ANNALS OF THE CHRYSENE FORMALISM

ORTHOGONAL REVIEW RECORD | VOL. 1, NO. 4 | ISSN 3144-2021 | NOVEMBER 2026

ORTHOGONAL REVIEW RECORD

The Source Duality: Endogenous and Exogenous Preparation of the Substituted-Chrysene Stock, Mixed Arm Words, and the Persistence Semigroup

Charles D. Schaper, Ph.D.
Annals of the Chrysene Formalism
Schaper Systems Press

Paper ID:	ACF-1.4-090
Issue:	Volume 1, Number 4, November 2026
ISSN:	3144-2021
Review architecture:	Earlier mathematical audit (M), followed by separate concordance (C) and chemistry (H) review passes and a revision response
Review focus:	Source laws, arm-word interpretation, persistence assumptions, thermodynamic scope and companion-file concordance
Record date:	4 November 2026
Document version:	v2.0
Editorial status:	Prepared for author review; no acceptance or online-publication date certified by this record

Declared review status

This is an AI-assisted simulated review record. The concordance and chemistry passes were performed separately from the production revision, but they are not independent human peer review, evidence of external referee appointment, or editorial acceptance. The mathematical stage below is preserved as a dated record of the preceding package; the new stages reconsider its remaining interpretive and dependency issues. The response records revisions prepared for the author's approval, rather than attributing unprovided statements to the author.

> **Scope of Review**
>
> The review evaluates the supplied Article 90 manuscript and its mathematical companions. It distinguishes exact finite results, analytic theorems under stated hypotheses, imported results and proposed chemical realizations. It does not certify unavailable predecessor proofs, compound-specific synthesis, experimental yields, kinetic measurements or the historical origin of a molecular system.

Combined recommendation

Retain the central conditional source-duality framework after the specified revisions. The earlier mathematical corrections remain necessary. The new concordance review identifies five clarification groups; the chemistry review identifies eight interpretation and physical-admissibility groups. The principal corrections specify the representative convention, repair dependency and rate hypotheses, distinguish ethers from esters, qualify contact energies and kinetic examples, and state the equilibrium premises of the work bound. These changes strengthen the conditional mathematical statement without replacing the paper's central architecture.

Contents

1 Stage summary **4**

2 Version and provenance **4**

3 Conclusions supported by the stated finite model **5**

3.1 Factorization and blindness 5

3.2 Arm classes and geometric counts 5

4 Corrections identified in the supplied source **6**

5 Verification recorded at the earlier mathematical stage **9**

6 Disposition of the earlier mathematical stage **9**

7 Concordance review **11**

7.1 C1 Specify arm representatives and registry relabeling 11

7.2 C2 Repair the hypothesis dependency register 11

7.3 C3 Separate preparation from zero-flip dynamics 11

7.4 C4 State the necessary absorption rates 12

7.5 C5 Harmonize identity, abstract and proof coverage 12

8 Chemistry and domain-admissibility review **13**

8.1 H1 Correct ether and ester terminology 13

8.2 H2 Distinguish proposed preparation from measured synthesis 13

8.3 H3 Give the contact score its proper physical scope 13

8.4 H4 Separate graph saturation from molecular immobilization 14

8.5 H5 Qualify kinetic language and time units 14

8.6 H6 Correct the proposed orientation readout 14
8.7 H7 State the thermodynamic premises of the selection-work bound 14
8.8 H8 Make the independence model experimentally testable 15

9 Revision policy and scope decisions **15**

10 Closure and current validation **15**

1 Stage summary

M — **Mathematical validity**

The earlier audit corrected source normalization and target limits, retention and return-hazard arguments, the cap-conversion mean, spectral-distance scope, and information interpretations of graph counts. Its detailed R1–R6 findings are retained below. Its verification figures refer to that earlier execution record, unless explicitly updated in the final status section.

C — **Concordance**

The new pass compares definitions, hypotheses, theorem dependencies, numerical and formal coverage, and package metadata. It requests a fixed representative/gauge convention, a corrected dependency register, separate placement and zero-flip hypotheses, complete absorption-rate assumptions, and harmonized companion metadata.

H — **Chemistry and domain admissibility**

The new pass preserves the abstract finite model but corrects chemical nomenclature, the status of proposed synthesis and yields, the meaning of contact scores, lattice locking, kinetic units, spectroscopic predictions, thermodynamic bounds and independence assumptions.

Reading the preserved mathematical stage

The following sections retain the substance of the prior mathematical report. Terms such as "cleavage" in that historical discussion are superseded by the precise arm-conversion distinction in H1. References to "revised text" in this preserved stage refer to the first publication package, not an assertion that the later C/H findings had already been resolved.

2 Version and provenance

The initial audit read the supplied manuscript source, `original/main.tex`, the prior report, `original/referee_report_a90.tex`, and the relevant verification source. The prior report carries October 2026 and provisional Volume 1, Number 5 metadata. It recommends "Accept with Minor Revisions" and claims clean execution of nine verification blocks. Those are statements in the supplied record; its source does not establish an independently identifiable referee or an editorial decision. It is retained unaltered as `provenance/Article90_Prior_Referee_Source.tex`.

The present audit does not inherit the prior report's blanket approval. In particular, its proposed desorption-delay argument is not a stochastic-order proof, and its approval of the spectral metric assertion is unsupported by the argument reproduced in Article 90. Its correct contact computations remain useful.

Supplied manuscript SHA-256

e2f33fe3bc430bcbfb7ba5caa6a30d1c

f2ccbbad091de13ba71d42472c91e426

Prior report SHA-256

252ba914865b3e65fd2902d0e8fa80d3

7c52b0442e66cb1cdb1eb3dbb52a33c9

Stable source labels below refer to the supplied manuscript. The revised-publication check is recorded separately at the end; its scope must not be confused with the initial source audit.

3 Conclusions supported by the stated finite model

3.1 Factorization and blindness

For a normalized product law with the conditional registry factor as defined, taking the negative base-two logarithm gives the additive selection cost. On a contact with equal coincidence scores $v_0 = v_1$, the common Boltzmann factor cancels and the registry bit is uniform. A column with $l - 1$ independent such contact bits therefore has registry cost $l - 1$. The zero partial derivatives hold for that specified factorized model and blind locus; they are not a universal assertion that real chemical kinetics cannot affect registry.

The explicitly listed contact patterns give

$$(v_0, v_1)_{D(eh),D(eh)} = (2, 0), \quad (v_0, v_1)_{D(eh),T(eeee)} = (1, 1), \quad (v_0, v_1)_{I(eee),D(eh)} = (1, 1).$$

The last example is a valid blind contact between two individually flip-visible units. Thus, under the acceptance rule that the substrate is $T(eeee)$ and the cap is $I(eee)$, cleavage of a neighboring D does not produce the proposed orientation bias. This does not exclude effects from arm words or interactions outside those assumptions.

3.2 Arm classes and geometric counts

For an m-letter arm alphabet, the canonical-support class counts are

$$N_D = \frac{m^2 + m}{2}, \qquad N_I = m^3, \qquad N_T = \frac{m^4 + m^2}{2}.$$

Their totals are 21 at $m = 2$ and 78 at $m = 3$. For two letters the blind counts are 2 on D, 0 on I and 4 on T. These counts are consistent; the probability normalization needs the distinct stabilizer clarification in R1.

The $3 \times 3 \times 6$ box has 54 nodes, 72 lateral bridges, 216 arm terminals, 72 free terminals and 45 vertical contacts. It has four sites complete in all three lattice directions. Those sites lie on the central column, whose free-arm set is empty and whose bridged-arm set is inversion invariant. The central-column and pre-closure blindness conclusions are conditional on the stated connector and assembly assumptions, not a count of experimentally observed states.

The alternating six-site species-cost formula is also correct as written:

$$-\log_2(f_D^3 f_T^3) = 3 \log_2 \left(\frac{(1 + \rho + x)^2}{\rho} \right).$$

It describes the six-site D, T, D, T, D, T cluster, not an additional initiator cap. Adding a cap adds its selection cost, its site and its contact; it must not silently retain the same six-site accounting.

4 Corrections identified in the supplied source

R1 Normalize the arm law and select every target coordinate

Location: `def:measure`, `lem:support`. The arm-word orbit factor must be taken under the stabilizer of the chosen support,

$$G_S = \{g \in \{1, \iota\} : gS = S\}, \qquad P([a] \mid S) = |\operatorname{Orb}_{G_S}(a)|\, \pi_h^{n_h(a)} (1 - \pi_h)^{|S| - n_h(a)}.$$

For the canonical support $I = \{3, 6, 12\}$, G_I is trivial. Its eight two-letter words therefore have multiplier one, not the multiplier two of a global support-changing inversion orbit. Summing over stabilizer orbits gives total probability one. If global supports are instead included as separate states, their distribution must be specified explicitly.

The support assertion must also be confined to the declared source alphabet. A two-letter $\{h, e\}$ law cannot have positive support on configurations containing the other arm letters admitted elsewhere in the corpus. Define the column product domain explicitly, using the supported D, I, T families and arm classes. If an additional compatibility restriction is imposed, its conditional normalization can destroy the independent factorization; that restricted law requires a separate analysis rather than an unqualified appeal to the same product proof.

The displayed external tilt selects κ^0 and η^0 but leaves $\mu_{\mathrm{dc}}(\epsilon; \gamma)$ unchanged. Its limit is consequently

$$\delta_{\kappa^0} \otimes \mu_{\mathrm{dc}} \otimes \delta_{\eta^0},$$

not δ_{c_0} unless the chirality law is already concentrated. A valid repair is to tilt the chirality factor as well, or, for Bernoulli parameters $\gamma_i = P(\epsilon_i = 1)$, impose $\gamma_i \to \epsilon_i^0$. The finite-space proof is then direct: each target-coordinate probability tends to one. The base measure must assign positive mass to the target; a finite or diverging exponential tilt cannot recover a state excluded by a zero base weight. Define the normalization of the tilted factors explicitly.

R2 Distinguish retention from mixing to uniform

Location: `prop:metastability`, its proof and verification description. For $m = l - 1$ independent symmetric contact flips at rate φ, started at a fixed field η^0, one has

$$p_t = P(\eta_i(t) = \eta_i^0) = \frac{1 + e^{-2\varphi t}}{2}, \qquad \|\mathcal{L}(\eta(t)) - \delta_{\eta^0}\|_{\mathrm{TV}} = 1 - p_t^m.$$

This follows because total variation from a point mass is one minus the mass at that point. A convenient bound is

$$1 - p_t^m \le m(1 - p_t) \le m\varphi t, \qquad \varphi \le \frac{\varepsilon_{\mathrm{reg}}}{(l-1)\delta T}.$$

The manuscript instead cites a bound on distance to the *uniform* law. That is a mixing statement and cannot be inverted to justify distance from the initial point mass. Both statements can coexist, but their reference measures and purposes must be named.

The old sufficient ceiling is not itself refuted: for $0 \le \varepsilon \le 1$,

$$\frac{m}{4}\log(1+4\varepsilon^2/m) \le \varepsilon^2 \le \varepsilon.$$

It therefore implies the union bound above, though by a different argument. Replacing it with the simpler retention ceiling changes the reported example at $l = 6$, $\delta T = 1$, $\varepsilon = 0.05$ to $\varphi \le 0.01$. This is a sufficient bound, not an exact threshold. A statement about the entire placed state also needs to account separately for chemical or frame-loss events.

R3 Prove the chemical hazard bound from the chain

Location: `prop:metastability`, `rem:d4` and the prior referee report's desorption recommendation. A larger mean hitting time alone does not imply stochastic domination of an exponential variable. That argument must be removed. The desired bound is nevertheless correct for the explicit chain

$$\kappa \xrightarrow{d} \sigma, \quad \sigma \xrightarrow{a} \kappa, \quad \sigma \xrightarrow{r} \beta,$$

started in κ, with β absorbing for the first-return calculation. Write $S(t) = P(\tau_\beta > t)$ and $y(t) = P(X_t = \sigma \mid \tau_\beta > t)$. The transient forward equations give

$$y' = d - (a+d+r)y + ry^2, \quad y(0) = 0, \quad S' = -ryS.$$

At $y = d/(a+d)$ the vector field equals $-rad/(a+d)^2 \le 0$. Hence $y \le d/(a+d)$, and, for $d > 0$,

$$S(t) \ge e^{-\bar{r}t}, \quad P(\tau_\beta \le t) \le 1 - e^{-\bar{r}t} \le \bar{r}t, \qquad \bar{r} = \frac{rd}{a+d}.$$

A union bound then applies without independence across sites. This is not a bound for arbitrary initial phase: from σ, the initial return hazard is r. The exact initial-κ mean is $(a+d+r)/(dr)$, or 101 020 for the supplied example rates, not 101 000.

The maximum over degrees $\{2, 3\}$ controls returns of those vulnerable degrees only. A claim about *any* bulk return of a column also containing degree-four sites must include r_4 in the bound or assume $r_4 = 0$. The event should be named accordingly.

R4 Correct the cap mean and state the observation premise

Location: `cor:cap_catastrophe`, abstract, synthesis and conclusion. A precise mean requires a declared transition graph. For

$$\kappa_3 \underset{a_3}{\overset{d_3}{\rightleftharpoons}} \sigma_3 \underset{e_3}{\overset{r_3}{\rightleftharpoons}} \beta_3 \xrightarrow{k_4} \beta_4 \underset{r_4}{\overset{e_4}{\rightleftharpoons}} \sigma_4 \xrightarrow{a_4} \kappa_4,$$

with κ_4 the first-hit target and positive necessary rates, first-step equations yield

$$E_{\kappa_3}\tau_{\kappa_4} = \frac{1}{d_3} + \frac{1+a_3/d_3}{r_3}\left(1+\frac{e_3}{k_4}\right) + \frac{1}{k_4} + \frac{1}{e_4} + \frac{1+r_4/e_4}{a_4}.$$

Indeed, if $C = (1+a_3/d_3)/r_3$, the degree-three equations give $T_{\sigma_3} = C + T_{\beta_3}$ and $T_{\beta_3} = 1/k_4 + (e_3/k_4)C$. The degree-four first-hit mean is $1/e_4 + (1+r_4/e_4)/a_4$. The original expression does not solve these equations. If a different graph is intended, supply it and solve its equations instead. A mean is an expectation, not a deterministic time at which every sample path has converted.

Cap conversion to inversion-invariant T removes that cap as a frame carrier. It does not, by global inversion alone, collapse all edge fields: if $\eta_i = w_i \oplus w_{i+1}$ then simultaneously replacing all w_i by $1 - w_i$ leaves every η_i unchanged. Erasure of all registry data therefore requires an additional observation map that forgets those data when every relevant frame carrier is gone. State this as a model premise, exclude another surviving cap or external reference, and distinguish loss of observability from actual flipping. Likewise, injectivity on a labelled configuration space retaining η does not imply injectivity on a physical quotient that identifies those labels after frame loss.

R5 Retain the spectral lower bound and withdraw unsupported upper claims

Location: `thm:connes`, `prop:jstr_injective` and summaries. Under the explicitly stated bound for the *entire* signed sum,

$$\left\| \left[D, \sum_{i=0}^{l-2} (-1)^{\eta_i^0} Z_i Z_{i+1} \right] \right\| \leq 2\gamma_z,$$

the normalized sum is an admissible test observable. Its expectation difference is $(l-1)/(2\gamma_z)$, proving the stated lower bound. This conditional argument is sound, but the commutator premise is not proved merely by calculating those expectation values.

Part (b) asserts bi-Lipschitz comparison, finite distances and an admissibility comparison. Testing one observable gives only a lower bound. Moreover, a bound on the full sum need not hold on an arbitrary partial sum because commutators can cancel. Neither scalar kernel nor boundedness of the whole Lipschitz ball is supplied. For example, a Dirac operator commuting with a nonconstant observable that separates two states allows arbitrary multiples of that observable in its seminorm ball and gives infinite distance. The displayed upper commutator bound by itself excludes no such case.

Remove part (b), or supply independently verified assumptions and proof establishing finiteness and a uniform upper estimate. On a fixed finite set, finiteness and positive separation would yield some bi-Lipschitz constants, but that is an additional conclusion requiring proof, not a consequence of the single lower-bound test. The rank-two site density is correctly a pure character when restricted to the *diagonal relative-registry algebra* $C^*(Z_i Z_{i+1})$; that algebra should not be confused with the full noncommutative algebra invariant under global flip.

R6 Keep graph counts separate from information assumptions

Location: `cor:lattice_cost`, `rul:contacts` and the summary table. There are 45 vertical contacts in nine disjoint six-site columns. Adding eight edges connecting those columns in a tree gives 53 edges. This is correct graph arithmetic. To call the additional eight relations eight more bits requires that their observable variables exist and be independently unbiased under the stated source law. It is not implied by the count of physical lateral bridges. Thus 45 is the vertical-edge count, 53 is the single-column or augmented-tree count, and their information interpretations have different premises. Uniformity on each contact alone also does not establish joint independence; use the product-source premise explicitly.

5 Verification recorded at the earlier mathematical stage

The revised Python source was copied outside the publication tree and rerun for this review with Python 3.12.14. It passed 153 checks per run: 71 finite mathematical or sample checks and 82 macro-uniqueness checks. Normal and optimized executions produced identical output. Five deliberate defects were each rejected in both modes, giving ten rejected mutation runs. The checks include normalized arm distributions, finite contact and lattice enumeration, the corrected return mean, the six-state first-step mean and retention at the new ceiling. At $\varphi = 0.01$, $l = 6$ and $t = 1$, the exact retention distance is approximately $0.04853274 < 0.05$.

The supplied Lean listing was also extracted and compiled for this review. Lean 4.24.0 rejected its four conjunction proofs using bare `rfl`; the earlier machine-check claim was therefore not supported by that supplied listing. The revised companion was independently rerun with the pinned Lean 4.24.0 executable: all nine declared finite certificates compiled, their printed axiom unions were empty, and two deliberately false targets were rejected. These statements certify the declared arithmetic, involution, finite enumerations and contact examples. They do not formalize the stochastic-limit arguments, the general source factorization theorem, the operator-algebraic embedding or the Connes-distance hypothesis.

Several unverified numerical literals in the supplied code were replaced by computed values. In particular, the three initiator-contact frequencies at the declared $\rho^* = 203.625$ are approximately 2.39×10^{-7}, 2.37×10^{-5} and 2.17×10^{-3}. The original finite word totals 21 and 78 and the geometric counts remain unchanged. Float-based checks are numerical evidence; exact fractions and exhaustive enumerations have their separate finite scope. The execution and source-hash records are `Verification-Code/execution_audit.json`, `Verification-Code/verification_report.json` and `Formal-Proof/formal_audit.json` in the supplements.

6 Disposition of the earlier mathematical stage

Revised mathematical text checked on 4 November 2026. The modular manuscript addresses R1–R6:

- The source domain, stabilizer probabilities, normalized tilts and full-coordinate target limit are explicit.
- Retention uses the initial point mass and a conditional-hazard proof. Its chemical event excludes degree-four returns unless added explicitly; phase and label laws are distinguished.
- The cap mean uses the declared six-state chain. Readout loss requires the separate observation quotient and absence of another reference.
- The connected-column embedding fixes unused factors. The trace statement is restricted to blind registry; only the conditional spectral lower bound remains.
- The 53-bit interpretation requires the additional independent unbiased observables. The verification appendix now matches actual executable coverage.

Version fingerprints and scope are recorded in `review_audit.json`.

Disposition. No unresolved objection to R1–R6 remains within this targeted audit. The recommendation concerns the corrected finite and explicitly conditional claims, not the withdrawn

assertions or experimental realizability.

7 Concordance review

Review C — scope and recommendation

The delivered manuscript, referee report, verification companion and formal-proof companion were compared. The Python harness was independently rerun during this pass: 153 checks per mode, identical normal/optimized outputs, and ten rejected mutation runs. The Lean source and recorded audit were inspected; that inspection is not a fresh compilation. Recommendation: retain the central conditional results after the following minor-to-moderate clarifications.

7.1 C1 Specify arm representatives and registry relabeling

The source coordinate κ uses an arm class $[a]$ modulo the support stabilizer, while v_η uses an oriented representative. Changing eh to he at one D site swaps v_0 and v_1 for its incident contact. An unoriented orbit alone therefore does not determine those labelled scores.

Fix representatives $r_S([a])$ once, for example with e at position 6 and h at position 12 for the representative $D(eh)$. Under an allowed nontrivial stabilizer change indexed by $s_i \in \{0,1\}$, relabel incident contact fields and target fields coherently:

$$\eta'_{ij} = \eta_{ij} \oplus s_i \oplus s_j, \qquad (\eta^0_{ij})' = \eta^0_{ij} \oplus s_i \oplus s_j.$$

Only support-preserving changes are permitted. The labels of the probabilities change, while blindness, entropy and corresponding target-event probabilities are invariant. This is a coordinate convention, not a physical rotation or a new symmetry assumption.

Revision response. Adopt the fixed-representative convention and coherent registry relabeling in the source-law definitions. Preserve the finite contact counts and distinguish representative-dependent labels from invariant conclusions.

7.2 C2 Repair the hypothesis dependency register

The earlier register says that the support lemma survives failure of the product/independent-draw premise. That is too strong: compatibility exclusions can assign zero probability to previously possible configurations. Only the finite arm-class counts survive without the source-law assumptions. Similarly, central-column blindness does not survive arbitrary failure of saturation. It can survive a change in accretion order when the same final saturated central pattern remains.

Revision response. Correct both dependency statements. Full support is asserted only on the declared product domain with positive factors; a restricted domain requires its own law and proof. Central blindness retains its final-pattern and saturation hypotheses.

7.3 C3 Separate preparation from zero-flip dynamics

The name H_{place} is used both for external preparation and for $\varphi = 0$. These are distinct propositions: a state can be prepared and then flip. Introduce a separate locking hypothesis $H_{\mathrm{lock}} : \varphi = 0$ wherever persistent registry invariance is intended.

Revision response. Separate placement from locking in the definitions, persistence theorem, explanatory text and hypothesis register. External preparation remains compatible with a positive subsequent flip rate.

7.4 C4 State the necessary absorption rates

The assertion of eventual conversion to T needs more than $r_n > 0$. If $d_n = 0$, a fixed site may never leave its initial occupied phase; if a required advance rate is zero, the path to T can be blocked. State the fixed-slot finite-chain setting and positivity of the necessary desorption, re-export and advancement transitions. With the target absorbing, every non-target state must have an enabled path to it and there must be no other closed class.

Revision response. Add the finite-chain and positivity assumptions to the eventual-absorption assertion. This repairs the stated theorem's domain rather than asserting absorption for an arbitrary chemical system.

7.5 C5 Harmonize identity, abstract and proof coverage

The manuscript and companions used different subtitles. The OJS abstract must match the manuscript's abstract while retaining math-mode source and removing purely typographic commands. Bibliographic identifiers and page metadata must follow the current source list and compiled manuscript.

The existing scope distinction is sound: 71 mathematical/sample checks plus 82 macro checks are not 153 analytic theorems; nine Lean certificates are finite certificates, not full formal proofs of source convergence, stochastic persistence or operator-algebraic results. Imported theorem numbers in unavailable Articles 40, 79, 80 and 84 are not independently certified by this review.

Revision response. Use the supplied full title throughout, regenerate the abstract from the manuscript, incorporate the supplied DOI list and maintain explicit certificate coverage. No expansion of Lean's formal scope is required merely to accommodate these clarification edits.

8 Chemistry and domain-admissibility review

Review H — scope and recommendation

The chemical interpretation was reviewed separately from the mathematical factorization and finite enumerations. Recommendation: substantive interpretive revision while preserving the finite source and persistence architecture. No laboratory result, compound-specific rate or physically complete interaction model is certified by this review.

8.1 H1 Correct ether and ester terminology

The declared e arm is $\mathrm{Ar{-}O{-}CH_2{-}CH_2{-}OH}$; it contains an aryl ether, not an ester. Conversion to $h = \mathrm{Ar{-}OH}$ cannot be called ester cleavage merely because an oxalate precursor appeared earlier in a proposed route. The IUPAC distinctions are explicit: ethers have an $\mathrm{R{-}O{-}R'}$ linkage, whereas esters derive from an oxoacid and an alcohol through loss of water [IUPAC-Ether, IUPAC-Ester].

Revision response. Use "arm conversion" for the abstract Bernoulli channel $e \to h$. Distinguish an ester-bearing precursor, if invoked, from the reduced ether arm. Do not assert that the abstract probability parameter establishes a selective or experimentally demonstrated ether-dealkylation protocol. The arm-word combinatorics is unchanged.

8.2 H2 Distinguish proposed preparation from measured synthesis

The bromination/oxalate/reduction sequence is presented too directly as a realized purifier, and the value $\eta = 0.998437$ can be read as a measured efficiency. The supplied theory supports an imported proposed coarse-grained scheme under positive-rate assumptions, with parameter-conditioned model output. A patent's disclosure of alternatives does not itself validate this exact kinetic chain, its selectivity or its isolated yield.

Revision response. Qualify the proposed route and numerical efficiency where they first appear and in the conclusion. Preserve the mathematical model while identifying which reaction transitions and rate ratios require compound-specific confirmation.

8.3 H3 Give the contact score its proper physical scope

The overlap counts v_η are not literal numbers of physical hydrogen bonds and do not exhaust contact free energy. Both h and e may participate in interactions; an uncounted mixed contact need not have zero physical energy. An admissible effective interpretation is

$$G_\eta = G_{\mathrm{ref}} - k_{\mathrm{B}} T \, \Delta_{\mathrm{rgy}} v_\eta,$$

with dimensionless Δ_{rgy} and equal orientation degeneracy. Geometry, solvent, ionization, entropy and other interactions must either be included in the effective model or stated as omitted. "Blind" means equality of the model descriptor $v_0 = v_1$, not full experimental energetic degeneracy.

Revision response. State this effective-energy interpretation and its assumptions. Retain the exact cancellation theorem under the declared model; avoid extending it to all real interactions.

8.4 H4 Separate graph saturation from molecular immobilization

The exact lattice counts—54 nodes, 72 lateral bridges, 45 vertical contacts and 72 free terminals—are graph-model results. Saturation and no bond exchange do not alone prove that every covalent node is rotationally immobile or that every relevant energy of a central column is orientation independent. Flexible bridges may move without the endpoint relabeling represented in the graph.

Revision response. Keep the exact graph counts and the saturated-pattern theorem. Separate fixed endpoint labels from physical conformational motions, and identify bridge activation chemistry, phosphorylation state where applicable, and molecular mobility as realization requirements.

8.5 H5 Qualify kinetic language and time units

Statements that intra-column rotation is "rapid," a retention condition is "modest," or placement requires a particular covalent tether are not supported by compound-specific measurements. Desorption, interphase transfer and re-entry are different transitions. Polarity alone does not determine their rates. The parameters a, d, e, r, k, φ have inverse-time units, so a mean of 101 020 is in the chosen model time unit, not automatically in seconds.

Revision response. Use explicitly conditional kinetic comparisons; identify the common time unit and distinguish transition mechanisms. Treat covalent locking as a possible implementation rather than the unique logical meaning of placement.

8.6 H6 Correct the proposed orientation readout

For the $D(eh), D(eh)$ example,

$$\theta_{DD} = (1 + e^{2\Delta_{\mathrm{rgy}}})^{-1}$$

is the unaligned fraction P_1, not the ratio of the two orientations. The model ratio is

$$\frac{P_1}{P_0} = e^{-2\Delta_{\mathrm{rgy}}}.$$

An experimental syn/anti assignment needs structural identification and control of exchange, aggregation and signal integration. A heterodimer's predicted $1:1$ population is conditional on equal degeneracy and this interaction model.

Revision response. Correct the fraction-versus-ratio description and state the structural and dynamical premises for an NMR or other population test. Preserve the formula and do not present the conditional prediction as a general chemistry law.

8.7 H7 State the thermodynamic premises of the selection-work bound

A universal work cost $k_{\mathrm{B}}T \ln 2\, c$ cannot be inferred from an arbitrary kinetic source probability μ. For an isothermal protocol with specified Hamiltonian, thermal reservoir, control and feedback accounting, the appropriate statement is the nonequilibrium free-energy bound $W \geq \Delta F_{\mathrm{neq}}$ [EVdB11]. If μ is the Gibbs reference for the unchanged Hamiltonian, then preparing a point mass gives

$$\Delta F_{\mathrm{neq}} = k_{\mathrm{B}}T D_{\mathrm{KL}}(\delta_{c_0} \| \mu) = k_{\mathrm{B}}T \ln 2\, [-\log_2 \mu(c_0)],$$

where D_{KL} uses natural logarithms and the target has positive reference mass. The equilibrium premise is additional to the kinetic law.

Revision response. Retain surprisal as a mathematical selection cost for the general source law. Restrict its direct work interpretation to the stated Gibbs/isothermal setting, with full protocol and resource accounting. This narrows the physical interpretation without altering the factorization identity.

8.8 H8 Make the independence model experimentally testable

Independent identical arm conversion and independent monomer draws are modeling assumptions. Correlated conversion, pH and ionization, solvent, concentration, counterions and aggregation may change the arm law or contact energies.

Revision response. Keep the binomial and product formulas under their explicit hypotheses. Identify compound characterization, an actual $e \to h$ conversion time course, source-composition measurements, and structural/kinetic assembly measurements as empirical tests. No unprovided experimental data are introduced.

9 Revision policy and scope decisions

All five concordance groups and all eight chemistry groups are accepted as targeted clarification or repair. The core finite framework, the two-letter arm model and the conditional source-duality argument remain the subject of the paper.

Scope boundary. Requiring an experimental synthesis campaign as a precondition for stating a conditional mathematical theorem would materially change the commission. That broader requirement is not adopted. Instead, chemical assumptions and testable realization requirements are made explicit. Similarly, expanding the interaction model to every solvent, conformational, electronic or cooperative contribution is deferred; the existing exact statements remain limited to the declared score and source law. These scope choices do not rebut the precise nomenclature, rate, energy or thermodynamic corrections above.

10 Closure and current validation

Focused source checks confirmed the targeted concordance and chemistry repairs. The reachability comparison is restricted to the source-supported product domain; placement is separated from locking; the graph's prohibited endpoint relabelling is distinguished from a physical bond-severing mechanism. All five concordance groups and eight chemistry groups are addressed without changing the finite architecture.

The unchanged Python source was rerun in normal and optimized modes: 153 checks per mode passed, outputs agreed, and all ten deliberate-defect runs were rejected. The Lean source is unchanged from the earlier nine-certificate build record. No new Lean compilation is claimed for this revision; the retained execution log documents the earlier build. Newly expanded analytic arguments are not claimed as Lean theorems. The current manuscript, report and companion documents were rebuilt and checked for layout and reference consistency before release.

Disposition: The conditional mathematical framework is suitable for author consideration following these revisions. This is not an experimental validation or editorial acceptance.

References

[IUPAC-Ether] IUPAC. Ethers. *Compendium of Chemical Terminology* (Gold Book). `https://doi.org/10.1351/goldbook.E02221`.

[IUPAC-Ester] IUPAC. Esters. *Compendium of Chemical Terminology* (Gold Book). `https://doi.org/10.1351/goldbook.E02219`.

[EVdB11] M. Esposito and C. Van den Broeck. Second law and Landauer principle far from equilibrium. *EPL*, 95:40004, 2011. `https://doi.org/10.1209/0295-5075/95/40004`.

Revision history

Version	Date	Scope
v1.0	4 November 2026	Targeted mathematical audit and first publication package.
v2.0	4 November 2026	Separate concordance and chemistry passes, recorded revision responses and updated production format.

Living Document Notice

This record identifies the scope and date of its checks. Later changes to the manuscript or to imported results require corresponding re-evaluation. Computational checks and finite proof certificates support only their stated coverage.

Editorial Disclaimer

This simulated review supports author-side assessment. It is not independent human peer review, an editorial acceptance, experimental confirmation or a claim that the full analytic paper has been formally verified in Lean 4.