Aging as irreversible information decay: a governing law, and the absence of a single-cell melting signature in human blood

Nathan Hangen[1]

[1] *Independent Researcher*

(Dated: July 23, 2026)

We model the aging epigenome as a block-structured binary codeword copied, at every cell division, through a noisy channel with no backup: DNA methyltransferase-1 re-templates the daughter strand from the parent strand, so a copy error is propagated, not corrected. Under this single assumption the mutual information between an organism's birth epigenetic state and its current state, $I(t) = \mathrm{MI}(\mathrm{birth}; \mathrm{state}_t)$, is a strict monotone that can only decrease—aging as the data-processing inequality applied to a cell lineage. We reduce the model to one relaxational law on an identity-contrast order parameter, $\tau_0 \dot{M} = JM - T\,\mathrm{artanh}(M) + h_{\mathrm{ext}}$, and show that its symmetry structure (a $\mathbb{Z}_2$ "no-reference" symmetry) forbids any internal quantity from restoring I: only an *external* reference—the mechanism partial reprogramming supplies—can raise it. A single dimensionless fate number $\mathcal{N} = J/T_{\mathrm{eff}}$ separates a maintained (ordered) regime from a decaying one and predicts an information wall, a regime where death is reached at any metabolic power. We then test the model's most consequential empirical prediction against public DNA-methylation data. A competing "melting" reading—that aging is a collective, in-principle reversible loss of an ordered epigenetic phase—makes two testable predictions the core model does not: a late-accelerating (plateau-then-collapse) age trajectory, and within-cell disordering at composition-inert, cross-lineage-saturated sites. In human blood, cross-sectionally, we find neither. The age-drift trajectory is front-loaded and decelerating (relaxation-to-floor; saturating/logistic fits beat linear by AIC), and a composition-inert probe—CpG sites saturated in every sorted blood lineage, for which cell-composition mixing is mathematically inert—shows no within-cell melting signature where it would be cleanest to see. We do *not* claim the epigenome universally cannot melt; the definitive test (read-level discordance stratified by cross-lineage saturation) is named as the pending falsifier.

I. INTRODUCTION

Aging biomarkers built on DNA methylation ("epigenetic clocks") [4, 5] are among the most accurate predictors of chronological and biological age, yet the field remains divided on what the underlying clock *is*. Two families of explanation compete. In the first, aging is the gradual accumulation of stochastic copy errors at individual loci—an information-decay picture in which the epigenome slowly forgets its birth state. In the second, aging is the loss of a collective, ordered epigenetic phase—a "melting" picture in which the youthful state is a low-entropy configuration held by cooperative interactions, and aging is the (in-principle reversible) collapse of that order. The two pictures carry very different therapeutic implications: the first says only an external reference can restore lost information; the second says the right collective nudge could re-freeze the melted state.

We make the information-decay picture precise, derive the law it obeys, and confront the melting picture with data.

a. The unbacked copy channel. The physical substrate forces the framing. At each mitosis the maintenance methyltransferase DNMT1 reads the parent (template) strand and writes the daughter strand to match. There is no stored, uncorrupted reference: if the parent strand's methylation state has drifted, the daughter inherits the drift. Formally, every division is one use of a binary symmetric channel (per-site flip probability p), and the lineage is a Markov chain. The data-processing inequality [1] then guarantees that $I(t) = \mathrm{MI}(\mathrm{birth}; \mathrm{state}_t)$ is non-increasing. The only way to raise it is to inject information from outside the chain—an external reference. This is the model's spine, and Sec. II turns it into an equation.

b. Reader, paper, and the palimpsest. The name is literal. A palimpsest is a manuscript scraped and overwritten, the earlier text still faintly legible beneath the new—a record kept with no clean master, only successive rewrites on one surface. The aging epigenome is such a document, and two failure sites degrade it. The *paper*— the methylation marks on the substrate—can be corrupted directly; and the *reader/writer*—DNMT1 transcribing parent strand to daughter—can copy with error, each division rewriting the page slightly wrong. In the law below these are the identity contrast M and the copy-noise term T. Because the reader holds no external exemplar to check against, an error in either, once written, becomes the new template. This fixes what each class of intervention can do: *prevention* means protecting the paper (raising the read–write coupling J so a domain reinforces its own marks) or improving the reader (lowering T)—both slow decay but, by the no-reference symmetry below, cannot recover what is already lost; *restoration* requires importing a reference from outside the lineage, which is precisely what partial reprogramming [11, 12] supplies. The result of this paper is that natural human blood aging is dominated by the reader-error term, not by a melting of the paper's ordered domains.

c. The question. Is measured bulk aging (a) per-locus information decay toward a maximum-entropy floor, (b) collective melting of an ordered phase, or (c) an artifact of the tissue's changing cell composition with age? We show the answer in human blood is dominated by (a) and (c), with no detectable contribution from (b) at the resolution the data allow.

II. A MINIMAL GOVERNING LAW

A. The order parameter

For a block/domain of the epigenome let m be its mean methylation and define the *identity contrast* $M \equiv 2(m-\frac{1}{2}) \in [-1,1]$. $M = \pm 1$ is a fully committed block; $M = 0$ is the maximum-entropy floor. Because the cell keeps no stored reference, the copy channel is symmetric under relabeling $M \to -M$—there is no ground truth to break the symmetry. The only quantity that points back to the birth state is $I(t) = \mathrm{MI}(\mathrm{birth}; \mathrm{state}_t)$, with $I \approx M^2/(2\ln 2)$ for small M.

B. The equation

The model reduces to one overdamped (model-A) relaxational law with an external source. We call it *minimal*, not *fundamental*: the $-T\,\mathrm{artanh}(M)$ term is the mean-field closure of the copy-channel dynamics, a well-motivated modeling choice rather than a first-principles necessity. Its value is that every term maps to an independently established result (Table I), not that it is derived from a microscopic Hamiltonian.

$$\tau_0 \frac{dM}{dt} = -\frac{\partial f}{\partial M} + h_{\mathrm{ext}}(t) = JM - T\,\mathrm{artanh}(M) + h_{\mathrm{ext}}(t), \qquad f(M) = -\tfrac{J}{2}M^2 - T\,s(M), \tag{1}$$

with s the binary entropy.

TABLE I. Every term of Eq. (1) maps to an independently established result.

term	meaning	result it carries
JM	collective read–write reinforcement (self-templating coupling)	ordered phase / reverse polarity
$-T\,\mathrm{artanh}(M)$	copy-channel decay toward the floor $m^* = \frac{1}{2}$	relaxation-to-floor drift (Sec. IV)
T	noise temperature, $T \propto p_{\mathrm{eff}} = p_{\mathrm{raw}}/(p_{\mathrm{raw}} + r)$, floored at $T(\rho)$	Eigen threshold; kinetic rate r; anchor ceiling
$h_{\mathrm{ext}}(t)$	external reference (e.g. OSK reprogramming); $\equiv 0$ internally	reversibility asymmetry

a. Disordered limit ($T > J$, near $M = 0$): $\tau_0 \dot{M} \approx -(T-J)M$, i.e. geometric relaxation-to-floor with decay constant $\kappa = (T-J)/\tau_0$. This is a front-loaded, decelerating age trajectory—the shape Sec. IV confirms—not a youthful plateau followed by late collapse.

b. Ordered limit ($T < J$): the free energy develops wells at $\pm M_0$ with a barrier Δf; the aligned youthful domain sits in a well and escaping it (aging) costs an Arrhenius time. This is the regime the reverse-polarity control law (companion paper) exploits; it is not the regime human blood aging occupies (Secs. IV–V).

C. The symmetry and the monotone

- The *noiseless* copy map is a bijection on codewords—measure-preserving and time-reversible—so I is *conserved*: a Liouville-type invariant (the frozen M_0 in the $T \to 0$ limit).
- *Noise breaks the time-reversal symmetry* of the copy map. The formerly conserved I becomes a strict monotone, $dI/dt \le 0$ (H-theorem / data-processing inequality).
- *Noether-type statement.* The $\mathbb{Z}_2$ no-reference symmetry forbids any *internal* charge that indexes the birth state. The sole candidate, I, is exactly the quantity the broken symmetry converts from conserved to dissipated. Only h_{ext}— external precisely because raising I requires a reference outside the system—can lift it. This is aging's arrow stated as a symmetry, and the reason reprogramming (an external reference) is the only information-restoring intervention the model admits.

a. Honest seam. The free energy $f(M)$ is a Lyapunov function ($df/dt = -(1/\tau_0)(f')^2 \leq 0$) but is *not* the same monotone as I. In the ordered phase f descends into a well while I stays flat—identity *committed*, not lost. The universal monotone is I (DPI); f governs the coarse dynamics. We do not equate them outside the disordered regime.

D. Phase boundaries and the fate number

1. *Maintained vs aged* (Eigen error threshold [2]): $r > p_{\text{raw}}(1-p^*)/p^* \Leftrightarrow p_{\text{eff}} < p^*$.
2. *Ordered / reverse-polarity*: $J > T(p_{\text{eff}})$ *and* spatial dimension ≥ 2 (no order in 1D; confirmed in simulation—1D neighbor consensus cannot sustain order, 2D can).
3. *Information wall* (power-independent death): even at maximal reinforcement $r \to r_{\max}(\rho)$, if $T(\rho) > J$ the ordered phase is unreachable at *any* metabolic input.

The single controlling ratio is the *fate number*

$$\mathcal{N} = \frac{J}{T_{\text{eff}}}, \qquad T_{\text{eff}} \propto p_{\text{eff}} = \frac{p_{\text{raw}}}{p_{\text{raw}} + r}, \text{ floored at } T(\rho). \tag{2}$$

$\mathcal{N} > 1$ means identity held; $\mathcal{N} < 1$, relaxation-to-floor aging; the ceiling $\mathcal{N}_{\max} = J/T(\rho)$—if $\mathcal{N}_{\max} < 1$, death at any metabolic power. $\mathcal{N}$ is nothing exotic mathematically: with $T_c = J$ it is the transition's *reduced temperature* T_c/T_{eff}, the standard order/disorder control ratio of any mean-field transition. The novelty we claim is not the ratio but its *mapping onto aging*: T_{eff} tied to copy noise and metabolic reinforcement, J to read–write chromatin coupling, and $\mathcal{N}$ crossing 1 as the maintained-to-aged transition. A disordered-regime cousin is the epigenetic Damköhler number $\text{Da} = r/p_{\text{raw}}$, critical at $\text{Da}_c = (1-p^*)/p^*$.

III. SIMULATION OF THE CORE MODEL

Simulating the block codeword through repeated binary-symmetric copying reproduces the closed-form results: $I(t)$ decays geometrically toward zero in the disordered regime; the per-locus methylation relaxes to the $m^* = \frac{1}{2}$ floor with the front-loaded shape the law predicts; and there is an Eigen error threshold—above a critical p_{eff}, reinforcement cannot hold identity and the codeword melts to noise regardless of copy fidelity elsewhere. The maintenance module realizes r (local reinforcement) as neighbor consensus with *no pristine reference*—the load-bearing invariant of the model: the only legitimate self-restoring force is local agreement, never comparison to a stored codeword. This is why the model predicts that genuine information *restoration* (as opposed to *retention*) requires h_{ext}.

IV. THE SHAPE DISCRIMINATES CORE FROM MELTING

The two pictures make opposite predictions about the *shape* of the age–drift curve. **Core** (per-locus relaxation): each locus relaxes toward the floor at a rate set by its distance from it, so drift is front-loaded and decelerating—fast early, saturating late. **Melting** (collective barrier): the ordered phase is protected by a cooperative barrier that holds until it yields, giving a youthful plateau then a late collapse.

We define drift $D(\text{age}) = \langle|\beta - \beta_{\text{young}}|\rangle$ on CpGs with a young baseline (baseline age ≤ 5 yr) in the blood-only subset of the aging blood cohort [14], and fit linear, saturating, and logistic forms by AIC.

a. Result. The trajectory is front-loaded: the youth slope ($\approx +1.33 \times 10^{-3}\,\text{yr}^{-1}$) exceeds the mid-life slope ($\approx +1.08 \times 10^{-3}\,\text{yr}^{-1}$); saturating ($k = 0.034$) and logistic ($x_0 \approx 22$ yr) fits both beat linear by AIC. This is the core model's relaxation-to-floor, not the melting picture's plateau-then-collapse.

b. Model-internal discriminator. To confirm the shape is genuinely diagnostic and not an artifact of the fitting, we simulate both mechanisms—core (relaxation) and melting (an Arrhenius barrier with T rising toward T_c)—and score each by the fraction of the trajectory that is front-loaded. The empirical curve is ∼78% front-loaded; simulated core is ∼74%, simulated melting ∼0%. The data selects core.

V. THE NEGATIVE: NO WITHIN-CELL MELTING WHERE COMPOSITION IS EXCLUDED

The age-associated rise in bulk methylation entropy is real and is the melting picture's headline evidence. But bulk methylation is a convex combination over cell types, $\beta_{\text{bulk}} = \sum_i f_i \beta_i$, and the blood cell-type fractions f_i shift

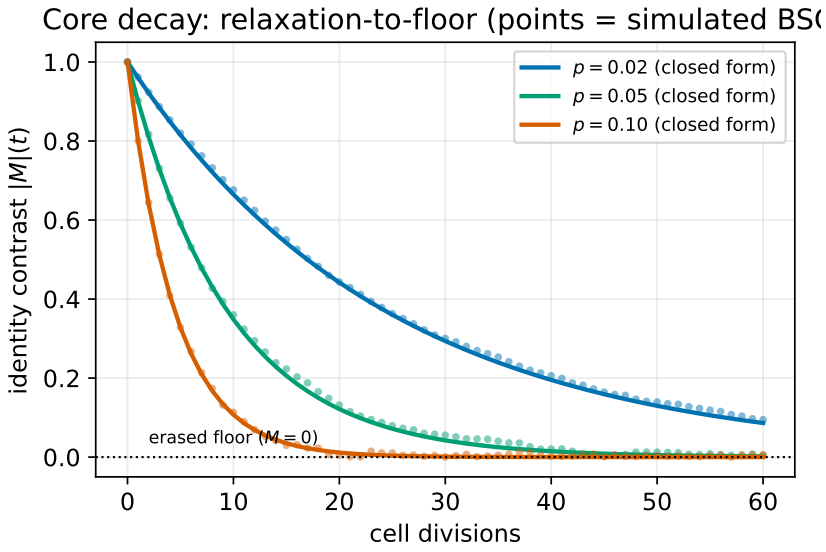

FIG. 1. Core-model relaxation-to-floor. Solid curves are the closed-form $|M|(t) = |1 - 2p|^t$; points are a simulated binary-symmetric lineage (20,000 loci). The decay is front-loaded and decelerating—most identity contrast is lost early—and saturates at the erased floor $M = 0$, with no plateau.

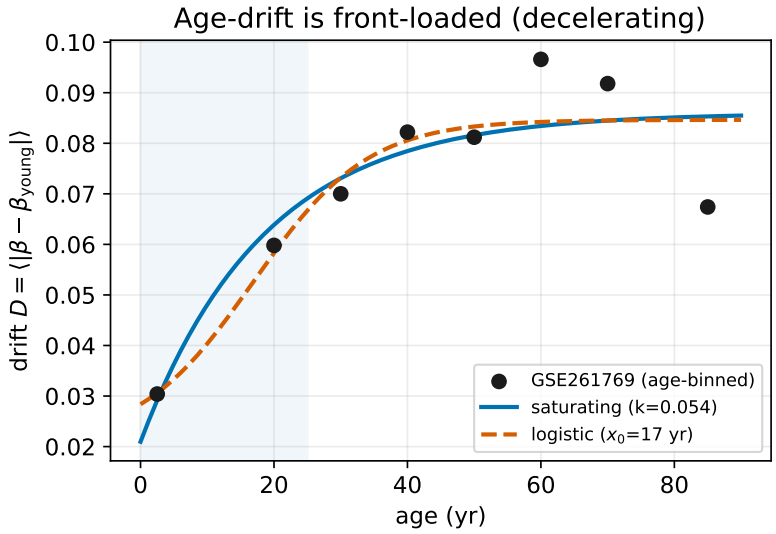

FIG. 2. Age-drift shape in human blood (46,517 CpGs with a trusted young baseline, GSE261769). Drift is front-loaded and decelerating; saturating ($k = 0.034$) and logistic ($x_0 \approx 22$ yr) fits beat linear by AIC.

systematically with age (immunosenescence), and standard reference-based deconvolution removes only part of the apparent effect [7]. A rise in bulk entropy can therefore arise with *no per-cell change at all*, purely from the mixing fractions moving toward intermediate values. Distinguishing within-cell disordering (melting) from between-cell mixing (composition) is the crux, and bulk data cannot do it directly.

a. The composition-inert probe. We exploit a mathematical fact: if a CpG is *saturated* in every sorted lineage—$\beta \geq 0.9$ in all of them ("pan-high") or $\beta \leq 0.1$ in all of them ("pan-low")—then the convex combination $\sum_i f_i \beta_i$ is inert to any change in the fractions f_i, being a weighted average of (near-)equal values. Composition mixing cannot move a pan-invariant site; therefore any age drift there must be genuine within-cell change, while drift at composition-sensitive "variable" sites is a mixture of both. We define pan-invariant and variable sites from six sorted blood lineages in a sorted-cell WGBS atlas [13] (granulocytes, monocytes, CD4-T, CD8-T, NK, B; three replicates each), drop CpGs

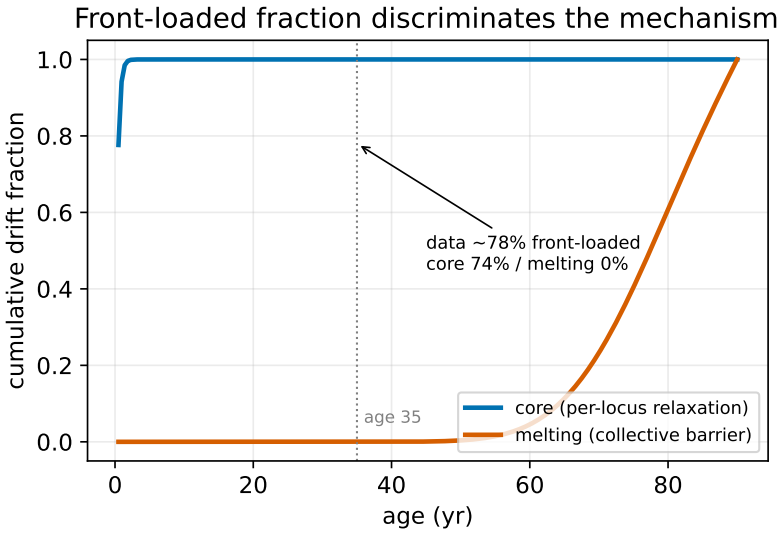

FIG. 3. The shape is diagnostic. Cumulative drift fraction vs age for the two mechanisms: core (per-locus relaxation) is ∼74% complete by age 35; melting (a collective barrier with T rising toward T_c) holds a youthful plateau (∼0% by 35). The data (∼78% front-loaded) matches core.

at or adjacent to common SNPs, and—because coverage noise is a symmetric artifact the pan-invariant contrast does *not* remove—coverage-match every donor by hypergeometric downsampling to a fixed depth and estimate β with a beta-binomial Jeffreys posterior.

b. Result. Pan-invariant sites show no significant age drift. The donor-level OLS slopes ($n = 23$ donors, ages 21–75, coverage-matched) are $+4.1 \times 10^{-5}\,\mathrm{yr}^{-1}$ (pan-high, 95% CI $[-2.5, +10.6] \times 10^{-5}$) and $-3.3 \times 10^{-5}\,\mathrm{yr}^{-1}$ (pan-low, CI $[-13.5, +6.9] \times 10^{-5}$); both CIs include zero. The decisive mixing control is the intermediate-band-fraction slope: pan-invariant $+6.2 \times 10^{-6}\,\mathrm{yr}^{-1}$ (CI $[-6.4, +7.7] \times 10^{-5}$) versus variable $+1.2 \times 10^{-3}\,\mathrm{yr}^{-1}$ (CI $[-1.3, +3.7] \times 10^{-3}$)—the variable point estimate is about 200× larger, the direction expected if composition mixing acts on variable sites, but at $n = 23$ the variable CI *also* spans zero. We are therefore explicit: the confound-immune result is the *pan-invariant flatness* (a tight null at a composition-inert probe), not a demonstrated significant rise at variable sites; the two-arm contrast lives in the point estimates. Stratifying pan-low sites by CpG-island context leaves even the island-like pan-low drift flat ($-1.9 \times 10^{-5}\,\mathrm{yr}^{-1}$, CI $[-9.2, +5.5] \times 10^{-5}$)— the class where an ordered-phase model would most expect melting. The verdict is keyed on the donor-level CI ($n = 23$ is the binding limit, not the ∼40k sites); the site-level bootstrap is tight by design and reported only as supplementary.

c. Interpretation. Where cell-composition mixing is mathematically excluded (pan-invariant sites), there is no detectable within-cell disordering—the melting picture's single-cell signature does not appear. The larger (point-estimate) drift at variable sites is where composition mixing *can* act, and where the literature independently locates genuine per-locus change: age-associated per-locus *mean* drift is well established to be largely composition-independent (≈ 50–83% of age-DMRs survive Houseman cell-fraction adjustment [6]) and to concentrate at *intermediate*, not already-saturated, sites—precisely the core model's relaxation at drift-prone loci. We therefore do *not* claim the entire bulk entropy rise is composition; we claim the *melting-specific* signal (within-cell disorder at composition-inert saturated sites) is absent, while the real per-locus change that does exist has the signature of core decay, not collective melting. Combined with the shape result, both melting-specific predictions fail in this tissue at this resolution.

d. Independent sorted-cell confirmation (mean level). A separate cohort of FACS-sorted classical monocytes—20 young (ages ∼24–30) and 20 old (∼57–70), eRRBS [19]—removes the composition confound by design (one purified cell type) and tests the same mean-level prediction directly. Across 2.8×10^6 CpGs, the methylated pole (young $\beta > 0.9$, 9.0×10^5 sites) shifts by only $\Delta\beta = -0.0009$ into age (5.1% desaturating below 0.9), the unmethylated pole (young $\beta < 0.1$) by $+0.0010$, and the intermediate-band mass barely moves (16.04% → 16.08%). Melting predicts the saturated poles regress toward $\frac{1}{2}$ and the intermediate mass grows; neither happens. Because the cells are sorted, this is a composition-free replication of the mean-level negative, at 20-versus-20 donors rather than the pan-invariant probe's $n = 23$.

e. What this does and does not establish. It establishes that, cross-sectionally in human blood, the composition-inert probe detects no within-cell melting signal above its noise floor. It does *not* establish that the epigenome universally cannot melt. Two gaps remain: (i) within-lineage sub-population mixtures could still hide per-cell change

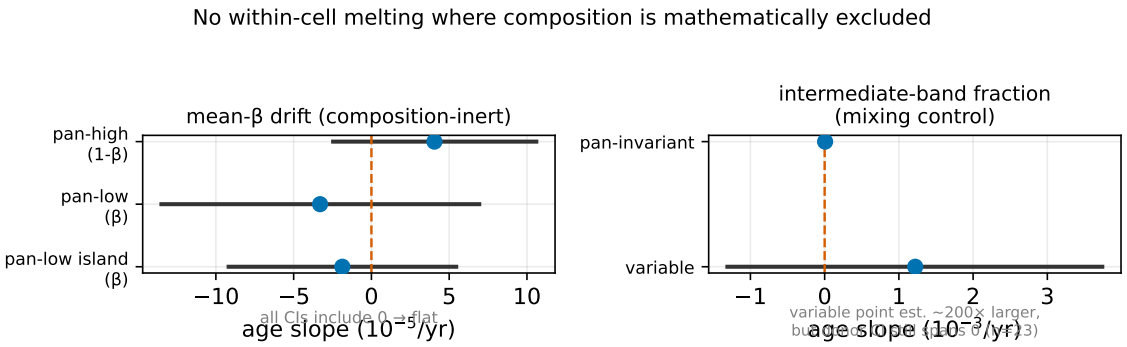

FIG. 4. The composition-inert negative (donor-level slopes, 95% CI, $n = 23$, coverage-matched to depth N). *Left:* mean-β drift at composition-inert pan-invariant sites (pan-high, pan-low, and island-context pan-low)—all CIs include zero. *Right:* the decisive mixing control, the intermediate-band-fraction slope: pan-invariant is $+6.2\times10^{-6}$/yr (flat), the variable point estimate is $+1.2\times10^{-3}$/yr—about 200× larger—but its donor CI still spans zero at $n = 23$. The contrast is in the point estimates plus the pan-invariant flatness, not in donor-level significance of the variable arm.

the sorted-atlas definition of "pan-invariant" cannot see; (ii) the probe is cross-sectional and pilot-scale ($n = 23$). The definitive test is read-level (Sec. VI).

VI. DISCUSSION

a. Therapeutic consequence. If aging is information decay under a no-reference channel, then no internal maintenance can restore lost identity information—the $\mathbb{Z}_2$ symmetry forbids it. Only an external reference can, which is exactly what partial reprogramming (OSK) supplies [11, 12]: h_{ext} transiently lifting M. Interventions that raise reinforcement r (or lower noise T) can *retain* information longer and even, in the ordered regime, hold it at near-zero cost—but retention is not restoration. The model draws a sharp line between the two.

b. The one result that would overturn the negative—and why the obvious version does not. Read-level metrics (proportion of discordant reads, PDR [8]; epipolymorphism [9]) score within-molecule methylation discordance, which cell-composition mixing cannot fake. It is already established that PDR *rises* with age in blood (a PDR-based age clock reaches $R^2 \approx 0.81$ on RRBS blood, ages 19–56 [10]), and that this rise is not fully explained by cell composition. This does not resurrect the melting picture, because a rising PDR is predicted equally by core per-locus decay: independent flips toward the $\frac{1}{2}$ floor produce reads carrying both methylated and unmethylated CpGs—discordant by construction. Generic PDR increase therefore does not discriminate core from melting. The discriminating test is a PDR rise *decoupled from mean-β drift*. The naive version fails: core decay *can* raise discordance even at a saturated site, by *desaturating* it. What separates the pictures is whether the discordance rise is accompanied by a mean-β shift: core decay raises PDR *and* moves the mean; collective melting can raise within-read discordance while the mean stays put (the ordered domain loses concordance before it loses its average commitment). The falsifier is therefore: at sites where the mean shows no age-associated drift (pan-invariant sites), does read-level PDR nonetheless rise with age? A yes is a melting signature the bulk mean-based probe cannot see. Its validity depends on the mean-null holding at the same sites, donors, and depth in the read-level data—so the two analyses must run on one cohort. This is runnable in principle on existing read-level blood data (RRBS/WGBS with epiread output) without new wet-lab work—the gate is a read-level methylation-calling pipeline plus a clean, wide-age cohort.

c. A pilot of the read-level test, run. We ran this test on an independent whole-blood cohort [15] (20 healthy donors, ages 21–52; RRBS-depth methylation haplotypes, hg19), binning reads into genomic windows. A first cut classified a window as composition-inert when its mean β was pinned (< 0.1 or > 0.9) in *every* donor; but that definition is a survivorship cut—conditioning saturation on the *old* donors discards exactly the sites melting would have desaturated in age, biasing the $\beta\approx1$ arm toward a flat null. We therefore use a bias-audited estimator. Because selecting saturated windows on the same young donors used to measure the change couples ceiling-selection noise into it, we define the saturated set on a *held-out* half of the young donors and measure the change on the other

half versus the old. Because a global slope extrapolated to $\Delta\beta = 0$ would read the falsifier off the thin tail of the data, we estimate it *directly* as the mean ΔPDR over the composition-inert subset ($|\Delta\beta| < 0.02$). And because tens of thousands of spatially autocorrelated, donor-sharing windows make naive confidence intervals fiction, we quote a donor- and 1 Mb-block bootstrap. Only the $\beta > 0.9$ arm discriminates the models: at $\beta \approx 0$, core (age hypermethylation of unmethylated sites), melting, *and* incomplete bisulfite conversion all add sporadic marks that couple disorder to the mean, so that arm is uninformative and we set it aside. On the discriminating, conversion-immune $\beta > 0.9$ arm ($\sim 3.0 \times 10^4$ inert windows), the mean ΔPDR at fixed composition is +0.0004 (donor+block bootstrap 95% CI $[-0.008, +0.006]$)—indistinguishable from zero: no melting-specific decoupling under a pinned mean. Tightening the inert window drives the effect *toward* zero and slightly negative (ΔPDR = −0.002 at $|\Delta\beta| < 0.01$)—the opposite of a melting signal, which would persist as $\Delta\beta \to 0$—while relaxing it lets coupled drift bleed back in (+0.003 at $|\Delta\beta| < 0.05$). The two residual estimator biases—measurement-error attenuation and ceiling selection—both push this estimate *upward*, toward a false melting signal, so the null is conservative, and the interval bounds any genuine melting decoupling below ≈0.006 in ΔPDR. Sequencing depth is uncorrelated with age ($r \approx 0$), ruling out a depth-driven artifact. A second, independent sorted population [16] (eight hematopoietic-progenitor samples, ages 20–49) is directionally concordant (inert-subset ΔPDR ≈ − 0.006) but rests on four-versus-four donors—pseudo-replication we report as a supporting anecdote, not an inferential replication. Two limits keep even the whole-blood result a corroborating pilot, not the definitive test: it caps near age 52 (no elderly), and it lacks the bisulfite-conversion spike-in controls that would license the $\beta \approx 0$ arm. Within these limits the read-level probe *corroborates* the core reading on the arm that can distinguish the models: no within-read melting when composition is held fixed.

d. Pushing to the widest public age gap. We repeated the read-level test on the most demanding cohort public data allows—sorted $CD4^+$ T cells (composition removed by design), WGBS (genome-wide, not island-biased), ages 18–25 versus 82–86, a ∼60-year gap (source WGBS [3]; mHapBrowser-recomputed per-locus PDR [18], n=3 vs 3). Here a small *melting-consistent* decoupling does appear: mean ΔPDR at fixed composition = +0.008 (extrapolating to +0.0025 as $|\Delta\beta| \to 0$), consistent across all three old donors (exact 6-donor permutation: the old/young split is the most extreme of $\binom{6}{3} = 20$, $p = 0.10$, the n=3 floor)—so it is not one aberrant donor. Two facts nonetheless argue it is an aged-DNA *artifact* rather than domain melting. First, deamination and sequencing error in 82–86-year-old DNA convert scattered methylated calls to apparent-unmethylated within otherwise-methylated reads, manufacturing within-read discordance at $\beta \approx 1$ with no mean change—exactly this signature. Second, and decisively for direction, the effect does *not* show the open-sea/late-replicating enrichment domain-melting predicts; it instead *rises with CpG density* (+0.0065 → +0.0122 from sparse to dense windows), the fingerprint of a per-CpG error whose per-read discordance probability $1 - (1 - \varepsilon)^k$ grows with the read's CpG count k. Third, and most directly, a library-preparation benchmark bounds how much of a +0.008 shift is assay rather than biology: in a panel where one whole-blood DNA sample per donor was prepared with three bisulfite kits across two sequencers [17], the across-preparation ΔPDR spread at the same composition-inert saturated sites is +0.022 to +0.049—three to six times the age gap—on *identical DNA at fixed age*; even nominally-identical technical replicates of one sample differ by up to +0.049, and this purely technical spread itself rises with CpG density (+0.036 → +0.070, sparse to dense), the same fingerprint. The extreme-age-gap signal therefore sits well inside the run-to-run technical envelope and carries its density signature. The clean veto—whether old-donor discordance is carried by singleton mismatched CpGs—needs read-level reads unavailable from precomputed metrics. We therefore report the extreme-age-gap decoupling as an aged-DNA/assay artifact rather than evidence for melting: the core reading is not overturned, and a conversion-controlled read-level cohort remains the deciding experiment.

e. Limits. Cross-sectional (no within-individual longitudinal trajectory; the wide-age WGBS range from newborn to centenarian is so far only cross-sectional and sparse [21]); one tissue (blood); pilot $n = 23$ for the pan-invariant negative; within-lineage mixtures not excluded. Each is a named gate on the corresponding claim, not a hidden caveat.

f. Relation to the companion control-law paper. A separate manuscript [22] develops the ordered-phase ("reverse-polarity") regime as a control law for a maintained epigenetic state. It is explicitly decoupled from this paper's negative, and the decoupling is *mechanical*: it is the same order parameter M and the same law (Eq. (1)), in two different regimes. This paper concludes human blood sits in the disordered regime ($T > J$, $\mathcal{N} < 1$), where the law reduces to relaxation-to-floor and there is no collective barrier to melt. The companion paper explores the ordered regime ($T < J$, $\mathcal{N} > 1$), where a barrier exists and can be engineered. It therefore does *not* claim the aging measured here is a melting it reverses—because in the $\mathcal{N} < 1$ regime there is no ordered phase to have melted. The honest link is negative: a reverse-polarity intervention would *engineer* an $\mathcal{N} > 1$ regime that natural blood does not occupy, not restore one it lost.

VII. METHODS

a. Datasets. An aging human blood methylation cohort [14] (blood-only subset) for the drift-shape and donor-level pan-invariant tests; a sorted-cell WGBS atlas [13] (six blood lineages, three replicates each, plus a 23-donor whole-blood cohort ages 21–75) fetched by HTTP range from the per-CpG count format in hg38 CpG-index order.

b. Site classification. Pan-high $\beta \geq 0.9$ / pan-low $\beta \leq 0.1$ in every lineage with a per-lineage standard-deviation robustness cap; variable = the composition-sensitive complement.

c. SNP masking. CpGs at or adjacent to common SNPs dropped. As a genetic-confound check, the pan-invariant signal was additionally tested against cis-mQTL CpGs from the goDMC catalog [20] (a specific set, unlike ubiquitous common-SNP proximity), so that a methylation-quantitative-trait-locus effect could be separated from an age effect.

d. Coverage matching. Hypergeometric downsample to a fixed depth per donor; beta-binomial Jeffreys posterior β—removing the symmetric coverage-noise channel the pan-invariant contrast does not.

e. Statistics. Donor-level OLS slope and Spearman with CI ($n = 23$) as the primary verdict; site-level bootstrap as supplementary only.

f. Drift shape. $D(\text{age}) = \langle|\beta - \beta_{\text{young}}|\rangle$ on young-baseline CpGs; linear/saturating/logistic by AIC; front-loaded fraction as the core-vs-melting discriminator.

g. Code and reproducibility. All analysis and simulation code is in the palimpsest repository; the automated test suite (76 tests) passes.

Appendix A: Assessment of claims

TABLE II. Claims and the strength of their support.

claim	evidence	strength
Core information decay (I monotone, relaxation-to-floor)	sim + closed form + data shape	strong
Minimal governing law (mean-field/model-A closure)	derivation, cross-checked	moderate (phenomenological)
Fate number = reduced temp T_c/T_{eff}; novelty is the aging mapping	mean-field standard + mapping	strong (mapping)
Relaxation-to-floor shape (not plateau-then-collapse)	blood cohort, cross-sectional	moderate
Melting signature absent at composition-inert sites	composition-inert probe	moderate (pilot n)
Per-locus mean drift real & largely composition-independent	this study + Ref. [6]	strong (external)
Generic PDR increase does not favor melting	literature + core-model logic	strong (logical)
Universal “epigenome cannot melt”	—	not claimed

Appendix B: Open experiments that gate the strong claims

1. Saturation-stratified read-level PDR (PDR rise decoupled from mean drift at pan-invariant sites, same cohort): the melting-specific falsifier. Generic PDR increase is already known and does not settle it.
2. Single-tissue longitudinal birth-to-old trajectory (clean the shape claim).
3. Independent-cohort / larger-n replication of the pan-invariant negative.
4. Regenerate all provisional figures and numbers from a frozen run.

[1] C. E. Shannon, *A mathematical theory of communication*, Bell Syst. Tech. J. **27**, 379 (1948); the source-channel framing and the monotonicity of mutual information under a Markov chain (data-processing inequality).

[2] M. Eigen, *Selforganization of matter and the evolution of biological macromolecules*, Naturwissenschaften **58**, 465 (1971); the quasispecies error threshold underlying the maintained-vs-aged boundary.

[3] G. Jenkinson, E. Pujadas, J. Goutsias, A. P. Feinberg, *Potential energy landscapes identify the information-theoretic nature of the epigenome*, Nat. Genet. **49**, 719 (2017); information-theoretic WGBS modeling, and the source of the sorted young/old WGBS samples (GSE79798) used in the extreme-age-gap read-level test.

[4] S. Horvath, *DNA methylation age of human tissues and cell types*, Genome Biol. **14**, R115 (2013).

[5] G. Hannum *et al.*, *Genome-wide methylation profiles reveal quantitative views of human aging rates*, Mol. Cell **49**, 359 (2013).

[6] T. Yuan *et al.*, *An integrative multi-scale analysis of the dynamic DNA methylation landscape in aging*, PLoS Genet. **11**, e1004996 (2015); age-DMRs largely survive Houseman cell-fraction adjustment and concentrate at intermediate, not saturated, sites.

[7] E. A. Houseman *et al.*, *DNA methylation arrays as surrogate measures of cell mixture distribution*, BMC Bioinformatics **13**, 86 (2012); the cell-fraction adjustment whose residual age effect motivates the composition-inert probe.

[8] D. A. Landau *et al.*, *Locally disordered methylation forms the basis of intratumor methylome variation in chronic lymphocytic leukemia*, Cancer Cell **26**, 813 (2014); origin of the proportion-of-discordant-reads (PDR) metric.

[9] G. Landan *et al.*, *Epigenetic polymorphism and the stochastic formation of differentially methylated regions in normal and cancerous tissues*, Nat. Genet. **44**, 1207 (2012); epipolymorphism, the read-level disorder measure PDR generalizes.

[10] D. I. Karetnikov *et al.*, *Age prediction using DNA methylation heterogeneity metrics*, Int. J. Mol. Sci. **25**, 4967 (2024); a heterogeneity-metric (PDR/epipolymorphism) age model on RRBS blood, ages 19–56 (PRJNA531784), $R^2 \approx 0.81$; PDR rises with age and is not fully explained by cell composition.

[11] A. Ocampo *et al.*, *In vivo amelioration of age-associated hallmarks by partial reprogramming*, Cell **167**, 1719 (2016).

[12] Y. Lu *et al.*, *Reprogramming to recover youthful epigenetic information and restore vision*, Nature **588**, 124 (2020); partial reprogramming as the external reference (h_{ext}) that raises I.

[13] N. Loyfer *et al.*, *A DNA methylation atlas of normal human cell types*, Nature **613**, 355 (2023); sorted-cell WGBS atlas (GSE186458) used to define pan-invariant sites.

[14] L. Mboning *et al.*, *BayesAge: a maximum likelihood algorithm to predict epigenetic age*, Front. Bioinform. **4**, 1329144 (2024); targeted bisulfite aging blood cohort (GSE261769) used for the drift-shape and pan-invariant analyses.

[15] Whole-blood WGBS/RRBS methylation-haplotype cohort, 20 donors ages 21–52 (NCBI BioProject PRJNA339792, "Potential epigenetic biomarkers of obesity-related insulin resistance in human whole blood"); the read-level pilot cohort.

[16] T. Schoofs *et al.*, *DNA methylation changes are a late event in acute promyelocytic leukemia and coincide with loss of transcription factor binding*, Blood **121**, 178 (2013); source (GSE42119) of the healthy $CD34^+$/promyelocyte progenitor WGBS used as the independent sorted population.

[17] L. Zhou *et al.*, *Systematic evaluation of library preparation methods and sequencing platforms for high-throughput whole genome bisulfite sequencing*, Sci. Rep. **9**, 10383 (2019); same-DNA multi-protocol panel (GSE128731) used to bound the technical PDR envelope.

[18] Y. Hong *et al.*, *mHapBrowser: a comprehensive database for visualization and analysis of DNA methylation haplotypes*, Nucleic Acids Res. **52**, D929 (2024); source of the precomputed per-locus PDR tracks.

[19] I. Shchukina *et al.*, *Enhanced epigenetic profiling of classical human monocytes reveals a specific signature of healthy aging in the DNA methylome*, Nat. Aging **1**, 124 (2021); sorted-monocyte young/old eRRBS used as an independent, composition-free mean-level check.

[20] J. L. Min *et al.*, *Genomic and phenotypic insights from an atlas of genetic effects on DNA methylation*, Nat. Genet. **53**, 1311 (2021); the goDMC cis-mQTL catalog used to test whether the pan-invariant signal is genetic.

[21] H. Heyn *et al.*, *Distinct DNA methylomes of newborns and centenarians*, Proc. Natl. Acad. Sci. USA **109**, 10522 (2012); the newborn-to-centenarian WGBS bounding the wide-age drift range (GSE31263).

[22] N. Hangen, *Reversed polarity: a control law for holding a maintained epigenetic state—and why it is not what bulk aging measures* (companion manuscript).