The VERSF History-Measure, Generation-Geometry and Branch-Closure Theorem

Operational Metric, Completion Geometry and Physical Branch from One Bath-Resolved History Transfer

▲ Programme Keystone Companion — Unified Substrate Dynamics and Standard Model Constitutive Closure Series

Gate HMGBC-1 / Companion to *The VERSF Master Substrate Action and Constitutive-Descent Theorem* (MASD-1)

Keith Taylor VERSF Theoretical Physics Programme Six sector gates → MASD-1 → HMGBC-1 → HFB-1 → Renormalised Parameter Closure → Final Audit Tier A conditional history-closure theorem, constitutive descent and finite-regulator execution protocol — July 2026

General Reader Summary

The Master Substrate Action paper built a single microscopic machine for the final Standard Model calculation. Its achievement was real: instead of giving the strong force, the weak force, particle masses and the Higgs sector their own separate rules, it made them different readings of one common underlying action. Its honesty was also real: it said plainly that the machine still had settings that had to be chosen by hand, and that choosing them differently produced different candidate universes. A machine that *can* be tuned to look like the Standard Model — and equally tuned not to — has capacity, not prediction.

This companion asks what could fix those settings without ever looking at the answer.

The starting point is something the programme already contains. Physical facts are irreversible and leave records. Earlier work shows that unresolved, almost-committed record modes form a physical environment — a bath. When those hidden modes are dropped from the description, the visible system remembers its past: its present behaviour depends on an entire weighted history, not only on its instantaneous state. That remembering is the missing ingredient.

The paper enlarges the microscopic state so that the visible variables and the unresolved record bath are described *together*. On this enlarged state the dynamics are ordinary, local, one step at a time, and positive — the cleanest kind of object there is. When the bath is integrated out, the visible system inherits a memory, and the curvature of that memory law says exactly how costly

each kind of microscopic failure is. The setting that previously had to be chosen — the operational metric — becomes an output.

Part of this is no longer a proposal. The smallest realistic case, a seven-state "closure wheel", has been carried through to exact numbers. The wheel's own rulebook is recovered rather than assumed; the six kinds of microscopic failure are given their natural directions; and the resulting cost matrix is computed exactly. It comes out full rank, with two cross-terms that arise on their own rather than being put in, and three cross-terms that are exactly zero for reasons of symmetry.

That calculation also taught the paper an important lesson about itself, which it now applies everywhere. **Exactly two features of the answer hold no matter which directions one picks for the six flaw types**: the exact zeros between different symmetry sectors, and the fact that the matrix is full rank. Those are the durable results, and both survive every test the paper applies — the full family of symmetric reweightings proved here, the restoration of parts initially left out, and an alternative underlying rule that was tried.

Everything else is less secure than it looks. The particular numbers are properties of the model. And everything *within* a symmetry sector — including the signs, and even whether two flaw types are linked at all — depends on the specific directions the paper freezes for them, and becomes physical only once those directions are derived from the theory rather than chosen. The paper says so explicitly, because a sign that looks like a prediction is exactly the sort of thing that later gets quietly matched to data.

The same history flow settles a question the programme had left as a philosophical fork: are different internal copies of a thing kept as separate protected accounts, or do they share one bath? Separate accounts allow only phase changes and swaps; a shared bath allows continuous transfer among copies, which is what the non-Abelian forces need. The paper turns that fork into a definite matrix audit on the same response that fixes the metric — the narrative choice is converted into a calculation, though the physical response algebra has not yet been evaluated.

Generations get the same treatment, subject to one inherited input. *Given* the rank-three, equal-layer structure taken from earlier papers in the programme, the way the six boundary states pair into three is not chosen — of all fifteen ways to pair them, four respect the dynamics and exactly one also treats the three families alike. That much is computed, not selected. It is not yet the physical projection: whether the record layer of the full theory pushes forward this way remains to be established. On its own, though, it leaves two of the three families indistinguishable. What separates them is a twist: a discrete winding that a committed fact deposits in the record structure. With any nonzero winding, the three families acquire three distinct depths. This is a genuine generation *geometry* obtained from the history law rather than an inserted diagram — though it is not yet a mass spectrum, and the numbers attached to it are wheel numbers, not physics.

Then comes the part needing the most care. The machine assigns a cost to every allowed branch and selects the cheapest. But comparing costs across branches is meaningful only if the branches are pieces of *one* underlying object measured on *one* scale — otherwise one is comparing prices in different currencies and calling the smaller number cheaper. The paper therefore makes it a required theorem, not an assumption, that all branches live inside one master operator with one common reference measure, and that they arise as its natural components rather than being placed side by side afterwards. If two branches genuinely tie, the theory must publish the tie; it may not quietly pick whichever looks more familiar, and it may not use measured data to break it.

One point is easy to get backwards, so it is worth stating directly: what is handed over is a recipe, not a result. It would seem natural to demand that every number be calculated before the execution stage opens. That is the wrong bar — it asks the derivation stage to perform the execution stage's job. What matters for honesty is not whether the answers are already known but whether any *choice* remains once the calculation starts. So the paper specifies a complete recipe **capable of being sealed**: the transfer operator, the way the bath is separated from the visible part, the definition of the metric, the projection onto the record layer, the branch decomposition, the tolerances, and the code manifest — every remaining numerical quantity defined as an output of one deterministic procedure, relative to a finite and explicitly declared set of conditional upstream inputs. On that basis the execution stage may open and run blind, and it may return success, failure, or a tie. **The seal itself takes effect only when the real manifest, its hashes and the two declared regulator inputs are published, and that publication has not happened.** So the accurate statement is: a sealable rulebook is specified; the seal is not yet in force; the run has not been made. No Standard Model number exists anywhere in this paper, and none may be filled in before the run.

The boundary of what is proved is drawn honestly, and it is worth stating in the paper's own voice rather than leaving a reader to assemble it. **This companion does not eliminate the Standard Model's freedom; it relocates it.** The freedom moves upstream — into one master generator, into inherited results from earlier papers in the programme, and into a short list of named assumptions. Every one of those is now written down in a single table, together with what fails if it turns out not to hold. Relocation is genuine progress, because it converts a vague "some rule must fix this" into something finite, named and checkable. It is not the same as elimination, and the paper does not claim it is.

Several things remain open, and they are stated as sharply as the results. The rule the exact calculations were performed on is not yet known to agree with the rule declared physical — their structures agree, their numbers do not, and that gap has its own gate. There are sixteen admissible ways for the microscopic history to be deformed that carry no known failure label at all, which is either a sign that a stated assumption is doing heavy lifting or a sign that the list of failure types is incomplete. And the way the failure coordinates are embedded into the space of rules is itself a convention that affects the answer, so it too has to be forced rather than picked.

When the calculation is finally run, it will return one metric, one generation geometry and either one physical branch or a declared degenerate set of them — or it will show that the framework does not return the Standard Model. Either outcome is a result. What this companion has done is make sure that when the run happens, no one gets to choose the answer.

Principal Claim

MASD-1 defines one finite-regulator machine whose scalar, gauge, fermion, colour and global-measure sectors descend from a common action. It also identifies, at full strength, the point at which that machine remains non-predictive: the operational metric W_prim, the generation/completion geometry and the physical branch are still constitutive inputs or candidate structures. Until one substrate history law returns all three from the *same* object — and until the branches among which that law selects are shown to be components of *one* history operator rather than separately constructed models — the Standard Model's freedom has been relocated into the microscopic wiring, not eliminated.

This companion supplies the missing history architecture, and states precisely which part of it is a theorem and which part is a finite-regulator evaluation not performed here.

At regulator level n, the resolved MASD configuration is

$$\mathfrak{X}_n = (\rho,\ U,\ J,\ \Phi;\ \Psi,\ \bar{\Psi}),$$

with primitive defect vector

$$\mathbf{D}_n = (\mathcal{D}_C,\ \mathcal{D}_J,\ \mathcal{D}_T,\ \mathcal{D}_R,\ \mathcal{D}_B,\ \mathcal{D}_H).$$

The history carrier is enlarged by the unresolved, sub-threshold record modes that the VERSF memory-kernel programme identifies as the physical bath:

$$\mathfrak{X}_{(n,k)} = (\mathfrak{X}_{(n,k)},\ q_{(n,k)}), \qquad q_{(n,k)} \in \mathcal{Q}_n.$$

The bath-resolved Euclidean history law is generated by one positive finite-regulator transfer operator

$$\mathbb{T}_n = \exp(-a_t\ \mathbb{H}_n),$$

where $\mathbb{H}_n$ contains the operational Dirichlet generator, the MASD defect potential, the positive unresolved-record bath operator and the defect–bath coupling returned by the same substrate dynamics. A Feynman–Kac transfer with a nonconstant potential is **not** itself stochastic — it does not preserve the constant function, and its rows do not sum to one (**verified**). The Markov evolution is obtained only after normalisation, in this order:

$$[\![\ \text{positive raw transfer } \mathbb{T} \longrightarrow \text{Perron data } (\lambda_0,\ \ell,\ r) \longrightarrow \text{Doob kernel } K \longrightarrow \text{Markov history process.}\ ]\!]$$

Only K, never the unnormalised $\mathbb{T}$, is the Markov evolution, and the history probability measure and its Radon–Nikodym derivative are defined with respect to K. With that understood: the reduced *resolved* process is generally non-Markovian while the enlarged process is Markovian. This is the pivotal structural fact: memory is not posited, it is what remains after integrating out modes that were themselves memoryless.

Integrating out the unresolved bath produces the finite-history Radon–Nikodym action

$$\mathcal{A}_{(n,N)} = -\log(d\mu_{hist}^{(n,N)} / d\mu_0^{(n,N)}),$$

and its long-history rate

$$\mathcal{I}_n = \lim_{(N\to\infty)} (1/N)\ \mathbb{E}_{(\mu_0)}\ \mathcal{A}_{(n,N)}.$$

The operational metric that MASD-1 leaves as an input is no longer selected independently. It is the quadratic defect curvature of that rate:

$$[\![\ (W_{prim}^{(n)})_{AB} = \partial^2\mathcal{I}_n / \partial\varepsilon_A\ \partial\varepsilon_B\ |_{(\varepsilon=0)}\ ]\!]$$

for $A, B \in \{C, J, T, R, B, H\}$, with all diagonal weights and cross blocks returned together, and, in the Gaussian finite-regulator sector, equal to one exact Schur complement of the bath block.

The same stationary history flow returns the universal record conductance. Its symmetric part defines one covariant record Laplacian,

$$\Delta_\Omega = B_\Omega^\dagger\ C_\Omega\ B_\Omega,$$

and the history-derived completion resolvent $R_\Omega = (W_H + \Delta_\Omega)^{-1} W_H$. Because $[W_H, \Delta_\Omega] \neq 0$ in general, R_Ω is **not** an ordinary positive operator; the manifestly positive object is the Hermitian contraction $\mathcal{S}_\Omega = W_H^{(1/2)}(W_H + \Delta_\Omega)^{-1} W_H^{(1/2)}$, $0 < \mathcal{S}_\Omega \leq I$, to which $R_\Omega = W_H^{(-1/2)}\ \mathcal{S}_\Omega\ W_H^{(1/2)}$ is similar. R_Ω is therefore a **W_H-self-adjoint contraction** with spectrum in (0, 1] (§33). The oriented completion operator is $\boldsymbol{\mathcal{R}}_\Omega = V_{(H,\star)}^\dagger R_\Omega$, so the general physical scalar–completion law becomes

$$[\![\ \mathcal{O}_{(H,\star)} = (\varphi_{(H,\star)}/\Lambda\star)\ \boldsymbol{\mathcal{R}}_\Omega\ P_H^R.\ ]\!]$$

The isotropic MASD-1 branch is recovered only when $R_\Omega = I$ on the scalar-access support, i.e. only when that support lies in $\ker \Delta_\Omega$. The generation and sector operators are compressions of the one universal object,

$$R_f = S_{fL}^\dagger\ \boldsymbol{\mathcal{R}}_\Omega\ S_{fR},$$

not independently selected graphs.

Finally, branch selection. This is the point the companion must handle with the most care, because a Perron root selects a stationary vector *within one irreducible operator* and says nothing, by itself, about a comparison *across* separately built operators. The companion therefore does two things that a naive "argmin over Perron roots" does not. First, it requires — as an explicit theorem obligation, not an assumption — that every admissible branch β be a component (invariant sector, boundary condition or reducing subspace) of *one* master history operator $\mathbb{T}_n^{hist}$ on one common history space, with all branch weights descending from one common reference measure (§43A). Only then are the branch free energies

$$f_\beta^{(n)} = -(1/a_t) \log \lambda_{(0,\beta)}^{(n)}$$

block Perron exponents of a single operator rather than incommensurable numbers. Second, given that common-history-space theorem, the physical branch set is

$$\mathfrak{B}_{phys}^{(n)} = \operatorname{argmin}_{(\beta \in \mathfrak{B}_{adm}^{(n)})} f_\beta^{(n)},$$

with a certified positive selection gap for uniqueness and a published degenerate set otherwise. No branch is preferred because it resembles the Standard Model.

One finite part of the history-content gate is executed. The canonical K = 7 closure-transition wheel supplies an explicit positive, irreducible and reversible one-step kernel, and the six MASD defects admit a minimal D_6-symmetry-adapted six-scalar tangent slice inside its admissible support. The resulting six-direction Fisher matrix can therefore be evaluated exactly. It is positive definite, has rank six, and returns two nonzero cross-defect blocks with definite signs and three symmetry-forced zeros — without inserting any of them. This is the first concrete history-derived metric calculation in the HMGBC programme (§§14A–14D, Corollary 16.4). Its *choice-independent* output — full Fisher rank and exact vanishing between inequivalent irreducible sectors — is a genuine candidate for the physical pattern. Its within-irrep cross blocks (values, signs, and even their non-vanishing) are properties of the declared generator family, and its rational magnitudes are specific to the canonical wheel model; neither is a physical prediction (§14C.3, Thm 14C.6). The one exception is the **J–H sign**, which is recovered whenever $\mathcal{D}_H$ is diagonal-supported with weights of one sign (Corollary 14C.4b); no such exception exists for W_{TB}.

The result is a canonical-wheel execution, not yet the full physical MASD return. Each MASD defect is vector-valued on a larger carrier, whereas the wheel calculation retains one symmetry-adapted scalar representative from each defect sector. The physical execution must prove that these six wheel directions are the faithful images of the actual MASD defect differentials (the **Defect-Lift Obligation**, whose four testable criteria and required proof content are itemised in HM10 and §14E), lift them through the multi-cell record, gauge, completion and fermionic carriers, determine whether the cross blocks that vanish on the seven-state wheel remain zero after that lift, and compute the physical magnitudes on the full operator rather than inheriting the wheel's.

HMGBC-1 has therefore progressed beyond a purely conditional history architecture. The canonical K = 7 transfer, its persistent/nonpersistent decomposition, a defect tangent family and its complete one-step Fisher metric are returned exactly. The remaining physical debt is narrower: derive the full typed lift of those six scalar representatives; derive the global and multi-cell history operator rather than choosing its spatial coupling; push the stationary flow canonically onto the terminal-record carrier; and exhibit all admissible branches as components of one master history operator. The canonical-wheel subcertificate is issued below; the full HMGBC-1 physical certificate remains unissued.

The honest one-line status. MASD-1 relocated the Standard Model's freedom into (W_{prim}, completion law, branch). HMGBC-1 relocates it one level further, into **one object — the complete bath-resolved master history generator $\hat{\mathbb{H}}_n^{hist}$** — whose blocks are the resolved local cost ($V_{loc} \to W_{core}$), the bath spectrum $\Omega^2_{(Q,n)}$ and the defect–bath coupling C_n, together with the proof that its invariant components *are* the admissible branches. It is essential that the gate be on the *whole generator*, not merely on (Ω^2_Q, C_n): because $W_{prim} = W_{core}(0) - C_n^\dagger (\Omega^2_{(Q,n)})^+ C_n$, a derived bath block with a *chosen* W_{core} would relocate the old freedom straight into the local history cost. If $\mathbb{H}_n^{hist}$ is itself *derived* from the substrate axioms and shown regular on the real K = 7 regulator, the freedom is removed; if any block is chosen, the relocation continues. This companion builds every mechanism that consumes $\mathbb{H}_n^{hist}$ and proves what it can at theorem level, but it does not evaluate $\mathbb{H}_n^{hist}$. That, together with the terminal-record projection (§31.1) and the one-operator branch structure (§43A), is HMGBC-1's set of **three** load-bearing gates, named as such throughout.

Headline Results and Conditionality

The following ledger uses exactly **four status labels**, and no others, because mixed vocabulary ("defined", "derived", "verified", "constructed", "pass", "closed") has been the main source of overstatement in this programme:

1. **Definition / protocol** — a construction, formula or procedure fixed by declaration. Carries no truth claim beyond internal consistency.
2. **Conditional theorem** — proved as finite-dimensional mathematics *given* a named premise that is not itself discharged here.
3. **Executed benchmark** — computed exactly on the canonical K = 7 wheel. Its *choice-independent* content — full Fisher rank and exact vanishing between inequivalent irreducible sectors — is a candidate physical pattern; its within-irrep cross blocks (values, signs, non-vanishing) belong to the declared generator family — except the J–H sign, recovered under support typing (Cor. 14C.4b) — and its magnitudes are model properties (§14C.3, F25).
4. **Physical return** — a numerical result on the actual MASD/CAR/Fock carrier. **No row in this paper carries this label as achieved**; the single row that would carry it is marked NOT ISSUED.

The operational heat-transfer formula and the central-algebra prescription are definitions or conditional theorems; the wheel calculations are executed benchmarks; every physical return remains open.

Object	Result	Status
Programme role	companion converting MASD-1 constitutive inputs into history-law outputs	Definition/protocol
Enlarged carrier	$\mathfrak{X} = (\mathfrak{X}, q)$ with unresolved record bath	Definition/protocol
Full transfer	$\mathbb{T}_n = e^{(-a_t \mathbb{H}_n)}$ on the enlarged carrier; requires the frozen finite-regulator generator (HM3)	Conditional theorem

Object	Result	Status
Reduced visible dynamics	non-Markovian after bath elimination (Nakajima–Zwanzig)	Conditional theorem
Finite-history action	$\mathcal{A}$_(n,N) = −log(dμ_hist/dμ_0)	Definition/protocol
Long-history rate	$\mathcal{J}$_n = lim N^{-1} $\mathbb{E}$ $\mathcal{A}$_(n,N); requires the ergodic limit (HP2)	Conditional theorem
Operational metric	W_prim = D^2_**D** $\mathcal{J}$_n(0) (Thm 16.1a); Gaussian Schur form W_core(0) − C†$(\Omega^2)^+$C under the range condition (Thm 16.1b); fixed only up to W ↦ J^TWJ until HM12 is discharged	Conditional theorem
— *Canonical K = 7 subtheorem* —		
Canonical K = 7 history kernel	explicit 7×7 stochastic T_0, π = (7,4,4,4,4,4,4)/31	Executed benchmark
Canonical Perron gap	1 − max_(λ≠1)|λ| = ½ (Thm 14A.1)	Executed benchmark
Persistent/nonpersistent split	one Perron mode + six-dimensional complement	Executed benchmark
Defect tangent family	support-preserving exponential deformation in C, J, T, R, B, H; linear independence proved (Thm 14B.1), uniqueness of the generators open (F26)	Definition/protocol
Canonical history metric	exact positive W_wheel (§14C, Cor. 16.4)	Executed benchmark
Canonical metric rank	full rank, no information-null direction; the minimum eigenvalue 3/31 is a model magnitude, not physical (F25)	Executed benchmark
Wheel within-irrep cross blocks	W_JH = −33/434, W_TB = 3/248 **for the declared generator family**; values, signs and non-vanishing depend on that frozen family and are **not** choice-independent, except as below (§14C.3, Thm 14C.6)	Executed benchmark
J–H sign under support typing	W_JH = −6π_h αsp − 6π_b βqr, so diagonal support with α, β of one sign forces sgn W_JH < 0 for every admissible weight (Cor. 14C.4b); reduces the residual freedom to two named questions about the physical $\mathcal{D}$_H and supplies HM10.4 with its discriminating content	Conditional theorem
Vanishing wheel cross blocks	W_BH = W_CH = W_RH = 0, forced by inter-irrep orthogonality for any representatives and any admissible weights (Lemma 14C.4)	Executed benchmark

Object	Result	Status
Wheel magnitudes as physics	prohibited: no rational entry, eigenvalue or condition number may be quoted as prediction or fed to HFB-1 (F25)	Definition/protocol
Supported Euclidean relaxation spectrum	spec(T_0^2) evaluated; magnitudes model-specific (§14D)	Executed benchmark
Restored D_6 doublets (8×8)	zeros survive; new W_(Tcos,Bsin) = $3\sqrt{3}/248$; T/B eigenvalues replaced by {9/62, 6/31} (§14F.4)	Executed benchmark
Tangent-space exhaustion	24-dimensional admissible tangent space, 8 defect directions, **16 unlabelled** — the entire $m = 2 \oplus m = 4$ sector carries no defect label (§14G); HM4 decoupling certificate outstanding (F36)	Executed benchmark
Canonical wheel recovered	Doob transform of $D^{(-1/2)}AD^{(-1/2)}$ returns T_0 exactly, 1/7 hub and 1/4 rim (Thm 11A.1); conditional on HM11 and on enumerating the move graph (F32)	Conditional theorem
Score-generator uniqueness	equivariance alone provably cannot fix the generators — every sector but C carries multiplicity 3 or 4 (Thm 14C.6); support and intertwining data required (F26)	Conditional theorem
Defect-lift obligation	four testable criteria on $D\mathcal{D}\_A \longrightarrow \Delta\_A\ T$ (HM10, §14E); closure sector passes the graph-lift analogues (Appendix I)	Definition/protocol
Tilt-family obligation	the exponential family with transported increments must be forced, not chosen (HM12, F34)	Definition/protocol
Kernel identification (G11a)	does L^op restricted to the closure catalogue equal $I - \mathcal{S}\_n$? an inherited attribution, §55A	Definition/protocol
Kernel relation (G11b)	given G11a, Doob(heat) = $\exp(-a\_t(I - T_0))$ exactly: same eigenvectors, same π, same D_6 content, eigenvalues mapped by $\lambda \mapsto e^{(-a\_t(1-\lambda))}$; no a_t reconciles the magnitudes (F35)	Conditional theorem
Polar orientation frame	$J\_A = V\_A\|J\_A\|$, equivariance propagated (Thm 14F.2), rank additive (14F.3); the magnitude-bearing lift Π^phys awaits evaluation of every J_A	Conditional theorem
Full physical W_prim	specification only: wheel structural pattern + multi-cell / typed-carrier lift, magnitudes recomputed on the physical operator	Definition/protocol

Object	Result	Status
— *Canonical projection and generation geometry* —		
Terminal-record map (canonical)	antipodal A = $\{b_0,b_3\}$, B = $\{b_1,b_4\}$, C = $\{b_2,b_5\}$: unique layer-equivalent partition among the four strongly lumpable matchings (Thm 32A.1); provenance inherited (Thm 32A.0); the physical pushforward and $\mathcal{U}\_\Omega$ descent are not established	Executed benchmark
Record pushforward (canonical)	equally weighted K_4, conductance 2/31 (§32A.2)	Executed benchmark
Three-generation Laplacian	$\Delta\_\Omega\^(0) = (8/31)(I_3 - ⅓J_3)$, spec {0, 8/31, 8/31} — degenerate doublet at zero holonomy	Executed benchmark
$\mathbb{Z}_7$ holonomy split	λ_m(k) distinct for every k ≠ 0 (k = 1,2,3 verified); nondegenerate generations follow from Fact occupancy	Executed benchmark
Completion contraction (canonical)	$\mathcal{S}\_\Omega\^(k) = (I_3 + \Delta\_\Omega\^(k)/w\_H)^{-1}$; spectra depend on w_H and are model-specific (F25)	Executed benchmark
Holonomy determinant ordering	D_(±1) < D_(±2) < D_(±3) for every w_H, g_Ω > 0 — ordering structural and w_H-independent (Thm 44A.1); a Gaussian *stiffness* comparison, not a block-Perron selection, and it does not discharge §43A	Conditional theorem
Holonomy descent	$q\_\star\gamma_6 = 2\gamma_3$ and k_Ω = 4k_vac (mod 7), so k_vac ≠ 0 ⟹ k_Ω ≠ 0 (G7d-1); physical basicness on a broken background open (G7d-2); labels permute (F33)	Conditional theorem
Gaussian determinant Perron-realisation	rank-one blocks give λ_(0,k) = π^(3/2)/√(det A_k); refutes an in-principle objection only, realises *any* positive λ_k, and does not satisfy Firewall H10	Conditional theorem
— *Structure and admission* —		
Cross-defect blocks	history covariances rather than inserted interactions, **once** the embedding, normalisation and tilt family are fixed (Cor. 16.2); the Fisher–Schur bridge itself is outstanding (F39)	Conditional theorem
Gauge/null quotient	kernel of history distinguishability plus an explicit classification audit; counting is not classifying (G5, F8)	Definition/protocol
Absolute action unit	per transfer step a_t, common to all sectors	Definition/protocol

Object	Result	Status
HP1 (cone positivity)	bath-resolved positivity with positivity improvement per branch; requires a Dirichlet/sub-Markov generator, not merely $L^{op} \geq 0$ (F37), plus irreducibility	Conditional theorem
HP2 (bi-Perron)	simple positive left/right Perron pair; the qualitative gap is automatic from positivity improvement, its magnitude and regulator-uniform bound are the open G3 return	Conditional theorem
HP2R (detailed balance)	reversible defect-free reference only; optional, never required	Definition/protocol
HP3 (refinement)	cylindrical block intertwining, reduced to three commuting-square identities (Appendix C)	Conditional theorem
History conductance / current	symmetric / antisymmetric stationary edge flow	Definition/protocol
Universal Laplacian	$\Delta_\Omega = B_\Omega^\dagger C_\Omega B_\Omega$	Definition/protocol
Completion resolvent	$R_\Omega = (W_H + \Delta_\Omega)^{-1} W_H$ is a W_H-self-adjoint contraction with spec $\subseteq (0,1]$ via $\mathcal{S}_\Omega$; sector compressions R_f do **not** inherit this (§33)	Conditional theorem
Terminal-record projection	C_Ω must be the representative-independent, cylindrically compatible, strongly lumpable pushforward with $\mathcal{U}_\Omega$ descending from polar transport (Thm 31.1); the residual tests validate but do not derive r_n (F40)	Definition/protocol
MASD-1 identity branch	$R_\Omega = I \Leftrightarrow$ scalar support $\subseteq \ker \Delta_\Omega$	Conditional theorem
Sector operators	$R_f = S_{fL}^\dagger \boldsymbol{\mathcal{R}}_\Omega S_{fR}$: typed, generally non-normal; physical content in singular values and relative frames	Definition/protocol
Bath–ledger verdict	commutant / reducing-projector audit of the bath response algebra (Thm 41.1), inheriting the range condition	Definition/protocol
One-operator branch structure	**all admissible branches components of one master history operator on one reference measure** (§43A); not established	Definition/protocol
Physical branch selector	Perron free-energy minimiser, conditional on §43A and on global dominance (§43A.5)	Conditional theorem
Selection uniqueness	positive branch gap $\Delta_{sel} > 0$, with the extensivity convention declared (F43)	Definition/protocol
Master generator $\hat{\mathbb{H}}_n^{hist}$	complete resolved-local (W_{core}) + bath (Ω^2_Q) + coupling (C_n) generator; not evaluated — the history-content gate	Definition/protocol

Object	Result	Status
Physical master transfer	$\mathbb{T}_0^{hist} = M \exp(-a_t L^{op}) M$ from the inherited operational generator; positive **given** a Dirichlet/sub-Markov L^{op} (§56A.2, F37)	Conditional theorem
Bath triple as compression	$\Omega^2_Q = Q\mathbb{H}Q$, $C = Q\mathbb{H}P\iota E_D$, $W_{core} = E_D^\dagger\iota^\dagger P\mathbb{H}P\iota E_D$ — no block independently chosen; requires the embedding ι and the P_kin/P_info separation (§56A.3)	Conditional theorem
Raw transfer from admissibility graph	$\mathfrak{B}_n = M \mathcal{S}_n M$ is cone-positive but carries negative inertia, so it is **not** $e^{(-a_t\mathbb{H})}$; the Euclidean object is $\mathbb{T}_E = \mathfrak{B}_n^\dagger\mathfrak{B}_n$ (§11A.2)	Conditional theorem
Branch decomposition (components)	$\mathfrak{B}_{adm} = \pi_0(G_n)$ gives commuting projectors by construction, but proves only that components are branches (§43A.6)	Conditional theorem
Branch algebra (central)	minimal central projections of $\mathfrak{M} = \{\mathbb{T}, \mathfrak{A}^{kin}\}''$; exhaustive **relative to** the frozen observable algebra, and exact superselection at finite regulator is itself open (F41)	Conditional theorem
Uniform gap (canonical family)	$\delta_{global} \geq \min\{½, \delta_X\}$ given HP3 intertwining and $\delta_X > 0$; two levels are a diagnostic, not a convergence certificate (F43)	Conditional theorem
Regulator closure	two-level intertwining plus an analytic bound, monotonicity theorem or controlled Cauchy estimate	Definition/protocol
Inherited-dependency ledger	every upstream attribution with what fails if it does not hold (§55A)	Definition/protocol
HMGBC structural theorem	complete on the stated premises (Thm 55.1)	Conditional theorem
HFB-1 execution admission	blind-execution protocol on the sealed package $\mathfrak{S}_{HMGBC}$ (Thm 56A.1); takes effect on publication of $\mathscr{M}_{hash}$ with real hashes and two declared regulator inputs — **manifest pending**	Definition/protocol
Renormalised Parameter Closure	consumes $\mathcal{P}_{SM}^{VERSF}(\mu_\star)$ from a passing HFB-1; architecturally specified, numerically locked	Definition/protocol
HMGBC physical certificate	**requires the frozen spectrum, couplings, one-operator proof, gap and returns of §§48–51**	**Physical return — NOT ISSUED**

Abstract

MASD-1 constructs a common finite-regulator action over independent substrate amplitude, link, current, record/completion and CAR/Fock matter variables. Its six primitive defects are priced by one operational metric W_{prim}; its scalar–completion law determines the fermion completion geometry; and its physical Standard Model branch must be selected before HFB-1 can execute the boundary calculation. MASD-1 states the decisive limitation at full strength: W_{prim}, the completion law and the branch label are not returned by a substrate history measure, so its candidate branches demonstrate capacity, not prediction.

This companion constructs the missing history architecture and proves what is provable at finite dimension, while marking three gates it does not close: the derivation of the complete master generator (history-content), the terminal-record projection, and the one-operator branch structure (branch-comparison).

At regulator level n, the MASD configuration $\mathfrak{X}_n$ is enlarged by a finite unresolved-record bath $\mathcal{Q}_n$, physically identified with sub-threshold commitment modes. A real, bounded-below, bath-resolved Hamiltonian $\hat{\mathbb{H}}_n$ defines the positive Euclidean transfer $\mathbb{T}_n = \exp(-a_t \hat{\mathbb{H}}_n)$, which is a positive weight operator and **not** a stochastic kernel; the **Doob-normalised** process on the enlarged carrier is Markovian. Eliminating the bath yields a non-Markovian resolved history law and a finite-history Radon–Nikodym action. The long-history relative-entropy rate $\mathscr{J}_n$ has positive-semidefinite defect curvature, and in the Gaussian finite-regulator sector that curvature is the exact Schur complement of the bath block. It returns every diagonal and cross block of W_{prim} at once, fixes the common per-step action normalisation, and makes exact gauge or terminal-null directions statistically invisible. This is a genuine improvement on positing a heat semigroup: positivity is carried microscopically by a *positive* transfer on the enlarged carrier, whose Doob normalisation defines the Markov history process, rather than being assumed.

The transfer positivity conditions are separated correctly. HP1 is cone positivity of the bath-resolved transfer with positivity improvement on each irreducible physical component. HP2 is the consequent simple positive left/right Perron structure — automatic at each finite level from positivity improvement, with its quantitative magnitude and regulator-uniform lower bound the separate open return; detailed balance is demoted to HP2R, an optional reversible-reference certificate rather than a requirement on the irreversible committed-history process. HP3 is cylindrical compatibility of the resolved, cohomological, bath and flat-holonomy blocks, under which the finite-history measures form a projective family.

The stationary edge measure of the reduced process decomposes into a symmetric conductance and an antisymmetric current. The conductance defines one covariant terminal-record Laplacian Δ_Ω and one completion resolvent R_Ω. The oriented history-derived completion map replaces the identity completion of MASD-1, recovered only when R_Ω is the identity on the scalar-access support. Quark and lepton generation operators are typed compressions of this common map. The bath-response algebra supplies a finite decision procedure for the programme's bath–ledger fork: nontrivial reducing projectors identify protected sub-accounts, while a trivial

commutant on a class certifies a saturated shared bath — the algebraic precondition of non-Abelian transport.

For branch selection, the companion makes explicit a precondition often left implicit: a leading Perron root selects a stationary vector *within one irreducible operator*, so comparing the Perron data of *separately constructed* branch operators compares incommensurable objects. The paper therefore states, as an HMGBC-1 theorem obligation (§43A, the **Common-History-Operator theorem**), that every admissible branch is a component of one master history operator on one common history space, with one reference measure; only under that theorem is the free-energy comparison $f_\beta = -(1/a_t) \log \lambda_{(0,\beta)}$ physical. Given it, compactness, lower semicontinuity and a positive selection gap yield a unique physical branch or a published degenerate set; no finite branch catalogue is assumed.

The paper closes the mathematical handoff MASD-1 → HMGBC-1 → HFB-1. Its strict physical certificate is conditional on the evaluated finite K = 7 bath spectrum, six-defect coupling map, null classification, one-operator branch structure, branch gap and regulator-convergence package. No Standard Model observable may be used to choose any of them.

Contents

Part IV — History Positivity and Refinement

Part V — Microscopic Generation Geometry

Part VI — Bath–Ledger Decision and Physical Branch

Part VII — Execution, Certificate and HFB-1 Handoff

Part VIII — Closure

Part I — Companion Position and Closure Target

1. Aim and Programme Position

MASD-1 defines the common microscopic parent required by the six final sector gates. It provides the independent variables, the six primitive defects, the common Hessian descent, the terminal-record address construction, the completion-strain bilinear, the physical evaluation protocol and the final boundary table. It also identifies, without hedging, why the machine is not yet predictive.

The action is written schematically as

$$\Gamma_{bos}^{(n)} = \Gamma_V^{(n)} + \tfrac{1}{2} \langle \mathbf{D}, W_{prim}\, \mathbf{D} \rangle,$$

but the history law that fixes W_prim is not evaluated there. The simplest scalar–completion branch uses an identity target on the scalar-access support, and the geometry that should distinguish generations is not returned. Candidate finite-cell branches (CFB-1/2/3) are available, but no physical history law selects among them; MASD-1 scores them as capacity, not prediction.

The programme dependency is

```
⟦ Six sector gates → MASD-1 → HMGBC-1 → HFB-1 → Renormalised Parameter Closure
→ Final Audit. ⟧
```

HMGBC-1 is not a competing master action. It is the companion that supplies the three objects MASD-1 deliberately leaves for a history theorem:

1. the operational metric W_prim;
2. the universal generation/completion geometry;
3. the physical branch selector.

HFB-1 then freezes those returns inside the MASD-1 execution protocol and calculates the microscopic Standard Model boundary. Two of the three — the metric and the geometry — are delivered here as theorems consuming a finite bath triple. The third — the selector — is delivered only after a precondition (§43A) that this paper elevates from an implicit assumption to a required theorem, because without it the selector is not even well posed.

2. The Exact Debt Inherited from MASD-1

The debt is not "find a better candidate metric." A selected positive matrix is still a selected matrix; a selected generation graph is still a selected graph; a branch chosen because its outputs look promising is still a fit. MASD-1's §47M.7 makes this the crux on which the entire derivational claim turns.

The required object is one finite-regulator probability law on microscopic histories whose different derivatives return all three deferred objects. With the six primitive defects $\mathbf{D} = (\mathcal{D}_C, \mathcal{D}_J, \mathcal{D}_T, \mathcal{D}_R, \mathcal{D}_B, \mathcal{D}_H)$, the history law must return

$$W_{prim} = D^2_{\mathbf{D}}\, \mathcal{I}(0),$$

```
Δ_Ω = (record-geometry reading of the same stationary flow),
𝔅_phys = (free-energy minimiser of the same raw transfer),
```

and it must do so from **one** law. A construction that uses one law for W_prim, a second for the generation geometry and a third for branch selection has not closed the programme — it has renamed the freedom. The single most important discipline of this paper is that all three descend from the one transfer $\mathbb{T}_n$ and its stationary structure.

To that debt this companion adds a fourth item that MASD-1's handoff left implicit and that a careful reading exposes as a genuine hole: the branches among which f_β selects must be components of one operator, or the comparison is meaningless. That is §43A, and it is treated with the same seriousness as the metric itself.

3. Claim Grades

H1 — Exact finite-dimensional mathematics. Gaussian integration, Schur complements, positive semigroups, Perron–Frobenius theory, stationary-flow decomposition, weighted Laplacians and operator compression. These hold regardless of any VERSF premise.

H2 — Conditional VERSF history theorem. The claim that the unresolved sub-threshold record sector is the complete same-order bath and that its frozen transfer generator is the physical one (premises HM1–HM4).

H3 — Finite-regulator execution. The actual K = 7 spectrum $\{\omega^2_{(n,a)}\}$, defect couplings C_n, transfer matrix $\mathbb{T}_n$, null classification, one-operator branch structure and branch weights. *A proper subset of H3 is executed:* the canonical K = 7 closure-wheel kernel, its Perron gap, its persistent/nonpersistent split, a six-scalar defect tangent family and the exact 6×6 W_wheel (§§14A–14D). The remainder of H3 — the full typed lift, the multi-cell/global operator, the record pushforward and the one-operator branch structure — is not executed.

H4 — HFB-1 boundary execution. The physical W_prim, R_Ω and branch are consumed by MASD-1 to return gauge, scalar, flavour, neutral, colour and strong-phase boundaries.

H5 — Empirical comparison. Renormalised observables compared with experiment after the microscopic boundary is frozen and transported.

This paper closes H1, specifies H2–H4, and now **executes a canonical-wheel fragment of H3** (§§14A–14D). It does not claim H5. It is explicit about the boundary: the mechanisms are H1; the canonical-wheel metric is an executed H3 prototype conditional on the Defect-Lift Obligation (HM10); and the full physical generator, the record pushforward and the one-operator branch structure remain H2/H3 obligations not discharged here.

4. Binding Notation and Type Discipline

The resolved finite-regulator configuration is

$\mathfrak{X}$_(n,k) = (ρ_(n,k), U_(n,k), J_(n,k), Φ_(n,k); Ψ_(n,k), Ψ⁻_(n,k)).

The history-step label k is a regulator ordering index; it is not assumed to be fundamental physical time. The transfer step is a_t, inherited from the refinement/continuum bridge.

The primitive defect carrier is $\mathcal{R}$_prim^(n) = ⊕*(A∈{C,J,T,R,B,H})* $\mathcal{R}$*_A^(n). The unresolved bath carrier is a finite real Hilbert space $\mathcal{Q}$_n = span{q*(n,a) : a = 1, …, N_Q(n)}. The complete history carrier is

$\mathcal{H}$_n = L²(X_n, μ_op^(n)) ⊗ L²($\mathcal{Q}$_n, dν_Q).

One distinction is binding and is the source of the whole construction's non-triviality:

⟦ reference operational geometry ≠ W_prim. ⟧

The first fixes coordinates, quotient volume and the kinetic Dirichlet generator. The second is the curvature of the full history rate, and is the object returned to MASD-1. Conflating them is exactly the circularity that MASD-1's §10.1.1 warned about — defining the metric by the very geometry it is supposed to produce. Here the reference geometry is upstream regulator data; W_prim is a computed Schur complement downstream of it. They are not permitted to be identified.

5. Premises and Firewalls

Premise HM1 — Complete unresolved-record bath. At the declared order, every unresolved mode capable of changing the defect rate is retained in $\mathcal{Q}$_n before reduction.

Premise HM2 — History locality on the enlarged carrier. The bath-resolved process is local in one history step on ($\mathfrak{X}$, q), even though the reduced resolved process is nonlocal in history.

Premise HM3 — Real bounded-below Euclidean generator. ℍ_n is self-adjoint or sectorially closed on the reversible reference branch and has a real, bounded-below potential on the physical quotient.

Premise HM4 — Defect-source completeness. The leading coupling of the resolved MASD variables to the bath factors through the six-defect tangent map. No Standard Model observable enters the coupling.

Premise HM5 — Common branch normalisation. Every branch is evaluated with the same reference measure, transfer step and determinant convention. Independent branchwise normalisation is forbidden. *(This premise is necessary but not sufficient for branch comparison; §43A supplies the sufficient condition.)*

Premise HM6 — Terminal records carry completion geometry. The same terminal-record carrier used for matter addresses carries the conductance from which the universal completion Laplacian is built.

Premise HM7 — Refinement compatibility. The resolved, cohomological, bath and flat sectors possess compatible lift/coarse maps satisfying the block conditions of Part IV.

Premise HM8 — Blind physical execution. All bath spectra, couplings, branch manifolds and regulator maps are frozen before any Standard Model output is inspected.

Premise HM9 — Single master history operator (branch closure). The admissible branches $\mathfrak{B}$_adm^(n) are the components of one operator $\mathbb{T}$_n^hist on one common history space $\mathscr{H}$_n^hist, with all branch weights induced by one common reference measure. This premise is what makes f_β a comparison of eigen-branches of one object rather than of separate models; it is the antecedent of §43A and is not discharged by HM5 alone.

Premise HM10 — Defect-lift faithfulness. The full vector-valued MASD defect differentials D$\mathcal{D}$_A compress to the finite symmetry-adapted history tangent directions used in any executed metric (on the canonical wheel, the six scalar Δ_A T of §14B). Equivalently, the wheel/internal tangent family is the faithful image of the physical defect carrier, not a lossy caricature of it. This premise is what upgrades the executed W_wheel (Corollary 16.4) to a claim about W_prim^physical; it is the **Defect-Lift Obligation** and is not discharged here. "Faithful" is not a slogan — it has four testable components, each of which is separately falsifiable (§14E, F23):

1. **Sector match.** Each physical defect carrier $\mathcal{R}$_A^(n) restricts to the wheel-supported subspace as a nonzero map; no MASD defect is annihilated by the compression.
2. **Rank preservation.** The compression π_wheel : $\mathcal{R}$_prim^(n) → T_(T_0) does not collapse two distinct defect labels onto the same wheel direction — dim of the image equals six, not fewer.
3. **Symmetry compatibility.** The compression intertwines the physical stabiliser action with the D_6 action on the wheel, so a defect that transforms in a given representation lands in the matching wheel representation (this is what selects the specific F_A of §14B rather than arbitrary scalars).
4. **Support typing (discriminating).** The inter-irrep vanishing relations and rank preservation follow automatically from equivariance and positivity, so requiring them alone would make this criterion unfailable. The discriminating requirement is on **support**: the physical defect differentials must carry the support types the benchmark assumes — in particular that $\mathcal{D}$_H is *diagonal-supported* (carried by self-transitions, not transport) and *uniformly signed across states*. Under those two conditions the J–H sign is forced for every admissible weight (Corollary 14C.4b), so the criterion has content and can fail. Within-irrep cross blocks are otherwise *outputs* of the evaluated J_A and are not required to match the benchmark family (Thm 14C.6).

A construction failing any of (1)–(4) means W_wheel is a symmetry caricature whose promotion to physics is unjustified. When (1)–(4) hold, what is claimed is: full Fisher rank; the inter-irrep vanishing relations; **and the J–H sign wherever support typing is established** (Corollary

14C.4b). The rational magnitudes are never claimed, and the remaining within-irrep blocks are outputs of the evaluated differentials rather than inherited patterns.

Premise HM11 — Canonical Admissibility Principle. Once the complete set of admissible elementary moves and the common operational measure have been fixed, no additional move-dependent weight is introduced unless it is derived from a substrate invariant; symmetry-equivalent elementary moves therefore carry equal reference conductance. This principle is what makes the reference kernel a *derived* object rather than a chosen one: it reproduces the canonical wheel exactly (Theorem 11A.1) and is consistent with the coarse-graining result that selects the lazy stochastic operator $\frac{1}{2}(I + D^{-1}A)$ within its declared family. **It is not proved on the complete MASD/CAR/Fock move graph**, and that is where the construction of §11A remains conditional (F32). The common operational-measure theorem independently fixes the quadratic measure and its common normalisation across admissible channels, so HM11 governs the *conductance weights*, not the measure.

Premise HM12 — Tilt-family obligation. The deformed kernel of §14B and §56A.5 is an *exponential* family, $K_\eta \propto \exp[\sum_A \langle \eta_A, \Xi_A \rangle] K_0$, built on covariantly transported increments Ξ_A. Fisher information is a metric: under a reparameterisation $\varepsilon \mapsto \varepsilon'$ with Jacobian J at the origin it transforms as $W \mapsto J^T W J$. So W_prim is defined only **relative to the chosen embedding of the defect coordinates into kernel space**, and a different admissible, support-preserving, symmetry-adapted embedding — for instance a non-exponential tilt with the same linearisation up to a defect-dependent rescaling — returns a *congruent but numerically different* W_prim. Because W_prim is consumed as a quadratic form on physically normalised defects in $\Gamma_{defect} = \frac{1}{2}\langle \mathbf{D}, W_{prim} \mathbf{D} \rangle$, that Jacobian is **not gauge**: it changes the physics. HM10.3 constrains the angular part of the embedding, not its radial/normalisation part. HM12 asserts that the exponential family with covariantly transported increments is *forced* by the substrate rather than chosen; it is the metric-side analogue of the branch-normalisation obligation of §43A and is not discharged here (F34).

Firewall H1 — No chosen W_prim. A positive matrix inserted before the history calculation is a benchmark only.

Firewall H2 — No real-time kernel as a probability covariance. The oscillatory retarded memory kernel may change sign. Positivity is carried by the Euclidean transfer and its positive history measure, not by treating the retarded kernel as a covariance.

Firewall H3 — No forced detailed balance. Irreversible committed histories need not be reversible. Detailed balance is optional HP2R, not mandatory HP2.

Firewall H4 — No selected generation graph. Generation adjacency, weights and phases must descend from the stationary history conductance and record transport.

Firewall H5 — No branchwise norm reset. Dividing each raw branch transfer by its own leading eigenvalue erases the physical free-energy comparison and is forbidden before branch selection.

Firewall H6 — No finite-branch assumption. The admissible branch space may be finite, countable or continuous. Compactness and a positive gap must be demonstrated, not inferred from a list of candidates.

Firewall H7 — No identity-completion default. The MASD-1 identity completion branch is a special case. It cannot be retained if the history law returns a nontrivial R_Ω.

Firewall H8 — No cross-operator branch comparison. Two branches may be compared by free energy only after they are exhibited as components of one master history operator with one reference measure (HM9, §43A). Comparing Perron roots of separately normalised or separately constructed operators is forbidden, because a Perron root is an invariant of a single operator and carries no cross-operator meaning.

Firewall H9 — No tuning of the executed wheel. The canonical kernel T_0, its stationary measure, the six defect score generators F_A and the resulting W_{wheel} are **frozen, substrate-grounded and symmetry-adapted** — not fitted to any downstream output; uniqueness of the F_A within each representation sector remains open and is part of HM10.3 (F26). In particular the wheel-vanishing cross blocks $W_{BH} = W_{CH} = W_{RH} = 0$ (Corollary 14C.1) may not be filled in by hand to match a target; any nonzero value for them must be *derived* from the multi-cell, internal-representation, terminal-record or completion lift. Choosing F_A off the symmetry-adapted family, or rescaling W_{wheel} entries to improve an output, voids the subcertificate.

Firewall H10 — No empty formal direct sum for branch comparison. Writing $\mathbb{T}_{formal} = \mathbb{T}_{CFB1} \oplus \mathbb{T}_{CFB2} \oplus \mathbb{T}_{CFB3}$ is mathematically legal but physically empty: the three summands currently use different configuration spaces and constitutive structures, so the relative measure assigned to them is arbitrary and simply *becomes* the new branch-selection parameter. The Common-History-Operator requirement (§43A) is satisfied only when the components arise as spectral or invariant projections $\{P_\beta \mathscr{H}^{hist} P_\beta\}$ of *one derived generator constructed before the branches* — $\mathscr{H}^{hist} \to \{P_\beta \mathscr{H}^{hist} P_\beta\}$ — not when separately built operators are placed side by side afterwards. A post-hoc direct sum with a free relative normalisation fails G9c and violates HM5/H5/H8.

Part II — Bath-Resolved Finite-Regulator Histories

6. The Resolved MASD Carrier

Let X_n be the finite physical quotient of the MASD configuration space after the exact gauge, record-coboundary and terminal-null identifications already known at the kinematic level. The operational-geometry programme supplies a finite-dimensional Hilbert structure and normalised operational measure $\mu\_op^{(n)}$. A point $x \in X\_n$ represents a physical resolved configuration, not a raw coordinate representative.

The six defects define a smooth finite-regulator map $\mathbf{D}\_n : X\_n \to \mathcal{R}\_prim^{(n)}$, whose zero locus is the admissible stationary set

$$\mathcal{M}\_n = \{ x \in X\_n : \mathbf{D}\_n(x) = 0 \}.$$

No choice of history metric is needed to define this locus. Its existence and, crucially, its *rigidity* (whether the reduced physical response is constant along it) are exactly MASD-1's open falsifier F28, and they reappear here as the null-classification obligation of §17 and Gate G5.

7. The Unresolved Record Bath

The memory-kernel programme identifies a physical environment: sub-threshold record modes that fail to persist as terminal facts but exchange energy, phase and distinguishability with the resolved commitment sector. At finite regulator, write the unresolved bath operator $\Omega^2\_(Q,n) \geq 0$ on $\mathcal{Q}\_n$ — positive on the modes it retains, but not assumed strictly positive, since under §56A.3 it is the compression $Q\_n\mathbb{H}^{hist}\ Q\_n$ and may be singular; the reduction then uses the pseudo-inverse under the range condition of Theorem 16.1b — with eigenpairs

$$\Omega^2\_(Q,n)\ v\_(n,a) = \omega^2\_(n,a)\ v\_(n,a).$$

The frequencies are not adjustable damping parameters; they are finite-regulator eigenvalues of the unresolved-record block. The defect–bath coupling is a real linear map $C\_n : \mathcal{R}\_prim^{(n)} \to \mathcal{Q}\_n$ with mode components $c\_(n,aA) = \langle v\_(n,a), C\_n\ \iota\_A\rangle$.

The complete numerical debt of the bath layer is therefore finite and explicit:

$$[\![\{\omega\_(n,a)\},\quad \{c\_(n,aA)\}.]\!]$$

Neither set may be selected from masses, couplings or mixing data (HM4, Firewall H1). **This finite pair, together with the resolved local block W_core (via V_loc, §10), constitutes the complete generator $\hat{\mathbb{H}}\_n^{hist}$ into which HMGBC-1 relocates the constitutive freedom** (§19). Whether the relocation is a reduction depends on whether the *whole generator* — not merely (Ω^2_Q, C_n) — is *derived* from the substrate memory-kernel and operational-geometry programmes and shown regular on the real regulator. This is the history-content gate, and it is on $\hat{\mathbb{H}}\_n^{hist}$ because $W\_prim = W\_core(0) - C\_n^\dagger\ (\Omega^2\_(Q,n))^+\ C\_n$ depends on W_core as an independent input (Theorem 16.1b).

8. The Enlarged Operational Carrier

A history state is $\hat{x} = (x, q) \in \hat{X}\_n = X\_n \times \mathcal{Q}\_n$, with reference measure

$$d\hat{\nu}\_n(\hat{x}) = d\mu\_op^{(n)}(x)\ d\nu\_Q^{(n)}(q),$$

where $d\nu_Q$ is the canonical Gaussian or Lebesgue volume fixed by the bath kinetic normalisation. The enlarged kinetic Dirichlet generator is

$$\hat{\Delta}_n = \Delta_{(X,n)} \otimes I + I \otimes \Delta_{(Q,n)}.$$

The first term is supplied by operational distinguishability geometry; the second is the standard positive bath-coordinate Laplacian. Neither contains Standard Model couplings.

9. Reference Measure and Finite History Space

For N steps, let $\hat{\gamma} = (\hat{x}_0, \ldots, \hat{x}_N) \in \hat{X}_n^{(N+1)}$. The canonical reference path measure is

$$d\hat{\mu}_0^{(n,N)} = d\hat{\pi}_{(0,n)}(\hat{x}_0) \prod_{(k=0)}^{(N-1)} \hat{K}_{(0,n)}(d\hat{x}_{(k+1)} \mid \hat{x}_k),$$

with $\hat{K}_{(0,n)}$ the defect-free reference kernel. The physical history measure is absolutely continuous with respect to it on the admitted support:

$$d\hat{\mu}_{hist}^{(n,N)} = \exp[-\mathcal{A}_{(n,N)}(\hat{\gamma})]\, d\hat{\mu}_0^{(n,N)}.$$

Absolute continuity is an admission condition. A branch whose support is singular relative to the declared reference requires a separate stratum and may not be hidden inside an infinite action. (This is one concrete way HM9/§43A can fail: two branches with mutually singular supports are not components of one operator on one measure.)

10. Bath-Resolved Transfer Generator

Define the Euclidean history Hamiltonian $\hat{\mathbb{H}}_n = -\hat{\Delta}_n + V_n(x, q)$, with

$$V_n(x, q) = V_{loc}^{(n)}(x) + \tfrac{1}{2} \langle q, \Omega^2_{(Q,n)} q\rangle + \langle q, C_n \mathbf{D}_n(x)\rangle.$$

The local potential $V_{loc}^{(n)}$ contains the MASD substrate potential and the direct local history cost returned before the unresolved bath is eliminated. The bilinear coupling is the most general declared-order real coupling between unresolved modes and the primitive defect vector. The full transfer is

$$[\![\ \mathbb{T}_n = e^{(-a_t \hat{\mathbb{H}}_n)}, \ ]\!]$$

with Trotter representation

$$\mathbb{T}_n = \lim_{(m\to\infty)} [e^{((a_t/m) \hat{\Delta}_n)} e^{(-(a_t/m) V_n)}]^m,$$

each factor of which preserves the positive cone when V_n is real and bounded below (HM3).

11. Positivity of the One-Step Kernel

The heat kernel of $e^{(a_t \hat{\Delta}_n)}$ is nonnegative; multiplication by $e^{(-a_t V_n)}$ is positive; the Trotter limit therefore has a nonnegative integral kernel $\hat{T}_n(\hat{x}', \hat{x}) \geq 0$. If the operational configuration graph is connected, the bath covariance is nondegenerate, and the allowed transition set communicates across the physical branch, the kernel is positivity improving:

```
f ≥ 0, f ≠ 0  ⟹  𝕋_n f > 0.
```

This is the correct microscopic route to terminal-cone positivity. It is stronger than what MASD-1's §10.1.1 could assert: there, cone positivity of the heat semigroup was inferred from a *posited* Dirichlet form; here it is a property of an explicitly positive transfer on the enlarged carrier, whose Doob normalisation is Markov. Positive-semidefinite Hessian matrices alone would never have sufficed — a distinction the paper enforces as Firewall H2 and re-states at §20.

11A. The Raw Master History Transfer from the Admissibility Graph

The history-content gate (§19) asks whether the master generator is derived or chosen. This section supplies a graph-based construction from the *complete admissible move graph* and the common operational measure. **Its role is diagnostic and derivational, not physical**: it proves that the canonical K = 7 wheel kernel is recovered rather than inserted (Theorem 11A.1) and it characterises admissible support, but the transfer used for physical execution is the spectrally positive *operational heat transfer* of §56A.2, built from the inherited Dirichlet generator L^{op}. The reason is the inertia obstruction of §11A.2 below. The single premise here is HM11.

11A.1 The full finite-regulator configuration graph

Let $\mathcal{C}_n = \{ (\rho, U, J, \Phi; \Psi, \Psi; q) \}$ be the complete finite-regulator configuration carrier, *including* the unresolved record variables q. The six typed defects are as fixed by MASD-1. Define the symmetric admissibility matrix

```
⟦ A_n(z, z′) = 1 if z ↔ z′ by one admissible elementary VERSF update, else 0,
⟧
```

including the persistence self-transition $A_n(z, z) = 1$ whenever the local carrier persists through one refinement step.

Scope caution: a finite regulator does not by itself give a finite move graph. The carrier (ρ, U, J, Φ; Ψ, Ψ; q) is finite-*dimensional* at finite regulator, but several of its variables remain **continuously valued**. A binary matrix $A_n(z, z')$ with sums over z′ is therefore well defined on the seven-state closure catalogue and **not** yet defined on the full MASD configuration space. Closing the construction requires *either* (i) a declared discretisation of every continuous carrier, with convergence under refinement, *or* (ii) a measure-valued proposal kernel $a_n(z, dz')$ with $d_n(z) = \int a_n(z, dz')$ and $K_{(0,n)}(z, dz') = a_n(z, dz')/d_n(z)$. Route (ii) requires the theory to **derive the distribution of update amplitudes**, not merely which update *types* are permitted: HM11's equal weighting of symmetry-equivalent move types does not determine a measure over their continuous amplitudes. HM11 accordingly closes the wheel and does **not** yet uniquely specify the complete MASD/CAR/Fock history law (F32).

With $d_n(z) = \sum_{(z')} A_n(z, z')$ and the normalised symmetric adjacency $\mathcal{S}_n = D_n^{(-1/2)} A_n D_n^{(-1/2)}$, the Perron vector of $\mathcal{S}_n$ is $r(z) = \sqrt{d_n(z)}$ and its Doob transform is

$$⟦\ K_0(z' \mid z) = \mathcal{S}_n(z, z')\, r(z') / r(z) = A_n(z, z') / d_n(z).\ ⟧$$

Theorem 11A.1 — The canonical wheel is recovered, not inserted

On the bare $K = 7$ closure catalogue the admissibility graph has degrees $d(h) = 7$ (hub persistence plus six rim moves) and $d(b_i) = 4$ (hub, self, two rim neighbours). Hence K_0 assigns **1/7 to every permitted hub transition and 1/4 to every permitted rim transition** — which is exactly the canonical kernel T_0 of §14A, entry for entry. **Verified exactly.** The published canonical wheel is therefore a *consequence* of the admissibility graph together with HM11, not an independently posited kernel. This closes, at the canonical level, the objection that T_0 was chosen; it does not close the full gate, because the complete MASD/CAR/Fock move graph has not been enumerated (F32).

11A.2 The raw transfer, and why it is *not* the Euclidean transfer

Let V_n^{ref} be the upstream finite-regulator reference cost — the already-derived substrate potential and the unresolved-record reference norm, and **not** the unknown physical W_{prim}:

$$⟦\ V_n^{ref}(z) = \Gamma_V^{(n)}(z) + \tfrac{1}{2} \|q\|^2_{(\mathcal{Q}_n,\ op)}.\ ⟧$$

Define the **raw cone-positive graph transfer**

$$⟦\ \mathfrak{B}_n = M_{(\exp(-a_t V_n^{ref}/2))}\ \mathcal{S}_n\ M_{(\exp(-a_t V_n^{ref}/2))}.\ ⟧$$

A spectral caution that must not be elided. $\mathcal{S}_n$ is similar to the Markov kernel $D_n^{-1}A_n$, so its spectrum is that of T_0 on the canonical wheel — which **contains negative eigenvalues** ($-\tfrac{1}{4}$ and $-3/28$, §14A). Congruence by a strictly positive diagonal preserves inertia (Sylvester), so $\mathfrak{B}_n$ likewise carries negative spectral inertia. **Therefore $\mathfrak{B}_n$ may not be identified with $e^{(-a_t \mathbb{H}_n)}$ for any real, self-adjoint, bounded-below $\mathbb{H}_n$**, whose spectrum would be strictly positive. Writing the graph transfer as a Euclidean exponential would contradict HM3 and §14D, which squares T_0 for exactly this reason. **Verified: the negative inertia persists under every positive-diagonal congruence.**

The two objects must therefore be kept distinct:

$$[\![\ \mathfrak{B}_n : \text{raw cone-positive transfer — supplies Perron data (HP1/HP2) and the Doob history law;}\ ]\!]$$

$$[\![\ \mathbb{T}_{(E,n)} = \mathfrak{B}_n^\dagger\, \mathfrak{B}_n \geq 0,\quad \mathbb{H}_{(E,n)} = -(1/2a_t) \log \mathbb{T}_{(E,n)} \text{ on } \ker(\mathbb{T}_{(E,n)})^\perp : \text{the positive Euclidean relaxation transfer and its generator.}\ ]\!]$$

This is the multi-cell generalisation of the $T_E = T_0^2$ construction of §14D, and it is the object that legitimately carries a real bounded-below Hamiltonian. (An equivalent route is the graph heat operator $\mathbb{T}_{(E,n)} = M\, e^{-a_t(I - \mathcal{S}_n)}\, M$, which is manifestly positive; the squared form is retained here for continuity with §14D.)

$\mathfrak{B}_n$ is nonnegative entrywise; self-adjoint in the declared reference measure; positivity improving on each connected component (HP1); **defined before any branch decomposition**; globally normalised by one action unit a_t; and independent of Standard Model measurements (Firewall H1, HM8). Its Doob transform returns the normalised stochastic history law *after* the raw Perron roots have been retained for branch comparison (Firewall H5). $\mathfrak{B}_n$ supplies a **candidate support-level implementation of part of §43A** — a common support, communicating components, and raw component Perron data *once the complete graph exists*. It does not supply the exhaustive branch algebra, continuous or metastable sectors, the full MASD/CAR/Fock graph, or the physical common operator. Everything requiring a Hamiltonian is supplied by $\mathbb{T}_{(E,n)}$.

11A.3 The metric as an output, not an input

With the typed transition scores $F_{(A,n)}$ of §14F, the deformed kernel is $K_{(n,\varepsilon)}(z'|z) \propto K_{(n,0)}(z'|z) \exp[\sum_A \varepsilon_A F_{(A,n)}(z,z')]$, the centred scores are $S_{(A,n)} = F_{(A,n)} - \mathbb{E}_{(K_{(n,0)})} F_{(A,n)}$, and

$$[\![\ (W_{prim}^{(n)})_{AB} = \lim_{N\to\infty} (1/N)\, \mathbb{E}[\, \partial_A \log d\mu_N \cdot \partial_B \log d\mu_N\,]_{(\varepsilon=0)}.\ ]\!]$$

W_{prim} is thus an output of the one operator and is nowhere used to construct it. That is the non-circularity §4 demands — with one qualification that must be stated rather than left implicit. The defect-source injection $E_{(D,n)}$ of §56A.3 is a *minimum-norm* pseudo-inverse, and minimum-norm is defined relative to the inner product supplied by the operational geometry. Hence W_{core} — and through it W_{prim} — is a **functional of the reference operational geometry**, not merely downstream of it. §4 permits this: upstream data may fix coordinates and quotient volume. But since F6 is precisely about normalisations leaking into the metric, the dependence is declared here rather than discovered later.

Status. ⟦ Canonical-wheel admissibility kernel and raw graph transfer: CONSTRUCTED under HM11. ⟧ This is deliberately *not* called a master-history operator: the complete continuous move law is undefined (F32), the graph object is diagnostic rather than the physical Euclidean transfer

(§56A.2), the operational-generator attribution is open (G11a), and the exhaustive branch algebra is open (G9c). What is constructed is an admissibility kernel and the raw graph transfer built on it.

12. Reduction to Non-Markovian Resolved Dynamics

Let P project the enlarged carrier onto resolved observables and Q = I − P. Exact projection gives a discrete Nakajima–Zwanzig equation

$$(d/dk)\ P\ \hat{\rho}_k = P\ L\ P\ \hat{\rho}_k + \sum_{j=0}^{k-1} \mathcal{K}_{k-j}\ P\ \hat{\rho}_j + \mathcal{N}_k,$$

with discrete memory kernel $\mathcal{K}_m = P\ L\ Q\ (Q\ L\ Q)^{m-1}\ Q\ L\ P$ and $\mathcal{N}_k$ depending on unresolved initial data. Thus

⟦ Doob-normalised bath-resolved process Markovian ⟹ resolved process generally non-Markovian. ⟧

The memory papers' Volterra kernels are continuum, reduced-sector manifestations of this exact finite-regulator statement. This is the load-bearing conceptual result of Part II: memory in the visible theory is a *derived* consequence of eliminating memoryless bath modes, not an independent postulate — which is precisely why the metric it produces cannot be freely chosen once the complete generator $\hat{\mathbb{H}}_n^{hist}$ (local block, bath block and coupling together) is fixed.

Part III — Radon–Nikodym Rate and Operational Metric

13. Finite-History Action versus Asymptotic Rate

The finite-history log-likelihood ratio is $\mathcal{A}_{(n,N)} = -\log(d\hat{\mu}_{hist}^{(n,N)} / d\hat{\mu}_0^{(n,N)})$. At finite N this is an action, not a rate. The rate functional is the normalised long-history object

$$⟦\ \mathcal{I}_n(\varepsilon) = \lim_{(N\to\infty)}\ (1/N)\ D_{KL}(\ \hat{\mu}_0^{(n,N)}\ \|\ \hat{\mu}_\varepsilon^{(n,N)}\ ).\ ⟧$$

This distinction prevents "rate functional" from being used loosely for an unnormalised finite path action. (The Principal Claim writes $\mathcal{J}_n = \lim N^{-1}\mathbb{E}_{(\mu_0)}\mathcal{A}_{(n,N)}$ with $\mathcal{A} = -\log(d\mu_{hist}/d\mu_0)$, while the display above writes the KL rate $D_{KL}(\hat{\mu}_0 \| \hat{\mu}_\varepsilon)$; these coincide, and both Hessians at ε = 0 equal the Fisher matrix, so the two forms are used interchangeably below.) Existence of the limit follows under appropriate stationarity, ergodicity and integrability conditions; a *uniform* spectral gap is not strictly necessary for existence but is the sufficient mechanism used here — it delivers exponential mixing, differentiability of the rate, resolvent control and regulator-uniform convergence together, which is why HP2's Perron gap is a required quantitative return (§21) rather than an assumed one.

14. Defect Deformation and Score Map

Choose a frozen orthonormal basis $\{e_A\}$ of the primitive defect tangent carrier. A defect deformation is $\mathbf{D}(\varepsilon) = \sum_A \varepsilon_A e_A$. With reduced history kernel $K_n^\varepsilon(y \mid x)$, the score is

$$S_A(x, y) = \partial/\partial\varepsilon_A \log K_n^\varepsilon(y \mid x) \,|_{(\varepsilon=0)},$$

and the information-rate matrix is

$$(W_{rate}^{(n)})_{AB} = \mathbb{E}_{(\pi_n K_n)}[S_A S_B].$$

Under regularity, $(W_{rate}^{(n)})_{AB} = -\mathbb{E}(\pi_n K_n)[\partial^2/\partial\varepsilon_A \partial\varepsilon_B \log K_n^\varepsilon \,|_{(\varepsilon=0)}]$. For a genuinely history-correlated reduced process the two definitions are **not** interchangeable, and the paper uses the path-score form

$$(W_{rate}^{(n)})_{AB} = \lim_{(N\to\infty)} (1/N)\, \mathbb{E}[\partial_A \log d\mu_\varepsilon^{(N)} \cdot \partial_B \log d\mu_\varepsilon^{(N)}]_{(\varepsilon=0)}$$

as primary. The one-step formula above presumes a transition law $K_\varepsilon(y \mid x)$; after bath elimination the resolved law is $K_\varepsilon(y_k \mid y_0, \ldots, y_{(k-1)})$, so the one-step expression is valid **only** under one of four conditions, one of which the execution must declare: (a) the Fisher information is computed on the *enlarged Markov* carrier, where the one-step form is exact; (b) the resolved state is enlarged to include sufficient memory variables to restore the Markov property; (c) conditional scores given the complete past are used, with the martingale-difference decomposition proved so that cross-time terms vanish in the Cesàro limit; or (d) a factorisation is proved that eliminates all cross-time terms. Routes (a) and (c) are one mechanism, (a) being its Markov instance: on the enlarged carrier the centred score satisfies $\sum_y K(y \mid x) S_A(x, y) = 0$ by construction, so the scores form a **martingale difference sequence**, all cross-time covariances vanish, and the asymptotic variance equals the one-step variance — which is exactly why the one-step formula is exact there. Route (a) is the natural choice here, since §§8–11 construct precisely that carrier. Absent such a declaration the one-step formula is a Markov-sector approximation, not an identity (F38).

14A. The Canonical K = 7 History Kernel

The canonical closure catalogue is

$$\mathcal{K} = \{ h, b_0, b_1, b_2, b_3, b_4, b_5 \},$$

with one hub state h and six cyclically ordered boundary states b_i, indices mod 6. The canonical refinement kernel T_0 is the row-stochastic 7×7 matrix that sends the hub uniformly to itself and all six rim states (each with weight 1/7), and sends each rim state b_i with weight 1/4 to the hub, to itself, and to its two rim neighbours $b_{(i-1)}$, $b_{(i+1)}$:

$$T_0(h, \cdot) = (1/7)(1, 1, 1, 1, 1, 1, 1) \quad \text{on } (h, b_0, \ldots, b_5),$$
$$T_0(b_i, h) = T_0(b_i, b_i) = T_0(b_i, b_{(i-1)}) = T_0(b_i, b_{(i+1)}) = 1/4, \quad \text{all other entries } 0.$$

It is row-stochastic, irreducible, aperiodic and reversible with respect to

$$⟦ \pi = (1/31)(7, 4, 4, 4, 4, 4, 4). ⟧$$

Its exact spectrum is

$$⟦ \mathrm{spec}(T_0) = \{ 1, ½, ½, -¼, -3/28, 0, 0 \}. ⟧$$

Consequently, its persistent projector on $L^2(\pi)$ is $P_{pers} f = \langle 𝟙, f\rangle_\pi 𝟙$, and its nonpersistent complement is

$$Q_{inc} = I - P_{pers}, \quad \dim \mathrm{Ran}\, Q_{inc} = 6.$$

This is an explicit finite-regulator realisation of the persistent/nonpersistent split that the record-current and memory-kernel programmes describe abstractly. It does not yet identify every vector in Ran Q_{inc} with a physical unresolved oscillator; it supplies the exact finite relaxation carrier on which that identification must be made.

Theorem 14A.1 — Canonical-Wheel History Conditions

The canonical wheel satisfies HP1 ($T_0 f \geq 0$ whenever $f \geq 0$, with positivity improvement after finitely many steps); HP2 ($\lambda_0 = 1$ is simple, and $1 - \max_{(\lambda\neq1)}|\lambda| = ½$); and HP2R ($\pi_x T_0(x, y) = \pi_y T_0(y, x)$). Thus the qualitative *and quantitative* finite-level Perron gap is not merely assumed on the canonical wheel:

$$⟦ \Delta_{Perron}^{wheel} = ½. ⟧$$

What remains open is a regulator-uniform lower bound for the full multi-cell and full-carrier history operator (§21, Gate G3).

14B. Six MASD Defects as Admissible Wheel Perturbations

Let $\varepsilon = (\varepsilon_C, \varepsilon_J, \varepsilon_T, \varepsilon_R, \varepsilon_B, \varepsilon_H)$ be six scalar defect coordinates. These are **not** asserted to exhaust the vector-valued MASD defect carriers; they are the minimal symmetry-adapted wheel representatives of those carriers. Define a support-preserving exponential deformation

$$T_\varepsilon(x, y) = T_0(x, y) \exp\Big[\sum_A \varepsilon_A F_A(x, y) \Big] / Z_x(\varepsilon), \quad Z_x(\varepsilon) = \sum_z T_0(x, z) \exp\Big[\sum_A \varepsilon_A F_A(x, z) \Big].$$

This family is row-stochastic for every finite ε, creates no transition absent from the canonical wheel, and remains strictly positive on every canonical allowed edge. Write $\theta_i = 2\pi i/6$. The six score generators are:

- **Closure defect.** $F_C(b_i, b_{i+1}) = +1$, $F_C(b_i, b_{i-1}) = -1$ — the unique rotation-invariant oriented circulation on the rim, the minimal wheel image of vacuum-relative closure holonomy.
- **Constitutive-current defect.** $F_J(h, b_i) = F_J(b_i, h) = 1$ — the isotropic radial conductance direction, the minimal wheel image of constitutive current realisation.
- **Transport-conservation defect.** $F_T(h, b_i) = \cos\theta_i$, $F_T(b_i, h) = -\cos\theta_i$ — one real axis of the lowest nonconstant radial harmonic (a local source/sink imbalance). Its sine partner is the symmetry-equivalent second basis vector of the same two-dimensional D_6 representation; selecting the cosine axis fixes a coordinate orientation and does not remove the partner from the full vector-valued carrier.
- **Record-composition defect.** For the undirected rim edge $e_i = \{b_i, b_{i+1}\}$, $F_R(b_i, b_{i+1}) = F_R(b_{i+1}, b_i) = (-1)^i$ — the alternating real triangle-composition mode carried by the six elementary wheel triangles.
- **Polar bond-compatibility defect.** $F_B(b_i, b_{i+1}) = \cos\theta_i$, $F_B(b_{i+1}, b_i) = -\cos\theta_i$ — a local oriented bond twist with zero total circulation, $\sum_{i=0}^{5} \cos\theta_i = 0$, so it distinguishes local link–record polar mismatch from the global closure circulation of F_C.
- **Scalar/completion-stretch defect.** $F_H(x, x) = 1$ for every $x \in \mathcal{K}$ — the D_6-scalar enhancement of local persistence relative to transport away from the same state.

All score values not specified above are zero. The tangent matrices are

$$(\Delta_A T)_{xy} = \partial T_\varepsilon(x,y)/\partial \varepsilon_A \big|_{\varepsilon=0} = T_0(x, y) \big[F_A(x, y) - \bar{F}_A(x) \big], \quad \bar{F}_A(x) = \sum_z T_0(x, z) F_A(x, z),$$

which obey $\sum_y (\Delta_A T)_{xy} = 0$ and $(\Delta_A T)_{xy} = 0$ whenever $T_0(x, y) = 0$.

Theorem 14B.1 — Canonical-Wheel Six-Defect Tangent Theorem

The six matrices $\{ \Delta_C T, \Delta_J T, \Delta_T T, \Delta_R T, \Delta_B T, \Delta_H T \}$ are six linearly independent admissible tangent directions of the canonical $K = 7$ stochastic kernel, and form a natural finite D_6-adapted scalar representative set for the MASD primitive defect labels. Two things they do **not** establish: (a) that the full MASD differentials compress to precisely these matrices (faithfulness); and (b) that this representative set is *unique* rather than one natural choice among symmetry-compatible alternatives. Both are the content of the **Defect-Lift Obligation**:

$$ D\mathcal{D}_A \longrightarrow \Delta_A\, T \quad \text{must be derived — faithfully and uniquely — on the full typed carrier (HM10, §14E).} $$

Faithfulness is HM10.1–4 (sector match, rank, equivariance, support typing); uniqueness is the equivariance item §14E(3), which is what would promote "a natural choice of F_A" to "the derived F_A." Until §14E(3) is proved, the specific score generators of §14B are well-motivated and symmetry-adapted, not forced (F26).

14C. Exact Canonical-Wheel History Metric

For the canonical stationary process, define the centred scores $S_A(x, y) = F_A(x, y) - \bar{F}_A(x)$. The one-step Fisher–Onsager matrix is

$$ (W_{wheel})_{AB} = \sum_{(x,y)} \pi_x\, T_0(x, y)\, S_A(x, y)\, S_B(x, y). $$

In the ordering (C, J, T, R, B, H), direct evaluation gives the symmetric matrix with diagonal (12/31, 75/434, 21/124, 12/31, 21/124, 75/434) and the only nonzero off-diagonal entries $W_{JH} = W_{HJ} = -33/434$ and $W_{TB} = W_{BT} = 3/248$:

$$ W_{wheel} = \begin{bmatrix} 12/31 & 0 & 0 & 0 & 0 & 0 \\ 0 & 75/434 & 0 & 0 & 0 & -33/434 \\ 0 & 0 & 21/124 & 0 & 3/248 & 0 \\ 0 & 0 & 0 & 12/31 & 0 & 0 \\ 0 & 0 & 3/248 & 0 & 21/124 & 0 \\ 0 & -33/434 & 0 & 0 & 0 & 75/434 \end{bmatrix}. $$

Its eigenvalues are

$$ \mathrm{spec}(W_{wheel}) = \{ 3/31,\ 39/248,\ 45/248,\ 54/217,\ 12/31,\ 12/31 \}, $$

so $W_{wheel} > 0$, rank $W_{wheel} = 6$, and condition number $\kappa(W_{wheel}) = 4$. There is no information-null direction in this six-scalar reduction.

Corollary 14C.1 — Calculated Cross Blocks

The only nonzero cross blocks are

$$ W_{JH} = -33/434, \quad W_{TB} = 3/248. $$

The first arises because radial current transport and local self-persistence compete for probability weight in the same rows of the kernel; the second because the radial continuity harmonic and the oriented bond harmonic overlap in the same first D_6 representation. All remaining cross blocks vanish in this scalar wheel reduction by symmetry and support orthogonality. In particular,

⟦ W_BH = W_CH = W_RH = 0 ⟧

on the bare seven-state wheel. **This is a substantive negative result.** The minimal wheel does not generate the scalar–bond, scalar–closure or scalar–record entanglement that the full nonfactorising MASD action may require. Such blocks must arise, if physical, from the multi-cell, internal-representation, terminal-record or completion-carrier lift. They may not be inserted into the wheel metric to improve downstream outputs (Firewall H1, H4).

Corollary 14C.2 — Reversible and Irreversible Directions

The symmetric score directions J, R, H preserve detailed-balance structure to first order after the corresponding stationary measure is updated. The oriented directions C, T, B generically produce stationary history currents and need not preserve detailed balance; they nevertheless remain positive stochastic deformations and retain bi-Perron stationarity for sufficiently small finite deformation. This explicitly supports the HMGBC distinction between HP2 (bi-Perron stationarity) and HP2R (reversible reference certificate): the irreversible defect directions are physical and present already at the wheel level.

Corollary 14C.3 — Choice-Independent, Generator-Family and Model-Specific Content

Not every feature of W_wheel has the same epistemic status, and the paper must not quote them on the same footing. The content partitions **three** ways, not two: a **choice-independent** part, a **generator-family** part, and a **model-specific magnitude** part. The separation is made by asking two distinct questions — which features survive an arbitrary D_6-symmetric, support-preserving change of the canonical kernel (a reweighting of the hub↔rim probability away from 1/7, say, or of the rim self/neighbour split) with the score generators held fixed; and which survive re-choosing the generators with the kernel held fixed.

There are in fact **two** independent invariance questions, and they must not be run together: invariance under *reweighting the kernel with the score generators held fixed*, and invariance under *re-choosing the generators with the kernel held fixed*. The second freedom is exactly HM10.3 / F26, which is open, and it is not small — see Theorem 14C.6.

Choice-independent invariants (survive both). There are exactly two, and they are narrower than a first reading suggests:

- **Inter-irrep pairings vanish**, and in particular W_BH = W_CH = W_RH = 0. This is forced by representation orthogonality for **any** representatives of the sectors involved, whatever the weights and whatever generators are chosen inside each sector.
- **Every linearly independent tangent family has full Fisher rank.** The proof is one line: the Fisher weights $\pi\_x/T_0(x, y)$ take exactly two values — 49/31 on hub edges and 16/31 on rim edges — both strictly positive, so the form is a positive diagonal in the edge basis and is nondegenerate on every subspace, the 24-dimensional tangent space included.

Rank is therefore a property of independence, not of the particular directions. (Consistently, the Gram spectrum on that space runs from 16/31 to 49/31.)

Properties of the frozen generator choice (survive reweighting only). Everything *within* an irrep sector belongs here — **including the non-vanishing of the cross blocks, not merely their signs**. Because J and H occupy two directions of a three-dimensional reflection-even m = 0 space (Theorem 14C.6), one may choose an admissible, same-sector H′ that is Fisher-orthogonal to J, giving **W_JH′ = 0 exactly** (**verified by explicit construction**); likewise a copy of B orthogonal to the selected T within the four copies of the m = 1 ⊕ m = 5 doublet. So:

```
⟦ within-irrep cross blocks — their values, their signs, and even whether they
vanish — are properties of the declared generator family, not of the wheel
(except as below). ⟧
```

They become physical **only once the actual defect differentials J_A select the physical representatives** (§14F, HM10.3, F26). Lemma 14C.4's sign results are invariance statements under *kernel reweighting with the generators held fixed*, and are not to be read more widely — except where support typing applies, which does extend the J–H case beyond reweighting (Corollary 14C.4b).

One exception, recoverable rather than conventional. The demotion above concerns *unrestricted* re-choice within a sector. Restricting the **support** of the persistence score restores the J–H sign: if $\mathcal{D}$_H is diagonal-supported with weights of one sign, then W_JH < 0 for every admissible weight assignment (Corollary 14C.4b). The J–H sign is therefore conventional only under re-choices that change the *support type* of H, and the residual freedom reduces to two named questions about the physical $\mathcal{D}$_H. No corresponding result is available for W_TB, whose doublet freedom is four-dimensional; that sign remains demoted.

Model-specific magnitudes (do NOT survive):

- the rational values themselves — 12/31, 75/434, 21/124, −33/434, 3/248 — and the eigenvalues {3/31, 39/248, 45/248, 54/217, 12/31, 12/31}, the condition number 4, and every other number. These depend on the canonical hub↔rim weight 1/7 and change under reweighting.

Explicit demonstration. Reweighting the hub→rim probability from the canonical 1/7 to 1/10 (holding graph, scores and D_6 symmetry fixed) leaves W_BH = 0 and sgn W_JH < 0 unchanged, but moves W_JH from −33/434 ≈ −0.076 to ≈ −0.115. The zero and the sign are invariant **under reweighting**; the magnitude is not. Neither statement makes the sign choice-independent — that is a separate question, settled negatively by Theorem 14C.6 and recovered only under support typing (Corollary 14C.4b).

Consequence for the claim. Independently of the generator representatives, the wheel fixes full Fisher rank and the inter-irrep vanishing relations — **and, given support typing, the J–H sign** (Corollary 14C.4b). Everything else within an irrep sector — block patterns, the T–B sign, and non-vanishing — **must be re-derived from the physical defect differentials** and may not be carried over. The rational magnitudes are properties of the canonical wheel model, quoted throughout as an executed benchmark and never as HFB-1 inputs. Anyone consuming W_wheel downstream may consume only its rank and inter-irrep zeros; everything else must be recomputed on the physical operator (Firewall H9, F24, F25).

Lemma 14C.4 — The structural partition, proved for arbitrary admissible weights

Corollary 14C.3's partition is not merely illustrated by one reweighting; the sign and vanishing statements hold for **every** admissible D_6-symmetric support-preserving weight assignment. Parameterise the hub row by (self s, six rim p each, s + 6p = 1) and each rim row by (hub q, self r, two neighbours n each, q + r + 2n = 1). Then, computing the one-step Fisher contributions row by row:

- **Hub row.** The centred scores of J and H are exactly **opposite** there. Since $\bar{F}_J(h) = 6p$ and $\bar{F}_H(h) = s$, one has S_J = (−6p on the diagonal, **s** off) and S_H = (+6p, −s), so $S_H = -S_J$ identically. Weighting by the row's stationary mass, the hub contribution to W_{JH} is $\mathbf{-6\,\pi_h\, s\, p < 0}$ (the unweighted row sum is −6ps, since 6p + s = 1). This anti-proportionality is the precise content of the informal statement that current and persistence "compete for probability weight in the same rows."
- **Rim rows.** $\sum_y T\, S_J\, S_H = q\, r\, (q + r + 2n - 2) = -q\, r < 0$ under the normalisation.

Hence **$W_{JH} < 0$ for all admissible positive weights.** ■ Similarly $\sum_y T\, S_T\, S_B = q\, n\, a(a - b)$ gives **$W_{TB} > 0$** for all admissible weights; and $\sum_y T\, S_H\, S_B = -n\, r\, (a - b)$ sums to zero over the rim because $\sum_i \cos\theta_i = 0$, so **$W_{BH} = 0$ exactly for arbitrary weights**, not merely canonical ones. **Verified symbolically.**

This upgrades §14C.3 from an illustrated partition to a proved one, which matters because F25 and every downstream "structure but not magnitude" consumer rest on it.

Corollary 14C.4b — The J–H sign is recovered by support typing

Theorem 14C.6 demotes the J–H sign because H may be re-chosen within a three-dimensional sector. That freedom is narrower than it appears once the *support* of the persistence score is typed. Let F_H be diagonal-supported with $F_H(h,h) = \alpha$ and $F_H(b_i,b_i) = \beta$, F_J unchanged. Then, for every admissible weight assignment,

$$W_{JH} = -6\,\pi_h\,\alpha\, s\, p - 6\,\pi_b\,\beta\, q\, r,$$

verified symbolically, and returning −33/434 exactly at the canonical weights with α = β = 1. Both terms are negative whenever α, β > 0, so

$$F_H \text{ diagonal-supported with } \alpha, \beta > 0 \;\Longrightarrow\; \operatorname{sgn} W_{JH} < 0 \text{ for every admissible weight.}$$

The sign is unforced **only** if α and β may differ in sign — i.e. only if persistence is weighted oppositely at the hub and on the rim. §14B declares $\alpha = \beta = 1$. So the residual J–H freedom reduces to two named questions about the *physical* $\mathcal{D}$_H, and they are exactly what HM10.3 must return for the H sector:

1. **Is $\mathcal{D}$_H diagonal-supported?** — carried by self-transitions rather than by transport (this is Lemma I.5 seen from the score side).
2. **Is it uniformly signed across states?** — α and β of the same sign, not opposite.

If both hold, the J–H sign is a genuine consequence of support and positivity rather than a convention. *(The corresponding question for W_TB is not settled here: the doublet freedom is four-dimensional and the analogous finite calculation has not been done — so the demotion of the T–B sign stands until it is.)*

Lemma 14C.5 — Kernel-family robustness of the structural invariants

The invariants survive not only reweighting but a change of kernel family. **Convention.** The score generators are held at their wheel supports while the heat kernel has support on *all* pairs; the comparison therefore evaluates the same declared generator family against a different kernel. Replacing T_0 by the Doob-normalised operational heat kernel $\exp(-a\_t(I - \mathcal{S}\_n))$ of §56A.2 and recomputing: **W_BH = W_CH = W_RH = 0 to machine precision, sgn W_JH < 0 and sgn W_TB > 0, at every a_t tested (½, 1, 2)**, while W_CC moves from $12/31 \approx 0.387$ to 0.140, 0.207, 0.245 respectively. **Verified numerically.** The **inter-irrep zeros are convention-independent** — they hold for either kernel and any representatives — while the **magnitudes are not**, and the within-irrep signs are properties of the declared family in both cases (§14C.3). This is the sharpest available statement of the partition, and it bears directly on Gate G11 below.

Theorem 14C.6 — Equivariance alone cannot fix the generators

Decompose the 24-dimensional admissible tangent space (§14G) under the C_6 action: **every sector m = 0, …, 5 carries exactly four real dimensions** (verified), and under the full D_6 action the m = 0 sector splits as three reflection-even plus one reflection-odd. Consequently:

- **C** is the unique reflection-odd m = 0 direction — **multiplicity one**, so symmetry genuinely forces it, exactly as §14B claims;
- **J and H** occupy two directions of a *three*-dimensional reflection-even m = 0 space — symmetry does not distinguish them from any other basis of that space;
- **T and B** are two of four copies of the $m = 1 \oplus m = 5$ doublet — symmetry does not say which two.

Therefore equivariance alone can never discharge §14E(3). Except for the closure direction, no defect generator is fixed by symmetry: a multiplicity-3 or multiplicity-4 freedom survives in every other sector, and this is the *generic* case rather than an edge case. §14F.2's remark that "residual freedom survives inside a repeated irreducible representation" is thus quantitative and pervasive. §14E(3) must accordingly be read as requiring **equivariance *together with* support and intertwining data**, and the execution must state what supplies the latter — on present evidence, only the actual J_A can (F26). This theorem and Corollary 14C.4b are complementary and should be read together: **symmetry alone does not fix the generators; symmetry *plus support typing* does fix the J–H sign.** The open question is therefore not whether the sign is conventional, but what support type the physical $\mathcal{D}$_H carries — a sharper and answerable form.

14D. Canonical-Wheel Relaxation Spectrum

Because T_0 contains negative and zero eigenvalues, the real logarithm $-\log T_0$ is not a positive history Hamiltonian on the entire seven-state space. Let $T_0^\sharp$ denote the adjoint in $L^2(\pi)$. Reversibility gives $T_0^\sharp = T_0$, so the positive two-step relaxation operator is

$$T_E = T_0^\sharp T_0 = T_0^2,$$

with spectrum

$$\mathrm{spec}(T_E) = \{ 1, ¼, ¼, 1/16, 9/784, 0, 0 \}.$$

On $\mathscr{H}$_supp = ker(T_E)^⊥, the supported Euclidean generator is

$$H_E = -(1/2a_t) \log T_E,$$

with finite relaxation energies 0, (1/2a_t) log 4 (multiplicity two), (1/2a_t) log 16, and (1/2a_t) log(784/9). The two kernel modes of T_E are erased in one canonical refinement step and carry infinite Euclidean cost in the logarithmic description. This is an exact canonical-wheel relaxation spectrum; it is **not** yet the physical oscillator spectrum {ω_(n,a)} of the complete unresolved-record bath, which requires the full typed bath generator and its embedding in the multi-cell history operator (§7, §19, history-content gate).

14E. The Defect-Lift Obligation — What a Proof Must Exhibit

The canonical-wheel execution stands or falls on one bridge: that the six scalar wheel directions are the faithful image of the physical MASD defects (HM10). This section makes the obligation checkable rather than asserted, so that a referee can see exactly what is owed and a future companion knows exactly what to prove. It is the metric-side analogue of Appendix F for branch selection.

A proof discharging HM10 must exhibit, for the physical multi-cell typed carrier (record, gauge, completion and CAR/Fock fermion sectors included):

1. **The compression map.** An explicit linear map $\pi\_lift : \mathcal{R}\_prim^{(n)} \to T\_(T_0^{full})$ from the full vector-valued defect carrier to the tangent space of the full history operator, together with its restriction to the wheel sector. On the bare wheel this is the identity on the six scalars; the obligation is to construct it on the *full* carrier.
2. **Sector match and rank (HM10.1–2).** A proof that π_lift annihilates no defect label and has image dimension equal to the defect count, so no two physical defects collapse to one history direction.
3. **Equivariance (HM10.3).** A proof that π_lift intertwines the physical stabiliser representation with the wheel/internal symmetry action, which is what forces the specific score generators F_A of §14B rather than arbitrary scalars — i.e. the choice of F_C, F_J, …, F_H is *derived* from the representation theory, not posited.
4. **Support typing (HM10.4).** Inter-irrep vanishing and rank preservation follow automatically from equivariance and positivity, so a criterion demanding only those could not fail. The discriminating requirement is on **support**: a proof that the physical defect differentials carry the support types the benchmark assumes — in particular that $\mathcal{D}$_H is *diagonal-supported* (carried by self-transitions rather than transport) and *uniformly signed across states*, α and β of one sign. Under those two conditions the J–H sign is forced for every admissible weight (Corollary 14C.4b), so the criterion has content and can fail. Remaining within-irrep blocks are outputs of the evaluated J_A, not targets inherited from the benchmark (Thm 14C.6).
5. **Magnitude derivation (separate, harder).** The physical *magnitudes* are then whatever the full typed operator returns; they are computed, not inherited from the wheel. A lift that reproduces the structural pattern but returns different rational values is a **success**, not a failure — the wheel never claimed the magnitudes (§14C.3).
6. **Cross-block completion audit.** Where the wheel returns an exact zero (W_BH = W_CH = W_RH = 0), the lift must either preserve the zero or *derive* a nonzero value from a specific multi-cell / internal-representation / terminal-record / completion mechanism, with that mechanism named. A nonzero value that appears without such a mechanism is an inserted coefficient (Firewall H9, F24).

Status. None of (1)–(6) is discharged in full here. But the obligation is no longer only a checklist: **Appendix I works the canonical closure *score* through an explicit two-cell history complex and verifies all four *graph-lift analogues* of the HM10 criteria on it exactly** — the physical sector-match and equivariance obligations for the actual $\mathbb{Z}_7$-valued closure differential remain open —, demonstrating that the list is achievable for at least one defect and extracting one new necessary condition (the support-preservation Lemma I.5). Item (3), global equivariance, remains the deepest open point: it is what would turn the score generators from a well-motivated symmetry-adapted *choice* into a *derivation*, and until it is proved for all sectors, §14B's tangent family is a canonical and natural representative set, not a proven-unique one. The honest reading of the whole canonical-wheel subtheorem is therefore: *the machinery is executed exactly, its choice-independent output is a genuine candidate for the physical pattern, its numbers and within-irrep blocks are properties of one specific model and generator family, and* *the lift to physics is verified only at the graph-internal level for one score and itemised-and-open for the rest.*

14F. The Polar-Isometric Defect Lift and the Restored D_6 Doublets

§14E itemises what a defect lift must exhibit. This section supplies a *canonical formula* for it, and works in the complete symmetry-adapted family of which §14B's six scalars are a reduced slice.

14F.1 The polar lift

At a frozen defect-zero background $z_\star$, let $J_A = D\mathcal{D}_A|(z\star) : T_{(z_\star)} \mathcal{C}_n \to \boldsymbol{\mathcal{R}}_A$ be the differential of each typed MASD defect, and take its polar decomposition on $(\ker J_A)^\perp$:

$$J_A = V_A\,|J_A|.$$

The partial isometry V_A is the canonical **orientation frame** for embedding the physical tangent directions into the typed residual carrier. It removes any arbitrary rescaling — but it does so by *discarding* $|J_A|$, which carries the singular values and hence the physical response magnitudes. The two objects must therefore be used in different places:

$$\Pi_A^{frame} = L_A \otimes V_A \quad \text{(benchmark orientation frame; equivariance, rank, typing)}$$
$$\Pi_A^{wheel} = L_A \otimes J_A = L_A \otimes V_A|J_A| \quad \text{(magnitude-bearing **wheel/internal benchmark**)}$$

Using Π_A^{frame} where Π_A^{wheel} is meant would silently normalise away the response magnitudes. **Neither object is the physical lift.** L_A is a *wheel support representative*, so both carry the diagnostic wheel factorisation; the authoritative physical lift is defined only through J_A, the tangent-to-history embedding ι and the transported increments Ξ_A (§56A.4–5), and never through L_A. With wheel support operators L_A (oriented rim circulation for C; symmetric hub–rim conductance for J; radial first-harmonic source–sink doublet for T; alternating triangle-composition mode for R; oriented polar rim-twist doublet for B; diagonal self-persistence for H), the local typed lift is $\Pi_A = L_A \otimes V_A$, and on a multi-cell substrate

$$\Pi_A^{(n)} = \sum_{(x \in X_n)} \Pi_{(A,x)} \otimes |x\rangle\langle x| + \sum_{(x\sim y)} \Pi_{(A,xy)} \otimes |y\rangle\langle x|,$$

the first term carrying local closure, composition and persistence, the second transported current, divergence and polar bond compatibility.

Theorem 14F.2 — Automatic equivariance of the polar lift

Suppose the physical symmetry acts unitarily with $J_A R_A(g) = S_A(g) J_A$. Then $R_A(g)^\dagger(J_A^\dagger J_A)R_A(g) = J_A^\dagger J_A$, so $|J_A|$ commutes with $R_A(g)$, and polar decomposition gives

```
⟦ V_A R_A(g) = S_A(g) V_A. ⟧
```

The polar factor therefore *propagates* equivariance: given that J_A intertwines, so does V_A. Two qualifications are essential. First, $J_A R_A(g) = S_A(g) J_A$ is a **hypothesis, not a result** — the theorem does not derive that the physical defect differential carries the required symmetry, only that polar decomposition preserves it once proved. Second, the conclusion concerns the orientation frame V_A, not the magnitude-bearing lift. So the correct status is:

```
⟦ polar orientation-frame and conditional equivariance lemma: DERIVED ⟧
⟦ complete magnitude-bearing six-defect lift: OPEN until every J_A is
evaluated and typed ⟧
```

Residual freedom survives inside a repeated irreducible representation: where a multiplicity exceeds one, a unitary remains and must be fixed by additional support/intertwining data (F26 narrows to this case).

Theorem 14F.3 — Rank additivity (benchmark factorisation)

Because $\mathcal{R}_{prim} = \bigoplus_A \mathcal{R}_A$ is a typed direct sum and $L_A \neq 0$ on each declared support, $\mathrm{rank}(\bigoplus_A \Pi_A) = \sum_A \mathrm{rank}\, J_A$ after classified nulls are removed. Two distinct defect labels cannot collapse merely because their wheel representations coincide: J and H, for instance, are separated by radial-edge versus diagonal support — which is Lemma I.5 seen from the lift side. Finally, in the **benchmark factorisation** — using the magnitude-bearing object $\Pi_A^{wheel} = L_A \otimes J_A$, **not** the frame Π_A^{frame}, which would normalise the response magnitudes away —

```
⟦ W_AB^wheel = (Π_A^wheel)† W_edge^(n) Π_B^wheel, ⟧
```

so every cross block is a calculated overlap (Corollary 16.2 at the level of the lift). This is a benchmark-factorised formula, **not** the authoritative physical definition, which is the path-law construction of §56A.5.

14F.4 The complete doublet calculation

The transport-conservation and polar-bond directions are two-dimensional D_6 representations. §14B works in the reduced slice spanned by one axis of each; the complete symmetry-adapted family retains both axes:

$$[\![L_T = (L_T^{\cos}, L_T^{\sin}), \quad L_B = (L_B^{\cos}, L_B^{\sin}),]\!]$$

giving eight symmetry-adapted directions (C, J, T_cos, T_sin, R, B_cos, B_sin, H) rather than six. The full 8×8 canonical-wheel Fisher matrix has been evaluated exactly. **What survives and what changes:**

Choice-independent, and unaffected by restoration:

- full rank — now **rank 8**, no information-null direction;
- the exact inter-irrep zeros **$W_{BH} = W_{CH} = W_{RH} = 0$**, for *both* the cosine and sine partners.

Survives restoration, within the declared generator family (not choice-independent):

- $W_{CC} = W_{RR} = 12/31$, $W_{JJ} = W_{HH} = 75/434$, $W_{JH} = -33/434 < 0$ — every six-slice diagonal and the J–H cross block are unchanged. (The *value* −33/434 is family-relative; the *sign* is recoverable independently under support typing, Corollary 14C.4b.)
- the doublet diagonals are equal, $W_{(T \cos)} = W_{(T \sin)} = W_{(B \cos)} = W_{(B \sin)} = 21/124$, as D_6 requires, and the cross-doublet entry $W_{(T \cos, T \sin)} = 0$.

Changes when the omitted doublet directions are restored:

- a **new nonzero cross-block appears with no six-slice counterpart**,

  $$[\![W_{(T \cos, B \sin)} = -W_{(T \sin, B \cos)} = 3\sqrt{3}/248,]\!]$$

  antisymmetric across the doublet pairing, alongside the symmetric $W_{(T \cos, B \cos)} = W_{(T \sin, B \sin)} = 3/248$;

- the T/B sector eigenvalues are **replaced**: the slice's {39/248, 45/248} become {9/62, 6/31}, each doubly degenerate, so

  $$[\![\mathrm{spec}(W_8) = \{ 3/31, 9/62, 9/62, 6/31, 6/31, 54/217, 12/31, 12/31 \}.]\!]$$

This is a clean test of the partition: restoring the omitted half of two irreducible representations **preserved the choice-independent content — full rank and the inter-irrep zeros — and altered magnitudes**, exactly as F25 requires. Note that the unchanged J–H entry and its sign survive *restoration within the same frozen generator family*; they are not thereby choice-independent invariants (§14C.3, Thm 14C.6). Downstream consumers must use the eight-direction object; the six-scalar W_wheel remains valid only as the executed benchmark of §14C, and its T/B eigenvalues may not be quoted at all.

Bookkeeping. The HFB-1 handoff (§52) and Appendix D continue to specify a 6×6 ledger. That is to be read as a **6×6 matrix of typed operator blocks**, in which the T and B entries carry two-dimensional irreducible carriers — not as a literal six-scalar matrix. The eight-direction wheel calculation is the correct scalar realisation of that block structure.

Status. ⟦ Wheel/internal benchmark lift formula: DERIVED (polar-isometric, equivariance propagated, rank-additive). ⟧ ⟦ Authoritative physical lift: OPEN, pending evaluation and typing of the full J_A, the embedding ι and the transported increments Ξ_A (§56A.4–5, F23). ⟧

14G. Tangent-Space Exhaustion — the Converse Audit

§17 imposes a hard obligation to classify every *null* direction of W_prim. The converse question is equally sharp and is not asked anywhere else: **are there admissible history deformations carrying no defect label at all?** On the canonical wheel this is a finite calculation, and it is performed here.

The admissible tangent space at T_0 consists of matrices supported on the allowed directed edges with vanishing row sums. There are 31 allowed directed edges (7 from the hub, 4 from each of the six rim states) and 7 row-sum constraints, so

```
⟦ dim 𝒯_(T₀)^adm = 31 − 7 = 24. ⟧
```

The eight symmetry-adapted defect directions of §14F.4 are linearly independent and lie inside this space (**verified: rank 8, contained**). Therefore

```
⟦ unlabelled admissible directions = 24 − 8 = 16. ⟧
```

The decomposition, computed. Per C_6 sector the tangent space carries four real dimensions: the hub row contributes 7 edge values minus 1 constraint = 6 total (one per sector), and the rim rows contribute four copies of the regular representation minus 6 constraints = three per sector, giving 1 + 3 = 4 per sector and 24 in all. Locating the eight defect directions within it (**verified**):

Sector	Real dim	Defect directions	Unlabelled
m = 0, reflection-even	3	J, H	1
m = 0, reflection-odd	1	C	0
m = 1 ⊕ m = 5 (doublet, ×4)	8	T_cos, T_sin, B_cos, B_sin	4
m = 2 ⊕ m = 4 (doublet, ×4)	**8**	**— none —**	**8**
m = 3	4	R	3
Total	**24**	**8**	**16**

Two things fall out. First, **the entire m = 2 ⊕ m = 4 sector — eight real dimensions, a third of the tangent space — carries no defect label whatsoever.** That is the single largest unlabelled block and the natural first target for the HM4 decoupling certificate. Second, a coincidence worth *checking* rather than asserting: m = 2 and m = 4 are also the modes carrying T_0's two zero eigenvalues (§14A). Those are functions on states, not tangent directions on edges, so the correspondence is not automatic — but if the same sector is both erased dynamically in one refinement step and defect-unlabelled, that is either a clean explanation of why HM4 holds there or a coincidence of labels. Either answer is a result, and neither is claimed here.

Why this matters. Premise HM4 asserts that the leading coupling of the resolved variables to the bath *factors through the six-defect tangent map*. Sixteen admissible deformations carrying no defect label means one of two things, and the paper does not get to leave the disjunction unstated: either (a) HM4 is doing very heavy lifting — the unlabelled directions exist but are dynamically decoupled from the bath — or (b) the defect basis is **incomplete**, and W_prim as constructed prices only a subspace of the admissible deformations. A full audit must decompose the 24-dimensional space irreducibly under the wheel symmetry and exhibit, sector by sector, which irreps carry a defect representative and which do not; a sector containing admissible deformations but no defect direction is precisely where an unpriced physical mode would hide.

Obligation. The physical execution must return the irrep-by-irrep table of $\mathcal{T}$^adm together with the defect image, and must certify for every unlabelled direction that it does not couple to the unresolved bath at the declared order. **Status: OPEN (F36).** This is the metric-side converse of the null-classification obligation, and like it, counting is not classifying.

---

15. Gaussian Bath Elimination

At quadratic order around a defect-zero branch, write the bath-resolved history Hessian in frequency space as

```
Ŵ_n(ω) = [ [ W_(core,n)(ω), C_n† ], [ C_n, A_(Q,n)(ω) ] ],  A_(Q,n)(ω) = ω² I
+ Ω²_(Q,n).
```

Integrating out q gives the exact Schur complement

```
⟦ W_(hist,n)(ω) = W_(core,n)(ω) − C_n† A_(Q,n)(ω)⁻¹ C_n. ⟧
```

The full block is positive semidefinite iff $A_{(Q,n)}(\omega) > 0$ and $W_{(hist,n)}(\omega) \geq 0$. For $\omega \neq 0$ the shift $\omega^2 I$ makes $A_{(Q,n)}(\omega)$ strictly positive and the plain inverse is legitimate; **the range condition of Theorem 16.1b bites precisely in the $\omega \to 0$ limit**, where $A_{(Q,n)}(0) = \Omega^2_{(Q,n)}$ may be singular and the inverse must be read as a pseudo-inverse. This is the same relaxation mechanism as MASD-1's auxiliary Schur complement, now at the history-law level. In bath eigenmodes,

```
(W_(hist,n))_AB(ω) = (W_(core,n))_AB(ω) − ∑_a c_(n,aA) c_(n,aB) / (ω² +
ω²_(n,a)).
```

The cross-defect blocks are therefore *generically generated* whenever one unresolved mode couples to more than one primitive defect — this is the microscopic origin of MASD-1's nonfactorising W_BH, W_CH, W_RH, and of the nonfactorising response the gauge and QCD gates required. A shared bath mode makes a definite contribution to a cross block, but the *total* block may still vanish through symmetry, opposing modes, or cancellation against W_core — which is exactly what the wheel's exact zeros W_BH = W_CH = W_RH = 0 (§14C.1) illustrate. The determinant term changes the branch weight but not the quadratic resolved Hessian (Appendix A).

16. The History-Metric Theorem

The metric result has two distinct grades, and they must not be run together into one universal identity: the **general** information-geometric statement, which is the definition of W_prim, and its **Gaussian-sector specialisation** to a Schur complement, which is a computational form valid only in the declared quadratic-bath regime.

Theorem 16.1a — History-Derived Operational Metric (general)

Assume: (1) the bath-resolved finite-history measure is absolutely continuous w.r.t. the frozen operational reference; (2) the defect deformation is differentiable in quadratic mean; (3) the long-history information rate exists. Then the MASD primitive metric is the defect curvature of the rate,

$$W_{prim}^{(n)} = D^2_{\mathbf{D}}\, \mathcal{J}_n(0),$$

positive semidefinite, containing every diagonal and cross block W_AB^(n), A, B ∈ {C, J, T, R, B, H}. **The identity is coordinate-dependent**: **D** must be the physically normalised defect coordinates, and which normalisation and tilt family those are is exactly what HM12 declares undischarged — so 16.1a fixes W_prim only up to the congruence $W \mapsto J^T W J$ induced by that choice. **Proof.** The KL rate is nonnegative and minimised at the reference branch, so its first derivative vanishes and its Hessian is PSD; the cross blocks are its mixed second derivatives. ■ This holds regardless of whether the bath elimination is Gaussian.

Theorem 16.1b — Gaussian-Sector Schur Form

Assume additionally (4) the full bath-resolved quadratic block is positive on the physical quotient; (5) the bath couples to the defects at most bilinearly (the declared Gaussian/quadratic-bath sector); and (6) the **range condition** ran C_n ⊆ ran Ω²_(Q,n) holds. Condition (6) is not decorative: when Ω²_(Q,n) arises as the compression Q_nℍ^hist Q_n (§56A.3) it need not be strictly positive, the inverse must be read as a pseudo-inverse $(\Omega^2_{(Q,n)})^+$, and the pseudo-inverse Schur complement is the exact reduction **only** under this range condition. Without it the Gaussian elimination is not exact and 16.1b does not apply. Then, and only then, the general curvature of 16.1a equals the Schur complement of §15 at zero frequency:

$$W_{prim}^{(n)} = W_{(hist,n)}(0) = W_{(core,n)}(0) - C_n^\dagger\, (\Omega^2_{(Q,n)})^+\, C_n.$$

Proof. Quadratic Gaussian elimination gives the resolved Hessian as the §15 Schur complement; evaluating the static one-step handoff at $\omega = 0$ returns the metric consumed by the local MASD action. No independent sector coefficient survives. ■

Scope caveat. Outside the Gaussian sector, integrating out the bath generates higher cumulants of the defect coupling, and the rate Hessian is *still* given by 16.1a but is **no longer** the simple Schur complement of 16.1b: the non-Gaussian corrections add to $D^2_{\mathbf{D}} \mathcal{J}_n(0)$ terms not captured by $C_n^\dagger(\Omega^2_{(Q,n)})^+ C_n$ alone. The physical execution must state which regime it is in; a Schur-form return is admissible only with a certificate that the bath–defect coupling is bilinear to the declared tolerance. This is also why the history-content gate is on the whole generator $\hat{\mathbb{H}}_n^{hist}$ and not on (Ω^2_Q, C_n): W_{core} enters 16.1b as an input, so a derived bath block with an underived W_{core} does not determine W_{prim} (§19).

Corollary 16.2 — Cross-Defect Non-Insertion (scope-limited).

Once the defect embedding, its normalisation and the tilt family are fixed (HM12), a nonzero W_{BH}, W_{CH}, W_{RH} or any other cross block is a history covariance rather than an inserted interaction: the cross blocks are *calculated overlaps* because they are second derivatives of one probability law. This discharges MASD-1's Theorem 9A at the level of provenance, and no further. It does **not** assert that the embedding is unique — pre-registering the exponential family prevents later fitting, but does not prove that family is derived (F34).

The Fisher–Schur bridge is a separate obligation, and it has a concrete obstruction. Theorem 16.1a defines W_{prim} as the Hessian of the path-space KL rate; Appendix A obtains a Schur complement from Gaussian elimination of the quadratic influence action. These coincide only if **Doob normalisation and Gaussian bath elimination commute** — and they need not. The Doob transform on the enlarged carrier uses the Perron function $r(x, q)$; marginalising the resulting kernel over q' equals the Doob transform of the marginalised transfer **only if r factorises as $r_X(x) \cdot r_Q(q)$**, which fails precisely when the defect–bath coupling $C_n \neq 0$ — that is, in every case of interest. So the obligation is sharp: **prove that Doob normalisation and Gaussian elimination commute, or bound the discrepancy.** HM12 additionally makes the coordinate-normalisation half non-cosmetic. **Status: obligation, not discharged here (F39).**

Corollary 16.3 — Monotonicity under Bath Enlargement.

Adding unresolved bath modes can only decrease the resolved quadratic stiffness in the Loewner order, provided the complete enlarged block remains positive. Screening is therefore generic, which is the structural explanation of MASD-1's CFB-2 finding that a $\sqrt{7002}$ Abelian amplitude was screened to $Z_q \approx 2$: the auxiliary/bath relaxation *must* reduce stiffness, so a large local amplitude does not survive as a large physical one. The electroweak deficit being "auxiliary-structural" is Corollary 16.3 in action.

Corollary 16.4 — First Executed History-Metric Branch.

For the canonical $K = 7$ kernel T_0 and the six scalar defect tangents of §§14B–14C,

⟦ $D^2\_\mathbf{D}\ \mathcal{I}\_\text{wheel}(0) = W\_\text{wheel},$ ⟧

where W_wheel is the exact positive matrix of §14C. Thus the History-Metric Theorem (16.1a) is *executed* on a nontrivial finite history law. **The headline is executability, not a physical pattern**: one transfer law returns every entry of the matrix from one measure with nothing inserted. Specifically: one transfer law can return all six diagonal coefficients simultaneously; cross-defect coefficients can arise as Fisher covariances rather than inserted terms; the resulting metric is full rank, which is choice-independent because the Fisher form is nondegenerate on the tangent space; and representation orthogonality forces exact zeros in the *inter-irrep* cross blocks (W_BH = W_CH = W_RH = 0), which is likewise choice-independent. The within-irrep entries — W_JH and W_TB, their signs and even their non-vanishing — are properties of the frozen generator family, with one exception: the **J–H sign** is recovered for every admissible weight once $\mathcal{D}$_H is typed as diagonal-supported with uniform sign (Corollary 14C.4b). No such exception holds for W_TB (§14C.3). It does **not** establish W_wheel = W_prim^physical, and it does **not** claim the rational magnitudes are physical: those are model-specific to the canonical hub↔rim weight (§14C.3) and would change under any admissible reweighting. Promotion of even the structural pattern to physics requires the Defect-Lift Obligation (§14B.1, §14E, HM10); promotion of magnitudes requires the full typed operator to be evaluated, which the wheel does not do. It is the first executed instance of Theorem 16.1a — landing in the Gaussian-independent general form, since the wheel Fisher matrix is computed directly from the exact kernel rather than through a Schur complement — and it is a benchmark of the machinery, not a physical return.

17. Gauge, Coboundary and Terminal-Null Directions — and the Null-Classification Obligation

Let v generate a deformation that changes only the presentation of a history, $K^\varepsilon = G_\varepsilon^{-1} K G_\varepsilon$. The physical path probabilities are unchanged up to relabelling, so the score vanishes, $S_v = 0$, hence $W_{prim}\, v = 0$. The information null space contains: (1) exact gauge transformations; (2) exact record coboundaries; (3) coordinate reparametrisations of the same history; (4) terminally invisible residue; (5) regulator boundary nulls certified by the blocking map.

The converse is an audit, not a theorem, and this paper treats it as a hard gate. A Fisher/history metric is positive semidefinite, and a PSD matrix may carry *additional statistically invisible directions* — defect variations that leave the one-step history law unchanged for reasons beyond gauge or frame. Any such direction lies in ker W_prim without being a gauge mode. Therefore:

⟦ physical rank of W_prim must be CALCULATED, not set equal to (dim − gauge count). ⟧

Every numerical null mode must be classified into categories (1)–(5) or retained as an unexplained physical degeneracy. Counting nulls is not classifying them. This is the history-measure face of MASD-1's 39 unclassified CFB-1 null modes, and it is coupled to F28: an unexplained null is either a genuine redundancy to be quotiented or a signal that the reduced response is not rigid along the defect-zero locus — in which case the boundary certificate is not a function of the frozen inputs alone. Define

```
N_null = [ n_1 ⋯ n_r ],  P_info = I − N_null N_null⁺,  W_prim^phys = P_info
W_prim P_info,
```

but P_info may enter a physical certificate only after every column of N_null is analytically identified (Gate G5). An algorithmic quotient against a merely numerical null space is **not** admissible for a physical return.

18. Absolute Normalisation and Dimensional Ledger

The history rate is measured per common transfer step a_t, so the quadratic action unit is fixed by 𝒥_n / a_t, not by independent colour, weak, Abelian, scalar or fermion conventions. The primitive dimensional ledger is (ℓ_n, a_t, ρ_c, λ_sub, κ_e, Λ⋆). HMGBC-1 fixes the action conversion and contributes a_t (frozen refinement step), κ_e (current-response compression of W_prim) and the history contribution to the completion residue entering Λ⋆. No dimensionful member may be fixed from a Standard Model observable (this closes MASD-1's F20 conditionally on the same bath evaluation as F19).

19. Relation to MASD-1 W_prim, and Honest Relocation of the Freedom

MASD-1 treats Γ_defect = ½ ⟨**D**, W_prim **D**⟩ as the physical local quadratic rate. HMGBC-1 supplies the interpretation and calculation:

```
⟦ W_prim = W_hist(0). ⟧
```

The full history theory additionally returns W_hist(ω), whose frequency dependence records memory and nonlocal response; MASD-1 consumes the static local boundary, while momentum-dependent sector calculations may retain the full kernel. A physical execution must publish both the static 6×6 defect-block matrix and the frequency-dependent residual needed for scalar and threshold audits.

Honest reduction and remaining relocation. HMGBC-1 removes one layer of arbitrariness at the canonical-wheel level. The one-step kernel T_0, its stationary measure and its six-dimensional nonpersistent complement are explicitly constructed and evaluated, and W_wheel is computed exactly from a declared tangent family (§§14A–14D, Corollary 16.4). The kernel and its spectral data are no longer free; the *tangent family* remains a declared choice (F26). The remaining freedom is not the entire generator being unknown; it is concentrated in four narrower identifications:

⟦ full MASD defect differentials ⟶ wheel and internal history tangents (Defect-Lift Obligation); ⟧
⟦ canonical local wheel ⟶ physical multi-cell history operator; ⟧
⟦ full stationary flow ⟶ terminal-record conductance (projection gate, §31.1); ⟧
⟦ admissible branch catalogue ⟶ components of one master operator (§43A). ⟧

The canonical-wheel benchmark is executed exactly. Its promotion to a history-content result remains conditional on the physical defect embedding (HM10.3, F26), the tilt family (HM12, F34) and the G11a generator attribution — so the history-content *subgate* has not passed; an executed benchmark has been produced. The full physical history-content gate remains open because the scalar wheel tangents have not yet been proved to be the complete images of the vector-valued MASD defects, and the spatial, record, gauge, completion and fermionic extensions of the kernel remain unevaluated. Structurally, the point still holds at full strength: HMGBC-1 does not remove the constitutive freedom, it moves it into **one object, the complete master generator $\hat{\mathbb{H}}_n^{hist}$**, together with the single-operator branch structure (HM9). It is essential that the object be the *whole generator* and not merely the bath pair $(\Omega^2_{(Q,n)}, C_n)$: Theorem 16.1b gives $W_{prim} = W_{core}(0) - C_n^\dagger (\Omega^2_{(Q,n)})^+ C_n$, in which W_{core} (via V_{loc}) is an independent input, so a derived bath pair with a *chosen* W_{core} would relocate the old freedom straight into the local history cost rather than remove it. The chain is

$$⟦\ \hat{\mathbb{H}}_n^{hist} \longrightarrow (W_{core}, \Omega^2_{(Q,n)}, C_n) \longrightarrow W_{prim},\ ⟧$$

and the gate is on the first arrow. The move is a genuine reduction **if and only if** the entire generator and the single-operator structure are forced by prior VERSF results — the memory-kernel programme, the operational-geometry programme, and the substrate axioms — rather than chosen. This paper builds every consumer of $\hat{\mathbb{H}}_n^{hist}$ and proves the consumer theorems (H1); it does not evaluate $\hat{\mathbb{H}}_n^{hist}$ (H3) or exhibit the single operator on the real regulator. **Three** open gates are named precisely:

- **History-content gate (F1, F6–F8, F19 below).** Is the *complete* generator $\hat{\mathbb{H}}_n^{hist}$ — W_{core} as well as $\{\omega^2_{(n,a)}\}$ and $\{c_{(n,aA)}\}$ — returned by the substrate programmes, and is W_{prim} independent of any normalisation not fixed upstream? If any block of $\hat{\mathbb{H}}_n^{hist}$ is chosen, the relocation is cosmetic.
- **Projection gate (§31.1, F10 below).** Is the terminal-record conductance C_Ω the pushforward of the full-configuration flow (representative-independent, cylindrically compatible, strongly lumpable), so that the *generation geometry* is a history return and not a projection chosen at the map onto Ω? A history-derived metric with an inserted projection would leave Firewall H4 violated one level down.
- **Branch-comparison gate (§43A, F16–F17 below).** Are the admissible branches components of one master history operator on one reference measure (comparability), *and* does dominance imply selection (global dominance, §43A.5)? If either fails, MASD-1's F29 is not resolved even with a per-branch Perron gap.

Until all three gates pass on the frozen K = 7 regulator, HMGBC-1 has relocated the Standard Model's freedom one level deeper and made it concrete and finite — a real advance in tractability — but it has not eliminated it.

Part IV — History Positivity and Refinement

20. HP1 — Bath-Resolved Cone Positivity

Let $\mathcal{C}_n = \{ f \in \mathscr{H}_n : f \geq 0 \}$.

HP1.

$$\mathbb{T}_n \mathcal{C}_n \subseteq \mathcal{C}_n,$$

and on each connected physical branch, $f \geq 0$, $f \neq 0 \Longrightarrow \mathbb{T}_n f > 0$.

Theorem 20.1.

If the kinetic generator is Dirichlet, the potential is real and bounded below, and the physical transition graph is irreducible, HP1 holds.

Proof. The heat semigroup is positivity preserving; the potential factor is positive multiplication; Trotter composition preserves positivity; irreducibility plus nondegenerate bath diffusion gives positivity improvement. ■

This theorem separates three ideas that must never be conflated: Hilbert-space positivity of a Hessian; positivity of a measure; order-cone positivity of a transfer kernel. HMGBC requires all three, at their proper places. The advance over MASD-1's §10.1.1 is exactly here: there, HP1 was asserted of a posited semigroup; here it is a theorem about an explicitly positive transfer — Markov after Doob normalisation — *conditional only on the irreducibility of the physical transition graph* — a property that is itself checkable on the frozen regulator rather than assumed.

20.2 Scope of HP1 — the Fermionic Phase Caveat

The Trotter/Feynman–Kac cone argument of Theorem 20.1 is clean for a **scalar Schrödinger-type transfer** $\mathbb{H} = -\hat{\Delta} + V$ with real, bounded-below scalar V. The full MASD carrier is not scalar: it also carries gauge connections, CAR/Fock fermionic variables and determinant-line phases, any of which can produce a sign or phase problem. The scalar heat-kernel argument therefore establishes HP1 for the **bosonic positive-pushforward transfer** and does *not*, on its own, establish positivity of the full MASD transfer. Extending it requires **one** of two explicit formulations, and the execution must declare which it invokes (Gate G2):

- **(a) Real nonnegative fermionic determinant.** HP1 applies to the bosonic transfer *after* the fermionic determinant has been shown real and nonnegative — the positive-pushforward branch — after which the bosonic cone result carries the full measure; or
- **(b) Full reflection/transfer positivity as a separate gate.** Full reflection/transfer positivity of the fermion-including transfer is adopted as an additional, separately certified physical gate.

This is not a gap in the bosonic theorem; it is the precise point at which HP1 couples to the determinant-line and strong-phase programme (§52A). If the fermionic reduction returns a complex phase, the positive-cone proof does not go through unchanged, and **HP1-failure and $\bar{\theta} \notin \{0, \text{global obstruction}\}$ (F20) become the same physical event**. Gate G2 must state whether it invokes (a) or (b); a bosonic HP1 pass quoted as if it certified the full transfer is inadmissible.

21. HP2 — Bi-Perron Stationarity, and What the Gap Does and Does Not Give

For a positivity-improving compact transfer operator, $\mathbb{T}_n r_n = \lambda_{(0,n)} r_n$ and $\ell_n^\dagger \mathbb{T}_n = \lambda_{(0,n)} \ell_n^\dagger$ with $r_n > 0$, $\ell_n > 0$.

HP2. The leading root is simple and isolated: $\lambda_{(0,n)} > |\lambda_{(1,n)}|$. The normalised reduced transition kernel is the Doob transform

$$K_n(y \mid x) = T_n(x, y)\, r_n(y) \,/\, (\, \lambda_{(0,n)}\, r_n(x)\,),$$

with stationary distribution

$$\llbracket\ \pi_n(x) = \ell_n(x)\, r_n(x) \,/\, \langle \ell_n, r_n \rangle,\ \rrbracket$$

so that $K_n \mathbb{1} = \mathbb{1}$ and $\pi_n K_n = \pi_n$. No detailed-balance assumption is used.

What is automatic, and what is not. The *qualitative* gap is automatic: for a finite positivity-improving matrix — or a compact strongly positive operator — the Perron–Frobenius/Jentzsch theorem already gives a simple dominant eigenvalue with $\lambda_{(0,n)} > |\lambda_{(j,n)}|$, $j \geq 1$. So HP1 positivity improvement delivers strict finite-level dominance without further work. What does **not** follow automatically, and what Gate G3 must return, is everything quantitative:

- a numerical **lower bound** on the gap $\lambda_{(0,n)} - |\lambda_{(1,n)}|$ — supplied on the intertwined canonical family by $\delta_{global} \geq \min\{½, \delta_X\}$ (§47A.1, verified);
- good **conditioning** of the reduced resolvent $(\lambda_0 I - \mathbb{T})^+$, which controls the first-order response of the Perron data to a defect deformation and hence the differentiability of the rate;

- **regulator-uniformity** — a gap bounded below uniformly in n, not merely positive at each level;
- **survival in the continuum limit**.

So the correct statement is: *HP1 gives strict finite-level dominance; G3 must calculate the gap's magnitude and prove regulator-uniform, continuum-stable control.* The warning that matters is preserved: the gap can become arbitrarily small near a bifurcation, and it is exactly there that the Fisher metric is ill-conditioned and the defect-zero locus fails to be isolated. This is the shared root of two MASD-1 falsifiers — the score resolvent diverges as the gap closes, so the metric is well-conditioned precisely when the gap is controlled and ill-conditioned when it is not. A proof of a *uniform* Perron gap therefore services both the metric (F19-analogue) and the zero-locus rigidity (F28-analogue) at once. HP2 existence is thus not conditional; its *quantitative control* is, and that control is the content of Gate G3, returned numerically with declared uncertainty rather than inferred.

Executed on the canonical wheel. For the K = 7 kernel T_0 the gap is exact:

$$\lambda_0 = 1, \quad \max_{\lambda\neq 1}|\lambda| = \tfrac{1}{2}, \quad \Delta_{\text{Perron}}^{\text{wheel}} = \tfrac{1}{2}$$

(Theorem 14A.1). Thus HP2 — both the qualitative Perron structure and a quantitative gap — is **passed on the canonical wheel**. Gate G3 therefore remains open only for the full history operator and for a regulator-uniform lower bound $\inf_n [\lambda_{(0,n)} - |\lambda_{(1,n)}|] > 0$; the wheel supplies the first level of that sequence.

22. HP2R — Reversible Reference Branch

The defect-free pre-commitment reference may satisfy $\ell_n = r_n$ in the operational inner product, giving detailed balance $\pi_n(x) K_n(y \mid x) = \pi_n(y) K_n(x \mid y)$. This is HP2R. It supplies a symmetric reference Dirichlet form and a reflection-positive Euclidean starting point, and it is valuable for that reason. It is **not** imposed on the irreversible committed-history process, whose antisymmetric current is physically meaningful and carries the CP/Fact-Momentum content of §30, §37 (Firewall H3).

23. HP3 — Cylindrical Refinement Compatibility

Let $V_n : \mathscr{H}_n \to \mathscr{H}_{(n+1)}$ be the normalised refinement lift.

HP3.

$$\mathbb{T}_{(n+1)} V_n = V_n \mathbb{T}_n,$$

equivalently, at generator level, $\mathbb{H}_{(n+1)} V_n = V_n \mathbb{H}_n$ on a common core, and at kernel level

$$K_{(n+1)}(x', q_n^{-1}(B)) = K_n(q_n x', B),$$

independent of the fine representative inside each coarse fibre — strong lumpability. This is HMGBC-1's replacement for the minimal-fibre-submersion premise that MASD-1's §10.1.1 had to leave open. Here it is decomposed into checkable block conditions (§§24–27) on inherited maps, rather than a single opaque geometric assumption, and Appendix C reduces it to three commuting-square identities. It remains a premise (HM7) discharged only numerically, but its content is now explicit and per-block auditable.

24. Scalar and Operational Block

The scalar operational quotient uses the normalised pullback $(V_n^X f)(x') = f(q_n x')$, with $(V_n^X)^\dagger V_n^X = I$ from projective measure consistency. The scalar/refinement generator must satisfy $L_{(X,n+1)} V_n^X = V_n^X L_{(X,n)}$. Toy causal-diamond random-walk constructions demonstrate the mechanism but do not substitute for the physical $K = 7$ execution.

25. Cohomological Transport Block

The refinement-cohomology programme supplies a canonical lift $L_n^{(1)} : H^1(\Lambda_n) \to H^1(\Lambda_{(n+1)})$ and coarse pullback $\Gamma_n^* : H^1(\Lambda_{(n+1)}) \to H^1(\Lambda_n)$ with $\Gamma_n^* L_n^{(1)} = 2I$. After multiplicity normalisation $\hat{\Gamma}_n^* = \frac{1}{2} \Gamma_n^*$, so

$$\llbracket \hat{\Gamma}_n^{*} L_n^{(1)} = I. \rrbracket$$

This is the exact HP3 component on the persistent H^1 carrier: it preserves transport classes while scalar gradients are quotiented out.

26. Bath Block

Let $L_n^Q : \mathcal{Q}_n \to \mathcal{Q}_{(n+1)}$ be the bath-mode refinement lift. Bath compatibility requires

$$\Omega^2_{(Q,n+1)} L_n^Q = L_n^Q \Omega^2_{(Q,n)} \quad \text{and} \quad C_{(n+1)} V_n^D = L_n^Q C_n,$$

where V_n^D is the defect-carrier lift induced from the resolved refinement. These state that refinement does not invent a new bath spectrum or a new defect coupling on the inherited coarse subspace — the precise sense in which the bath triple, once fixed at one level, is not silently re-chosen at the next.

27. Flat Holonomy Block

The flat history sector is carried by non-bounding record cycles; its refinement lift must preserve the holonomy class. Two independent certificates are required: **occupancy** (the physical branch has a nonzero flat class) and **pinning** (the surviving class remains in the inherited discrete root structure rather than relaxing to arbitrary U(1) holonomy). Failure of occupancy removes the flat contribution without invalidating the curvature sector; failure of pinning weakens exact discrete identification but does not destroy continuum U(1) transport.

28. Projective History Measure and Continuum Compatibility

Theorem 28.1 — Projective History Extension.

If HP3 holds and the finite-history measures are normalised and cylindrically consistent, then $(q_n)_* \mu_{hist}^{(n+1,N)} = \mu_{hist}^{(n,N)}$ for every finite N, and standard projective extension yields a measure on the inverse-limit history space.

Regulator convergence requires more than projective existence. The execution must publish: a uniform lower spectral gap on the physical transfer; stable null dimension and classification; convergence of $W_{prim}^{(n)}$; convergence of $R_\Omega^{(n)}$; and convergence of the branch free-energy ordering **and its gap** (§46). Projective existence is necessary; it is not the convergence certificate.

Part V — Microscopic Generation Geometry

29. Stationary Edge Flow

Let K_n be the resolved stationary kernel and π_n its stationary measure. Define the stationary directed edge weight

$$m_n(x, y) = \pi_n(x)\, K_n(y \mid x),$$

the amount of history flow carried by the transition $x \to y$.

30. Conductance and Irreversible Current

Decompose $m_n(x, y) = c_n(x, y) + j_n(x, y)$, with

$$\llbracket\, c_n(x, y) = \tfrac{1}{2} [\, m_n(x, y) + m_n(y, x) \,] \,\rrbracket \quad \text{and} \quad \llbracket\, j_n(x, y) = \tfrac{1}{2} [\, m_n(x, y) - m_n(y, x) \,]. \,\rrbracket$$

Then $c_n(x, y) = c_n(y, x) \geq 0$ and $j_n(x, y) = -j_n(y, x)$. The conductance defines the positive generation geometry; the current carries the history arrow, entropy production and Fact-

Momentum-type irreversible flow. Detailed balance is the special case $j_n = 0$ (HP2R). The current's survival on the committed-history branch is what makes a nonzero CP invariant *possible* (§37) — a reversible reference branch alone could not source it.

31. Covariant Record Incidence

Let Ω_n be the universal terminal-record carrier. For each transition $x \to y$ let $\mathcal{U}_{yx}^{\Omega} : \Omega_x \to \Omega_y$ be the inherited polar record transport. Define the covariant incidence operator

$$(B_\Omega \psi)(x, y) = \psi_y - \mathcal{U}_{yx}^{\Omega} \psi_x,$$

and let C_Ω multiply each oriented edge by the symmetric conductance $c_n(x, y)$. The covariant Dirichlet form is

$$\mathcal{E}_\Omega(\psi) = \tfrac{1}{2} \sum_{(x,y)} c_n(x, y) \, \| \psi_y - \mathcal{U}_{yx}^{\Omega} \psi_x \|^2 .$$

Theorem 31.1 — Terminal-Record Pushforward (obligation, the projection gate)

There is a step here that is easy to pass over and is in fact a third load-bearing gate. The stationary flow $m_n(x, y) = \pi_n(x) K_n(y \mid x)$ lives on the **full resolved configuration space** X_n; the completion Laplacian Δ_Ω lives on the **terminal-record carrier** Ω. Introducing $\mathcal{U}^\Omega$, B_Ω, C_Ω does not, by itself, prove that the record geometry is *induced by* the full history process rather than being a fresh chosen projection. The obligation is:

$$[\![\, C_\Omega(a, b) = \sum_{(x : r(x)=a)} \sum_{(y : r(y)=b)} c_n(x, y), \,]\!]$$

the pushforward of the full-configuration conductance under the terminal-record map $r : X_n \to \Omega$, subject to three conditions: (i) **independence from fine representatives** within each terminal class; (ii) **cylindrical compatibility** under refinement; and (iii) **strong lumpability**, so that C_Ω is a function of the terminal class alone. The record transport $\mathcal{U}_{yx}^{\Omega}$ must likewise *descend from the polar refinement transport*, not be assigned.

Status: obligation, arithmetically discharged and provenance derived-upstream on the canonical wheel (§§32A, 44A), holonomy descent open. If (i)–(iii) hold, Δ_Ω is genuinely history-derived and Firewall H4 is honoured all the way down to Ω. On the canonical wheel they *do* hold exactly: the antipodal map is the unique strongly lumpable equal-layer partition (Theorem 32A.1), the pushforward is K_4, and the generation Laplacian and $\mathbb{Z}_7$-split completion spectra are exact. The provenance — why three two-state layers, why antipodal — is inherited from the anomaly-descent/census chain (Theorem 32A.0), conditional on those upstream papers rather than on an HMGBC premise. What remains open is the **holonomy descent** onto the generation triangle (Gate G7d) and the kernel-identification question of Gate G11. The projection gate therefore stands as: arithmetic PASS (G7a), provenance PASS conditional on upstream (G7b), typed-lift form PASS (G7c1/c2), holonomy descent OPEN (G7d).

31.2 Terminal transport as the polar part of the conductance-weighted microscopic transport

Theorem 31.1 requires $\mathcal{U}_\Omega$ to *descend* rather than be assigned. The canonical construction supplies the descent formula. With $c_n(x,y) = ½[\pi_n(x)K_n(y|x) + \pi_n(y)K_n(x|y)]$ and the terminal-record map $r_n : \mathcal{C}_n \to \Omega_n$, the only admissible terminal conductance is the pushforward $C_\Omega(a,b) = \sum_{(x:\, r_n(x)=a)} \sum_{(y:\, r_n(y)=b)} c_n(x,y)$, and for typed polar transport one sets

$$[\![M_{(ba)}^\Omega = \sum_{(x:\, r_n(x)=a)} \sum_{(y:\, r_n(y)=b)} c_n(x,y)\, S_b^\dagger\, \mathcal{U}_{(yx)}^{rec}\, S_a, \quad \mathcal{U}_{(ba)}^\Omega = \mathrm{Pol}(M_{(ba)}^\Omega).]\!]$$

The terminal transport is thus the polar part of the conductance-weighted microscopic transport, not an independent assignment. This is the formula whose absence made Gate G7 open; supplying it converts the remaining obligation from "construct a descent" to "verify this descent on the physical background," which is the narrower statement the gate now carries.

32. Universal Generation Laplacian

The record Laplacian is

$$[\![\Delta_\Omega = B_\Omega^\dagger\, C_\Omega\, B_\Omega,]\!]$$

positive semidefinite. Its kernel consists of covariantly constant record sections; its nonzero spectrum measures the cost of generation-depth variation under the history law. No quark, charged-lepton or neutrino graph is selected — there is one universal object (Firewall H4). This is the derived replacement for MASD-1's CFB-2/CFB-3 hand-chosen generation graphs: those were the honor-system pre-registered objects that §47M.1 flagged; here the graph is the conductance of the one stationary flow.

32A. Conditional Canonical K = 7 Holonomy-Sector Calculation

This section instantiates the universal Laplacian of §32 on the canonical K = 7 wheel and carries the calculation through to the completion-contraction spectra. **Its status is fixed at the outset and governs everything below:**

```
⟦ Exact canonical-wheel calculation, CONDITIONAL on the terminal-record
pushforward (§31.1). ⟧
```

It does **not** close Gate G7. The §31.1 obligation requires representative independence, cylindrical compatibility, strong lumpability *and* descent of $\mathcal{U}_\Omega$ from the polar refinement transport; this section discharges the quotient arithmetic and (via the inherited chain) the layer provenance, but the physical pushforward — in particular the $\mathcal{U}\_\Omega$ descent — remains open. Nothing here may be quoted as a physical generation return.

A typing obligation that must not be elided. The calculation supplies three *conjugate nonzero* $\mathbb{Z}_7$ *sector pairs*, {±1}, {±2}, {±3}. These are **holonomy sectors, not physical generations.** The three physical generations are the three eigenvalues λ_m(k) *within* a given sector k. No result below proves that the sector pairs are the generations, and the inherited corpus does not yet supply that identification; conflating the two is a typing error and fires F31.

§31.1 leaves the terminal-record map r as an obligation, of which the *provenance* — why three two-state layers, and why antipodal — is the constitutive half. That provenance is located upstream, in the inherited anomaly-descent and census results, so it is a theorem conditional on those papers rather than an independent assumption of HMGBC-1. **The relocation is only as strong as the upstream attributions; it does not make the layer structure unconditional, and the upstream theorems carry their own named premises — and it does not by itself discharge the §31.1 pushforward conditions.**

Theorem 32A.0 — Terminal-Layer Provenance (conditional on the QN-3 / SC-1 chain)

Take as inputs, from the inherited corpus: (i) the Substrate Anomaly-Descent result, that the physical support quotient has **rank three**, with each direction *bidirectional* — carrying equal descent and return branches; and (ii) the census result, that each support-return layer is independently terminal and no terminal carrier constitutes more than one layer, so the three support-return directions require exactly **three** complete generation packages and no fourth belongs to the minimal ledger. The canonical K = 7 wheel independently realises the directional boundary catalogue as six cyclic boundary states plus one hub anchoring their common closure condition. Then:

```
⟦ 3 physical directions × 2 orientations = 6 wheel boundary states, ⟧
```

and identifying the two orientations of each physical direction is the antipodal quotient b_i ∼ b_(i+3). **This is the unique D_6-equivariant such quotient**, because the antipodal map is the unique nontrivial involution (rotation by π) of the six-state cycle — the only order-2 element of C_6 — so no other identification of orientations is D_6-compatible. It yields the three antipodal terminal classes A = $\{b_0, b_3\}$, B = $\{b_1, b_4\}$, C = $\{b_2, b_5\}$, and layer-local terminality with layer-exclusive attachment identifies these with the three independent support-return completion ledgers.

Status: DERIVED within the inherited QN-3 / SC-1 premise chain — not a new HMGBC-1 premise. The three-layer count and the antipodal identification are inherited from the rank-three, bidirectional physical-support structure; they are no longer asserted here. The qualification is essential: QN-3 and SC-1 are themselves conditional on their named anomaly-descent premises, so this is a *relocation* of the assumption upstream, not its elimination. HMGBC-1's own contribution is the D_6-uniqueness step (verified above) and the lumpability selection (Theorem

32A.1); the physics inputs (i)–(ii) are taken from the cited papers and are only as secure as those papers. What HMGBC-1 no longer needs is a standalone premise asserting "three two-state layers."

Theorem 32A.1 — Uniqueness of the antipodal terminal-record map

Given the three-layer, layer-equivalent structure of Theorem 32A.0, the only partition of the six D_6-symmetric boundary states into three two-state classes that is *strongly lumpable* with respect to T_0 **and** treats the three layers equivalently is, uniquely up to generation relabelling, the **antipodal** partition

$$A = \{b_0, b_3\}, \quad B = \{b_1, b_4\}, \quad C = \{b_2, b_5\},$$

with the hub as the shared closure anchor H. **Proof (exhaustive enumeration, verified exactly).** There are 15 perfect matchings of the six rim states into three pairs. Evaluating strong lumpability against T_0 for each (with the hub as its own class) returns **four** strongly lumpable matchings, not one:

- $\{b_0b_1, b_2b_5, b_3b_4\}$, $\{b_0b_3, b_1b_2, b_4b_5\}$, $\{b_0b_5, b_1b_4, b_2b_3\}$ — each strongly lumpable, but with **inequivalent** quotient classes: one generation class has self-return probability ½, another ¼, and some class-to-class transitions vanish. These violate the inherited requirement (Theorem 32A.0) that the three pre-hierarchy layers be representation-equivalent.
- $\{b_0b_3, b_1b_4, b_2b_5\}$ (antipodal) — strongly lumpable **and** the unique matching for which all three generation rows are identical, (¼, ¼, ¼, ¼).

(The fully-adjacent matching $\{b_0b_1, b_2b_3, b_4b_5\}$ is *not* even lumpable — one representative reaches a neighbouring class with probability ¼, the other with 0 — so it is excluded at the first step.) Therefore strong lumpability *alone* does not select antipodal pairing; strong lumpability **plus the inherited layer-equivalence** does, and does so uniquely up to relabelling. ■

This is a sharper theorem than "lumpability picks antipodal": it isolates exactly which hypothesis (layer-equivalence, now inherited via Theorem 32A.0 rather than assumed) does the selecting. Together with 32A.0 it closes the **provenance** of the canonical projection within the inherited premise chain: the record classes are the unique lumpable equal-layer partition, and the equal-layer structure is inherited, not posited. What remains genuinely open is narrower and stated in §32A.3.

32A.2 The exact quotient geometry

The induced quotient kernel on (H, A, B, C) is row-stochastic with hub row (1/7, 2/7, 2/7, 2/7) and each generation row (1/4, 1/4, 1/4, 1/4), stationary measure

$$\bar{\pi} = (1/31)(7,\ 8,\ 8,\ 8).$$

Every distinct pair of quotient classes carries the same stationary conductance $c\_uv = \bar{\pi}\_u\ \bar{T}\_uv = 2/31$, so the pushed-forward record geometry is the **equally weighted complete graph K_4**. Eliminating the hub — the shared closure anchor, not a generation layer — by the exact Laplacian Schur complement gives the three-generation graph with effective pair weight $g\_\Omega = 8/93$ and Laplacian

```
⟦ Δ_Ω^(0) = (8/31)(I₃ − ⅓ J₃),  spec = {0, 8/31, 8/31}. ⟧
```

The representative-independence, strong-lumpability, layer-equivalence and conductance-pushforward conditions of §31.1 pass exactly on the canonical wheel. **Cylindrical compatibility and physical polar-transport descent are not established by a one-level matching calculation and remain open** (G7d, G7h). The zero-holonomy branch therefore, *given the inherited three-layer structure (Theorem 32A.0)*, fixes the pairing and returns one singlet and one degenerate doublet: **the wheel does not distinguish all three generation depths on its own; the three-layer count is inherited from the anomaly-descent/census chain, not derived by the wheel.**

32A.3 What is closed and what remains

Closed within the inherited premise chain: the three-layer count and antipodal identification (Theorem 32A.0, conditional on QN-3/SC-1); the pairing is antipodal, unique, strongly lumpable (Theorem 32A.1); the quotient is K_4; the generation Laplacian is $(8/31)(I_3 - \tfrac{1}{3}J_3)$. The layer structure is inherited rather than assumed here.

Gate G7 nevertheless remains OPEN overall, and this section may not be cited as closing it. The correct split is:

```
⟦ G7 canonical wheel quotient arithmetic: PASS ⟧
⟦ G7 physical terminal-record pushforward: OPEN ⟧
```

The second is more than the layer provenance: §31.1 requires representative independence, cylindrical compatibility, strong lumpability, **and** descent of $\mathcal{U}\_\Omega$ from the polar refinement transport onto the record carrier. The layer provenance (32A.0) and the pairing uniqueness (32A.1) address the *partition*; they do not establish the *transport descent*, which is the part that would make the generation geometry a history return rather than an identification. The upstream attributions are also taken on trust from the cited papers and carry their own premises. Holonomy descent (Gate G7d) and terminal-map derivation or output-invariance (Gate G7h, F40) are open separately.

A typing caution on w_H. The isotropic local form $W\_H\text{^local} = w\_H\ I_3$ is fixed in *form* by the inherited structure — equal support directions force equal diagonal entries, layer-local terminality forbids local off-diagonal borrowing, and cross-layer coupling enters only later through Δ_Ω. But the *magnitude* w_H is the HMGBC history-metric stiffness (canonically 75/434 on the wheel) and **must not be identified with the anomaly-descent localization number 8/3**, which is a differently typed object. The prior papers close the form and positivity of the lift, not its history-metric magnitude; the magnitude remains part of the history-content gate (G4), not the projection gate.

32A.4 Holonomy-Split Generation Geometry

The degeneracy of $\Delta\_\Omega^{(0)}$ is broken by closure holonomy. What follows is a **complete finite calculation under six explicit conditions**, listed here so that no step of it is load-bearing without being visible:

1. **Coefficient group.** The orientation-invisible closure residue is the $\mathbb{Z}_7$ kernel of the derived D_7 transport group — not a chosen coefficient system. *(Typing caution: the primitive fold register and the $\mathbb{Z}_7$ transport register are* ***differently typed****; a value in one may not be read as a value in the other, exactly as w_H may not be read as the anomaly-descent 8/3.)*
2. **Definedness of the class.** Gate-2 orientation coherence detwists the coefficient system and the relevant closure sector is flat on elementary plaquettes, so $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$ is defined at all.
3. **Protection.** The protected primitive-Fact cycle is non-bounding under refinement.
4. **Descent.** The terminal-record projection carries the Gate-3 Wilson period to the quotient triangle with total phase $\varphi\_k = 2\pi k/7$ (this is Gate G7d, **open** — see below).
5. **Occupancy.** The closure phase records primitive commitment faithfully, so the committed-Fact branch has $k \neq 0$. This excludes the globally inert $\mathbb{Z}_7$ branch but does **not** fix the primitive charge to ± 1.
6. **Isotropic stiffness.** The local completion stiffness is isotropic, $W\_H = w\, I_3$ with $w > 0$; the value $w = 75/434$ is used **only** as the canonical-wheel benchmark, never as a physical prediction (§14C.3, F25).

The closure-transport programme supplies (1); the Gate-3 protection results supply (3). Conditions (2), (4) and (5) are named gates below.

Holonomy-descent obligation (Gate G7d). There is a bridge here that §32A/§44A computes across but does not derive: Gate-3 supplies the class on the *vacuum transport complex* Γ_vac, whereas the spectrum below places the flux on the *generation triangle* A → B → C → A. That requires the pushforward $H^1(\Gamma\_vac; \mathbb{Z}_7) \to H^1(\Gamma\_gen; \mathbb{Z}_7)$ — equivalently, a proof that the descended polar record transport $\mathcal{U}\_\Omega$ carries the Gate-3 class onto the terminal-record cycle with total Wilson phase $\varphi\_k = 2\pi k/7$. The general projection obligation (§31.1) already requires $\mathcal{U}\_\Omega$ to descend from the polar refinement transport rather than be assigned; the holonomy-descent obligation is the flat-sector instance of exactly that requirement, and it is **not** discharged here. The eigenvalue calculation below is therefore exact *conditional on the descended class*; the placement of the flux on the generation graph remains an identification until Gate G7d is proved.

For total quotient holonomy $\varphi\_k = 2\pi k/7$, the covariant three-cycle Laplacian has eigenvalues

$$[\![\lambda\_m(k) = 2\, g\_\Omega\, [\, 1 - \cos(\, (2\pi m + \varphi\_k)/3\,)\,], \quad m = 0, 1, 2, \quad g\_\Omega = 8/93.]\!]$$

For every nonzero k the three eigenvalues are distinct (verified for k = 1, 2, 3): e.g.

⟦ spec Δ_Ω^(1) ≈ {0.00764339, 0.21032618, 0.29815946}. ⟧

Theorem 32A.5 — Nonzero occupancy gives nondegenerate generations

The latest Gate-3 result establishes that a primitive committed Fact carries a **nonzero** $\mathbb{Z}_7$ residue, $k \neq 0$, provided the closure phase faithfully records the closure transition; it does *not* by itself fix $|k| = 1$. Since every nonzero k yields three distinct λ_m(k), it follows that

⟦ Fact existence ⟹ $k \neq 0$ ⟹ nondegenerate three-level generation geometry. ⟧

Unit winding is not required for hierarchy — nonzero occupancy alone suffices. The charge status must be stated in exactly two parts, and they are not the same claim:

⟦ $k \neq 0$ under faithful non-inert $\mathbb{Z}_7$ recording ⟧
⟦ $|k| = 1$ as the primitive Fold value: OPEN ⟧

Faithful recording excludes the globally inert sector; it does **not** fix the primitive charge to unit winding. The determinant calculation of §44A may show that $|k| = 1$ is the lowest Gaussian sector among the admitted nonzero sectors, but that energetic ordering may **not** be quietly converted into a proof that every primitive Fact deposits unit winding (Gate G7e, F30).

This is the first point in the chain that produces a genuinely nondegenerate three-generation operator from the K = 7 wheel + antipodal record quotient + $\mathbb{Z}_7$ holonomy — with the three-layer count supplied by the inherited provenance (Theorem 32A.0), not derived by the wheel, and with the sector-versus-generation typing obligation of §32A's preamble still open (F31). It is a generation *geometry*, not yet a mass spectrum: the Yukawa Hessian, sector selectors and kinetic normalisations act downstream, and these are canonical-wheel values, not physical mass ratios.

32A.6 Completion contraction

With the isotropic canonical scalar lift W_H^(3) = w_H I_3 (cross-generation structure supplied by Δ_Ω, never inserted into the local H-block), the positive Hermitian contraction is $\mathcal{S}$_Ω^(k) = $(I_3 + \Delta\_\Omega\text{\^{}}(k)/w\_H)^{-1}$. **Its spectrum depends on w_H and is therefore model-specific in exactly the sense F25 forbids treating as physical** — unlike the determinant *ordering* of §44A.1, which is provably w_H-independent. Evaluated at the canonical benchmark w_H = 75/434 **(a wheel magnitude, not a prediction)**, the spectrum is

⟦ spec $\mathcal{S}$_Ω^(0) = {1, 0.40106952, 0.40106952}, spec $\mathcal{S}$_Ω^(1) ≈ {0.95764367, 0.45104219, 0.36692543}, ⟧

and similarly for k = 2, 3 — three nondegenerate suppression levels for every nonzero winding. The canonical history law thus returns a genuine three-level completion hierarchy rather than the identity or a hand-selected graph.

32A.7 Exact holonomy descent and the labelling correction

The holonomy-descent obligation (Gate G7d) can be discharged on the canonical quotient by an exact chain-level computation. Let $\gamma_6 = [b_0b_1] + [b_1b_2] + \cdots + [b_5b_0]$ be the oriented rim cycle and $\gamma_3 = [AB] + [BC] + [CA]$ the quotient triangle. Under the antipodal quotient each quotient edge is covered twice, so

$$[\![\; q_\star \gamma_6 = 2 \gamma_3. \;]\!]$$

If the flat closure cochain is **basic** with respect to the antipodal quotient, $\rho\_vac = q^\star \rho\_\Omega$, then $k\_vac = \langle \rho\_vac, \gamma_6 \rangle = \langle \rho\_\Omega, q\_\star\gamma_6 \rangle = 2\langle \rho\_\Omega, \gamma_3 \rangle = 2\, k\_\Omega \pmod 7$. Since $2^{-1} = 4 \pmod 7$,

$$[\![\; k\_\Omega = 4\; k\_vac \pmod 7), \quad \text{equivalently } k\_vac = 2\; k\_\Omega. \;]\!]$$

Therefore $k\_vac \neq 0 \Longrightarrow k\_\Omega \neq 0$: the quotient cannot annihilate the Gate-3 class, because the covering degree 2 is invertible in $\mathbb{Z}_7$. Nonzero occupancy upstream survives descent — which is exactly what Gate G7d required, on the canonical branch.

Labelling correction (mandatory). Multiplication by 4 permutes the six nonzero classes, so the *set* of spectra is unchanged but the *labels* are not. Verified explicitly, the conjugate pairs map as

$$[\![\; \text{quotient } \pm1 \leftrightarrow \text{vacuum } \pm2, \quad \text{quotient } \pm2 \leftrightarrow \text{vacuum } \pm3, \quad \text{quotient } \pm3 \leftrightarrow \text{vacuum } \pm1. \;]\!]$$

Every occurrence of k in the generation-triangle spectrum (§32A.4–6) and in the determinant ordering (§44A) denotes the **quotient triangle class k_Ω**. In particular the dominant sector is

$$[\![\; |k\_\Omega| = 1 \quad \Leftrightarrow \quad |k\_vac| = 2, \;]\!]$$

so a result quoted as "|k| = 1" in vacuum labels would be **wrong**. Any statement about the primitive Fold charge is a statement about k_vac and must be converted (F33).

Residual condition. The descent above assumes antipodal *basicness* of the physical cochain, $\rho(r_3\, e) = \rho(e)$. That holds on the canonical D_6-symmetric closure branch; it remains to be checked on a fully symmetry-broken physical background. **Status: Gate G7d PASS on the canonical branch; basicness on a broken background OPEN.**

33. History-Derived Completion Resolvent

Let W_H > 0 be the H-defect block of the history-derived metric on the scalar-compatible completion carrier. The quadratic completion functional is

$$\Gamma\_comp^{(2)} = \tfrac{1}{2} \,\| X - \varphi\_H\, P\_H^R \|^2\_(W\_H) + \tfrac{1}{2} \langle X, \Delta\_\Omega\, X \rangle.$$

Stationarity gives $(W_H + \Delta_\Omega) X = W_H \varphi_H P_H^R$, hence

$$\llbracket\ X_\star = \varphi_H R_\Omega P_H^R, \quad R_\Omega = (W_H + \Delta_\Omega)^{-1} W_H.\ \rrbracket$$

Positivity, stated correctly. R_Ω is *not* Hermitian in the ordinary inner product when $[W_H, \Delta_\Omega] \neq 0$, so "$0 < R_\Omega \leq I$" in the Loewner sense is false. The manifestly positive object is the symmetrised contraction

$$\llbracket\ \mathcal{S}_\Omega = W_H^{1/2} (W_H + \Delta_\Omega)^{-1} W_H^{1/2} = \left(I + W_H^{-1/2} \Delta_\Omega W_H^{-1/2} \right)^{-1}, \quad 0 < \mathcal{S}_\Omega \leq I,\ \rrbracket$$

where the second equality shows $\mathcal{S}_\Omega$ is the inverse of I plus a positive semidefinite operator, hence a genuine Hermitian contraction. The solution map is related by similarity,

$$\llbracket\ R_\Omega = W_H^{-1/2} \mathcal{S}_\Omega W_H^{1/2},\ \rrbracket$$

so R_Ω has the same spectrum as $\mathcal{S}_\Omega$, namely $\mathrm{spec}(R_\Omega) \subseteq (0, 1]$, and is **self-adjoint in the W_H-weighted inner product**: $\langle x, R_\Omega y\rangle(W_H) = \langle R_\Omega x, y\rangle(W_H)$. R_Ω is therefore a *W_H-self-adjoint contraction*, not an ordinary positive operator. The suppression-factor reading uses $\mathrm{spec}(R_\Omega) = \mathrm{spec}(\mathcal{S}_\Omega) \subseteq (0, 1]$, which is the physically intended statement; the hierarchy is a *return* through these eigenvalues, not a graph choice. Downstream, the sector compressions $R_f = S_{fL}^\dagger \boldsymbol{\mathcal{R}}_\Omega S_{fR}$ do **not** inherit W_H-self-adjointness: for arbitrary left and right selectors, and after insertion of the left–right polar map $V_{(H,\star)}^\dagger$, R_f is a typed and generally non-normal map between two *different* induced weighted carriers. Self-adjointness would require additional intertwining conditions relating S_{fL} to S_{fR}, the left and right weights, V_H and the selected subspaces — none of which is assumed. The physical content of R_f is therefore carried by its **singular values and relative frames**, as §35 and §37 in fact use it, not by any self-adjointness property.

A second caution follows. Similarity to a contraction bounds the **eigenvalues** of R_Ω, not its ordinary singular values or Euclidean operator norm. Suppression statements must therefore be made either in the W_H-energy norm or through the Hermitian operator $\mathcal{S}_\Omega$; a positivity or contraction certificate must be issued for $\mathcal{S}_\Omega$, never for R_Ω in the naive product.

34. Oriented Completion Operator

The positive resolvent lives on the right completion carrier; the physical left–right operator includes the polar orientation:

$$\llbracket\ \boldsymbol{\mathcal{R}}_\Omega = V_{(H,\star)}^\dagger R_\Omega.\ \rrbracket$$

The oriented Higgs completion is

$$\llbracket\ \mathcal{O}_{(H,\star)} = (\varphi_{(H,\star)}/\Lambda_\star)\ \boldsymbol{\mathcal{R}}_\Omega P_H^R,\ \rrbracket$$

and the general history-derived scalar defect is consequently

$$\mathcal{D}_H^{\mathrm{hist}} = \mathcal{O}_H - (\varphi_H/\Lambda\star)\, \boldsymbol{\mathcal{R}}_\Omega\, P_H^R.$$

This replaces a freely chosen completion target (MASD-1's MA9/MA9′ candidate premise) with a history return. Where MASD-1 had to declare the generation-resolved law MA9′ as a candidate premise awaiting derivation, HMGBC-1 supplies $\mathcal{D}_H^{\mathrm{hist}}$ as that derivation, conditional on the complete generator $\mathbb{H}_n^{\mathrm{hist}}$, the projection gate (§31.1) that fixes Δ_Ω, and the W_H-self-adjoint contraction structure of R_Ω (§33) — the orientation $V_{(H,\star)}^\dagger$ carries the polar phase, so $\boldsymbol{\mathcal{R}}_\Omega$ is well typed even though R_Ω is only weighted-self-adjoint.

35. Sector Compressions

Let $S_{fR} : \mathbb{C}^3 \to \Omega^R$ and $S_{fL} : \mathbb{C}^3 \to \Omega^L$ be the frozen terminal-record occupancy selectors. The sector completion operators are

$$R_f = S_{fL}^\dagger\, \boldsymbol{\mathcal{R}}_\Omega\, S_{fR}, \quad f \in \{u, d, e, \nu\}.$$

These are maps between typed generation carriers, not four separately selected matrices. Mass hierarchy comes from the singular values of the resulting canonical completion maps; mixing comes from their relative left frames; CP structure comes from nontrivial oriented holonomy and failure of the sector compressions to admit a common real frame.

36. Identity Branch and MASD-1 Amendment

MASD-1's leading isotropic branch uses $\mathcal{O}_{(H,\star)} = (\varphi_{(H,\star)}/\Lambda\star)\, V_{(H,\star)}^\dagger P_H^R$. HMGBC recovers this exactly when $R_\Omega P_H^R = P_H^R$, equivalently $\Delta_\Omega P_H^R = 0$. Thus the identity branch is not rejected but classified:

$$\text{MASD-1 isotropic completion} \iff \text{scalar support} \subseteq \ker \Delta_\Omega.$$

A physical HMGBC return with nonzero completion Laplacian requires MASD-1's $\mathcal{D}_H$ to be replaced by $\mathcal{D}_H^{\mathrm{hist}}$ before HFB-1 executes (Firewall H7). This is the precise sense in which MASD-1's degenerate-mass CFB-1 result is not a failure of the architecture but the $\ker \Delta_\Omega$ corner of it.

37. Mixing, Hierarchy and CP Holonomy

Diagonal suppression alone does not generate mixing. The physical generation geometry requires three ingredients simultaneously: (1) nondegenerate eigenvalues of R_Ω or its sector compressions; (2) sector-dependent selector orientation; (3) nontrivial record holonomy. If every S_f selects the same reducing subspace and every holonomy is trivial, the completion matrices are simultaneously diagonalisable and the mixing matrices are trivial — which is exactly MASD-1's degenerate CFB-1.

A nonzero CP invariant requires an obstruction to a common real basis. In HMGBC this obstruction may come from complex polar record transport, nontrivial flat holonomy, the antisymmetric history current (§30), or noncommuting sector compressions. The paper does not assert these produce the observed CKM or PMNS values; it identifies the only admissible source of those values inside the history-selected geometry, and forbids any other (Firewall H4). Whether the actual bath triple yields the *measured* mixing is the H4 boundary execution, not a claim of this paper.

38. Relation to CY3′, and Binding It to the Selected Branch

MASD-1 compares Y_f^{mixed} (from the independent completion bilinear) and Y_f^{Hess} (from the Schur-reduced bosonic Hessian). HMGBC supplies the upstream completion operator and the history metric used by *both* routes; it does not make them equal by definition. The physical CY3′ test remains

$$[\![\ P_Y^L \left[\mathcal{O}_{(H,\star)} - (v_{cl}/\sqrt{2})(\mathcal{D}_{(\varphi_H)}\, \mathcal{O}_H)_\star \right] P_Y^R = 0\ ]\!]$$

together with $Y_f^{\mathrm{mixed}} \stackrel{?}{=} Y_f^{\mathrm{Hess}}$.

Binding condition (strengthened). This comparison is discriminating only on the *physically selected* hierarchical branch $\beta_\star$ returned by §45 under the §43A theorem — not on an arbitrary hierarchical candidate, and never on an identity branch where both sides are I and the equality is vacuous. A pass on some hand-chosen hierarchical branch would test the machinery on a non-physical wiring and would not license the certificate. As of this paper, CY3′ has zero discriminating evidence (MASD-1 F9): every executed instance is identity-trivial or computes Y_f^{mixed} alone. HMGBC-1 does not change that status; it makes the *object on which the test must run* — $\beta_\star$ with its Δ_Ω and R_Ω — precise, so that HFB-1's execution of CY3′ is on the selected branch or it is not admissible (§52A, Gate binding).

Part VI — Bath–Ledger Decision and Physical Branch

39. The Bath Response Algebra

The static unresolved response on the primitive defect carrier is $\boldsymbol{\mathcal{B}}_n = C_n^\dagger (\Omega^2_{(Q,n)})^+ C_n$, the pseudo-inverse being required because $\Omega^2_{(Q,n)}$ arises as a compression and may be singular (§56A.3); the bath-ledger verdict therefore inherits the range condition ran $C_n \subseteq$ ran $\Omega^2_{(Q,n)}$ of Theorem 16.1b. On a class carrier $\mathscr{H}_\mathcal{C}$, define the compressed response $\boldsymbol{\mathcal{B}}_\mathcal{C} = P_\mathcal{C} \boldsymbol{\mathcal{B}}_n P_\mathcal{C}$. Let

```
𝔄_𝒞 = alg^*{ 𝓑_𝒞, class phase generators, refinement transports }
```

be the generated finite-dimensional *-algebra, with commutant $\mathfrak{A}_\mathcal{C}'$. This algebra contains exactly the projectors and observables protected by all bath-mediated transport at the declared order. Note $\boldsymbol{\mathcal{B}}_n$ is the $\omega = 0$ bath response — the same object whose Schur complement is W_{prim} (§15) — so the ledger verdict and the metric are two readings of one response, not two constructions.

40. Protected Sub-Accounts and Partial Baths

A nontrivial reducing projector $0 < P < I$ with $[P, \mathfrak{A}_\mathcal{C}] = 0$ defines a protected sub-account. If all one-dimensional sector projectors commute with the response algebra, the branch is a **ledger**. If the minimal central projections have ranks $k_1, \ldots, k_r$, the branch is a **partial bath** with block structure $U(k_1) \times \cdots \times U(k_r)$ after the appropriate phase and continuity tests. This classification is calculated from the same response that contributes to W_{prim}.

41. Saturated-Bath Criterion

Theorem 41.1 — Bath Saturation by Commutant Triviality.

Let $\mathfrak{A}_\mathcal{C}$ act on a class carrier $\mathbb{C}^k$. If $\mathfrak{A}_\mathcal{C}' = \mathbb{C} I$, then the response acts irreducibly and no proper protected sub-account survives. Conversely, any nontrivial protected sub-account produces a non-scalar commutant element.

Proof. Finite-dimensional Schur theory: a reducing projector lies in the commutant; trivial commutant forbids one; a nontrivial invariant decomposition supplies its orthogonal projector. ■

This converts bath saturation from a narrative premise into an operator audit — the resolution of the programme's bath–ledger fork.

42. Connection to U(k) and SU(k) Transport

The bath-response audit alone establishes irreducible sharing. To reach the non-Abelian transport theorem, the class must also possess unitary reversible transport, continuity, and sector-resolved phase. Under those inherited conditions, saturated irreducible bath response gives U(k) on the class carrier; removing the shared central phase leaves SU(k). Thus HMGBC converts the bath premise into a finite numerical audit, which HFB-1 must execute on the physical carrier — no response algebra or commutant has been evaluated here:

```
⟦ trivial commutant + continuous phase-bearing transport  ⟹  saturated bath.
⟧
```

If the response algebra is block-reducible, the full non-Abelian claim fails or reduces to the returned blocks. This is where colour SU(3) and weak SU(2) must appear as verdicts rather than inputs; a ledgered colour class would be a definite falsification (F12).

43. Admissible Branch Space

Let $\mathfrak{B}_{adm}^{(n)}$ be the set of branches satisfying: defect-zero stationarity; physical positivity; anomaly and representation constraints; terminal-record and carrier typing; HP1–HP3; and frozen regulator and reference conventions. The set is not assumed finite (Firewall H6). For a continuous branch manifold, require compactness or coercivity sufficient to ensure a free-energy minimum exists.

43A. The Common-History-Operator Theorem — the Branch-Comparison Gate

This section supplies the precondition without which the branch selector of §§44–46 is not merely uncertain but ill posed. It is the companion's answer to the sharpest structural objection to the programme: **a Perron root is an invariant of one irreducible operator, so comparing the Perron data of separately constructed branch operators compares incommensurable quantities.** Minimising over such numbers is not selection; it is a category error dressed as an argmin.

43A.1 The obligation

Common-History-Operator Obligation (HM9 made precise). There must exist one master substrate history operator $\mathbb{T}_n^{hist}$ on one common history space $\mathscr{H}_n^{hist}$, and a common reference measure $\hat{\nu}_n^{hist}$, such that:

1. **Component structure.** Every admissible branch $\beta \in \mathfrak{B}_{adm}^{(n)}$ is a component of $\mathbb{T}_n^{hist}$ — an invariant subspace, superselection sector, boundary condition or reducing projector P_β with $\mathbb{T}_n^{hist} P_\beta = P_\beta\, \mathbb{T}_n^{hist} P_\beta$ on the physical support.
2. **Common normalisation.** All branch weights $Z_{(n,\beta)}$ and the transfer step a_t descend from the one reference measure $\hat{\nu}_n^{hist}$; no branch is separately normalised (HM5, Firewall H5, Firewall H8).
3. **Block-Perron identity.** The branch free energy is the block Perron exponent of the *one* operator restricted to its component,

⟦ $f\_\beta^{(n)} = -(1/a\_t) \log \lambda\_0(P\_\beta \mathbb{T}\_n^{hist} P\_\beta)$, ⟧

so that comparing f_β across β compares eigen-branches of one operator, not eigenvalues of unrelated operators.

43A.2 What follows, and what does not

If the obligation holds, then the free-energy comparison of §44–§45 is a comparison of the block spectra of a single positive operator, the argmin is physical, and a positive selection gap (§46) certifies a unique dominant component. Standard Perron–Frobenius block theory then gives: the dominant component is the one whose restricted operator carries the largest block Perron root; ties are genuine superselection degeneracies; and refinement stability (§47) is the statement that the dominant *component* persists under the intertwining of §23.

If the obligation fails — if the branches are separately constructed operators, or carry mutually singular supports (§9), or are separately normalised — then no per-branch Perron gap rescues the comparison. Each branch may have a perfectly isolated Perron root and the cross-branch minimisation still means nothing. This is the exact sense in which MASD-1's F29 is *not* resolved by a per-branch gap, and it is why HMGBC-1 separates the two:

```
⟦ per-branch gap (HP2/G3)  ≠  cross-branch comparability (§43A/G9). ⟧
```

43A.3 Status — partially advanced, core still open

The obligation splits into parts of unequal difficulty, and the inherited corpus closes some of them:

- **Common operational carrier and measure (G9a): PASS conditional on the operational-measure theorem.** That theorem supplies one common quadratic probability measure across the substrate's admissible coherence channels, with cross-channel bridging forcing one normalisation, and it makes clear the channels are transport blocks permitting cross-channel superposition — so there is *no literal superselection obstruction* to global normalisation. This removes one way the obligation could have failed.
- **Canonical $\mathbb{Z}_7$ holonomy-sector comparison (G9b): conditional pass** (Theorem 44A.1), as a Gaussian determinant ordering, not yet a block-Perron calculation.
- **Full-branch embedding (G9c): OPEN.** The operational-measure theorem does **not** place CFB-1/2/3 inside one operator. MASD-1 expressly states those constructions use different arenas, metrics and completion laws, that the remaining branch freedom is *continuous* (no finite isolated branch set has been returned by the history measure), and that their increasing Standard Model resemblance is evidence of increased *tuning*, not convergence of one derivation. A formal direct sum does not help (Firewall H10): its arbitrary relative normalisation just becomes the new selection parameter. The components must arise as spectral/invariant projections of one generator constructed *before* the branches.

- **Raw block-Perron evaluation and global dominance (G9d): OPEN.**

So the obligation is **not discharged**: G9c/G9d are the genuinely hard, still-open core, and Appendix F specifies exactly what their proof must exhibit — a construction of $\mathbb{T}_n^{hist}$ whose invariant decomposition *is* the admissible branch set, over the branch-labelling observables (vacuum, support, polar, bundle, orientation, refinement), verified on the K = 7 regulator. The *arena* question (one carrier, one measure, no superselection wall) is inherited-closed; what remains is the single large structural task of constructing and evaluating the one master history operator. Both open conditions are logged as falsifiers (F16 comparability, F16b global dominance).

Theorem 43A.4 — Reduction of Branch Selection to One Operator.

Assume the Common-History-Operator Obligation and HP1–HP2 on each component. Then $\mathfrak{B}_{phys}^{(n)} = \mathrm{argmin}_\beta f_\beta^{(n)}$ is the set of dominant components of $\mathbb{T}_n^{hist}$, is nonempty when the component set is compact or f_β is coercive, and is a singleton iff the dominant block Perron root is non-degenerate across components. **Proof.** Under the obligation, $\{f_\beta\}$ are the block Perron exponents of one operator; minimising $-(1/a_t)\log \lambda_0$ selects maximal-λ_0 components; existence follows from compactness/coercivity and lower semicontinuity of the top eigenvalue; uniqueness is non-degeneracy of the dominant block root. ■

The theorem is genuine H1 mathematics — *given the obligation.* Its entire physical content is therefore transferred to the obligation, which is the H3 gate.

43A.5 Global Dominance versus Superselection — What the Perron Minimum Does and Does Not Prove

One further distinction must be drawn, because §43A.4 as stated proves less than it may appear to. Even with the exact block decomposition $\mathbb{T}_n^{hist} = \bigoplus_\beta \mathbb{T}_\beta$, the largest block Perron root identifies the **thermodynamically dominant** component — the one that dominates the long-history weight — and this is *not* the same as proving that the other invariant components cease to be physical. Dominance is not, by itself, ontological selection.

Dominance becomes selection only under a global-support-and-normalisation theorem.

Global Dominance Condition. Let the reference positive density ρ_0 have nonzero support on every admitted component, $P_\beta \rho_0 \neq 0$ for all β, and define the **globally** normalised evolved density

$$\rho_N = (\mathbb{T}^{hist})^N \rho_0 / \| (\mathbb{T}^{hist})^N \rho_0 \|_1.$$

If one block has strictly maximal Perron root $\lambda_{(0,\beta_\star)} > \lambda_{(0,\beta)}$ for all $\beta \neq \beta_\star$, then

$$\rho_N \longrightarrow P_{\beta_\star}\, r_{\beta_\star} / \| P_{\beta_\star}\, r_{\beta_\star} \|_1 \quad \text{as } N \to \infty.$$

The four hypotheses are load-bearing, and each is its own failure mode: (i) the reference state must have nonzero support on the dominant component (else it can never be reached); (ii) the evolution must be normalised **globally across the full direct sum**, not block-by-block — block-wise renormalisation silently violates Firewall H5/H8 and destroys the comparison; (iii) the long-history limit must be taken **before** any per-block conditioning; (iv) the blocks must not be independently retained as permanently coexisting superselection worlds, since a genuine superselection rule would keep every component physical regardless of its Perron weight.

Consequence for the certificate. Absent the Global Dominance Condition, the free-energy argmin of §45 identifies the dominant branch but does **not** prove the others are non-physical: it is a thermodynamic selection, not an ontological one. A physical certificate must state which it claims. The distinction is not academic — if VERSF's branch structure is genuinely superselecting (case iv), then "the Standard Model is the dominant branch" is a statement about weight, and the coexistence of other branches would itself be a physical prediction to be confronted, not an artefact to be minimised away. The branch-comparison gate therefore has **two** independent parts: *comparability* (§43A.1–4, that the branches are one operator's components) and *global dominance* (§43A.5, that dominance implies selection). Neither is discharged here; both are logged (F16 comparability, F16b global dominance).

43A.6 Branches as communicating components — a first decomposition, superseded by §56A.8

*Forward reference: the definition below is **not** the paper's final position. It is refined into the central decomposition of §56A.8, which is exhaustive relative to the frozen observable algebra; where the two differ, §56A.8 governs.*

The obligation's hardest part (G9c) is not to *assemble* a direct sum but to have one arise. The canonical construction of §11A supplies it. Let $G_n = (\mathcal{C}_n, E_n)$ be the complete admissible-move graph of §11A.1 and define

$$\mathfrak{B}_{adm}^{(n)} = \pi_0(G_n),$$

the set of its **communicating components**. For each component β let P_β be multiplication by its characteristic function. Then automatically

$$[\mathbb{T}_n^{hist}, P_\beta] = 0 \quad \text{and} \quad \mathbb{T}_n^{hist} = \bigoplus_{\beta \in \pi_0(G_n)} P_\beta\, \mathbb{T}_n^{hist}\, P_\beta.$$

This is not a post-hoc formal direct sum (Firewall H10): the operator and its move graph are constructed *first*, and the components are *discovered* afterwards. The corresponding central branch algebra is $\mathfrak{z}_n = \ell^\infty(\pi_0(G_n))$, whose minimal projections are exactly the P_β; the freeze-manifest labels — vacuum, support, polar class, bundle class, fermion orientation, refinement class, flat holonomy — become **coordinates on this centre** rather than independently assembled models. Appendix F's items 1–3 are addressed at the support level by this construction rather than by assumption; items 4–6 — one reference measure, block-Perron comparability with proven gaps, and refinement covariance — are not, and the exhaustiveness question below remains.

Consequence for the CFB family. A proposed CFB configuration counts as a physical branch **only if it embeds into one of these communicating components**. Otherwise it is a benchmark model, not an admissible component of the physical operator — which is the precise sense in which MASD-1's differently-normalised CFB constructions cannot be compared until they are so embedded. The construction is also neutral on the open Reversible Connectedness question: if RC holds the relevant configurations lie in one component; if RC fails, the sealed components are *returned* as genuine separate branches. Either way the answer is computed, not assumed.

What this proves, and what it does not. It proves that **every communicating component is a branch**. It does *not* prove the converse — that every physical branch is a communicating component. Additional invariant or reducing sectors may arise from conserved representation labels *inside* a connected support, phase or determinant-line sectors, boundary conditions, repeated irreducible representations, continuous moduli, or central projections not visible as disconnected support. Closing G9c therefore requires proving

$$\llbracket\ \mathfrak{z}_n = Z(\mathrm{Comm}(\mathfrak{B}_n)) = \ell^\infty(\pi_0(G_n))\ \rrbracket$$

(or its direct-integral analogue), **and** that every MASD/CFB branch label is represented in that algebra.

Status. ⟦ Communicating-component branch decomposition: DERIVED CANDIDATE. ⟧ ⟦ Exhaustive physical branch embedding: OPEN. ⟧ Remaining: the centre-identification above and the finite enumeration of the admissible move graph (F32).

44. Raw Transfer and Branch Free Energy

Let $\mathbb{T}_{(n,\beta)}$ be the raw, commonly normalised transfer on branch β — i.e. the restriction $P_\beta \mathbb{T}_n^{hist} P_\beta$ of §43A, not a separately built operator. Its leading Perron root is $\lambda_{(0,\beta)}^{(n)}$, and the **one-step** free energy is

$$\llbracket\ f_\beta^{(n)} = -(1/a_t) \log \lambda_{(0,\beta)}^{(n)}.\ \rrbracket$$

This is a one-step free energy, not yet a density, and the distinction is not cosmetic. On a multi-cell regulator $\log \lambda_0$ is generally **extensive**, so f_β grows with the cell count and cross-level comparison is meaningless until an extensivity convention is declared. The execution must state which of the following holds: $\mathbb{T}_n$ is already normalised per cell; f_β is divided by the spatial volume or cell count; a common extensive vacuum term is subtracted; or the limiting object is a pressure/free-energy density rather than a one-step exponent. The *within-level* argmin of §45 is unaffected by a common additive extensive term, but §47's cross-level comparison and the selection gap Δ_{sel} are not — both require the declared convention (F43).

Independent normalisation $\mathbb{T}_{(n,\beta)} \mapsto \mathbb{T}_{(n,\beta)}/\lambda_{(0,\beta)}$ is permitted only after the f_β comparison has been recorded (Firewall H5). The Gaussian bath determinant contributes to $\lambda_{(0,\beta)}$ (Appendix A) and cannot be discarded as a branch-independent constant unless verified branch-independent — a verification that is itself part of the §43A obligation, since a branch-dependent determinant that is not a block of one operator would break comparability.

44A. Common-Gaussian Holonomy-Subsector Determinant Comparison

This section executes, on one explicitly declared canonical measure, the *only* branch-type comparison HMGBC-1 is currently able to perform: a comparison of Gaussian **stiffness blocks** across the $\mathbb{Z}_7$ holonomy subsector. It is placed here, beside the branch free energy of §44, because that is what it is a candidate contribution to — not in the generation-geometry sections, where it would read as a physical selection.

Mandatory status statement. This is a comparison of Gaussian stiffness blocks on one explicitly declared canonical measure. **It does not discharge the full Common-History-Operator Obligation of §43A, and it is not yet a block-Perron physical-branch selection theorem.** A common Gaussian stiffness direct sum is not the §43A object (Firewall H10); promotion requires Gate G7f.

Theorem 44A.1 — Canonical Gaussian Holonomy-Sector Determinant Ordering (|k| = 1, degenerate in sign)

The seven holonomy sectors can be placed in **one** common block-diagonal stiffness operator with one common Gaussian measure and no branchwise renormalisation,

$$\mathbb{A} = \bigoplus_{k \in \mathbb{Z}_7} \left(w_H I_3 + \Delta_\Omega^{(k)} \right).$$

This is a *stiffness* operator, not yet the master history *transfer* operator of §43A; the theorem below is accordingly a common-measure Gaussian comparison across the holonomy sectors, and it supplies the candidate **determinant contribution** to the eventual block-Perron free energies — it does **not** yet discharge the one-operator comparability requirement (see the promotion note). The Gaussian static free energy is $f_k = f_{common} + (1/2a_t) \log D_k$ with $D_k = \det(w_H I_3 + \Delta_\Omega^{(k)})$, which reduces exactly to

$$D_k = w_H^3 + 6\, g_\Omega\, w_H^2 + 9\, g_\Omega^2\, w_H + 2\, g_\Omega^3\, [1 - \cos(2\pi k/7)] = D_{common} + 2\, g_\Omega^3\, [1 - \cos(2\pi k/7)].$$

Structural ordering (model-independent). Because only the last term depends on k, and $1 - \cos(2\pi k/7)$ is strictly increasing in $|k|$ over $\{0, 1, 2, 3\}$ while $g_\Omega^3 > 0$,

$$D_{(\pm 1)} < D_{(\pm 2)} < D_{(\pm 3)} \quad \text{for every } w_H > 0,\ g_\Omega > 0.$$

The absolute-winding ordering is therefore **structural** — it does *not* depend on the canonical value $w_H = 75/434$ (the completion eigenvalues of §32A.6 are canonical-wheel magnitudes, but the *ordering* is not). The $k = 0$ sector is the global minimum but is **excluded on the committed-Fact branch** by Gate-3 nonzero occupancy (Theorem 32A.5) — an exclusion that remains conditional on faithful nonzero primitive-Fact occupancy and must stay visible. Among the admitted nonzero sectors the minimum is therefore

$$[\![\mathfrak{B}\_\text{dominant}^\text{wheel} = \{ k = +1, k = -1 \},]\!]$$

with gap to the next pair $\Delta f_{(1\to 2)} = (1/2a_t) \log(D_2/D_1) \approx 0.0162703 / a_t$ (this *magnitude* is canonical-wheel-specific; the *sign* $\Delta > 0$ is structural). **Nonzero occupancy plus a common isotropic positive H-block orders the sectors so that unit absolute winding minimises the Gaussian determinant free energy**, structurally, even though Gate-3 topology alone forced only a nonzero residue. The two signs $\{+1, -1\}$ remain exactly degenerate as orientation/CP conjugates; a separately derived commitment orientation may select the sign, and it **may not** be selected from observed CP data (HM8, §53; F18). The proper certificate is: canonical holonomy-sector determinant ordering selects $|k| = 1$, DEGENERATE $\{+1, -1\}$.

Promotion to a §43A comparison (open). To make this a genuine Common-History-Operator comparison, one must construct a single positive history transfer $\mathbb{T}\_\text{hol}^\text{hist} = \bigoplus_{(k \in \mathbb{Z}_7)} e^{(-a_t \mathbb{H}_k)}$ on one history space and prove that the branch-dependent part of its block Perron exponent $-(1/a_t) \log \lambda_0(P_k \mathbb{T}\_\text{hol}^\text{hist} P_k)$ equals the determinant expression above. Until that is done, $\mathbb{A}$ is a stiffness comparison, not the transfer-operator comparison §43A requires.

Scope. Even once promoted, this would be a one-operator comparison only across the $\mathbb{Z}_7$ **holonomy sectors** with the canonical isotropic w_H lift and the Gaussian weight. It is *not* the full branch-comparison gate: the CFB/MASD branches beyond holonomy, the non-isotropic physical w_H, the non-Gaussian corrections and the global-dominance condition (§43A.5) are untouched. What it shows is that the branch-selection machinery, on the one subsector where all inputs are canonical and exact, returns a definite, non-fitted, structurally-ordered answer — a proof of concept for §43A, not its discharge.

44A.2 Two Readings of k — Branch Label versus Phase-Assignment Law

A second distinction is mandatory, and it is easy to elide. The determinant ordering says something different depending on what kind of object k is:

- **If k is a branch label** — a degree of freedom over which the common history measure compares admissible sectors — then the determinant ordering *conditionally selects* $|k| = 1$ among the committed-Fact sectors, in the sense of Theorem 44A.1.
- **If the microscopic phase-assignment law fixes a particular nonzero k *before* branch comparison**, then the calculation **evaluates that sector but does not override it.** The ordering is then a consistency observation, not a selection: one does not get to re-select a quantity the substrate has already fixed.

The paper does not claim to know which reading is correct — that is a fact about the phase-assignment law, not about this calculation — and the honest statement of Theorem 44A.1 must therefore carry the disjunction. Relatedly, the still-open **Primitive-Fold Unit-Step** statement $\Delta\theta = \pm 1$ would independently select the same pair $\{+1, -1\}$; but it is *not required* for the determinant ordering, and the two must not be quoted as mutual support. Faithful nonzero occupancy (condition 5) excludes $k = 0$; it does **not** fix the primitive charge to ± 1. That separation is Gate G7e.

So the finite result is exactly:

```
⟦ conditional nonzero-holonomy generation spectrum: calculated exactly ⟧
⟦ Gaussian ordering of the nonzero sectors: |k| = 1, with ± degeneracy ⟧
```

while holonomy-to-generation typing (G7d), faithful occupancy and unit charge (G7e), continuum pinning, and Gaussian-to-Perron promotion (G7f) remain explicit physical gates.

44A.3 The Gaussian Determinant Perron-Realisation Lemma

Gate G7f asks that the determinant ordering be shown to *be* a block-Perron free energy of one positive transfer, not merely to resemble one. On the holonomy subsector this can be done exactly. Let $A_k = w_H I_3 + \Delta_\Omega^{(k)} > 0$ and $u_k(x) = \exp(-\tfrac{1}{2} x^\dagger A_k x)$ for $x \in \mathbb{R}^3$. On the single space $L^2(\mathbb{Z}_7 \times \mathbb{R}^3$, counting $\times\, dx)$ define one positive operator by

$$⟦\ (\mathbb{T}_{hol} f)(k, x) = u_k(x) \int_{\mathbb{R}^3} u_k(y)\, f(k, y)\, dy.\ ⟧$$

This is one formula, on one space, with one measure. Holonomy is conserved, so the k-sectors are its invariant components; each block is rank one and positivity improving, with Perron eigenfunction u_k and eigenvalue

$$⟦\ \lambda_{(0,k)} = \|u_k\|_2^2 = \int_{\mathbb{R}^3} e^{-x^\dagger A_k x}\, dx = \pi^{3/2} / \sqrt{\det A_k}.\ ⟧$$

Hence the block-Perron free energy is

$$⟦\ f_k = -(1/a_t) \log \lambda_{(0,k)} = -(3/2a_t) \log \pi + (1/2a_t) \log D_k, \quad D_k = \det A_k,\ ⟧$$

verified exactly. The branch-dependent part is precisely the Gaussian determinant term. This *proves*, rather than assumes, that the determinant contribution is the block-Perron free energy of a positive common transfer, and since $D_{(\pm 1)} < D_{(\pm 2)} < D_{(\pm 3)}$ the committed nonzero quotient sectors give $\mathfrak{B}_{dominant}^{hol} = \{k_\Omega = +1, k_\Omega = -1\}$ — in upstream vacuum labels $\{k_{vac} = +2, k_{vac} = -2\}$ (§32A.7, F33).

Scope — how little this establishes. The content must be stated at its true strength, which is small. Each block of $\mathbb{T}$_hol is **rank one**, so its Perron root is simply $\|u_k\|^2$; consequently, for *any* prescribed positive numbers λ_k whatever, some common positive operator has exactly those block Perron roots — choose u_k with the matching norm. The lemma therefore establishes only that the determinants D_k are positive, and it **refutes one objection and no more**: the objection that a determinant ordering could not *in principle* be a block-Perron ordering. It establishes nothing about the physical operator.

It does not satisfy Firewall H10. H10's criterion is that components arise as projections of a generator constructed *before* the branches. Here each A_k was constructed per-k first and the components assembled afterwards, with the relative normalisation supplied by the assembler; writing the result on $L^2(\mathbb{Z}_7 \times \mathbb{R}^3)$ does not change that. The lemma is therefore explicitly **not** an instance of the §43A object, and it is not counted toward G7f or G9b.

Global dominance (§43A.5) applies blockwise: with initial support in every admitted component the normalised evolution converges to the dominant component, but the two signs are exactly degenerate, so the return is the pair $\{+1, -1\}$ and the limiting mixture depends on their initial weights. The sign is not selectable from measured CP data (HM8, §53; F18).

Status. ⟦ Refutation of the in-principle objection: established. ⟧ ⟦ Any statement about the physical operator: NOT ESTABLISHED (does not satisfy H10). ⟧

45. Physical-Branch Selector

Theorem 45.1 — Bi-Perron Physical-Branch Selection.

Assume: (1) the Common-History-Operator Obligation (§43A); (2) every admitted component satisfies HP1 and HP2 with a proven gap; (3) the component set is compact or f_β is coercive; (4) f_β is lower semicontinuous; (5) one common normalisation (HM5). Then the physical branch set is nonempty and equals

$$[\![\ \mathfrak{B}\_\text{phys}^{(n)} = \operatorname{argmin}_{(\beta \in \mathfrak{B}\_\text{adm}^{(n)})} f_\beta^{(n)}.\ ]\!]$$

A unique physical branch exists iff the minimum is unique. **Proof.** Theorem 43A.4 plus the direct method of the calculus of variations. ■

Hypothesis (1) is not decorative: dropping it makes the argmin ill posed, as §43A.2 shows.

46. Selection Gap and Degeneracy Discipline

For a unique minimiser $\beta_\star$, define $\Delta\_\text{sel}^{(n)} = \inf_{(\beta \neq \beta_\star)} [\ f_\beta^{(n)} - f_{(\beta_\star)}^{(n)}\ ]$. The selection is certified only if $\Delta\_\text{sel}^{(n)} > 0$ with numerical uncertainty below $\Delta\_\text{sel}^{(n)}/2$. If Δ_sel $= 0$, the physical return is a degenerate branch set, and no secondary Standard Model comparison may break the tie (HM8, §53; F18). A degenerate return is a legitimate physical output —

possibly a genuine multiplicity — and is published as such, never resolved by resemblance to data.

47. Regulator Stability of the Branch

A physical branch must persist under refinement. Let ι_n map branches at level n into level n+1; require $\iota_n(\mathfrak{B}_{phys}^{(n)}) \cap \mathfrak{B}_{phys}^{(n+1)} \neq \emptyset$, and, for a unique branch, convergence of $W_{prim}^{(n)}$, $R_\Omega^{(n)}$, $f_{(\beta_\star)}^{(n)}$ and $\Delta_{sel}^{(n)}$. Under §43A this is the statement that the dominant *component* of $\mathbb{T}_n^{hist}$ maps to the dominant component of $\mathbb{T}_{(n+1)}^{hist}$ under the HP3 intertwining — a single-operator statement, hence well posed. A branch that changes because of finite-regulator ordering rather than a controlled bifurcation is not yet physical.

Part VII — Execution, Certificate and HFB-1 Handoff

47A. Regulator Convergence on the Intertwined Canonical Family

The coarse-graining programme supplies the stochastic spatial update $P_{(X,n)} = ½(I + D_n^{-1} A_n)$, the factor ½ being selected within its declared family by constant-mode stochasticity and exact extinction of the alternating fine mode. Let V_n^X be the normalised refinement lift with $P_{(X,n+1)} V_n^X = V_n^X P_{(X,n)}$, and define the full lift $V_n = I_{wheel} \otimes V_n^X \otimes I_{int} \otimes I_{bath}$. If the typed scores and reference potential are cylindrically compatible then

$$\mathbb{T}_{(n+1)}^{hist} V_n = V_n \mathbb{T}_n^{hist},$$

whence $W_{prim}^{(n+1)} V_n = V_n W_{prim}^{(n)}$, $\Delta_\Omega^{(n+1)} V_n^\Omega = V_n^\Omega \Delta_\Omega^{(n)}$, and $P_{(\iota_n(\beta))}^{(n+1)} V_n = V_n P_\beta^{(n)}$: metric, terminal geometry and branch projectors all intertwine.

47A.1 Uniform gap

For the factorised reference kernel $K_{(0,n)} = T_0 \otimes P_{(X,n)}$ the spectrum is the product set. With spatial gap $\delta_X = 1 - \sup_n |\lambda_2(P_{(X,n)})| > 0$,

$$\delta_{global} \geq \min\{ ½, \delta_X \},$$

verified exactly (the canonical wheel contributes its own gap ½ from Theorem 14A.1). If the refinement lift preserves the inherited spectrum and sends new fine alternating modes to eigenvalue zero, the bound is uniform at every regulator level — which is precisely the regulator-uniform lower bound Gate G3 requires, on this family.

47A.2 Approximate intertwining

If instead $\|\mathbb{T}_{(n+1)} V_n - V_n \mathbb{T}_n\| \le \varepsilon_n$ with uniform gap $\delta > 0$, standard spectral perturbation gives

```
⟦ ‖P_(n+1)^Perron V_n - V_n P_n^Perron‖ = O(ε_n/δ),  ‖W_prim^(n+1) - V_n
W_prim^(n) V_n†‖ = O(ε_n/δ²), ⟧
```

and the same form controls R_Ω, the sector compressions and the branch free-energy ordering.

Status. ⟦ Canonical regulator-convergence theorem: PASS *given* exact intertwining (HP3) and a uniform spatial gap $\delta_X > 0$. ⟧ ⟦ Physical two-level convergence execution: OPEN. ⟧ Both antecedents are assumptions until two actual regulator levels are assembled and checked, so **Gate G10 remains open**; what the section supplies is the error model that the check will use.

Two levels are a diagnostic, not a convergence certificate. Agreement between levels n and n+1 within tolerance is necessary and far from sufficient: it cannot by itself establish continuum convergence or a regulator-*uniform* gap, since any finite pair is consistent with drift that only appears later. A genuine certificate requires one of: an **analytic a priori bound** on ε_n; a **monotonicity theorem** in n; or a **controlled Cauchy/error estimate** giving summable ε_n. The two-level run is the minimum admissible evidence, and the paper does not treat it as more (F43).

48. Frozen Input Package

The physical HMGBC execution must publish one immutable package. **It freezes carriers, builders, algorithms and conventions — never quantities that Stage A is meant to calculate** (§52.1, Appendix G). It contains: (1) regulator complex and orientation data; (2) resolved physical quotient X_n; (3) operational measure $\mu_{op}^{(n)}$; (4) resolved kinetic Dirichlet generator L_n^{op}, with its positivity-preserving certificate (F37); (5) the **unresolved bath carrier and basis** — *not* the evaluated $\Omega^2_{(Q,n)}$, which is a compression of the assembled generator (§56A.3); (6) the **defect map and its tangent embedding ι** — *not* the evaluated coupling C_n, likewise a compression; (7) defect tangent basis, typed embeddings and the declared tilt family (HM12); (8) branch parameterisation, admissibility tests, the **observable algebra $\mathfrak{A}_n^{kin}$ and the branch-decomposition algorithm** — *not* the returned central components, which are a Stage-A output; (9) refinement lifts and coarse maps, which must preserve local self-transition (persistence) support per Lemma I.5, or the scalar-persistence defect $\mathcal{D}_H$ collapses onto the current defect; (10) the **microscopic record transport $\mathcal{U}^{rec}$ and the terminal-record map r_n** — *not* the pushed-forward terminal transport $\mathcal{U}_\Omega$, which is calculated from them by polar decomposition (§56A.7); (11) the numerical tolerance block $\mathcal{T}_{num}$; (12) code, dependency and hash manifest, with admission hashes covering builders and inputs only (Appendix G). A paper that publishes only W_prim without the carrier, builder and algorithm package does not permit the history claim to be audited.

49. Required Numerical Outputs

At each of at least two regulator levels, export: $\{\omega^2_{(n,a)}\}$; C_n; $\mathbb{T}_n$; $(\lambda_{(0,n)}, \ell_n, r_n)$; $W_{hist}^{(n)}(\omega)$; $W_{prim}^{(n)} = W_{hist}^{(n)}(0)$; the complete block ledger $[W_{AB}^{(n)}]$; the null basis and classification; $C_\Omega^{(n)}$, $\Delta_\Omega^{(n)}$, $R_\Omega^{(n)}$; the sector compressions R_u, R_d, R_e, R_ν; the bath-response algebra and commutant $\mathfrak{A}_{\mathcal{C}}$, $\mathfrak{A}_{\mathcal{C}}'$; **the master-operator component decomposition and its reference measure (§43A)**; and the branch table $\{\beta, f_\beta, \lambda_{(0,\beta)}, \Delta f_\beta\}$.

50. HMGBC-1 Admission Tests

- **Gate G1 — Source completeness.** Every same-order unresolved mode is represented or bounded below the published error threshold (HM1).
- **Gate G2 — History positivity.** The full bath-resolved quadratic block is nonnegative on the physical quotient and the transfer satisfies HP1. **The gate must declare which route of §20.2 it invokes for the fermion sector: (a) real nonnegative fermionic determinant, or (b) certified reflection/transfer positivity. A bosonic HP1 pass quoted as certifying the full transfer fails this gate (F21).**
- **Gate G3 — Bi-Perron gap magnitude.** Strict finite-level dominance $\lambda_{(0,n)} > |\lambda_{(j,n)}|$ is already given by HP1 (§21); this gate must return the gap's **magnitude**, a **regulator-uniform lower bound**, resolvent conditioning, and continuum-limit stability, each with declared uncertainty. Existence is not the deliverable; quantitative, uniform control is (F4).
- **Gate G4 — Metric return.** The complete 6×6 defect-block metric is published with no inserted sector coefficient.
- **Gate G5 — Null classification.** Every null mode is classified analytically into the categories of §17 or retained as a declared physical degeneracy. A merely counted null space fails this gate (F8).
- **Gate G6 — Refinement compatibility.** HP3 passes numerically and analytically on the inherited coarse subspace, per block (§§24–27).
- **Gate G7 — Terminal-record pushforward and generation geometry. Status: OPEN overall**, with one conditional sub-result:
- `⟦ G7 canonical wheel quotient arithmetic: PASS ⟧`
- `⟦ G7 physical terminal-record pushforward: OPEN ⟧`

  The §31.1 requirement is representative independence, cylindrical compatibility, strong lumpability **and** descent of $\mathcal{U}_\Omega$ from the polar refinement transport. The quotient arithmetic and layer provenance address the *partition*; they do not supply the *transport descent*. Sub-gates:

- **Gate G7a — Canonical quotient arithmetic.** Antipodal quotient, quotient kernel, stationary law, K_4 conductance, hub Schur complement and zero-holonomy Laplacian reproduced exactly (§32A). **Status: PASS under the inherited three-equivalent-two-state-layer structure.**

- **Gate G7b — Terminal-generation provenance.** The three-layer count and antipodal identification are derived within the inherited anomaly-descent/census chain (Theorem 32A.0). **Status: PASS conditional on QN-3/SC-1** — this does not close G7 overall.
- **Gate G7c1 — Generation carrier and selector typing.** Sector selectors inherited from the terminal-record ledger ($\mathscr{H}\_\Omega = \mathbb{C}^3\_\mathrm{gen} \otimes \mathscr{H}\_\mathrm{role}$), not chosen from masses or mixing. **Status: PASS (form).**
- **Gate G7c2 — Isotropic local form $W\_H = w\_H\, I_3$. Status: PASS (form).** The *magnitude* w_H (canonically 75/434, not the anomaly-descent 8/3) belongs to G4/history-content.
- **Gate G7d-1 — Canonical chain arithmetic and non-annihilation.** $q\_{\star}\gamma_6 = 2\gamma_3$ with $2^{-1} = 4 \pmod 7$ gives $k\_\Omega = 4k\_\mathrm{vac}$, so the degree-two quotient cannot annihilate the Gate-3 class and $k\_\mathrm{vac} \neq 0 \Longrightarrow k\_\Omega \neq 0$ (§32A.7). **Status: PASS.**
- **Gate G7d-2 — Physical cochain basicness and polar-transport descent.** The decisive premise $\rho\_\mathrm{vac} = q^\wedge\star\rho\_\Omega$ (antipodal basicness) must be established on the physical, symmetry-broken background, together with descent of $\mathcal{U}\_\Omega$ from the polar refinement transport (§31.2). The chain arithmetic does not prove it. **Status: OPEN (F29).**
- **Gate G7e — Occupancy and charge.** The closure phase must be shown to record primitive commitment faithfully, excluding k = 0 on the committed-Fact branch. Separately, if the unit charge |k| = 1 is claimed as a *substrate* fact rather than a determinant ordering, the Primitive-Fold Unit-Step statement $\Delta\theta = \pm 1$ must be returned independently — the two may not be quoted as mutual support (§44A.2). **Status: faithful nonzero occupancy conditional; unit value OPEN.**
- **Gate G7f — Gaussian-to-Perron promotion.** The common-measure Gaussian determinant ordering (Theorem 44A.1) must be shown to equal the branch-dependent contribution to the leading Perron exponents $-(1/a\_t)\log \lambda_0(P\_k\, \mathbb{T}^\wedge\mathrm{hist}\, P\_k)$ of *one globally normalised positive history operator* before it may be called physical branch selection. A common Gaussian **stiffness** direct sum is not, by itself, that comparison (Firewall H10). **Status: OPEN** (this is the subsector instance of G9b/G9c).
- **Gate G7h — Terminal-map derivation or output-invariance.** Either the upstream construction **uniquely determines** r_n, or **every** admissible map passing the projection residuals of §56A.7 returns equivalent Δ_Ω, R_Ω and physical boundary. The residual tests validate or reject a candidate; they do neither of these. **Status: OPEN (F40).**
- **Gate G7g — Sector/generation typing.** The three conjugate nonzero $\mathbb{Z}_7$ sector pairs {±1}, {±2}, {±3} must be shown to correspond — or not — to the three physical generations. The calculation supplies *sectors*; the generations are the eigenvalues *within* one sector (F31). A proof must exhibit four things: (i) **a map**, an explicit assignment from sector labels to generation labels (or a proof that none exists); (ii) **typing**, that the assignment respects the carrier types — sectors index components of the branch algebra, generations index the $\mathbb{C}^3\_\mathrm{gen}$ factor of $\mathscr{H}\_\Omega$, and these are different objects; (iii) **invariance**, that the assignment is stable under the label permutation $k\_\Omega = 4k\_\mathrm{vac}$ (§32A.7) and under regulator refinement; and (iv) **independence from cardinality**, an argument that does not rest on both sets having three elements. **Status: OPEN.**
- **Gate G8 — Bath–ledger verdict.** The response commutant is published for each relevant class (§§39–42).

- **Gate G9 — Branch selection. Status: OPEN overall**, with one executed subsector result:
- `⟦ canonical Gaussian ℤ₇ subsector comparison: PASS ⟧`
- `⟦ full one-history-operator branch comparison: OPEN ⟧`

  The master history operator, its invariant component decomposition, one reference measure, the **raw block-Perron** free-energy table and the selection gap must be published together. A common Gaussian **stiffness** direct sum is not, by itself, this gate (Firewall H10). Sub-gates:

- **Gate G9a — Common operational carrier and measure.** One common quadratic probability measure across the substrate's admissible coherence channels, with cross-channel bridging forcing one normalisation and no literal superselection obstruction to global normalisation (inherited from the operational-measure theorem). **Status: PASS conditional on that upstream theorem.**
- **Gate G9b — Canonical $\mathbb{Z}_7$ holonomy-sector comparison.** The common-measure Gaussian determinant ordering (Theorem 44A.1) is executed and structural. **Status: CONDITIONAL PASS** — a stiffness/determinant comparison, not yet a block-Perron calculation until the corresponding one positive transfer is constructed.
- **Gate G9c — Exhaustive full-branch embedding in one history operator.** CFB-1/2/3 and the full MASD branch family must be exhibited as components $\{P_\beta \mathscr{H}^{hist} P_\beta\}$ of *one derived generator* — arising from spectral/invariant projections, not placed beside each other after construction. **Status: OPEN, with a structural candidate.** §43A.6 defines branches as the communicating components $\pi_0(G_n)$, so the decomposition and commuting projectors arise by construction rather than assembly (Firewall H10 satisfied). But that proves only that components *are* branches, not that every branch is a component: closing the gate requires $\mathfrak{z}_n = Z(\mathrm{Comm}(\mathfrak{B}_n)) = \ell^\infty(\pi_0(G_n))$ and the representation of every MASD/CFB label in that algebra, plus the enumeration of G_n (F32).
- **Gate G9d — Raw block-Perron evaluation and global dominance.** The raw block Perron exponents on a common action scale, plus the global-dominance condition (§43A.5). **Status: OPEN.**
- **Gate G10 — Regulator convergence.** The physical outputs, including Δ_{sel}, are stable across at least two regulator levels within a declared error model.
- **Gate G11a — Generator attribution.** Does L^{op} restricted to the closure catalogue equal the graph Laplacian $I - \mathcal{S}_n$? This is an *inherited-attribution* question rather than a computation, and it is recorded as such in the §55A ledger. **Status: OPEN, upstream.**
- **Gate G11b — Kernel relation, given G11a.** Given G11a, the relation is a theorem, not an open question: **the Doob-normalised graph heat transfer on the closure catalogue is exactly $\exp(-a_t(I - T_0))$**, the continuous-time Markov semigroup generated by T_0 (**verified to machine precision**). It therefore has the *same eigenvectors, the same stationary measure* $\pi = (7,4,4,4,4,4,4)/31$ *and the same D_6 decomposition*; only the eigenvalue map $\lambda \mapsto e^{-a_t(1-\lambda)}$ differs. This **explains** Lemma 14C.5 rather than merely recording it: the structural invariants survive because the two kernels share eigenvectors and symmetry, so the representation-orthogonality zeros are automatic rather than fortunate. **Status: PASS.**

Failure of any gate prevents the physical certificate.

51. The Physical Certificate

HMGBC-1 Physical Certificate — fields: regulator levels; package hashes; history support; HP1 / HP2 / HP2R / HP3 status; W_prim block matrix; null dimension and classification; universal conductance; Δ_Ω; R_Ω; R_u, R_d, R_e, R_ν; colour bath verdict; weak bath verdict; Abelian ledger verdict; master-operator component structure and reference-measure hash; physical branch; selection gap; convergence norm; unresolved residual; certificate verdict ∈ {PASS, FAIL, DEGENERATE, NOT ISSUED}. **The current paper issues no filled numerical physical-result certificate, because HFB-1 Stage A has not yet been executed.**

Certificate versus admission. The absence of a filled physical-result certificate does not prevent admission of the sealed numerical execution defined in §56A. This certificate records the *outcome* of HFB-1 Stage A; it is not the prerequisite for beginning that stage. Before execution the status is **NOT EXECUTED**; after execution it becomes PASS, FAIL or DEGENERATE. The separate prerequisite for beginning is publication of the immutable construction manifest 𝓜_hash and the declared regulator inputs.

52. HFB-1 Two-Stage Execution Contract

HFB-1 is admitted by the sealed deterministic package of §56A, **not** by the prior existence of its numerical outputs. Its execution has two consecutive stages. Requiring the handoff table to be populated *before* opening would demand as a precondition exactly what Stage A is admitted to compute.

52.1 Stage A — HMGBC numerical execution

Stage A applies the immutable map 𝓕_HMGBC : 𝔖_HMGBC ⟶ 𝓗_HMGBC^out at no fewer than two regulator levels and calculates

```
⟦ { W_prim, [W_AB], P_info, C_Ω, Δ_Ω, R_Ω, 𝓡_Ω, R_u, R_d, R_e, R_ν, 𝔅_phys,
Δ_sel }, ⟧
```

together with the Perron, bath-response, null-classification, terminal-pushforward, component-decomposition and regulator-convergence certificates. **The numerical HMGBC return is therefore an output of HFB-1 Stage A, not a precondition for opening HFB-1.** Stage A returns PASS, DEGENERATE or FAIL. A PASS or admissible DEGENERATE populates the handoff table below; a FAIL terminates HFB-1 and does not permit downstream boundary execution.

52.2 Stage-A handoff table

The completed Stage-A return contains: $\beta_\star$, or a regulator-stable degenerate component set, with its **central-component identifier**; W_prim including the complete 6×6 **typed operator-block** ledger [W_AB] (T and B entries carrying restored doublet indices, §14F.4); the classified information-null projector P_info (distinct from the upstream kinematic quotient P_kin, §56A.3); C_Ω, Δ_Ω, R_Ω **and its positive contraction $\mathcal{S}_\Omega$** (§33); the oriented completion map $\mathcal{R}_\Omega$; the sector compressions R_u, R_d, R_e, R_ν; colour, weak and Abelian bath–ledger verdicts; the common transfer Perron package (λ_0, ℓ, r); the central component decomposition and common-reference-measure hash; the selection gap Δ_{sel} or the certified degeneracy; and the two-or-more-level regulator error and convergence package.

No row may be filled with a canonical-wheel magnitude, CFB benchmark, identity completion, measured Standard Model parameter or fitted substitute (Firewalls H1, H4, H9; F18, F25).

52.3 Stage B — Physical boundary execution

Only after Stage A returns PASS — or a DEGENERATE result for which **every** surviving component gives the same physical boundary within the declared tolerance — does Stage B execute the frozen MASD-1 boundary protocol (§§48–53 there), returning

$$\mathcal{P}_{SM}^{VERSF}(\mu_\star) = \{ g_s, g, g', \mu_H^2, \lambda, Y_u, Y_d, Y_e, \mathcal{B}_\nu, \bar\theta \}.$$

If distinct members of a degenerate branch set produce **distinct** physical boundaries, HFB-1 returns DEGENERATE and Renormalised Parameter Closure remains locked (§56A.3). No experimental result may break such a degeneracy (HM8, §53; F18).

52A. The Strong-Phase Decomposition — Typed and Non-Double-Counting

The strong phase $\bar\theta$ in the boundary set is a single physical invariant, not a free sum of three phases, and HMGBC-1 must hand HFB-1 a typing that makes it checkable rather than a formula that risks counting the same information twice. This section supplies that typing; it is the companion's fix to the determinant-line double-counting hazard.

52A.1 The invariant

The physical, basis-independent strong phase is the two-term object

$$\bar\theta = \theta_{top} + \arg\det(\mathbb{M}_u\, \mathbb{M}_d) \pmod{2\pi},$$

where θ_{top} is the coefficient of the topological density in the substrate history measure — the VERSF analogue of the QCD vacuum angle, returned by the substrate-cohomology / global-measure structure — and $\mathbb{M}_f = (v_{cl}/\sqrt{2})\, Y_f$ are the physical mass operators built from the history-derived Yukawas. This two-term combination is the only physical content. Any further additive term must be shown to vanish or to be a rewriting of these two.

52A.2 What the determinant line is, and is not

The determinant-line construction is the *framework* in which arg det($\mathbb{M}_u$ $\mathbb{M}_d$) is globally well defined across the faithful-group orbit stack: it fixes the reference section against which the determinant phase is measured. It is therefore **not** an independent additive third phase. The three-term writing $\bar{\theta} = \theta_{gauge} + \arg\det(Y_u Y_d) + \theta_{line}$ is admissible only under the identification $\theta_{gauge} \equiv \theta_{top}$, $\arg\det(Y_u Y_d) \equiv \arg\det(\mathbb{M}_u \mathbb{M}_d)$ (up to the real positive Jacobian $v_{cl}/\sqrt{2}$, which carries no phase), and $\theta_{line} \equiv$ the global-section bookkeeping that makes the determinant term single-valued — not a residual holonomy to be added. On a trivial determinant line, $\theta_{line} = 0$ by construction and the invariant reduces to the two terms.

52A.3 The three admissible fates of θ_line

θ_{line} must be classified into exactly one of:

1. **Zero** — the determinant line is globally trivialisable; $\bar{\theta} = \theta_{top} + \arg\det(\mathbb{M}_u \mathbb{M}_d)$ with no third term. Expected generic case.
2. **A global anomaly obstruction** — the line is nontrivial and admits no global section; θ_{line} is then not a number to add but the statement that arg det is not globally single-valued, and $\bar{\theta}$ is defined only relative to a declared section, with the obstruction class published. This is the "global obstruction" alternative of MASD-1 F24.
3. **A genuine residual physical holonomy** — distinct from both θ_{top} and the determinant phase. **Forbidden unless accompanied by a proof that the holonomy is not already contained in θ_{top} or arg det($\mathbb{M}_u$ $\mathbb{M}_d$).** Absent that proof, case 3 is a double-counting error, not a physical phase.

52A.4 Non-double-counting certificate (mandatory)

The strong-phase row is populated only when the branch publishes: (i) θ_{top} from the topological sector of the measure alone, with the Yukawa sector set to a real reference; (ii) arg det($\mathbb{M}_u$ $\mathbb{M}_d$) from the Yukawa determinants alone, on the fixed determinant-line section; (iii) an explicit demonstration that θ_{line} falls into fate 1 or 2 — i.e. that no phase information is counted in both (i) and (ii). A returned $\bar{\theta}$ requiring a nonzero fate-3 θ_{line} without the independence proof fails the certificate and does not populate the boundary table.

52A.5 Consistency with positive pushforward

If the positive-pushforward ontology holds, it constrains the *invariant*: a positive pushforward excludes both generic complex phases and the signed $\bar{\theta} = \pi$ case, so

```
⟦ positive pushforward  ⇒  θ¯ ∈ { 0, global obstruction }. ⟧
```

The typing of §52A.1–§52A.4 is exactly what makes this checkable: without separating θ_{top}, the determinant phase and the section bookkeeping, "$\bar{\theta} = 0$" cannot be verified without the risk that a hidden phase has merely moved between terms. HMGBC-1 supplies the source (θ_{top} from the history measure; the oriented completion $\boldsymbol{\mathcal{R}}_\Omega$ from which arg det $\mathbb{M}_f$ descends); HFB-1 performs the combination under this certificate.

53. Non-Insertion and Blind Evaluation Protocol

Prohibited before the HMGBC package is frozen: choosing bath frequencies to fit a mass hierarchy; changing defect–bath couplings to repair hypercharge; choosing a generation graph from CKM or PMNS; deleting null modes because they spoil rank; normalising branches separately before comparison; changing the branch parameterisation after seeing f_β; constructing the master-operator component structure to place a preferred branch at the minimum; using measured Standard Model numbers as priors; selecting an attachment channel from the observed atmospheric octant; forcing the bath verdict because non-Abelian forces are known to exist. Experimental data may enter only after the physical certificate and HFB-1 boundary are frozen.

Part VIII — Closure

54. Falsification Conditions

HMGBC-1 fails physically if any of the following occurs.

- **F1** — No finite unresolved-record bath reproduces the declared memory response (HM1 fails).
- **F2** — The bath-resolved transfer is not cone positive (HP1 fails).
- **F3** — The physical quotient retains a negative history-rate mode.
- **F4** — The leading Perron root, though *strictly dominant at each finite level by HP1* (§21), fails to be **quantitatively controlled**: no numerical lower bound, ill-conditioned reduced resolvent, non-uniform in regulator level, or vanishing in the continuum limit. **Status: existence is automatic; magnitude and regulator-uniform control are the required G3 returns.**
- **F5** — The long-history information rate does not exist (couples to F4: no *uniform* gap, no Cesàro limit).
- **F6** — W_{prim} depends on a coordinate or defect normalisation not fixed upstream (reference geometry leaks into the metric — the §4 circularity), **or W_{core} is chosen rather than derived, so $\hat{\mathbb{H}}_n^{hist}$ is not fully fixed (§19, F19).**

- **F7** — Cross-defect blocks require inserted coefficients (contradicts Theorem 16.1a; would fire only if the bath response failed to source them).
- **F8** — Numerical nulls cannot be classified and alter physical rank; **status: the null-classification obligation of §17 is a hard gate (G5), and this is the history-measure face of MASD-1 F28.**
- **F9** — HP3 fails on the inherited coarse subspace (a block condition of §§24–27 fails).
- **F10** — The universal completion geometry does not define typed sector maps, **or the terminal-record conductance C_Ω is not the representative-independent, cylindrically compatible, strongly lumpable pushforward of the full-configuration flow (Theorem 31.1 fails). Status: this is the projection gate; on the canonical wheel the arithmetic is discharged (antipodal map unique and lumpable, §§32A, 44A, Gate G7a PASS) and the provenance is derived upstream (Theorem 32A.0, Gate G7b PASS conditional on QN-3/SC-1), leaving the holonomy descent (Gate G7d) open. Fires if the upstream layer provenance does not hold, i.e. if the antipodal partition is imposed rather than inherited.**
- **F11** — The history law returns only the identity completion and hence no route to hierarchy, unless HFB-1 nevertheless returns the physical boundary (the ker Δ_Ω corner, §36).
- **F12** — The bath response is ledgered where the non-Abelian gauge-origin branch requires a saturated bath (colour/weak class fails §41–§42).
- **F13** — The branch free-energy minimum does not exist.
- **F14** — The minimum is degenerate but a single branch is nevertheless selected (violates §46).
- **F15** — The selected branch changes discontinuously under regulator refinement (§47).
- **F16** — *(Comparability.)* The determinant or normalisation convention differs across branches, **or the branches are not exhibited as components of one master history operator on one reference measure (HM9/§43A.1–4 fails)**; the free-energy comparison is then not physical. **Status: half of the branch-comparison gate, open — a per-branch Perron gap does not close it (§43A.2).**
- **F16b** — *(Global dominance.)* Even with one operator, the Global Dominance Condition (§43A.5) fails — the reference state misses the dominant component, normalisation is per-block rather than global, or the blocks are permanently coexisting superselection worlds — so thermodynamic dominance does not imply ontological selection. **Status: the other half of the branch-comparison gate, open.**
- **F17** — HFB-1 cannot reproduce the same W_prim and $\mathcal{R}$_Ω from the frozen package.
- **F18** — A Standard Model observable was used before package freeze.
- **F19** — *(History-content gate.)* The **complete generator ℍ̂_n^hist** — its local block W_core as well as (Ω²_(Q,n), C_n) — is not returned by the substrate memory-kernel and operational-geometry programmes but chosen in any block; then HMGBC-1 has relocated the freedom, not removed it (§19). **Status: OPEN. The canonical-wheel *benchmark* is executed — T_0, its stationary measure and W_wheel are computed, not chosen (§§14A–14D) — but the history-content *gate* has not passed: it awaits the G11a generator attribution, the physical defect embedding (F26) and the tilt derivation (F34), on top of the lift to the full typed, multi-cell generator.** Because W_prim = W_core(0) − C_n†(Ω²_Q)⁺C_n, a derived bath pair with a chosen W_core still fires this at full scope.

- **F20** — *(Strong phase.)* $\bar{\theta}$ requires a fate-3 residual θ_line without an independence proof, i.e. the determinant-line term double-counts the topological or determinant phase (§52A).
- **F21** — *(Fermionic positivity.)* The full MASD transfer fails cone positivity because the fermionic determinant is complex or signed, and neither route (a) real-nonnegative determinant nor (b) certified reflection positivity of §20.2 is supplied; then the bosonic HP1 does not extend to the physical measure and F20/F21 are the same event.
- **F22** — *(Gaussian-sector overreach.)* A Schur-complement return $W\_prim = W\_core - C^\dagger(\Omega^2)^+C$ is quoted outside the certified bilinear-coupling regime, where non-Gaussian bath elimination adds higher cumulants that Theorem 16.1b omits (§16, scope caveat).
- **F23** — *(Defect-lift faithfulness.)* The six wheel/internal history tangents are not the faithful compression of the full vector-valued MASD defect differentials — any of HM10's four criteria (sector match, rank, equivariance, support typing) fails, per §14E; then W_wheel is a symmetry caricature and even its structural pattern cannot be promoted (§14B.1, §14E, Corollary 16.4). **Status: open. Appendix I verifies all four *graph-lift analogues* for the wheel closure score on a two-cell lift, so the checklist is not empty; the physical sector-match and equivariance obligations for the actual $\mathbb{Z}_7$-valued closure differential, the remaining five defects, and full equivariance are all open.**
- **F24** — *(Illegitimate cross-block completion.)* A nonzero W_BH, W_CH or W_RH is reported without deriving it from the multi-cell / internal-representation / terminal-record / completion lift — i.e. the wheel-exact zeros of Corollary 14C.1 are filled in by hand (violates Firewall H9). Equally, any rescaling of W_wheel entries to fit a target fires this.
- **F25** — *(Wheel-magnitude and within-irrep misuse.)* Any rational magnitude of W_wheel (a diagonal entry, a cross-block value, an eigenvalue, the condition number) is quoted as a physical prediction or fed to HFB-1 as an input, rather than only the choice-independent content — full Fisher rank and exact inter-irrep vanishing. Equally fires if a within-irrep cross block (W_JH, W_TB), its sign, or its non-vanishing is quoted as structural without qualification: these depend on the declared generator family (§14C.3, Thm 14C.6). **Exception:** the J–H *sign* may be quoted wherever support typing is established, since diagonal support with uniform sign forces it for every admissible weight (Corollary 14C.4b). No such exception exists for W_TB. The magnitudes are properties of the canonical hub↔rim weight 1/7 and change under admissible reweighting; treating them as physical is a category error, the metric-side analogue of a cross-operator branch comparison.
- **F26** — *(Score-generator non-uniqueness.)* The six score generators F_A of §14B are treated as uniquely derived when in fact HM10.3 (equivariance, §14E) has not been proved; then §14B's tangent family is one natural symmetry-adapted representative set presented as *the* physical lift. Until §14E item (3) is discharged, the choice is well-motivated but not unique.
- **F27** — *(Projection provenance.)* The three-layer count or antipodal identification is treated as unconditionally derived when in fact it rests on the anomaly-descent/census (QN-3/SC-1) attributions (Theorem 32A.0), which carry their own premises and are taken from papers not re-derived here. Fires if those upstream results do **not** in fact establish a rank-three bidirectional support quotient with layer-exclusive terminality, or if the D_6-uniqueness step is misapplied — i.e. if the provenance is quoted as absolute rather than chain-conditional.

- **F28** — *(Holonomy-selection overreach.)* The $|k| = 1$ ordering of Theorem 44A.1 is quoted as the full branch-comparison result rather than a common-measure Gaussian *determinant* comparison on the $\mathbb{Z}_7$ holonomy subsector (a stiffness comparison, not yet a transfer-operator Perron comparison). Extending it to CFB/MASD branches, non-isotropic w_H, non-Gaussian corrections or ontological (global-dominance) selection, or asserting it discharges §43A comparability without the promotion construction, fires this. The sign degeneracy $\{+1,-1\}$ being resolved from observed CP data (rather than a derived commitment orientation) also fires this (HM8, §53; F18).
- **F29** — *(Holonomy descent.)* The $\mathbb{Z}_7$ flux is placed on the generation triangle without proving the descent $H^1(\Gamma\_vac; \mathbb{Z}_7) \to H^1(\Gamma\_gen; \mathbb{Z}_7)$ — i.e. that the descended polar record transport carries the Gate-3 class onto the terminal-record cycle with Wilson phase $2\pi k/7$ (Gate G7d), or without certifying Gate-2 detwisting and elementary flatness so that κ is defined at all. The §32A/§44A spectra are exact given the descended class; asserting the placement without the descent fires this.
- **F30** — *(k-reading conflation.)* The determinant ordering of Theorem 44A.1 is quoted as *selecting* $|k| = 1$ when the microscopic phase-assignment law in fact fixes a nonzero k **before** branch comparison — in which case the calculation evaluates that sector and does not override it (§44A.2). Equally, quoting the Primitive-Fold Unit-Step $\Delta\theta = \pm 1$ and the determinant ordering as mutual support fires this: they are independent, and faithful occupancy excludes only $k = 0$, not $|k| \neq 1$ (Gate G7e).
- **F31** — *(Sector/generation typing.)* The three conjugate nonzero $\mathbb{Z}_7$ sector pairs $\{\pm1\}$, $\{\pm2\}$, $\{\pm3\}$ are identified with the three physical generations. They are **holonomy sectors, not generations**: the three generations are the eigenvalues $\lambda\_m(k)$ *within* one sector k. The corpus does not yet prove any correspondence between the sector pairs and the generation triple, and asserting one — or letting the coincidence of the number three do the work — fires this.
- **F32** — *(Admissibility-graph closure.)* The Canonical Admissibility Principle (HM11) fails on the complete MASD/CAR/Fock move graph — some admissible elementary move carries a substrate-derived weight unequal to that of a symmetry-equivalent move — or the move graph is not finitely enumerable at the frozen regulator. Then §11A's raw transfer is not the derived object it claims to be, and the recovery of T_0 (Theorem 11A.1) is a canonical coincidence rather than a derivation.
- **F33** — *(Holonomy label conflation.)* A winding statement is quoted without declaring whether k denotes the quotient-triangle class k_Ω or the upstream vacuum-cycle class k_vac. Under the degree-2 antipodal covering $k\_\Omega = 4k\_vac \pmod 7$, so $|k\_\Omega| = 1 \Leftrightarrow |k\_vac| = 2$ and the conjugate pairs permute (§32A.7). Reporting "$|k| = 1$" in vacuum labels, or transferring a primitive-Fold statement about k_vac onto the generation spectrum without the factor, fires this.
- **F34** — *(Tilt-family choice.)* The exponential tilt of §14B/§56A.5 is treated as canonical when HM12 is undischarged. Because Fisher information transforms as $W \mapsto J^TWJ$ under reparameterisation, and W_prim is consumed as a quadratic form on physically normalised defects, a different admissible embedding returns a congruent but numerically different metric — so the radial/normalisation part of the embedding changes the physics and is not gauge. Fires if the tilt family is asserted forced without proof.
- **F35** — *(Kernel identification.)* The executed results on T_0 are quoted as results on the physical transfer without discharging Gate G11a. Given G11a the relation is exactly the

semigroup statement of G11b — same eigenvectors, same π, same D_6 content, eigenvalues mapped by $\lambda \mapsto e^{(-a_t(1-\lambda))}$ — so the structural invariants transfer but the magnitudes do not. **The magnitude gap cannot be closed by tuning the transfer step:** across the whole heat family W_{CC} rises from 0 as $a_t \to 0$ to a maximum ≈ 0.2457 near $a_t \approx 2.21$ and falls to $192/961 \approx 0.19979$ as $a_t \to \infty$, so the discrete-step value $12/31 \approx 0.387$ lies **outside the attainable range for every a_t** (verified). Quoting any T_0 magnitude as physical fires this, as does describing the subcertificate as executed on the physical kernel rather than on a diagnostic one.

- **F36** — *(Tangent-space incompleteness.)* An admissible history deformation exists that couples to the unresolved bath at the declared order and is **not** in the image of the six-defect tangent map — i.e. one of the 16 unlabelled admissible directions of §14G is physically active. Then HM4 fails and W_prim prices only a subspace of the admissible deformations.
- **F37** — *(Cone positivity from Hilbert positivity.)* The physical transfer is asserted cone-positive on the ground that $L_n^{op} \geq 0$ as a quadratic form. This is invalid: $L = [[1,1],[1,1]]$ is positive semidefinite while $\exp(-tL)$ has negative off-diagonal entries for all $t > 0$. Fires unless L_n^{op} is certified to generate a positivity-preserving (Dirichlet / sub-Markov) semigroup — nonpositive off-diagonal entries, the Beurling–Deny criterion — and, separately, unless positivity *improvement* is obtained from an irreducibility theorem rather than assumed.
- **F38** — *(One-step Fisher on a non-Markovian law.)* The one-step formula $W_{AB} = \mathbb{E}_{(\pi K)}[S_A S_B]$ is used for the reduced process without declaring one of the four conditions of §14. Routes (a) and (c) are the same mechanism — centred scores are a martingale difference sequence, so cross-time covariances vanish and asymptotic variance equals one-step variance — and (a) holds on the enlarged carrier but not on the reduced one. Since bath elimination makes the resolved law depend on the whole past, the one-step expression is otherwise a Markov-sector approximation quoted as an identity.
- **F39** — *(Unbridged Fisher–Schur identification.)* The Gaussian influence action of Appendix A is identified with the log-likelihood ratio whose path-space KL Hessian defines $\mathcal{J}_n$, without proving that **Doob normalisation and Gaussian bath elimination commute**. They do so only if the enlarged Perron function factorises as $r_X(x) \cdot r_Q(q)$, which fails whenever $C_n \neq 0$; absent a commutation proof or a bound on the discrepancy, the two constructions are distinct objects (§16). HM12 makes the coordinate-normalisation half non-cosmetic.
- **F40** — *(Terminal map validated but not derived.)* The residual tests of §56A.7 are treated as *deriving* r_n. They can only reject a candidate: they do not show the upstream construction uniquely determines r_n, that no other admissible map passes, that two passing maps return the same Δ_Ω, or that the search is exhaustive. Absent a uniqueness-or-invariance theorem on the full carrier, the projection has been moved into the sealed package rather than derived.
- **F41** — *(Exact superselection assumed.)* Branches are taken to be exact reducing components at finite regulator when the physically significant sectors are nearly-invariant metastable subspaces that separate only in an infinite-volume or continuum limit. Fires unless exact finite-regulator superselection is proved, or the framework is extended to

asymptotic ergodic components / low-lying metastable subspaces with a declared separation criterion.

- **F42** — *(Sealed read as derived.)* A frozen, hashed convention is presented as physically forced. Pre-registration prevents tuning; it does not confer derivation. Each convention of §56A.11 must be proved forced by upstream axioms, proved immaterial across all admissible alternatives, or retained as an explicit conditional premise — and the "no further constitutive choice" claim must be read as "no further choice beyond those premises."
- **F43** — *(Free-energy normalisation and convergence overreach.)* Branch free energies are compared across regulator levels without declaring the extensivity convention (per-cell normalisation, volume division, vacuum subtraction, or a density limit), since log λ_0 is generally extensive on a multi-cell regulator. Equally, agreement at two regulator levels is quoted as a convergence certificate without an analytic a priori bound, a monotonicity theorem, or a controlled Cauchy estimate.

55. The HMGBC-1 Theorem

Theorem 55.1 — VERSF History-Measure, Generation-Geometry and Branch-Closure Theorem.

Let MASD-1 supply the finite-regulator resolved configuration space, six primitive defects, terminal-record carrier and common physical quotient. Assume HM1–HM12, the terminal-layer provenance of Theorem 32A.0 (inherited from the QN-3/SC-1 chain), and the admission conditions of Parts II–IV. Then:

1. the unresolved sub-threshold record modes define an enlarged bath-resolved finite-regulator history process;
2. the Euclidean transfer is cone positive and, on each irreducible physical component, has a simple positive left/right Perron pair (automatic from positivity improvement), whose *quantitative* gap magnitude and regulator-uniform control are the separate required return (§21);
3. eliminating the bath produces a non-Markovian resolved history law;
4. the long-history Radon–Nikodym rate has positive-semidefinite defect curvature;
5. that curvature returns the complete MASD operational metric W_{prim} — including every cross-defect block and the common action unit — **relative to the fixed, physically normalised defect embedding and tilt family**; the return becomes unique only once HM12 is discharged or output-invariance across all admissible embeddings is proved (F34);
6. exact gauge, record-coboundary and terminal-null directions lie in the information kernel — and any further kernel direction must be classified, not assumed absent;
7. the symmetric stationary flow defines one covariant terminal-record Laplacian Δ_Ω;
8. the resolvent $R_\Omega = (W_H + \Delta_\Omega)^{-1} W_H$ defines the universal generation/completion geometry;

9. every sector completion operator is a typed compression of the one oriented map **$\mathcal{R}$**_Ω = V_H† R_Ω;
10. the commutant of the bath-response algebra decides whether each identical-sector class is ledgered, partially bathed or saturated;
11. **given the Common-History-Operator Obligation (§43A)**, the commonly normalised block Perron roots of one master operator define a physical branch set by free-energy minimisation;
12. blockwise refinement intertwining gives a projectively consistent history family and a regulator-stability test.

No measured Standard Model parameter is required by any construction in the theorem. A physical HMGBC pass additionally requires the finite-regulator spectrum, couplings, matrices, the master-operator component structure, the branch gap and the convergence package of §§48–51. **Items 1–10 and 12 are H1 mathematics given their premises; item 11 is H1 given §43A, whose antecedent is an H3 obligation not discharged here.**

55A. Inherited-Dependency Ledger

Three gates have been advanced by the same move: **relocation upstream**. That is legitimate and it is progress, but it means the paper's conditionality is real and distributed, and it should be declared in one place in the paper's own voice rather than reconstructed by a reader from nine sections. Each row records an object taken from an upstream result, and what fails if the attribution does not hold.

Object taken from upstream	Used in	If the attribution fails
Rank-three bidirectional physical-support quotient; layer-exclusive terminality (QN-3 / SC-1 anomaly-descent and census chain)	Theorem 32A.0 — three-layer count and antipodal identification	The terminal-record provenance reverts to an assumed premise; G7b reopens; Theorem 32A.1's uniqueness loses its selecting hypothesis (F27)
One common quadratic operational measure across admissible channels, with cross-channel bridging and no superselection wall (operational-measure theorem)	G9a; §43A.3 arena closure; the common normalisation of §44	Branch free energies are not commensurable; §43A's comparability half reopens (F16)
$\mathbb{Z}_7$ coefficient group as the rotation kernel of the D_7 transport group (closure-transport programme)	§32A.4 condition 1; the entire holonomy split	The coefficient group becomes a choice; §32A.4–6 and §44A lose their derived character (F29)
Non-bounding protection of Fact-isolating cycles under refinement; faithful nonzero occupancy (Gate-3 results)	§32A.4 conditions 3 and 5; Theorem 32A.5; the exclusion of k = 0	k ≠ 0 is no longer forced; the nondegenerate generation geometry becomes conditional on an unproved occupancy (G7e)

Object taken from upstream	Used in	If the attribution fails
Degree-corrected stochastic coarse-graining operator $\frac{1}{2}(I + D^{-1}A)$ with its selection within the declared family	§47A; the uniform-gap bound	The regulator family is not canonical; §47A's intertwining hypothesis loses its source (G10)
L_n^op, v̂_n^op, Π_pers^(n), 𝔄_n^kin genuinely inherited and **complete on the full MASD/CAR/Fock carrier**	§56A.1–56A.2; the entire sealed package	The sealed map 𝓕_HMGBC is not fully specified; admission under Theorem 56A.1 is not established (§56A.11)
L_n^op restricted to the closure catalogue equals the graph Laplacian I − 𝓢_n	Gate G11a; the bridge from the executed benchmark to the physical kernel	The semigroup relation of G11b has no antecedent, and the executed wheel results lose their connection to the physical transfer entirely (F35)
L_n^op generates a **positivity-preserving** (Dirichlet / sub-Markov) semigroup, not merely L^op ≥ 0	§56A.2 cone positivity; HP1 on the physical carrier	The physical transfer is not cone-positive; HP1–HP2 fail on the full carrier and the Perron/Doob construction collapses (F37)
MASD-1 defect definitions, terminal-record address construction, physical evaluation protocol	throughout	HMGBC-1 has no companion to serve

The pattern, stated plainly. HMGBC-1 does not eliminate the Standard Model's constitutive freedom; it relocates that freedom upstream — into the master generator, into the inherited operational data, and into the attributions above — and converts what remains into deterministic maps and residual tests. Relocation is progress precisely because it makes the freedom finite, named and checkable. It is not elimination, and the paper does not claim it is.

56. Closure Certificate

Structurally closed (H1, given premises): enlarged history carrier; bath-resolved transfer architecture; Markovian-to-non-Markovian reduction; finite-history Radon–Nikodym action; long-history information-rate definition; the general metric identity W_prim = D²_**D** 𝒥(0) (Theorem 16.1a) and its Gaussian-sector Schur form (Theorem 16.1b), *stated as two results*; corrected HP1 with its fermionic-phase scope (§20.2) and the HP2/HP2R/HP3 hierarchy *with the gap magnitude, not existence, logged as the G3 return*; conductance/current decomposition; universal record Laplacian; the completion resolvent as a W_H-self-adjoint contraction via 𝓢_Ω (§33) and the oriented map; sector compression rule; bath–ledger commutant audit; the reduction of branch selection to one operator (Theorem 43A.4); the Global Dominance Condition separating thermodynamic from ontological selection (§43A.5); branch free-energy selector; strong-phase typing (§52A); HFB-1 handoff and non-insertion protocol.

Still open (H2/H3, three named gates):

- *History-content gate (§19, F19):* the complete master generator $\mathbb{H}$_n^hist — the resolved local block W_core as well as the K = 7 bath spectrum {ω²_(n,a)} and coupling C_n — must be derived, not chosen, in every block; the physical transfer matrix; the complete null classification (F8); a nontrivial R_Ω; the bath verdicts on colour and weak classes (F12). **The canonical-wheel *benchmark* is executed (see subcertificate below); the history-content *gate* itself remains open, awaiting the G11a generator attribution, the physical defect embedding (F26), the tilt derivation (F34), the Defect-Lift Obligation (F23) and the multi-cell/typed lift.**
- *Projection gate (§31.1, §32A, F10):* **OPEN overall**, with the quotient arithmetic passed. Of the four strongly lumpable rim matchings, the antipodal one is the unique layer-equivalent partition (Theorem 32A.1); its provenance is derived within the inherited anomaly-descent/census chain (Theorem 32A.0, G7b, conditional on QN-3/SC-1); the quotient is K_4, the generation Laplacian is $(8/31)(I_3 - ⅓J_3)$, and the typed-lift *form* passes (G7c1/G7c2). **But the physical terminal-record pushforward remains OPEN** — §31.1 additionally requires descent of $\mathcal{U}$_Ω from the polar refinement transport, which the partition results do not supply — as do holonomy descent (G7d, F29), unit charge (G7e), and sector/generation typing (G7g, F31). The canonical calculation of §32A is exact *conditional on* that pushforward and may not be quoted as closing G7.
- *Branch-comparison gate (§43A, F16/F16b):* the **arena** sub-question is now inherited-closed — one common operational carrier and measure across substrate channels, no literal superselection obstruction to global normalisation (G9a, conditional on the operational-measure theorem). On the canonical $\mathbb{Z}_7$ holonomy subsector a common-measure Gaussian *determinant* comparison is executed (Theorem 44A.1, G9b): nonzero occupancy + isotropic positive H-block orders the sectors so |k| = 1 minimises, structurally (independent of the model-specific w_H), degenerate in sign {+1,−1} — a stiffness comparison, not yet a transfer-operator Perron comparison. **The core remains OPEN (G9c/G9d):** CFB/MASD branches must be exhibited as spectral/invariant components of *one derived generator* (a formal direct sum is empty — Firewall H10), with raw block-Perron evaluation and global dominance (§43A.5). Comparability, the selection gap (F13, F14), and regulator convergence (F15) all still depend on that one construction.

Plus the fermionic-positivity route declaration (F21) and the non-double-counting strong-phase certificate (F20).

Canonical Six-Gate Construction — Certificate.

```
⟦ admissibility-graph recovery of the canonical wheel kernel: EXACT PASS under
HM11 ⟧
⟦ full MASD/CAR/Fock proposal measure and spectrally positive Euclidean master
transfer: OPEN ⟧
⟦ polar orientation-frame and conditional equivariance: DERIVED ⟧
⟦ complete magnitude-bearing six-defect lift: OPEN pending evaluation of all
J_A ⟧
⟦ canonical terminal quotient and chain-level holonomy non-annihilation: EXACT
PASS ⟧
⟦ physical cochain basicness and polar terminal-transport descent: OPEN ⟧
```

⟦ communicating-component branch decomposition: STRUCTURAL CANDIDATE ⟧
⟦ exhaustiveness of the branch algebra and full CFB/MASD embedding: OPEN ⟧
⟦ Gaussian determinant Perron-realisation lemma: refutes an in-principle objection only; does NOT satisfy Firewall H10 and is not a pass ⟧
⟦ identification with the physical master-transfer blocks and global dominance: OPEN ⟧
⟦ canonical refinement theorem: PASS conditional on exact intertwining and a uniform spatial gap ⟧
⟦ physical two-level regulator execution: OPEN ⟧
⟦ eight-direction wheel Fisher calculation (D_6 doublets restored): EXACT PASS ⟧

Accordingly the load-bearing work is concentrated in four actual computations: derive the continuous/full proposal kernel; assemble a spectrally positive physical Euclidean transfer; evaluate the complete J_A and the terminal pushforward; and run two regulator levels while enumerating the full branch algebra.

Canonical K = 7 Wheel Subcertificate — ISSUED.

⟦ kernel T_0 (7×7 stochastic): executed [structural] ⟧
⟦ stationary measure π = (7,4,4,4,4,4,4)/31: exact [model-specific] ⟧
⟦ spectrum {1, ½, ½, −¼, −3/28, 0, 0}: exact [model-specific magnitudes] ⟧
⟦ Perron gap ½: exact [model-specific magnitude; positivity structural] ⟧
⟦ persistent/nonpersistent split (1 ⊕ 6): exact [structural] ⟧
⟦ six scalar defect tangents (C,J,T,R,B,H): constructed, D_6-symmetry-adapted six-scalar slice, independent [structural, uniqueness open — F26] ⟧
⟦ history metric W_wheel: exact, positive, rank 6, κ = 4 [positivity+rank structural; κ and entries model-specific] ⟧
⟦ within-irrep cross blocks W_JH = −33/434, W_TB = 3/248: exact [values and non-vanishing depend on the frozen generator choice; the T−B sign likewise; the **J−H sign** is recovered under support typing (Cor. 14C.4b) — F26] ⟧
⟦ inter-irrep zeros W_BH = W_CH = W_RH = 0: exact [choice-independent — the strongest *surviving* claim, but with the force of a symmetry selection rule rather than a history return: it holds for any D_6-invariant form, including a hand-written one; F24 guards against hand-completion] ⟧
⟦ supported Euclidean relaxation spectrum: exact [model-specific magnitudes] ⟧
⟦ closure-score two-cell lift: executed; all four *graph-lift analogues* verified (Appendix I) [physical sector-match and equivariance for the $\mathbb{Z}_7$ closure differential remain open] ⟧
⟦ support-preservation condition (Lemma I.5): extracted as a necessary constraint on physical refinement lifts [added to §48 package] ⟧
⟦ canonical quotient arithmetic: antipodal unique among 4 lumpable matchings by layer-equivalence (Thm 32A.1); K_4; Δ_Ω^(0)=(8/31)(I_3−⅓J_3) — G7a PASS ⟧
⟦ terminal-layer provenance: three-layer count + antipodal identification DERIVED upstream (Thm 32A.0, QN-3/SC-1) — G7b PASS conditional ⟧
⟦ physical terminal-record pushforward (incl. $\mathcal{U}$_Ω descent from polar refinement transport): OPEN — Gate G7 NOT closed ⟧
⟦ three-generation Laplacian (8/31)(I_3−⅓J_3), $\mathbb{Z}_7$-split completion geometry: exact, nondegenerate for k≠0 [conditional on inherited layer-count (Thm 32A.0), w_H lift, and holonomy descent G7d] ⟧

⟦ canonical holonomy determinant ordering: nonzero occupancy + isotropic positive H-block orders sectors so |k|=1 minimises, sign-degenerate {+1,−1} (Theorem 44A.1) — ordering is STRUCTURAL (w_H-independent); a Gaussian stiffness comparison, not yet a §43A Perron comparison; and a *selection* only if k is a branch label rather than fixed upstream by the phase-assignment law (§44A.2, F30) ⟧
⟦ subcertificate verdict: PASS as an executed benchmark; the choice-independent content (full Fisher rank; exact inter-irrep vanishing) is a candidate physical pattern — within-irrep blocks, the T–B sign and non-vanishing are generator-family properties, while the J–H sign is recovered under support typing (Cor. 14C.4b) (§14C.3) — conditional on the Defect-Lift Obligation (HM10, §14E, F23), which is now demonstrated achievable for the closure sector; the projection arithmetic and canonical holonomy branch are exact given the inherited layer provenance (Theorem 32A.0) and the isotropic lift; magnitudes are model properties, not predictions (§14C.3, F25) ⟧

Programme verdict.

⟦ HMGBC-1 canonical-wheel subtheorem: EXECUTED / PASS ⟧
⟦ HMGBC-1 sealed execution rulebook: COMPLETE ⟧
⟦ HFB-1 admission theorem: PASS ⟧
⟦ HFB-1 run effectivity: PENDING publication of $\mathcal{M}$_hash and two regulator inputs ⟧
⟦ HFB-1 Stage-A and Stage-B outcomes: NOT YET EXECUTED ⟧
⟦ Renormalised Parameter Closure: LOCKED pending HFB-1 PASS ⟧

Admission is by **sealed algorithm**, not by completed numerics (§56A) — a criterion on the construction, not on results. **The three constitutive gates are not left as undefined choices: each is converted into a deterministic execution map and a residual test. Their numerical outcomes remain unknown and may return PASS, FAIL or DEGENERATE.**

The **three** gates on which everything turns are named as such, and their status is uneven — which is itself the point. The **projection gate** has its arithmetic and its provenance passing (the latter conditional on the inherited anomaly-descent/census chain), and **two principal debts remain**: physical transport and holonomy descent (G7d), and terminal-map derivation or output-invariance (G7h, F40) — the residual tests validate a frozen map without deriving it. The **history-content gate** has its canonical-wheel *benchmark* executed, but the gate itself is open, awaiting the G11a generator attribution, the physical defect embedding and the tilt derivation as well as the full typed lift. The **branch-comparison gate** has its *arena* sub-question inherited-closed (one carrier, one measure, no superselection wall) and its $\mathbb{Z}_7$ subsector executed, but its *core* — exhibiting the full MASD/CFB family as spectral components of one derived generator, and evaluating the raw block-Perron free energies with global dominance — remains open, and a formal direct sum cannot substitute for it (Firewall H10). So the programme is left with **one central computational construction — the complete master history operator — together with a small number of independent proof obligations that constructing it would not by itself settle**: uniqueness or output-invariance of the terminal map (F40); completeness of the observable algebra (§55A); exact versus metastable branch structure (F41); the physical defect-coordinate and tilt normalisation (HM12, F34); the Fisher–Schur bridge (F39); and regulator convergence beyond the two-level diagnostic (F43). These are derivation obligations, not numerical readings of the assembled operator. Every other mechanism is either proved (H1), executed on the canonical wheel (H3 fragment), inherited-closed from prior papers (conditional on them), or reduced to that one construction.

56A. Sealed-Algorithm HFB-1 Admission and End-to-End Execution

The proper distinction is not between a numerical HMGBC package and an HFB-1 calculation. It is between a **frozen, non-adjustable construction** and the **numerical execution** of that construction. HMGBC-1 must close the first; HFB-1 performs the second.

Requiring every numerical intermediate to be published before HFB-1 begins would duplicate the execution stage inside HMGBC-1 — it would make HFB-1 wait for the calculation it exists to perform. Conversely, allowing HFB-1 to choose an unresolved operator, projection, branch weight or regulator rule would undo the constitutive closure. **The correct admission object is therefore a sealed deterministic algorithmic package, not a table whose numerical outputs are already known.**

56A.1 The sealed HMGBC construction package

At two or more regulator levels $n \in \mathcal{N}\_\mathrm{run}$, define

$$[\![\mathfrak{S}\_\mathrm{HMGBC} = \{ \mathfrak{X}\_n, \hat{\nu}\_n^{op}, L\_n^{op}, \Gamma\_V^{(n)}, \Pi\_\mathrm{pers}^{(n)}, \mathcal{D}\_n, r\_n, \mathfrak{A}\_n^{kin}, V\_n^{ref}, I\_n^{(n+1)}, C\_n^{(n+1)}, \mathcal{T}\_\mathrm{num}, \mathcal{M}\_\mathrm{hash} \}\_{(n \in \mathcal{N}\_\mathrm{run})}.]\!]$$

Its members: $\mathfrak{X}\_n = X\_n \times \mathcal{Q}\_n$, the complete physical finite-regulator carrier (resolved MASD variables, the complete CAR/Fock carrier, unresolved record modes); $\hat{\nu}\_n^{op}$, the one common operational measure, with no branch-dependent coefficient or normalisation; $L\_n^{op} \geq 0$, the inherited operational Dirichlet/transport generator; $\Gamma\_V^{(n)}$, the upstream substrate reference potential, excluding any independently chosen W_prim; $\Pi\_\mathrm{pers}^{(n)}$, the persistent-record projector with $Q\_n = I - \Pi\_\mathrm{pers}^{(n)}$; the full typed defect map $\mathcal{D}\_n$; the terminal-record/address map $r\_n : \mathfrak{X}\_n \to \Omega\_n$ fixed by the occupied-address and terminal-projector construction; $\mathfrak{A}\_n^{kin}$, the frozen physical observable algebra carrying the vacuum, support, polar, bundle, determinant-line, orientation, refinement and flat-holonomy labels; V_n^{ref}, built only from upstream substrate and unresolved-record data; the refinement lift $I\_n^{(n+1)}$ and coarse map $C\_n^{(n+1)}$; $\mathcal{T}\_\mathrm{num}$, all tolerances, differentiation rules, eigensolver conventions, rank thresholds and convergence norms; and $\mathcal{M}\_\mathrm{hash}$, the immutable source, input, code and dependency manifest.

Every unknown physical quantity below is an output of one deterministic map

$$[\![\mathcal{F}\_\mathrm{HMGBC} : \mathfrak{S}\_\mathrm{HMGBC} \longrightarrow \mathcal{H}\_\mathrm{HMGBC}^{out}.]\!]$$

No measured Standard Model quantity occurs in 𝔖_HMGBC.

56A.2 The spectrally positive master transfer

The physical master transfer used for execution is the **operational heat transfer**, not the normalised adjacency operator of §11A:

```
⟦ 𝕋_(0,n)^hist = M_(exp(−a_t V_n^ref/2)) exp(−a_t L_n^op) M_(exp(−a_t V_n^ref/2)). ⟧
```

Positivity requires more than L_n^op ≥ 0. Nonnegativity as a quadratic form gives *Hilbert-space* positivity of exp(−a_t L_n^op), which is not the same as *cone* positivity — the distinction §20 makes and which must not be relaxed here. The standard counterexample settles it: L = [[1,1],[1,1]] is positive semidefinite, yet exp(−tL) has strictly negative off-diagonal entries for every t > 0 (**verified**). What is required is that **L_n^op generate a positivity-preserving (Dirichlet / sub-Markov) semigroup** — equivalently, in the discrete setting, that L_n^op have nonpositive off-diagonal entries, the Beurling–Deny criterion. The inherited operational generator must be certified to satisfy this; it is an additional hypothesis on L_n^op, not a consequence of L_n^op ≥ 0 (F37). The canonical graph Laplacian I − 𝓢_n does satisfy it (**verified**), which is why the wheel construction goes through. Positivity *improvement* is a further step requiring an irreducibility theorem on the physical transition graph, not merely positivity preservation.

Given that certification, exp(−a_t L_n^op) preserves the positive cone, hence 𝕋_(0,n)^hist > 0 and, on its supported subspace,

```
⟦ ℍ_n^hist = −(1/a_t) log 𝕋_(0,n)^hist ⟧
```

is real, self-adjoint and bounded below — **verified**: congruence by a positive diagonal preserves positive-definiteness, so the negative-inertia obstruction of §11A.2 does not arise here. This is the correct resolution of that obstruction: the physical transfer descends from the *operational generator*, not from the normalised adjacency. The §11A construction remains an exact derivation of the canonical K = 7 wheel kernel and a diagnostic of admissible support; **it is not used as the physical Euclidean transfer.**

For the fermionic carrier the execution must invoke exactly one route of §20.2 — real nonnegative determinant, or independently certified reflection/transfer positivity. **Failure of both is an HFB-1 FAIL, not a freedom to alter the transfer.**

56A.3 The bath triple as a compression of the one generator

With P_n = Π_pers^(n), Q_n = I − P_n, decompose ℍ_n^hist into PP, PQ, QP, QQ blocks. Let J_n = D𝓓_n|*(z⋆*,n) be the complete six-defect differential on the physical tangent quotient.

Two typing requirements, both load-bearing. *(i) An explicit embedding.* J_n and its pseudo-inverse act between the configuration tangent space $T_{(z_\star)}\mathcal{C}_n$ and the defect carrier $\boldsymbol{\mathcal{R}}_{prim}^{(n)}$, whereas P_n, Q_n and $\mathbb{H}_n^{hist}$ act on the L^2 history carrier. Expressions such as $Q_n\mathbb{H}^{hist} P_n E_D$ are therefore **not typed as written** without a declared embedding $\iota : T_{(z_\star)}\mathcal{C}_n \to \mathrm{Dom}\, \mathbb{H}_n^{hist}$ together with its adjoint convention; the execution must supply ι and every expression below is to be read with ι and $\iota^\dagger$ inserted. *(An alternative that avoids the issue entirely is to take the second variation of the history quadratic form on the tangent space, $\mathbb{H}^{(2)}(z_\star) = D^2\mathfrak{E}_{hist}(z_\star)$, and project with tangent-space projectors P_T, Q_T; that route yields genuinely typed Hessian blocks and is the preferable construction where it is available.)*

(ii) Two distinct projectors, not one. The injection must use the **upstream kinematic quotient** P_{kin} — removing already-proved gauge, coordinate and record-coboundary redundancies — and **not** the downstream information-null projector P_{info} derived from $\ker W_{prim}$ (§17). Using P_{info} here would be circular: §56A.6 classifies the null space *after* W_{prim} is computed, while these blocks are what W_{prim} is computed *from*. The two are separated notationally and logically throughout:

$$\llbracket\ E_{(D,n)} = P_{kin}^{(n)} J_n^\dagger \left(J_n P_{kin}^{(n)} J_n^\dagger \right)^+ \quad \text{(upstream, kinematic)},\ \rrbracket$$
$$\llbracket\ P_{info}^{(n)} = I - N_{null} N_{null}^+ \quad \text{(downstream, from ker } W_{prim}\text{, §17)}.\ \rrbracket$$

Then the bath objects are **not separately supplied**:

$$\llbracket\ \Omega^2_{(Q,n)} = Q_n\, \mathbb{H}_n^{hist}\, Q_n,\quad C_n = Q_n\, \mathbb{H}_n^{hist}\, P_n\, \iota\, E_{(D,n)},\quad W_{(core,n)} = E_{(D,n)}^\dagger\, \iota^\dagger\, P_n\, \mathbb{H}_n^{hist}\, P_n\, \iota\, E_{(D,n)}.\ \rrbracket$$

So $\mathbb{H}_n^{hist} \longrightarrow (W_{(core,n)}, \Omega^2_{(Q,n)}, C_n)$ is a **compression of one previously constructed operator**; no block is selected independently. This bears on the history-content gate's central worry (§19, F19), and the strength of the claim matters. Defensible: **no independent W_{core} block is inserted once the sealed premises are fixed** — W_{core} is a compression of the same generator that supplies the bath, not a separately supplied matrix. Not defensible: that "a chosen W_{core} is impossible". Its *value* still moves with choices upstream of the compression — the tangent-to-history embedding ι, the Moore–Penrose injection convention, the operational inner product defining minimum norm, the selected background $z_{(\star,n)}$, and the persistent projector Π_{pers} — every one of which is an acknowledged conditional convention (§56A.11). Where the Gaussian/bilinear certificate of Theorem 16.1b passes, $W_{prim}^{(n)} = W_{(core,n)} - C_n^\dagger(\Omega^2_{(Q,n)})^+ C_n$; outside that sector HFB-1 evaluates the full information-rate Hessian of Theorem 16.1a.

56A.4 The full defect lift without loss of magnitude

The physical lift is the complete differential $J_n = \bigoplus_A \mathcal{J}_{(A,n)}$, **not** merely the partial isometries of its polar decomposition. Writing $J_{(A,n)} = V_{(A,n)}|J_{(A,n)}|$, the frame $V_{(A,n)}$ carries orientation while the positive factor carries the physical singular-value magnitudes; the frame-only object may audit equivariance but **may not replace** the magnitude-bearing object in the physical metric.

Which object is authoritative. Two distinct lifts must be kept apart, or an unproved wheel factorisation is smuggled into the physical calculation:

```
⟦ Π_(A,n)^wheel = L_(A,n) ⊗̂ J_(A,n)   — the wheel/internal **benchmark** lift only; ⟧
⟦ physical lift = J_(A,n) together with the embedding ι and the transported increments Ξ_(A,n) (§56A.5) — **authoritative**. ⟧
```

$L_{(A,n)}$ was introduced as a *wheel support representative*; tensoring the complete physical differential with it would import the diagnostic wheel factorisation into the physical operator. The physical metric is therefore defined directly from $J_{(A,n)}$, ι and $\Xi_{(A,n)}$ by the path-law construction of §56A.5, and never through L_A. The T and B defect rows retain their complete two-dimensional irreducible carriers, so the published 6×6 ledger is a matrix of six **typed operator blocks**, with TT, TB, BT and BB carrying restored doublet indices (§14F.4).

56A.5 Direct path-law definition of the physical metric

For a physical transition $z \to z'$ let $\Xi_{(A,n)}(z, z')$ be the covariantly transported increment of the full typed defect $\mathcal{D}_A$, evaluated with the inherited link, record and completion transport. With source $\eta \in \mathcal{R}_{prim}^{(n)}$ and the normalised exponential deformation $K_{(n,\eta)}(dz'|z) \propto \exp[\sum_A \langle \eta_A, \Xi_{(A,n)}(z,z')\rangle]\, K_{(n,0)}(dz'|z)$,

```
⟦ (W_prim^(n))_AB = lim_(N→∞) (1/N) 𝔼[ S_(A,n)^(N) ⊗ S_(B,n)^(N) ]_(η=0),
S_(A,n)^(N) = ∂_(η_A) log dμ_(n,η)^(N)|_(η=0). ⟧
```

This uses the actual MASD defects on the actual carrier. The canonical wheel scores remain a symmetry-reduced *test* of the construction and are not substituted for the full $J_{(A,n)}$.

56A.6 Executable null classification

After computing $W_{prim}^{(n)}$, HFB-1 computes its numerical kernel and tests every null vector against the five predeclared subspaces $\mathcal{N}_{gauge}$, $\mathcal{N}_{cob}$, $\mathcal{N}_{coord}$, $\mathcal{N}_{terminal}$, $\mathcal{N}_{regulator}$. With $P_{(\mathcal{N}_{known})}$ the projector onto their sum, the classification residual is $\varepsilon_{null}(v) = \|v - P_{(\mathcal{N}_{known})} v\|$, and a null is classified only when $\varepsilon_{null}(v) \le \tau_{null}$. Any remaining null is **published as an unexplained physical degeneracy** and forces DEGENERATE or FAIL according to whether the boundary observables are constant along it. It may not be silently removed (G5, F8).

56A.7 Terminal-record pushforward as an execution test

From the Doob-normalised stationary process, $m_n(x,y) = \pi_n(x)K_n(y|x)$ and $c_n = \tfrac{1}{2}(m_n + m_n^{T})$. The terminal conductance is fixed by $C_{(\Omega,n)}(a,b) = \sum_{(x:\, r_n(x)=a)} \sum_{(y:\, r_n(y)=b)} c_n(x,y)$, and with frozen embeddings S_a, S_b,

$$M_{(\Omega,n)}(b,a) = \sum\sum c_n(x,y)\, S_b^\dagger\, \mathcal{U}_{(yx)}^{rec}\, S_a, \quad \mathcal{U}_{(\Omega,n)}(b,a) = \mathrm{Pol}\, M_{(\Omega,n)}(b,a).$$

HFB-1 then **evaluates rather than assumes** the lumpability residual $\varepsilon_{lump}^{(n)} = \|E_{(r,n)}K_n - K_{(\Omega,n)}E_{(r,n)}\|$, the representative residual $\varepsilon_{rep}^{(n)}$ = max over x, x' with $r_n(x) = r_n(x')$ of $\|K_n(\cdot|x)_\Omega - K_n(\cdot|x')_\Omega\|$, and the holonomy-basicness residual $\varepsilon_{basic}^{(n)} = \|\rho_{vac}^{(n)} - r_n^\star \rho_\Omega^{(n)}\|$. **Gate G7 passes precisely when these fall below their frozen tolerances and the refinement square commutes; otherwise HFB-1 returns FAIL.** The flux is never placed on the terminal triangle by hand (F29). This converts G7d-2 and the §31.2 descent from open assertions into computed residuals.

What the residuals do and do not establish. They can *reject* a candidate r_n. They do not establish that the upstream construction **uniquely determines** r_n, that no other admissible map also passes, that two passing maps return the same Δ_Ω, or that the search over admissible maps is exhaustive. On the canonical wheel the corresponding uniqueness *is* available — the finite classification of all fifteen matchings (Theorem 32A.1) — but the full-carrier protocol needs its own uniqueness-or-invariance theorem. Without it the projection choice has been **moved into the sealed input package rather than derived** (F40).

56A.8 Exhaustive one-operator branch decomposition

Communicating components (§43A.6) are included but do not exhaust every reducing sector. The complete decomposition is therefore **central**. Let

$$\mathfrak{M}_n = \{ \mathbb{T}_n^{hist}, \mathfrak{A}_n^{kin} \}'', \quad \mathfrak{z}_n = Z(\mathfrak{M}_n),$$

and let the physical branches be the minimal central projections $\{P_{(\beta,n)}\} = \mathrm{MinProj}(\mathfrak{z}_n)$, or the direct-integral fibres when the centre is non-atomic: $\mathcal{H}_n = \int^\oplus \mathcal{H}_{(\beta,n)}\, d\nu_n(\beta)$, $\mathbb{T}_n^{hist} = \int^\oplus \mathbb{T}_{(\beta,n)}^{hist}\, d\nu_n(\beta)$. **This is exhaustive relative to the frozen observable algebra**: disconnected move-graph components reappear as central projections, while conserved representation, determinant-line, boundary and holonomy sectors inside a connected support are also retained — which is exactly the gap §43A.6 left open. A CFB configuration is a physical branch only if its state and observable algebra embed in one of these central fibres. No post-hoc direct sum is permitted (Firewall H10).

Two limits on this exhaustiveness, both stated rather than assumed away. First, it is exhaustive *relative to* $\mathfrak{A}_n^{kin}$, so physical exhaustiveness rests on the completeness of that observable algebra on the full carrier — an inherited attribution, listed in §55A and not re-derived. Second, and more structurally: **at finite regulator, physically significant phases are often not exact reducing components.** They may be nearly-invariant metastable sectors that become genuinely disjoint only in an infinite-volume or continuum limit. If the framework requires exact finite-regulator superselection, that requires a theorem; absent one, the branch framework must also admit asymptotic ergodic components or low-lying metastable spectral subspaces, with a declared separation criterion. **Status: OPEN (F41).**

56A.9 Perron evaluation and global dominance

For every admitted central component $\lambda_{(0,\beta)}^{(n)} = \rho(P_{(\beta,n)}\, \mathbb{T}_n^{hist}\, P_{(\beta,n)})$ and $f_\beta^{(n)} = -(1/a_t) \log \lambda_{(0,\beta)}^{(n)}$; no block is normalised separately before this table is written (Firewall H5). With frozen reference density satisfying $P_{(\beta,n)}\rho_{(0,n)} \neq 0$ for every admitted component and $\rho_{(N,n)}$ the globally normalised evolution (§43A.5), HFB-1 reports **PASS** when one regulator-stable component has a strictly maximal Perron root and all physical gates pass; **DEGENERATE** when a regulator-stable set shares the maximal root; **FAIL** when positivity, projection, null classification, branch comparability, global dominance or convergence fails. No experimental result may break a degeneracy (HM8, §53; F18; §46).

56A.10 Regulator execution

The same sealed construction is applied **without modification** at no fewer than two regulator levels, returning $\varepsilon_{\mathbb{T}}^{(n)} = \|\mathbb{T}_{(n+1)}^{hist}\, I_n^{(n+1)} - I_n^{(n+1)}\, \mathbb{T}_n^{hist}\|$, $\varepsilon_W^{(n)} = \|W_{prim}^{(n+1)} - I\, W_{prim}^{(n)}\, I^\dagger\|$, $\varepsilon_\Omega^{(n)} = \|\Delta_\Omega^{(n+1)} - I_\Omega\, \Delta_\Omega^{(n)}\, I_\Omega^\dagger\|$ and $\varepsilon_f^{(n)} = \max_\beta |f_{(\iota_n(\beta))}^{(n+1)} - f_\beta^{(n)}|$. Two levels constitute the **minimum execution diagnostic**: the gate at this level passes only when these residuals lie below frozen tolerances, the null classification is stable, and the dominant component or degenerate set is preserved. **A physical regulator-convergence certificate requires more** — an analytic a priori bound on ε_n, a monotonicity theorem in n, or a controlled Cauchy estimate giving summable ε_n (§47A). Passing the two-level diagnostic is therefore necessary and not sufficient, and the execution must report which of the three error-control results it supplies (F43).

Theorem 56A.1 — Blind-Execution Protocol and Admission Condition

HFB-1 is admitted for blind physical execution when: (1) every member of $\mathfrak{S}_{HMGBC}$ is frozen and hash-covered before any Standard Model output is inspected; (2) every formerly open physical quantity **other than the conditional upstream inputs and conventions listed exhaustively in §56A.11** is defined as an output of the deterministic map $\mathscr{F}_{HMGBC}$. Those listed items remain *conditional inputs*, and the protocol is blind with respect to them rather than deriving them; §56A.11 is the single authoritative list, so that this theorem and the ledger cannot drift apart; (3) the full master transfer is spectrally positive and constructed from the inherited operational generator and reference potential (§56A.2); (4) the bath triple is obtained by persistent/unresolved compression of that one generator (§56A.3); (5) the complete magnitude-bearing defect differential is used in the history metric (§56A.4–5); (6) the terminal conductance, polar transport and holonomy are computed by the pushforward tests of §56A.7; (7) the branch family is obtained from the central decomposition of one operator algebra (§56A.8); (8) the same construction is run at two or more regulator levels (§56A.10); (9) HFB-1 is forbidden to modify any rule after seeing a Standard Model output (HM8, §53); and (10) the branch construction either **proves that exact finite-regulator superselection is exhaustive**, or includes a **frozen metastable-sector criterion with a regulator-stable spectral-separation threshold** (§56A.8, F41) — without one of these the protocol can be blind but incomplete, since the physically significant sectors may separate only in the continuum limit.

The numerical values W_prim, P_info, C_Ω, Δ_Ω, R_Ω, R_f, 𝔅_phys, Δ_sel need not be known before admission. They are the outputs HFB-1 is admitted to calculate.

What this theorem is, stated precisely. It establishes a **blind-execution protocol**: a construction frozen and hashed in advance so that no discretionary choice can be made during execution. It is **not** a physical closure theorem, and it does not establish that the frozen conventions are *derived*. A convention does not become physically forced by being pre-registered — pre-registration prevents tuning, which is a different and lesser guarantee. The conventions still outstanding are named in §56A.11: the numerical tolerances, the Moore–Penrose injection, the background z_(⋆,n), the tilt family (HM12), and the completeness of the inherited L^op, v̂^op, Π_pers and 𝔄^kin. Each must eventually be (i) proved forced by upstream axioms, (ii) proved immaterial because all admissible alternatives give equivalent physical outputs, or (iii) retained as an explicit conditional premise. Until then, the claim that execution proceeds "without making any further constitutive choice" is to be read as **"without any further choice beyond the declared conditional premises of §56A.11"** — which is what the protocol actually delivers (F42).

Proof. The purpose of an admission gate is to prevent discretionary choices during execution. A sealed deterministic construction achieves this whether or not its numerical outputs have been evaluated; requiring those outputs beforehand confuses *admission* with *successful completion* and duplicates the execution stage inside HMGBC-1. Conversely a merely symbolic formula containing a free operator, branch prior, projection, relative normalisation or regulator rule does not qualify, because HFB-1 could still choose the answer. The package above removes **discretionary choice during execution**: the transfer descends from the operational generator; the bath blocks are compressions of that transfer; the metric is its path-information curvature; the terminal geometry is computed from its stationary flow *given the frozen map r_n*; the branches are the central components *of the frozen observable algebra*; and convergence is tested by repeating the same construction. It does **not** remove the declared conditional freedom in the terminal map, the tilt family, the observable algebra or the residual conventions of §56A.11 — that is precisely the distinction between a sealed protocol and a physical derivation. The package is therefore sufficient to begin a blind execution, while retaining the possibility that the execution returns FAIL or DEGENERATE. ■

56A.11 The residual discretionary surface

Sealing is not the same as having no conventions, and the ones that remain must be visible because they are the only surface on which an execution could still be steered:

- **𝒯_num tolerances** (τ_null, the ε-thresholds of §§56A.7, 56A.10) are genuine free parameters. They are legitimate only if **pre-registered and hash-covered before the run**; a tolerance adjusted after seeing a residual is a fit (F18).
- **The injection convention** E_(D,n) is the Moore–Penrose minimum-norm choice — canonical, but a declared convention rather than a derived object.
- **The background z_(⋆,n)** at which J_n is evaluated presumes a determinate defect-zero locus; if 𝓜_n is non-isolated the choice is not unique, which is MASD-1's F28 reappearing at the execution layer.

- **The terminal-record map r_n** is a frozen input, not an output: the §56A.7 residuals validate or reject a candidate but do not derive it, exhaust the alternatives, or show that all passing maps give the same Δ_Ω (Gate G7h, F40).
- **The exponential tilt family** and the defect-coordinate normalisation fix W_{prim} only up to $W \mapsto J^T W J$ (HM12, F34).
- **The metastable-sector separation criterion**, where exact finite-regulator superselection is not proved (Theorem 56A.1 condition 10, F41).
- **The completeness of the terminal-map search**, or equivalently the invariance class within which a passing map is deemed equivalent to any other.
- **The upstream attributions** — that L_n^{op}, $\hat{\nu}_n^{op}$, $\Pi_{pers}^{(n)}$ and $\mathfrak{A}_n^{kin}$ are genuinely inherited and complete on the full MASD/CAR/Fock carrier — are taken from the cited corpus and are not re-derived here.

None of these is a branch prior or a free operator; all are declarable in advance. They are recorded so that the seal is auditable rather than merely asserted.

Proposition 56A.2 — Programme Verdict under the Sealed Package

With the sealed package and immutable manifest published, the correct status is

```
⟦ HMGBC-1 rulebook and execution map: CLOSED FOR NUMERICAL EXECUTION ⟧
⟦ HFB-1 execution admission: PASS ⟧
⟦ HFB-1 physical boundary certificate: PENDING EXECUTION ⟧
⟦ Renormalised Parameter Closure: LOCKED UNTIL HFB-1 RETURNS PASS ⟧
```

This admission does **not** assert that the Standard Model boundary has been obtained, nor that the frozen conventions are derived. It asserts that the calculation may be run blind, with no discretionary choice available during execution beyond the declared conditional premises of §56A.11. A failed HFB-1 run is a scientific result and may not be repaired downstream.

Effectivity condition. The admission theorem is *established* here; the admission *takes effect* only on publication of $\mathscr{M}_{hash}$ with **real hashes rather than placeholders**, together with the two declared regulator inputs. Until those exist the correct wording is **"admission theorem established; execution manifest pending"** — never "execution completed." At the state of this paper the manifest is pending.

Corollary 56A.3 — Renormalised Parameter-Closure Lock

The Renormalised Standard Model Parameter-Closure stage consumes $\mathcal{P}_{SM}^{VERSF}(\mu_\star) = \{ g_s, g, g', \mu_H^2, \lambda, Y_u, Y_d, Y_e, \mathscr{B}_\nu, \bar{\theta} \}$. Because HFB-1 has been *admitted* but has not *returned*, this frozen boundary set does not exist. Therefore

```
⟦ Renormalised Parameter Closure: architecturally specified, numerically
LOCKED pending an HFB-1 PASS. ⟧
```

That stage may be written in advance as an immutable transport and evaluation protocol — scheme bridge, renormalisation-group evolution, threshold matching, electroweak breaking, running-mass calculation, pole extraction. **No physical parameter row may be populated** until a passing HFB-1 certificate exists. **Admission to run HFB-1 is not permission to populate the boundary table before the run.** Canonical-wheel values, CFB branch outputs, identity Yukawa matrices, benchmark gauge responses and measured Standard Model values may not fill the missing boundary (Firewalls H1, H4; F18, F25).

56A.12 Programme Status

```
⟦ HMGBC-1 sealed execution rulebook ⟶ HFB-1 Stage A (numerical HMGBC
execution) ⟶ HFB-1 Stage B (physical boundary) ⟶ Renormalised Parameter
Closure. ⟧
```

The first arrow is open in principle and awaits only the published manifest; no later arrow has been traversed. "Sealed rulebook" and "execution seal" denote the construction; "physical certificate" is reserved throughout for the numerical output certificate of §51.

---

57. Conclusion

MASD-1 built the common machine. HMGBC-1 supplies the rule that must fix its settings, proves the rule's consumer mathematics, and marks with precision the three things that the rule still needs.

The central move is to treat the unresolved record sector as part of the microscopic state rather than as an afterthought. On the enlarged carrier the transfer is local in the history step and positive; on the reduced carrier the same dynamics become non-Markovian. The history law remembers because the degrees of freedom that carried the memory have been integrated out. That one construction performs several jobs: its Radon–Nikodym curvature returns the operational defect metric, including all cross-sector response and the absolute action unit; its symmetric stationary flow returns the weighted terminal-record geometry, whose sector compressions generate hierarchy and mixing; its response algebra decides whether identical sectors retain protected ledgers or share a saturated bath; and its block Perron roots compare the admissible branches — once, and only once, they are shown to be components of one operator.

The identity completion used in the first MASD branch is now placed correctly. It is not the default law of nature; it is the special case in which the scalar-access support lies in the kernel of the history-derived completion Laplacian. A nontrivial history geometry replaces it automatically.

The paper also holds a tight standard of proof. A positive Hessian is not a positive transfer — and a *bosonic* positive transfer is not yet a positive *fermion-including* transfer (§20.2). A retarded oscillatory memory kernel is not a probability covariance. Stationarity does not require detailed balance. A list of candidate branches does not prove the branch space is finite. A null count is not a null classification. A gap's *existence* (automatic) is not its *uniform magnitude* (a required return). A per-branch Perron gap is not cross-branch comparability, and cross-branch comparability is not ontological selection (§43A.5). A history-derived metric is not yet a history-derived generation geometry until the projection onto the record carrier is itself a certified pushforward (§31.1). A solution map similar to a positive operator is not itself positive in the naive inner product (§33). A three-term strong phase is not three independent phases. And a derived bath spectrum with a chosen local cost is not a derived metric. Each of those distinctions is now built into a gate or a firewall.

The remaining calculation is difficult, but it is no longer conceptually diffuse. It is the frozen finite-regulator evaluation of the complete generator and its transfer

$$[\![\ \mathbb{H}_n^{hist} = (W_{core},\ \Omega^2_{(Q,n)},\ C_n),\quad \mathbb{T}_n = e^{(-a_t\ \mathbb{H}_n^{hist})},\ ]\!]$$

together with three certifications: that every block of $\mathbb{H}_n^{hist}$ is *derived* rather than chosen (history-content); that the record conductance is a genuine pushforward of the full-configuration flow (projection); and that the branches are components of one master operator on one measure with dominance implying selection (branch-comparison). Those objects return

$$[\![\ W_{prim},\quad R_\Omega,\quad \mathfrak{B}_{phys}.\ ]\!]$$

When they are evaluated and those certifications pass, HFB-1 will return a single microscopic Standard Model boundary, a declared degenerate boundary set, or a definite failure. That is the proper endpoint of the programme: not another candidate architecture, but a calculation that is no longer free to choose its answer. HMGBC-1's contribution is to have made that freedom finite, named and testable — reduced from "invent a rule" to "evaluate one generator, certify one projection, and prove one operator" — and to have refused, at every step, to call a relocated freedom an eliminated one.

Appendix A — Gaussian Influence-Functional Derivation

Let the quadratic bath-resolved action be

$$S[D, q] = \tfrac{1}{2}\, D^\dagger\, W_{core}\, D + D^\dagger\, C^\dagger\, q + \tfrac{1}{2}\, q^\dagger\, A_Q\, q.$$

Complete the square:

$$q^\dagger A_Q\, q + 2\, D^\dagger C^\dagger q = (q + A_Q^{-1} C\, D)^\dagger\, A_Q\, (q + A_Q^{-1} C\, D) - D^\dagger C^\dagger A_Q^{-1} C\, D,$$

so

$$S[D, q] = \tfrac{1}{2} (q + A_Q^{-1} C D)^\dagger A_Q (q + A_Q^{-1} C D) + \tfrac{1}{2} D^\dagger (W_{core} - C^\dagger A_Q^{-1} C) D.$$

Gaussian integration gives

$$\int dq\, e^{-S[D,q]} = (\det A_Q)^{-1/2} \exp\left[-\tfrac{1}{2} D^\dagger W_{hist} D \right], \quad W_{hist} = W_{core} - C^\dagger A_Q^{-1} C.$$

The determinant contributes to branch weights and cannot be discarded before the common-normalisation audit (§44) — and, under §43A, must be a block of the one master operator, not a separately computed factor, or the branch comparison is broken.

Appendix B — Bi-Perron and Doob Formulas

Let T be a finite positivity-improving matrix. Perron–Frobenius gives $T r = \lambda_0 r$ and $\ell^\dagger T = \lambda_0 \ell^\dagger$ with $r, \ell > 0$. Define $K_{xy} = T_{xy} r_y / (\lambda_0 r_x)$. Then $\sum_y K_{xy} = (T r)_x / (\lambda_0 r_x) = 1$. Define $\pi_x = \ell_x r_x / \sum_z \ell_z r_z$. Then

$$(\pi K)_y = \sum_x (\ell_x r_x / \langle \ell, r\rangle)(T_{xy} r_y / (\lambda_0 r_x)) = r_y (\ell^\dagger T)_y / (\lambda_0 \langle \ell, r\rangle) = \pi_y.$$

Detailed balance holds only when $\ell_x T_{xy} r_y = \ell_y T_{yx} r_x$; it is not required for stationarity (HP2 vs HP2R). For the block-Perron use in §43A, the same formulas apply to each restricted operator $P_\beta T^{hist} P_\beta$, and the block roots $\lambda_0(P_\beta T^{hist} P_\beta)$ are the comparable free energies — this is the only sense in which cross-branch comparison is defined.

Appendix C — Refinement Intertwining Proof (HP3)

Write the enlarged lift as $\hat{V}_n = V_n^X \otimes L_n^Q$. Assume the three commuting-square conditions

$$L_{(X,n+1)} V_n^X = V_n^X L_{(X,n)}, \quad \Omega^2_{(Q,n+1)} L_n^Q = L_n^Q \Omega^2_{(Q,n)}, \quad C_{(n+1)} V_n^D = L_n^Q C_n.$$

Then every kinetic and potential block of $\hat{\mathbb{H}}_{(n+1)} \hat{V}_n$ agrees with the corresponding block of $\hat{V}_n \hat{\mathbb{H}}_n$, so $\hat{\mathbb{H}}_{(n+1)} \hat{V}_n = \hat{V}_n \hat{\mathbb{H}}_n$, and functional calculus gives $e^{-a_t \hat{\mathbb{H}}_{(n+1)}} \hat{V}_n = \hat{V}_n e^{-a_t \hat{\mathbb{H}}_n}$. Therefore HP3 holds. The three conditions are exactly the per-block checks of §§24–27; HP3 is not opaque but a conjunction of three auditable identities on inherited maps.

Appendix D — Full 6×6 History-Metric Ledger

The physical execution must publish the Hermitian block matrix

```
W_prim = [ W_CC W_CJ W_CT W_CR W_CB W_CH ;
           W_JC W_JJ W_JT W_JR W_JB W_JH ;
           W_TC W_TJ W_TT W_TR W_TB W_TH ;
           W_RC W_RJ W_RT W_RR W_RB W_RH ;
           W_BC W_BJ W_BT W_BR W_BB W_BH ;
           W_HC W_HJ W_HT W_HR W_HB W_HH ].
```

For each block publish: carrier dimensions; basis labels; units; Hermiticity residual; regulator dependence; statistical or differentiation uncertainty; bath-mode decomposition (which $\omega_{(n,a)}$ source it, via §15); null overlap; and whether the block vanishes by symmetry or only numerically. No block may be omitted because a separable approximation is preferred. The bath-mode decomposition column is what certifies Corollary 16.2: a nonzero off-diagonal block must be traceable to a shared bath mode $c_{(n,aA)}\, c_{(n,aB)}$, or it is an inserted coefficient and F7 fires.

Appendix E — Universal Completion and Sector Typing

Let $\mathcal{S}_\Omega = W_H^{(1/2)}(W_H + \Delta_\Omega)^{-1} W_H^{(1/2)}$ be the positive Hermitian contraction of §33, and let $R_\Omega = W_H^{(-1/2)}\, \mathcal{S}_\Omega\, W_H^{(1/2)}$ be the completion solution map. R_Ω is *W_H-self-adjoint* and similar to the positive contraction $\mathcal{S}_\Omega$ (spectrum in (0,1]); it need not be positive or Hermitian in the unweighted inner product. Let $V_H^\dagger : \Omega^R \to \Omega^L$ be the polar orientation. Then $\boldsymbol{\mathcal{R}}_\Omega = V_H^\dagger R_\Omega : \Omega^R \to \Omega^L$. For sector selectors $S_{fR} : \mathbb{C}^3 \to \Omega^R$ and $S_{fL} : \mathbb{C}^3 \to \Omega^L$, the compression $R_f = S_{fL}^\dagger\, \boldsymbol{\mathcal{R}}_\Omega\, S_{fR}$ is a 3×3 typed left–right map. The physical canonical mass operator additionally includes the address kinetic factors R_{fX} and scalar normalisation returned by MASD-1; HMGBC supplies the universal completion geometry, not a shortcut around canonical normalisation. In particular Y_f^{mixed} and Y_f^{Hess} of MASD-1 §§35, 37 both consume this $\boldsymbol{\mathcal{R}}_\Omega$, so CY3′ tests two readings of one HMGBC object — which is why a genuine (non-identity) CY3′ pass is evidence about the history geometry, not merely about MASD-1's bookkeeping.

Appendix F — The Common-History-Operator Construction: What a Proof Must Exhibit

This appendix specifies the content of the §43A obligation so that its discharge is checkable rather than rhetorical. A proof must exhibit:

1. **One operator.** A single positive transfer $\mathbb{T}_n^{hist} = e^{(-a_t \hat{\mathbb{H}}_n^{hist})}$ on one history space $\mathscr{H}_n^{hist}$, with $\hat{\mathbb{H}}_n^{hist}$ real and bounded below (HM3), whose construction consumes only the *complete derived generator* (resolved local block W_{core}, bath block Ω^2_Q, coupling C_n) and inherited operational data — no block chosen.
2. **A branch-labelling observable algebra.** A commutative von Neumann algebra $\mathfrak{z}_n \subseteq$ commutant($\mathbb{T}_n^{hist}$) generated by the six branch labels of the freeze manifest (vacuum, support, polar, bundle, orientation, refinement). The admissible branches are its minimal spectral projections $\{P_\beta\}$.
3. **Direct-integral / superselection decomposition.** $\mathscr{H}_n^{hist} = \bigoplus_\beta P_\beta \mathscr{H}_n^{hist}$ (or the direct-integral analogue for a continuous label), with $\mathbb{T}_n^{hist} = \bigoplus_\beta P_\beta \mathbb{T}_n^{hist} P_\beta$, so each branch is literally an invariant component and $\mathbb{T}_n^{hist}$ commutes with every P_β.
4. **One reference measure.** A single $\hat{\nu}_n^{hist}$ on $\mathscr{H}_n^{hist}$ inducing all $Z_{(n,\beta)}$; no per-branch normalisation.
5. **Block-Perron comparability.** Each $P_\beta \mathbb{T}_n^{hist} P_\beta$ positivity-improving on its support with a proven gap, so $f_\beta = -(1/a_t) \log \lambda_0(P_\beta \mathbb{T}_n^{hist} P_\beta)$ are the comparable free energies of Theorem 43A.4.
6. **Refinement covariance.** The label algebra and its decomposition intertwine with HP3, so the dominant component maps to the dominant component under refinement (§47).

The natural candidate for $\mathfrak{z}_n$ is the algebra of conserved history charges (flat holonomy class, terminal-support type, bundle sector). Proving that these exhaust the branch freedom, commute with $\mathbb{T}_n^{hist}$, and carry one reference measure is the H3 work HMGBC-1 defers. **Absent this appendix's six items, "branch selection" is not defined**, and no amount of per-branch Perron structure supplies it.

Appendix G — Freeze Manifest and Machine-Readable Handoff

Manifest fields are split by stage, because hashing an object that Stage A is meant to *calculate* would demand the result before admission.

Admission hashes (frozen before execution; these are hashes of builders, schemas and inputs, never of results): hmgbc_version; masd_companion_version; regulator_id; regulator_hash; regulator_input_hash; reference_measure_hash; resolved_basis_hash; bath_basis_hash; defect_basis_hash; observable_algebra_input_hash; **master_operator_builder_hash**; **branch_decomposition_algorithm_hash**; refinement_map_hash; tolerance_block_hash; code_commit; dependency_lock_hash; random_seed_or_deterministic_flag.

Stage-A output hashes (produced by the run, not required for admission): **assembled_master_operator_hash**; **central_decomposition_hash**;

physical_branch_return_hash; metric_return_hash; null_classification_hash; convergence_report_hash.

Minimum HFB-1 handoff fields: physical_branch_id; physical_branch_set; master_operator_component_label; reference_measure_hash; selection_gap; perron_root; left_perron_vector; right_perron_vector; W_prim_file; W_block_ledger_file; null_basis_file; null_classification_file; Delta_Omega_file; R_Omega_file; R_u_file; R_d_file; R_e_file; R_nu_file; bath_ledger_verdicts; theta_top_file; det_line_section; theta_line_fate; hp1_status; hp2_status; hp2r_status; hp3_status; regulator_convergence_report; uncertainty_report. These are **Stage-A output** fields: a missing field prevents Stage B from opening, not admission of Stage A (§52, Theorem 56A.1). Admission is gated instead by the **admission hashes** of Appendix G — builder, algorithm, schema and input hashes, never hashes of results — which must be **real, not placeholders**, together with the two declared regulator inputs. The `master_operator_component_label` and `reference_measure_hash` fields enforce the §43A obligation at the manifest level: a handoff that names a branch without naming the operator it is a component of is refused.

Appendix H — Programme Source Map

This companion draws on: *The VERSF Master Substrate Action and Constitutive-Descent Theorem* (MASD-1); *Probability as Admissible Measure on the VERSF Operational Hilbert Geometry*; *Operational Geometry in the VERSF Framework*; *Unified Refinement Hilbert Geometry in VERSF*; *The Memory Kernel from First Principles: A Dynamic Derivation of K(τ) in VERSF*; *History-Dependent Dynamics from Persistent Commitment Memory*; *Coherence Scale, Memory Kernel Reduction, and Non-Markovian Gravity in VERSF*; *Fact-Momentum: Commitment-Capacity Field Dynamics and the Propagation of Irreversible Events*; *Refinement Cohomology and the Spectrum of Γ* on Admissible Substrates*; *Admissible Coarse-Graining and Continuum Emergence in VERSF*; *Continuum-Limit Regularity and Cone Convergence in VERSF*; *Deriving Regularity Hypotheses from Admissible TPB Dynamics in VERSF*; *The Admissible Lift Theorem*; *The Gate-2 Inspection over the Transport Construction*; *Gate 3 and Fact Momentum*; *The Non-Abelian Bath-Transport Theorem, Continuous Colour Gauge Dynamics, and Full Gauge-Product Closure in VERSF*; *The Substrate Cohomology, Global Gauge Measure, and Strong-Phase Trivialisation Theorem in VERSF*; and the scalar, neutral, Yukawa, gauge-response and QCD final-completion gates consumed by MASD-1.

Appendix I — Worked Closure-Defect Two-Cell Lift

§14E states what a defect-lift proof must exhibit but does not show the list is *achievable* on even one defect. This appendix works the canonical-wheel closure score F_C through an explicit two-cell history complex, computes the lifted Fisher pairings exactly, and tests the **four internal**

graph-lift analogues of the HM10 criteria against it. It is a single-sector, two-cell demonstration on the *wheel score* — not the full physical lift, and specifically **not** yet a lift of the physical closure object, which in the Gate-3 corpus is a $\mathbb{Z}_7$-valued inter-hub connection on a transport 2-complex. What it returns is nonetheless genuine: the wheel closure score is stable under a nontrivial two-cell extension, remains Fisher-independent of the bridge, and exposes a support-preservation condition that a naive lift violates.

I.1 The two-cell complex

Take two canonical cells A and B, each a hub with a six-state rim, joined by one bridge edge between rim state a_0 of A and rim state b_0 of B. States are labelled 0 = h_A, 1…6 = A-rim, 7 = h_B, 8…13 = B-rim. Build the kernel from symmetric edge conductances — hub–rim and rim–rim ring edges of unit weight in each cell, a bridge conductance g on $\{a_0, b_0\}$, and a self-loop of unit weight on every state to preserve the local persistence support that the single wheel carried through T_0(b_i, b_i) = ¼ — via

```
⟦ T(x, y) = c(x, y) / ∑_z c(x, z),  π(x) = deg(x) / ∑_w deg(w),  deg(x) = ∑_z
c(x, z). ⟧
```

This T is row-stochastic, irreducible, and reversible with respect to π by construction (a conductance kernel is always reversible). With g = 1/3 it has a simple Perron root 1 and a positive spectral gap, so HP1–HP2 hold on the two-cell complex.

I.2 The lifted closure score

The closure defect lifts as the oriented rim circulation in **each** cell with the same chirality:

```
F_C(a_i, a_(i+1)) = +1,  F_C(a_i, a_(i−1)) = −1,  and identically on the B-
rim.
```

The bridge carries its own auxiliary direction F_BR (score +1 on the bridge edge), and the persistence direction F_H (self-persistence) and the **symmetric** current direction F_J — the radial-conductance score F_J(h, b_i) = F_J(b_i, h) = 1 of §14B, *not* an oriented variant — lift cell-wise as on the single wheel.

I.3 Exact lifted pairings

Carrying the bridge conductance g symbolically, the closure Fisher weight is an **exact rational function of g**. Every rim state — bridged or not — contributes 2/Σ with Σ = ∑deg = 62 + 2g, because the bridge conductance cancels against the raised degree at a_0 and b_0. Hence

```
⟦ W_CC^(2-cell)(g) = 24/(62 + 2g) = 12/(31 + g), ⟧
```

verified symbolically, which at g = 1/3 gives 18/47 ≈ 0.383 and as g → 0 returns the single-wheel value 12/31 exactly. So the proximity to 12/31 is not a numerical coincidence at one parameter value but a **continuity statement**: the closure sector interpolates analytically from the single-wheel value, so its pattern is reproduced *analytically* rather than checked at a point (the support-typing criterion HM10.4 is tested separately in §I.4, with Lemma I.5 as its graph-side instance). The key structural test also passes exactly:

```
⟦ W_(C,BR) = 0. ⟧
```

The closure circulation does **not** mix with the bridge coupling: the oriented rim mode is orthogonal to the inter-cell edge under the Fisher pairing. Meanwhile the bridge does generate genuinely new structure absent on the single wheel, W_(J,BR) and W_(H,BR) ≠ 0, exactly the multi-cell cross-block generation that §14E(6) anticipated as the mechanism by which wheel-zero blocks can become nonzero after lifting.

I.4 The four HM10 criteria, tested

- **Sector match (graph-lift analogue of HM10.1).** W_CC ≠ 0 on the lift: the wheel closure score is not annihilated by the two-cell compression. This is the *graph-internal* analogue; the *physical* sector-match — that the MASD closure differential D𝓓_C maps onto F_C at all — is **not** established here (see §I.6). **Pass, graph-internal only.**
- **Rank preservation (HM10.2).** This is where the worked example earns its place. With the self-loops of §I.1 present, and using the §14B symmetric J-score throughout, the full lifted (C, J, H, BR) block is rank 4 with strictly positive spectrum {≈ 0.00968, 0.0971, 0.246, 0.383} — every lifted direction, including the new bridge mode, is independent. **Without** the self-loops — a naive conductance-only lift — the persistence direction H loses its independent support and collapses onto the current direction J, dropping the rank to 3. The self-loop is not decoration: it is the support that keeps H independent, inherited from T_0(b_i, b_i) = ¼ on the single wheel. **Pass with self-loops; fails spuriously without them — hence Lemma I.5.**
- **Equivariance (graph-lift analogue of HM10.3).** The same-chirality rim circulation in both cells is the unique lift of the *wheel score* F_C compatible with the per-cell D_6 action and the bridge's residual symmetry; the opposite-chirality choice is a different (bridge-odd) direction. This is equivariance of the wheel score under the graph symmetry — **not** the physical equivariance §14E(3), which would require mapping the MASD closure cochain to F_C. **Pass at the graph level only; physical equivariance open.**
- **Support typing (HM10.4).** This appendix does not merely rename the fourth criterion — it *tests* it, and Lemma I.5 is its graph-side twin. HM10.4 requires that the physical differentials carry the benchmark's support types, in particular that 𝓓_H is diagonal-supported. Lemma I.5 is precisely the graph-level instance: a lift that discards diagonal support collapses the persistence direction onto the current direction, because with only the hub self-loop surviving S_H = −S_J identically (Lemma 14C.4). So the two-cell calculation exhibits, on an explicit complex, both what the criterion demands and what fails when it is violated. **The support requirement is verified here at graph level; the physical question — whether the actual $\mathbb{Z}_7$-valued 𝓓_H is diagonal-supported and uniformly signed — remains open.**

The closure sector's own reproduction is also stronger than a spot check: $W_{CC}(g) = 12/(31 + g)$ exactly (§I.3), returning the single-wheel 12/31 as $g \to 0$, so the pattern is reproduced *analytically* rather than at one parameter value, and the paper can say precisely how far a given bridge moves the magnitude. The magnitude drift is more visible in the auxiliary sector, where the bridge generates entirely new nonzero blocks $W_{J,BR}$ and $W_{H,BR}$ with no single-wheel counterpart. **Pass on support typing at graph level and on analytic pattern reproduction; magnitude correctly treated as model-specific.**

I.5 Lemma I.5 — Support-Preservation Condition

Any faithful lift must preserve the local self-transition support of the wheel (the rim self-loops). A lift that maps the wheel into a pure inter-state conductance complex without self-loops collapses the scalar-persistence direction onto the current direction and fails HM10.2 for a spurious, avoidable reason. Equivalently: the persistence defect $\mathcal{D}_H$ is carried by the *diagonal* of the kernel, so any lift that discards diagonal support cannot represent it. The mechanism is the exact anti-proportionality of Lemma 14C.4: with only the hub self-loop surviving, $S_H = -S_J$ identically on the hub row, so the two directions are not merely close but *collinear*. That makes Lemma I.5 a theorem about score supports rather than an empirical observation. **This is a concrete, checkable necessary condition on the physical lift, extracted from the worked example** — it belongs in the §48 frozen-input package as a constraint on the refinement lifts.

I.6 What this appendix does and does not establish

It establishes, exactly and by computation, that the canonical-wheel closure *score* is stable under a nontrivial two-cell history complex: it survives with the right sign, decouples from the bridge, reproduces the single-wheel pattern without its magnitude, and passes the four graph-internal analogues of the HM10 criteria — with Lemma I.5 serving as the graph-side instance of the support-typing criterion HM10.4 itself, not merely a precondition for the others. It thereby demonstrates that the §14E obligation is *structurally achievable* for the wheel score of one defect, and it upgrades one criterion (rank) into a specific named condition on the lift.

It does **not** establish physical closure-defect faithfulness. The physical closure object is a $\mathbb{Z}_7$-valued inter-hub connection on a transport 2-complex; this appendix does **not** construct the map from the MASD plaquette-holonomy differential (or the face labels $\rho_f \in \mathbb{Z}_7$) to the real six-rim circulation score F_C. So the physical **sector-match** and **equivariance** obligations (HM10.1, HM10.3) remain open — the appendix passes their graph-internal analogues only. It also works one defect, two cells, one bridge; the full lift is all six defects through the full multi-cell, record, gauge, completion and fermionic carriers, and whether the wheel-zero blocks W_{BH}, W_{CH}, W_{RH} acquire *physical* nonzero values there is open (F23, F24). What the appendix removes is the worry that §14E is an empty checklist: the wheel closure score passes its graph-internal analogues, and the support-preservation lemma is a genuine, reusable result.