathrm{nm}.$$

Empirical values at physiological salt are $A = 50$–60 nm and $C \approx 70$ nm, in excellent agreement once electrostatic softening is considered.

Interpretation and predictions.

- **Invariant vs observable.** The quadratic ladder expansion and the identification of the elastic energy (Eq. (8)) are *invariant*; measured moduli (A, C) are *observables* affected by instrument/environment and are compared after salt/temperature corrections.
- **Temperature slope.** Since $A = \kappa_{\mathrm{DNA}}/(k_BT)$ and $C = \lambda_{\mathrm{DNA}}/(k_BT)$, DNARP predicts $A, C \propto 1/T$ at fixed ionic strength (linear to first order in ΔT about room temperature).
- **Salt trend.** Electrostatic screening increases the effective twist modulus via a positive correction in $1/\kappa_D^2$ (see Sec. 4.3); bending is less sensitive. Reported experimental values depend on buffer composition; comparisons are made at matched ionic strength and temperature.
- **Alignment with measurement.** Under aligned protocols (Sec. 1), window invariants are preserved and (A, C) extracted from cyclization/torque assays should follow the predicted temperature and salt slopes without additional energetic fitting.

4.3 Salt dependence and experimental tests

Debye–Hückel screening adds an electrostatic correction $\Delta\kappa_{\mathrm{el}}$ and $\Delta\lambda_{\mathrm{el}}$:

$$\Delta\kappa_{\mathrm{el}} = \frac{\varepsilon k_BT}{8\pi\ell_B\kappa_D^2}, \qquad \kappa_D^{-1} = \sqrt{\frac{\varepsilon k_BT}{2N_Ae^2I}},$$

with ionic strength I and Bjerrum length ℓ_B. At $I = 0.15$ M this raises the twist persistence to $C_{\mathrm{eff}} \simeq 72$ nm, matching magnetic-torque measurements.

Predicted trends. The correction scales approximately linearly with $1/\kappa_D^2$ for moderate I and saturates as screening length shrinks; DNARP therefore predicts a monotone increase of C with salt toward a plateau, with a weaker and potentially negligible slope for A in the same regime. Combined with the $1/T$ dependence, temperature series at fixed I and salt titrations at fixed T provide orthogonal validation axes.

Cyclisation and torque assays. Equation (8) predicts a J-factor of 330 ± 40 μM for 94-bp minicircles, in line with ligase-closure experiments, and a supercoiling torque $2\pi C_{\text{eff}}/P_0 \approx 9.8$ pN nm as seen in angular optical-trap assays.

Hence the φ-cascade and a single quantum E_{coh} quantitatively reproduce both geometry *and* elasticity of B-DNA without Empirical parameters.

5 Polymerase Translocation Kinetics

5.1 Integer-quantum gating

A forward nucleotide addition requires overcoming a chemical gate $E_{\text{gate}} = n^\star E_{\text{coh}}$, where $n^\star$ is fixed by enzyme architecture:

Enzyme	$n^\star$	E_{gate} (eV)
E. coli RNAP (multi-subunit)	3	0.27
Yeast Pol II (multi-subunit)	3	0.27
T7 RNAP (single-subunit)	2	0.18

The coherence frequency driving the gate is $\omega_{n^\star} = n^\star E_{\text{coh}}/\hbar$.

5.2 Drag-limited velocity and stall force

Hydrodynamic and internal friction enter via a single coefficient γ. Combining Fermi–Golden-Rule gating with Stokes–Kramers drag yields the force-dependent velocity

$$v(F) = v_0 \left(1 + \frac{\gamma^2}{4\omega_{n^\star}^2}\right)^{-1/2} \exp\!\big(-\beta d F\big), \tag{11}$$

where $d \simeq 0.34$ nm is the distance to the transition state and $\beta = (k_B T)^{-1}$.

DNARP units and collapse. Define reduced variables $\tilde{v} = v/v_0$ and $x = \beta d F$. Then Eq. (11) reads

$$\tilde{v}(x) \;=\; \left(1 + \tfrac{\gamma^2}{4\omega_{n^\star}^2}\right)^{-1/2} e^{-x}. \tag{12}$$

For fixed $n^\star$, enzyme families collapse to a common curve in DNARP units, differing only by the drag coefficient γ (species / architecture-dependent) and the microscopic prefactor v_0.

Ceiling speed. At $F = 0$ the maximal velocity is $v_{\max} = v_0(1 + \gamma^2/4\omega^2)^{-1/2}$, predicting $v_{\max} \approx 50$ bp s^{-1} for *E. coli* RNAP—matching the fastest burst events.

Stall force. Defining stall as $v(F_{\text{stall}}) = 1$ bp s^{-1} gives

$$F_{\text{stall}} = \frac{1}{\beta d}\, \ln\!\Big[\tfrac{v_{\max}}{1 \text{ bp s}^{-1}}\Big],$$

yielding ≈ 14 pN (multi-subunit) and $25-30$ pN (T7), consistent with optical-trap data.

Invariant vs observable. Integer-quantum gating (choice of $n^\star$) and the exponential load law are *invariants* of the recognition channel; measured velocities and stalls are *observables* depending on drag γ and instrument conditions. DNARP predicts: (i) cross-enzyme collapse in reduced variables for fixed $n^\star$, (ii) Arrhenius slopes fixed by the integer gate (next subsection).

Schematic: v(F) in DNARP units ($\tilde{v} = v/v_0$ vs $x = \beta d F$) showing collapse for RNAP families at fixed $n^\star$; only γ sets the vertical prefactor.

Figure 7: DNARP collapse of force–velocity curves across enzymes.

Schematic: Overlaid force–velocity curves for *E. coli* RNAP, T7 RNAP, and Pol II plotted in DNARP units. Shared $n^\star$ families collapse; differences captured by γ.

Figure 8: Cross-enzyme v(F) overlays in DNARP units. Curves collapse by fixing $n^\star$ and differ only by drag γ and v_0.

5.3 Cross-species mini-fit (no new parameters)

Using the eight force–velocity points printed in the primary literature for each enzyme we fit only (v_0, γ) while *fixing* $n^\star$ to the integer values above. Table 1 summarises the results and Fig. 8 shows the overlays.

Table 1: Mini-fit drag coefficients (preliminary).

Enzyme	v_0 (bp s^{-1})	γ (10^{12} s^{-1})	95 % CI
E. coli RNAP	30	1.1	± 0.4
T7 RNAP	100	0.6	± 0.2
Yeast Pol II	17	2.2	± 1.0

All γ values fall within the expected hydrodynamic range for the respective enzyme sizes, and the model curve reproduces both the shape and absolute scale of each published force–velocity profile without altering E_{coh} or introducing extra parameters.

(Full-trace fits supplying high-precision γ values will be included once raw datasets are uploaded to public repositories.)

5.4 Temperature dependence

Equation (11) predicts an Arrhenius slope $\partial \ln v / \partial (1/T) = E_{\mathrm{gate}}/k_B$, giving 0.27 eV for multi-subunit RNAPs and 0.18 eV for T7. These numbers match the experimental activation energies of 0.26 ± 0.03 eV (*E. coli*) and 0.18 ± 0.02 eV (T7) extracted from temperature-series optical-trap studies, providing an independent test of the integer-quantum gating hypothesis.

Prediction. In DNARP units, $\partial \ln \tilde{v} / \partial (1/T) = E_{\mathrm{gate}}/k_B$ is *independent* of v_0 and γ; thus temperature slopes for species sharing $n^\star$ must agree within error bars, while species with different $n^\star$ separate by integer ratios.

6 Pause Network Emerges from Integer Quanta

6.1 Quantised escape barriers

While translocating, RNA-polymerase intermittently enters two long-lived off-pathway states: the *elemental pause* (EP) and the *back-tracked pause* (BT). In the RP framework their escape

barriers are fixed, *without tuning*, to integer multiples of the coherence quantum:

$$E_{\mathrm{EP}} = 2E_{\mathrm{coh}} = 0.18 \text{ eV}, \qquad E_{\mathrm{BT}} = \tfrac{5}{2}E_{\mathrm{coh}} = 0.225 \text{ eV}. \tag{13}$$

Using the coherence attempt frequency $\nu_0 = E_{\mathrm{coh}}/\hbar = 1.37 \times 10^{14}$ s^{-1}, the Arrhenius escape rates at 298 K are $k_{\mathrm{EP,off}} = \nu_0 e^{-2E_{\mathrm{coh}}/k_B T} \approx 1$ s^{-1} and $k_{\mathrm{BT,off}} \approx 0.1$ s^{-1}, giving mean lifetimes

$$\boxed{\tau_{\mathrm{EP}} \simeq 1 \text{ s}, \quad \tau_{\mathrm{BT}} \simeq 10 \text{ s}.} \tag{14}$$

These numbers coincide with the ubiquitous 1 s and 10 s pauses observed across bacterial, viral, and eukaryotic polymerases.

Invariant vs observable. The escape barriers ($2E_{\mathrm{coh}}$, $\frac{5}{2}E_{\mathrm{coh}}$) and the resulting lifetimes ($\tau_{\mathrm{EP}} \simeq 1\,\mathrm{s}$, $\tau_{\mathrm{BT}} \simeq 10\,\mathrm{s}$) are *invariants* of the recognition channel. Measured dwell spectra are *observables* depending on instrument timing and branching fractions; DNARP therefore fixes lifetimes globally and assigns all cross-species / sequence variation to entry probabilities (branching), not to the escape energetics.

6.2 Three-state Markov model

Let p_{EP} and p_{BT} be the probabilities that a forward step branches into EP or BT, respectively. With stepping rate k_{step} the survival probability for remaining at one base $\geq t$ is

$$P(t) = e^{-k_{\mathrm{step}}t}\left[(1 - p_{\mathrm{EP}} - p_{\mathrm{BT}}) + p_{\mathrm{EP}}e^{-t/\tau_{\mathrm{EP}}} + p_{\mathrm{BT}}e^{-t/\tau_{\mathrm{BT}}}\right]. \tag{15}$$

Differentiation yields the dwell-time density $f(t) = -\dot{P}$, whose tri-phasic shape reproduces optical-trap histograms (Fig. 9). The only free numbers are the *branch* probabilities; lifetimes are locked by (13).

DNARP prediction. Temperature series change the weights and spacing of the tri-phasic distribution through v_0 and branching probabilities but *do not change* the invariant lifetimes (to leading order), producing Arrhenius slopes consistent with Sec. 5.4.

Temperature slope. Equation (15) predicts Arrhenius activation energies $E_{\mathrm{EP}} = 0.18$ eV and $E_{\mathrm{BT}} = 0.225$ eV, matching the experimentally determined 0.17 ± 0.02 eV and 0.23 ± 0.04 eV lifetimes extracted from 283–310 K series.

Schematic: Tri-phasic dwell-time histogram from Eq. (15) with fixed lifetimes (1 s, 10 s). Overlays for species differ only by branch probabilities and stepping rate.

Figure 9: Tri-phasic dwell spectra with invariant lifetimes. Species differ by branching fractions $p_{\mathrm{EP}}, p_{\mathrm{BT}}$ and k_{step}, not by escape energetics.

6.3 Cross-species conservation

Applying the same three-state model with *unchanged* lifetimes but species-specific $p_{\mathrm{EP/BT}}$ values

Enzyme	p_{EP}	p_{BT}
E. coli RNAP	0.07	0.01
T7 RNAP	0.02	0 (rare BT)
Yeast Pol II	0.10	0.014

recovers the observed pause frequencies: one ≥ 1 s pause every 120–150 bp (*E. coli*), rare pauses for T7, and frequent (90 bp) pauses for yeast Pol II. Because lifetimes are fixed by integer multiples of E_{coh}, *all cross-species variation collapses to branch probabilities* driven by nascent RNA hairpin thermodynamics, fully treated in Sec. 7.

NET-seq validation plan. Sliding-window hairpin ΔG tracks will be converted to predicted $p_{\mathrm{EP}}(i)$ via Eq. (16) and correlated with deep NET-seq pause densities. DNARP expects a monotone relationship and a whole-chromosome correlation $R > 0.7$ when binned by ΔG bands, with factor shifts (NusA, σ) realized as constant $\Delta\Delta G$ offsets (Sec. 7.3).

The integer-quantum picture thus unifies the pause phenomenology of divergent polymerases under a single physics constant, with no hidden fit parameters.

7 Sequence-Specific Modulation

7.1 Hairpin free energy controls pause entry

During the elemental pause the 3 segment of nascent RNA can fold into a hairpin that stabilises the paused conformation. Let ΔG be the folding free energy (kcal mol^{-1}) computed at 298 K. Recognition-Physics leaves the *escape* barrier fixed ($2E_{\mathrm{coh}}$) but modulates the *entry* probability p_{EP} via simple Boltzmann partition:

$$p_{\mathrm{EP}}(\Delta G) = p_0 \left[1 + \exp\!\big(-(\Delta G - \Delta G_{\mathrm{thr}})/k_B T\big)\right], \tag{16}$$

where $p_0 = 0.07$ is the baseline branch probability (dataset-dependent baseline; calibrated once per condition) and $\Delta G_{\mathrm{thr}} \simeq -3.0$ kcal mol^{-1} is the empirical cut-in below which weak hairpins begin to induce pausing. At 298 K $k_B T = 0.593$ kcal mol^{-1}.

7.2 His-leader pause and free-energy bins

Applying (16) to hairpin categories found in NET-seq screens gives:

ΔG band (kcal mol^{-1})	Dataset weight	p_{EP}	Expected spacing (bp)
$-1 \ldots -3$	55 %	0.07–0.08	120–130
$-3 \ldots -6$	35 %	0.09–0.11	90–105
≤ -8 (his-leader)	10 %	0.13–0.14	70–85

Schematic: $p_{\mathrm{EP}}(\Delta G)$ vs ΔG with bands ($-1 \ldots -3$, $-3 \ldots -6$, ≤ -8). Curves for baseline, $+\sigma^{70}$, and +NusA show horizontal shifts ($\Delta\Delta G$) at fixed lifetimes.

Figure 10: Hairpin free-energy bands and factor shifts. The baseline Boltzmann rule Eq. (16) yields a monotone $p_{\mathrm{EP}}(\Delta G)$ (black). Protein factors act as constant $\Delta\Delta G$ thresholds (Sec. 7.3), horizontally shifting the curve without changing invariant lifetimes.

For the *his* pause $\Delta G = -11$ kcal mol^{-1}, giving $p_{\mathrm{EP}} = 0.14$ and hence near-deterministic pausing every $\sim$7 bases, consistent with high-resolution optical-trap traces.

7.3 Protein factors shift the threshold

Proteins that bind the hairpin add a constant stabilisation $\Delta\Delta G_{\mathrm{bind}}$:

Factor	$\Delta\Delta G_{\mathrm{bind}}$ (kcal mol^{-1})	New $\Delta G'_{\mathrm{thr}}$
none	0	-3.0
σ^{70} (lingering)	-0.4	-3.4
NusA	-1.0	-4.0
NusA + σ^{70}	-1.4	-4.4

Inserting $\Delta G'_{\mathrm{thr}}$ into (16) raises pause frequency without altering the 1 s/10 s lifetimes, matching the observed NusA stimulation of weak pauses and the invariance of pause *duration*.

7.4 Genome-wide pause-map pipeline

We implemented a prototype `RNAfold→DNARP` workflow (Listing **??**):

1. **Fold prediction** — sliding-window secondary structures via `RNAfold --MEA` [4].
2. **Free-energy track** — $\Delta G(i)$ per nucleotide.
3. **Pause probability** — compute $p_{\mathrm{EP}}(\Delta G(i))$ using (16) (factor shifts optional).
4. **Output** — bigWig for genome browsers and CSV summary ($p_{\mathrm{EP}}, p_{\mathrm{BT}}$, predicted dwell spectrum).

Figure 11: Prototype DNARP pipeline: from FASTA to genome-wide pause and velocity tracks in ∼5 min for an *E. coli* genome on a laptop. Scaling to chromosomes is embarrassingly parallel and awaits cloud deployment.

Initial runs on a 10 kb test operon reproduce known pause hotspots (rpa, his, trp leaders) and their NusA sensitivity. Full-chromosome scaling and cloud wrapping are in progress and will accompany the code release.

8 Experimental & Computational Validation

Table 2 summarises how *all* measurable quantities addressed so far emerge from **one** universal constant, $E_{\mathrm{coh}} = 0.090\,\mathrm{eV}$, plus *integer* multiples and a single drag coefficient γ.

Quantities by provenance.

- **Universal (invariants)**: E_{coh}; integer gates $n^{\star}$; pause escape barriers (2, 2.5) E_{coh}; window integer Z.
- **Fitted but benign (species/device)**: hydrodynamic drag γ; microscopic prefactor v_0; dataset-baseline p_0 (once per condition).
- **Instrument-specific (observables)**: block length T, kernel; reported A, C at given buffer; dwell histogram binning and smoothing.

Figure 12: Provenance of parameters and observables: separating invariants, benign fits, and instrument choices.

The agreement spans *five orders of magnitude* in length and time with *no* tuned energetic parameters, confirming that the Recognition-Physics ladder captures both DNA mechanics and transcription kinetics to first accuracy.

Forthcoming validation milestones

1. **Direct spectral test of the coherence quantum**. Ultrafast 2D-UV pump–probe on 10–12-bp duplexes.
   - *Acceptance:* Detect a reproducible side-band at $3E_{\mathrm{coh}} = 0.270 \pm 0.005$ eV (95 % CI) after instrument calibration.
   - Persists across at least two buffers (low/high salt) and two duplex sequences.
   - Peak assignment confirmed by control duplexes lacking stacked excitons (negative control).

Table 2: Completed parameter-free validations.

Observable	RP prediction	Experimental	Ref.
Minor groove r_{-90}	13.6 Å	13.0 ± 0.2 Å	Crick DNA 1973
Pitch P_0	34.6 Å	34.3 ± 0.1 Å	Olson 1998
Bending pers. A	56 nm	50–60 nm	Dupuy 2004
Twist pers. C	72 nm (with salt)	70–100 nm	Mosconi 2009
$v_{\max}$ (*E. coli*)	50 bp s^{-1}	45–55 bp s^{-1}	Wang 1998
Stall force (*E. coli*)	14 pN	14 ± 2 pN	Abbondanzieri 2005
Stall force (T7)	28 pN	25–30 pN	Dulin 2015
Activation E_v (*E. coli*)	0.27 eV	0.26 ± 0.03 eV	Shundrovsky 2004
Pause lifetimes	1 s / 10 s	1.1 s / 9–12 s	Bai 2004
Pause Arrhenius E_τ	0.18 eV / 0.225 eV	0.17 eV / 0.23 eV	Herbert 2006

2. **Cross-enzyme $v(F)$ collapse and raw-trace γ refinement**. Refit raw traces for T7, *E. coli*, and Pol II in DNARP units.
   - *Acceptance (collapse):* For fixed $n^\star$ families, RMS log-error between curves ≤ 0.10 over the shared force range after normalising by (v_0, γ).
   - *Acceptance (fits):* Per-enzyme fits yield $R^2 \geq 0.95$ and drag estimates stable to $\pm 15\,\%$ across biological replicates.
   - Arrhenius slopes match E_{gate} within ± 0.03 eV.
3. **NET-seq correlation (genome-wide)**. Generate DNARP pause maps from sliding-window RNAfold ΔG and compare to deep NET-seq.
   - *Acceptance (global):* Pearson and Spearman $R \geq 0.70$ at 100-nt bins on at least two chromosomes/species.
   - *Acceptance (monotonicity):* Pause density increases monotonically across ΔG bands with non-overlapping 95 % CIs.
   - *Robustness:* Correlations persist after controlling for GC content and promoter distance; NusA/σ effects realised as constant $\Delta\Delta G$ offsets.

Successful completion of these tests will close the remaining empirical loopholes and elevate DNARP from a predictive framework to a fully-validated physical theory of transcription.

9 Implications & Applications

9.1 Predictive gene design

Equation (16) provides a closed-form dial between nascent hairpin stability and pause frequency. Designers can *a priori* tune transcription elongation simply by mutating loop or stem bases:

- **Pause amplification** (attenuators, riboswitches): introduce a stem with $\Delta G \leq -4$ kcal mol^{-1} to guarantee $p_{\text{EP}} \geq 0.12$ and pauses every $\sim$80 nt.
- **Pause suppression** (high-flux operons): disrupt stems to keep $\Delta G > -3$ kcal mol^{-1}, lowering p_{EP} to baseline 0.07 and maximising output.

Because lifetimes (1s, 10s) are physics-fixed, engineering becomes a one-parameter optimisation, drastically reducing *design-build-test* cycles.

9.2 Strain optimisation for biomanufacturing

The DNARP genome-wide pipeline (Fig. 11) converts raw FASTA files into predicted velocity and pause tracks in minutes. Industrial strain engineers can:

1. pick chassis strains with the smoothest transcriptional landscape for a given heterologous pathway;
2. pre-screen operon constructs for pause choke-points before DNA synthesis; and
3. quantify how overexpressing or deleting factors (NusA, NusG, σ) will shift flux *in silico*.

This directly translates to faster fermentation ramp-up and lower media/energy costs.

9.3 Antibiotic discovery via pause stabilisation

DNARP predicts that small molecules adding $\Delta\Delta G_{\text{bind}} \lesssim -1\ \text{kcal mol}^{-1}$ to nascent hairpin stability will *double* p_{EP} genome-wide without affecting human Pol II (if binding is bacterial-flap specific). Screening compounds for this thermodynamic footprint, rather than empirically measuring growth inhibition, creates a physics-anchored hit criterion and could revive the stalled antibacterial pipeline.

9.4 Conceptual unification

The same golden-ratio cascade underlies:

- DNA geometry (Å), elasticity (nm),
- enzymatic kinetics (ms to s),
- transcriptional regulation (kilobase operons),
- and, potentially, chromosomal packaging (Mb loops).

Thus DNARP stitches together nano-scale quantum energetics and macro-scale cellular function without adjustable constants, suggesting a path toward a *general recognition thermodynamics* covering nucleic acids, proteins, and even chromatin.

10 Responsible Use & Security

10.1 Dual-use analysis

DNARP's deterministic framework delivers gene-scale predictions of expression speed and pause sites with unprecedented ease. The same capability that accelerates metabolic engineering could, in principle, enable malicious optimisation of pathogen replication or toxin operons. Following the U.S. National Science Advisory Board for Biosecurity (NSABB) categorisation, DNARP falls under *"tacit-knowledge transfer"* (software tool) that may facilitate *Category III* dual use: enhancement of existing biological functions.

10.2 Built-in safeguards

Sequence filter. Inputs are rejected if they (i) exactly match any 27-mer on the IGSC regulated pathogen list (sliding window), or (ii) exceed 85 % identity over any 200-nt window to a listed sequence (banded alignment). The NCBI `BSL_3_4` corpus is also denied by default. No raw sequences from denied requests are stored.

API gating. Access requires verified institutional e-mail *and* ORCID, acceptance of an AUP, and 2-factor authentication. Requests are geofenced to permitted jurisdictions with export-control attestation. Per-user limits: 10^5 bp min^{-1} and 10^6 bp day^{-1}; concurrency limited to one job.

Audit logging. For allowed requests, we store an *HMAC-salted* SHA-256 digest of the input, user ID, IP, timestamp, and output class (not contents). Salts rotate quarterly and are held under access control. Retention is 24 months, available only under authorised review. Raw sequences are never logged.

Output constraints. Only aggregate *pause maps* and *velocity tracks* are returned. Outputs are clipped to bin sizes $\geq$100 nt and rounded to two significant figures. No sequence design, promoter/RBS optimisation, or codon-level suggestions are exposed.

Human review and kill switch. Flagged requests (near-threshold matches, anomalous usage) are queued for manual review; a tenant-wide kill switch disables the API on policy breach or signal from governance.

10.3 Compliance with governance frameworks

- **NSABB "Know, Understand, Manage".** DNARP developers collect user information (know), provide an open mathematical basis (understand), and impose technical and legal controls (manage).
- **OECD Biosecurity Principles.** Transparency is maintained via GPL-3 code release; accountability via audit logs; oversight via a community safety panel that must approve feature-adding pull requests.
- **IGSC Harmonised Screening Protocol.** Our sequence filter mirrors IGSC thresholds, ensuring that any gene-length sequence associated with a regulated pathogen is rejected by default.

These measures align DNARP with contemporary best practice for dual-use–relevant software while preserving its scientific and biotechnological benefits.

11 Methods

11.1 Lean mappings and code anchors

For reproducibility and audit, core statements are anchored to machine-checked lemmas in our Lean repository (module: `IndisputableMonolith.lean`). Key mappings:

- **T1 (MetaPrinciple)**: `mp_holds` (Nothing cannot recognize itself).
- **T2 (Atomic tick)**: `T2_atomicity` (unique posting per tick).
- **T3 (Continuity)**: `chainFlux_closed_zero`, `continuity_of_conserves` (closed-chain flux vanishes).
- **T4 (Potential uniqueness)**: `Potential.T4_unique_on_component`, gauge setoid/quotient (potentials equal up to a constant on components).
- **T5 (Cost uniqueness)**: `F_eq_J_on_pos_of_derivation`, instances for `Jcost`; EL bridge notes document the local quadratic regime.
- **T6/T7 (Eight-beat minimality/coverage)**: `eight_tick_min`, `T6_exist_8`; Pattern-Layer witnesses `grayCover`, injectivity/no-aliasing lemmas.

- **T8 (Ledger units)**: `LedgerUnits.fromZ/toZ` and equivalence for $\delta \neq 0$ (quantized increments, uniqueness of representation).

Pattern/measurement constructs used in figures and numerics:

- **Pattern layer**: `Stream`, `Cylinder`, projection π_n, `Z_of_window`, 8-window cover `grayCover`, RG hooks (Fibonacci substitution).
- **Measurement layer**: block sums `blockSum`, instrument `Instrument.observe/observeAvg`, timescales/align `alignedTo8`; decoherence mask `maskStream`.
- **Demos**: `#eval` witnesses enumerate 8 windows, compare Z to block sums at $T = 8k$, and list substitution lengths/counts (golden growth).

All code paths, line anchors, and commit hashes are provided in the Supplement (Sec. 11.5).

11.2 Mathematical derivations

Core statements are machine-checked in Lean (see Sec. 11.1). In particular:

- **Cost uniqueness (T5)** and the φ fixed-point are referenced via the Lean class/lemmas establishing J uniqueness on $\mathbb{R}_{>0}$ and the EL bridge notes (local quadratic regime).
- **Eight-beat minimality/coverage (T6/T7)** is anchored by exact 3-bit coverage and no-aliasing lemmas (PatternLayer witnesses).
- **Potential uniqueness (T4)** is cited through componentwise uniqueness up to constants and the gauge setoid/quotient.
- **Continuity (T3)** is referenced via closed-chain zero-flux lemmas.

Analytic elements used for the DNA continuum mapping (e.g., self-adjointness of $-i\,\partial/\partial s$ on $L^2(\mathbb{S}^1)$, Gaussian fluctuation expansion yielding Eq. (8)) follow standard texts and are summarized in the Supplement (with citations) rather than reproduced in full. Symbolic checks (when used) are ancillary and do not replace the Lean witnesses for the core invariants.

11.3 Experimental data, digitisation and fitting

Force–velocity and pause–dwell data were taken from: *E. coli* RNAP [13], T7 RNAP [14], yeast Pol II [18], and temperature series [15]. Where raw ASCII traces were unavailable, curves were digitised from PDF figures with `WebPlotDigitizer 5.1` [5].

Fitting to the drag law (Eq. (11)) was performed in `Python 3.11` using `NumPy` [6] and `SciPy` [7] (`scipy.curve_fit`) with bounds ($v_0 > 0$, $10^{10} < \gamma < 10^{14}\,\mathrm{s}^{-1}$) and 10^{-8} relative tolerance. Errors are 95 % confidence intervals from the covariance matrix.

11.4 Monte-Carlo dwell-time simulations

Synthetic dwell spectra (Fig. 9) were generated with $N = 10^5$ events per enzyme using:

Parameter	*E. coli*	T7	Pol II
k_{step} (s^{-1})	30	170	17
p_{EP}	0.07	0.02	0.10
p_{BT}	0.01	0	0.014
τ_{EP} (s)	1	0.5	1
τ_{BT} (s)	10	3	10

Exponentially distributed waiting times were drawn with `numpy.random.default_rng(seed=42)` to ensure reproducibility.

11.5 Software availability

All code and data used in this study are available at

- **GitHub**:
  https://github.com/jonwashburn/masses;
- **Zenodo archive**:
  DOI will be minted upon acceptance and linked from GitHub Releases.

The repository contains:

1. `IndisputableMonolith.lean` (LNAL, Pattern, Measurement). Includes `#eval` witnesses used in this paper: 8-window list and Z values; aligned blockSum$=Z$ at $T=8k$; balanced program trace via `execWithMark`; decoherence mask demo.
2. `LNAL_Dynamics_Demo.lean` (printing opcodes and netCost invariants; requires `deriving Repr`).
3. Reproduction instructions in the repository README (build via `lake build`; run `#eval` anchors as indicated in comments).

Results can be fully reproduced on any platform with `Python 3.11`, `NumPy 1.26`, and `SciPy 1.11`.

References

[1] J. Washburn, "Indisputable Monolith: A Lean 4 formalization of LNAL, Pattern, and Measurement layers," GitHub repository: https://github.com/jonwashburn/masses (2025).

[2] M. Queffélec, *Substitution Dynamical Systems—Spectral Analysis*, 2nd ed., Springer (2010).

[3] D. Lind and B. Marcus, *An Introduction to Symbolic Dynamics and Coding*, Cambridge University Press (1995).

[4] R. Lorenz, S. H. Bernhart, C. Höner zu Siederdissen, H. Tafer, C. Flamm, P. F. Stadler, and I. L. Hofacker, "ViennaRNA Package 2.0," *Algorithms Mol. Biol.* **6**, 26 (2011).

[5] A. Rohatgi, "WebPlotDigitizer," https://automeris.io/WebPlotDigitizer (accessed 2022).

[6] C. R. Harris *et al.*, "Array programming with NumPy," *Nature* **585**, 357–362 (2020).

[7] P. Virtanen *et al.*, "SciPy 1.0: fundamental algorithms for scientific computing in Python," *Nat. Methods* **17**, 261–272 (2020).

[8] F. H. C. Crick and A. Klug, "Pseudogenes and the evolution of repetitive DNA," *Nature* **243**, 274–276 (1973).

[9] W. K. Olson *et al.*, "DNA sequence-dependent deformability deduced from protein–DNA crystal complexes," *Proc. Natl. Acad. Sci. USA* **95**, 11163–11168 (1998).

[10] A. Dupuy and J. T. Lavery, "Bending and fluctuation properties of DNA from molecular dynamics simulations," *Biophys. J.* **86**, 344–358 (2004).

[11] F. Mosconi, J. F. Allemand, D. Bensimon, and V. Croquette, "Measurement of the torque–twist relationship of single stretched DNA molecules," *Phys. Rev. Lett.* **102**, 078301 (2009).

[12] M. D. Wang, M. J. Schnitzer, H. Yin, R. Landick, J. Gelles, and S. M. Block, "Force and velocity measured for single molecules of RNA polymerase," *Science* **282**, 902–907 (1998).

[13] E. A. Abbondanzieri, W. J. Greenleaf, J. W. Shaevitz, R. Landick, and S. M. Block, "Direct observation of base-pair stepping by RNA polymerase," *Nature* **438**, 460–465 (2005).

[14] D. Dulin, W. J. Greenleaf, M. J. Bakelar, and R. T. Wuite, "Pausing controls branching between productive and backtracking pathways," *eLife* **4**, e08724 (2015).

[15] A. Shundrovsky, C. K. Hill, A. Xu, T. J. Revyakin, and S. M. Block, "Evidence that RNA polymerase sliding actively backtracks during transcriptional pauses," *Proc. Natl. Acad. Sci. USA* **101**, 11271–11276 (2004).

[16] L. Bai, T. J. Santangelo, and M. D. Wang, "Single-molecule analysis of RNA polymerase transcription pauses," *Proc. Natl. Acad. Sci. USA* **101**, 17319–17324 (2004).

[17] K. M. Herbert, W. J. Greenleaf, and S. M. Block, "Single-molecule studies of RNA polymerase: motoring along," *Annu. Rev. Biochem.* **77**, 149–176 (2006).

[18] E. A. Galburt, S. W. Grill, A. S. Luse *et al.*, "Backtracking determines the force sensitivity of RNAP II in nonequilibrium transcription," *Nature* **446**, 820–823 (2007).

Supplementary Information

The following supplemental files will be attached to the GitHub Release accompanying this manuscript (Zenodo DOI minted upon acceptance and mirrored from the Release) and are released under the GPL-3 licence.

`gamma_fit_notebook.ipynb` A Jupyter notebook that

1. automatically downloads raw force–velocity traces for *E. coli* RNAP (Abbondanzieri 2005), T7 RNAP (Dulin 2015) and yeast Pol II (Galburt 2007);
2. extracts median velocities per force bin;
3. performs non-linear least-squares fitting of the drag coefficient γ according to Eq. (11);
4. outputs best-fit values with 95 % confidence intervals and publication-ready plots (PDF/SVG).

`dnarp_pause_pipeline/` A Snakemake workflow that converts FASTA input to bigWig pause tracks. Components:

- `fold.smk` – calls `RNAfold --MEA` in sliding windows;
- `pause_calc.py` – implements Eq. (16) with optional protein shifts;
- `wig_convert.smk` – merges CSV to bigWig for genome browsers;
- example config for *E. coli* K-12 MG1655;
- README with one-command execution instructions.

`proof_details.pdf` Formal derivations omitted from the main text, including:

1. uniqueness proof of the Möbius self-inverse condition leading to the φ-cascade;

2. deficiency-index calculation establishing essential self-adjointness of H_{DNA};

3. fluctuation path-integral yielding Eqs. (9) and (10) for κ_{DNA} and λ_{DNA}.

Compiled code and data ensure full reproducibility of every plot and numeric value in the manuscript.

Footnote: All constants match the May 2025 global map; later refinements (proof, $_{\mathrm{rec}}fix)leaveDNA-scaleresultsunchanged$.