Deriving the Commitment Barrier from Closure Entropy and the Coherence Scale

VERSF Theoretical Physics Program

For the General Reader

Every measurement in physics produces a definite outcome --- a fact about the world. But quantum mechanics describes systems that exist in superpositions until something irreversible happens. This paper addresses the question: what is the minimum cost, in terms of physical entropy, of producing one such irreversible fact?

We work within the Void Energy-Regulated Space Framework (VERSF), which proposes that space, time, and physical distinctions emerge from an underlying substrate through a process called *commitment* --- the irreversible selection of one outcome from competing possibilities. The central quantity we derive is the *commitment barrier* Φc, a dimensionless number that governs how sharp this transition must be.

Our main finding is that Φc is not a free parameter of the theory. It equals r, where r is an independently constrained coefficient of the closure-stiffness term in the theory's equations of motion. Under a physically motivated identification of the theory's entropy with thermodynamic entropy --- which is supported by four structural arguments but stated as an explicit conditional assumption --- $r = 1$, giving $\Phi c = 1$. This corresponds to a commitment energy scale of approximately 2.5 millielectronvolts --- well within the reach of low-temperature experiments. The value $r = 1$ is conditional; the energy scale formula $Ec \sim r \cdot \hbar c/\xi$ holds unconditionally.

The derivation rests on a chain of named results, clearly separated into two tiers. The *unconditional* results follow from combinatorics, graph theory, and the causal structure of spacetime: they do not depend on any interpretation of the theory's entropy as physical entropy. The *conditional* results --- including $\Phi c = r$ and $r = 1$ --- additionally require identifying the theory's abstract distinguishability classes with the microstate classes of statistical mechanics. This identification is strongly motivated by four independent structural arguments, but is stated as an explicit assumption rather than a theorem, so that the reader can see exactly where the physics enters.

The paper is self-contained. Readers unfamiliar with VERSF can treat the framework axioms as given and follow the derivations from Section 3 onward. Within the broader VERSF framework, this paper is specifically about the threshold for fact formation, not the later stability of records once formed; that question belongs to a companion treatment of persistence capacity.

Table of Contents

1. Introduction and Motivation
2. Closure-Entropy Framework
   - 2.1 Closure-Stiffness and the Wilson Penalty
   - 2.2 The K = 7 Closure Cell
3. Closure Entropy: Forced Logarithmic Form and Binary Minimality
   - 3.1 Theorem A --- Admissible Closure Entropy Form
   - 3.2 Lemma B --- Binary Minimality
   - 3.3 Corollary C --- Primitive Entropy Quantum in Closure Units
4. Primitive Commitment: Admissible Class and Spectral Theorem
   - 4.1 The Admissible Primitive Class
   - 4.2 Theorem D --- Spectral Minimum for Primitive Commitment
   - 4.3 Spectral Lemma --- Primitive Closure Cost
5. Bridge Lemma: Connecting Spectral Structure and Binary Refinement
6. Wheel Uniqueness --- Theorem E
7. The Barrier Formula and Physical Normalization
   - 7.1 Proposition --- Barrier Formula (Unconditional)
   - 7.2 Conditional Theorem G --- Physical Entropy Identification
   - 7.3 Conditional Corollary H
   - 7.4 Falsification Paragraph
8. The Commitment-Capacity Chain
   - 8.1 The Physical Normalization of Closure Entropy
   - 8.2 Theorem I --- Quartic Commitment-Capacity Invariant
   - 8.3 Corollary --- The Coherence Scale
   - 8.4 Theorem J --- Minimal Admissible Wheel Cell
   - 8.5 Consistency Theorem --- r = 1
9. Experimental Implications
   - 9.1 Predictions at Two Strengths
   - 9.2 The Two-Observable Test
   - 9.3 Physical Intuition for K = 7
   - 9.4 The Unified Derivation Chain
10. Summary: Logical Structure
    - 10.1 The Full Theorem Chain
    - 10.2 Status of Each Result
    - 10.3 Unified Statement

Appendix: Spectral Rigidity of the K = 7 Closure Cell

Abstract

We derive the dimensionless commitment barrier Φc from the closure-entropy structure of the Void Energy-Regulated Space Framework (VERSF). The derivation is structured as a sequence of named results separating unconditional from conditional claims at every step.

Unconditional results. Theorem A proves that any admissible closure entropy measure satisfying monotonicity, additivity, and the null-singleton condition must take the form $\tilde{S} = \ln N$, forcing the logarithmic form from axioms rather than assuming it. **Lemma B** establishes binary minimality: the smallest nontrivial irreversible partition refinement has cardinality change $1 \to 2$. **Corollary C** gives the primitive entropy quantum $\tilde{\Theta}_0 = \ln 2$ in closure units. **Theorem D** (Spectral Minimum) establishes $\lambda^* = 2$ by Rayleigh--Ritz over an independently defined admissible class $\mathcal{A}$prim. A **Spectral Lemma** gives the dimensionless primitive closure cost $\Delta\tilde{S}$prim = r with no free factor. **Theorem E** (Wheel Uniqueness) proves the cancellation $(1/2)\cdot\lambda^*(K) = 1$ holds if and only if $K = 7$. **Bridge Lemma F** connects the spectral and counting sides: the $\lambda^* = 2$ eigenspace is the unique lowest-roughness spectral sector capable of realizing primitive binary refinement. **Theorem I** derives the quartic commitment-capacity invariant $\chi(L) = \rho L^4/(\hbar c)$ from first principles --- the L^4 scaling is forced by the action budget of a bounded causal region, not assumed --- and establishes the coherence scale $\xi = (\hbar c/\rho)$^(1/4) from $\chi(\xi) = 1$.

Conditional results (on the physical entropy identification of Theorem G). Under the identification $S = k_B \tilde{S}$, the primitive entropy quantum becomes $\Theta_0 = k_B \ln 2$, implying $\eta = 1$ and hence $\Phi c = r$. The Consistency Theorem then selects $r = 1$ as the primitive consistency value. The identification is supported by four distinct lines of argument within the programme: binary irreducibility establishing one bit as the minimal record; structural counting giving $\tilde{\Theta}_0 = \ln 2$ unconditionally; the Matching Lemma fixing $\alpha = k_B$ once the state classes are identified; and the Bekenstein threshold route requiring one-bit threshold behaviour at the coherence scale. **Theorem J** proves that within the natural wheel-cell family, $K = 7$ is the unique minimal admissible closure cell; full uniqueness among all connected graphs remains an open graph-spectral exclusion problem.

The strongest unconditional result is $\Delta\tilde{S}$prim = r. The physical results $\Phi c = r$ and $Ec \approx r \times 2.5$ meV are conditional on the entropy identification; $r = 1$ is the primitive consistency value under that identification.

1. Introduction and Motivation

In the fact-formation framework of VERSF, a *committed fact* is an irreversible distinction produced once a reversible distinguishability exceeds a minimum entropy threshold. The dimensionless parameter governing the sharpness of this transition is the commitment barrier Φc. Within the broader VERSF programme, this paper concerns commitment capacity --- the capacity for fact formation; coherence capacity governs reversible pre-commitment distinguishability, and persistence capacity governs the later stability of committed records. In prior work, $\Phi c = 1$ appeared as a threshold when the amplification variable is normalised by $A_{min} = k_B \ln 2$. The present paper argues that Φc is not a free parameter by providing a complete proof architecture with every assumption explicitly labelled.

The derivation separates into two tiers:

Unconditional tier: Theorems A, D, E, I; Lemmas B, F; Corollary C; Spectral Lemma; Proposition (barrier formula $\Phi_c = \eta r$). These depend only on the combinatorial structure of the $K = 7$ closure cell, the VERSF axioms of finite distinguishability and irreversible commitment, the admissibility conditions for partition refinements, and the action budget of a bounded causal region.

Conditional tier: Conditional Theorem G, Conditional Corollary H, and the Consistency Theorem $r = 1$. These are conditional on the identification of VERSF closure entropy with physical thermodynamic entropy --- an internal programme assertion of VERSF whose precise locus is identified in Section 8.1.

The logical spine runs as follows. Define the dimensionless closure action cost of a mode with eigenvalue λ as

$$\Delta\tilde{S}(\lambda) := (r/2) \cdot \lambda,$$

and the physical entropy cost as $\Delta S_c = \Theta_0 \cdot \Delta\tilde{S}_c$, where Θ_0 is the entropy quantum. The commitment barrier is

$$\Phi_c = \Delta S_c / (k_B \ln 2) = \eta \cdot (r/2) \cdot \lambda^*, \quad \eta := \Theta_0 / (k_B \ln 2).$$

The unconditional tier pins $\lambda^* = 2$, gives $\Delta\tilde{S}_{prim} = r$, and establishes $\tilde{\Theta}_0 = \ln 2$. The conditional tier yields $\eta = 1$, $\Phi_c = r$, and $r = 1$.

Notation convention. Throughout this paper, tilde quantities ($\tilde{S}$, $\tilde{\Theta}_0$, $\tilde{\Phi}$) denote dimensionless closure-unit quantities. Untilded quantities (S, Θ_0, Φ) denote the corresponding physical thermodynamic quantities. The two sets are related by $S = k_B \tilde{S}$ under the physical entropy identification of Theorem G; all results above the conditional gate are stated and proved in closure units, with the physical interpretation applied only below it.

2. Closure-Entropy Framework

2.1 Closure-Stiffness and the Wilson Penalty

In *From Schrödinger to Dirac: First-Order Closure Flow in VERSF*, the Wilson term is derived as the unique local quadratic entropy penalty associated with closure roughness. The entropy functional is

$$S_{ent}[\Psi] = \lambda \langle \Psi, D_G^2 \Psi \rangle,$$

with the corresponding Hamiltonian correction

$$H_W = (r/2) \cdot (\hbar c/\xi) \cdot (D_G^2 \otimes \beta),$$

where r is the dimensionless closure-stiffness coefficient and ξ is the closure coherence scale. The prefactor 1/2 is structural --- it arises from the quadratic nature of the entropy penalty in the second variation of the closure-entropy functional --- and is independent of the cell spectrum. The closure-stiffness paper argues that $r = \mathcal{O}(1)$ on independent stability grounds: a too-soft closure is unstable against fluctuations.

2.2 The K = 7 Closure Cell

The K = 7 closure cell has graph Laplacian Lcell with explicit spectrum

Spec(Lcell) = { 0, 2, 2, 4, 4, 5, 7 }.

Eigenvalue Multiplicity Sector

---

0 1 Constant (trivial) --- global phase

2 2 Minimum transport --- primitive commitment

4 2 Second transport

5 1 Near-UV

7 1 UV --- suppressed by HW

The choice K = 7 has so far been motivated by the closure-flow derivation in companion papers. Section 8.4 (Theorem J) provides a structural argument that K = 7 is the *minimal* admissible closure cell under the constraints of finite distinguishability, binary primitive commitment, closure-stiffness stabilization, and primitive commitment-capacity threshold matching --- making K a derived minimum rather than a free parameter.

3. Closure Entropy: Forced Logarithmic Form and Binary Minimality

3.1 Theorem A --- Admissible Closure Entropy Form

Theorem A (Forced Logarithmic Form). *Let $\tilde{S} : \mathbb{N} \to \mathbb{R}_{\geq 0}$ be any function assigning a closure entropy to a system with N distinguishable states, satisfying:*

(i) ***Monotonicity.*** *If $N' > N$ then $\tilde{S}(N') > \tilde{S}(N)$: finer distinguishability means higher entropy.*

(ii) ***Additivity.*** *For two independent subsystems with N_1 and N_2 distinguishable states respectively, $\tilde{S}(N_1 \cdot N_2) = \tilde{S}(N_1) + \tilde{S}(N_2)$.*

(iii) ***Null singleton.*** *$\tilde{S}(1) = 0$: a system with a single distinguishable state has zero entropy.*

Then $\tilde{S}(N) = C \ln N$ for some constant $C > 0$. In closure units, set $C = 1$, giving $\tilde{S}(N) = \ln N$.

Proof.

Step 1: Multiplicative-to-additive structure. Define $f(N) := \tilde{S}(N)$. Condition (ii) gives $f(N_1 \cdot N_2) = f(N_1) + f(N_2)$ for positive integers N_1, N_2.

Step 2: Monotonicity forces logarithmic form on integers. Condition (i) gives $f(N) < f(N')$ whenever $N < N'$, so f is strictly monotone on $\mathbb{N}$. By the integer version of the Cauchy functional equation (Aczél 1966, §2.1): any monotone solution of $f(N_1 \cdot N_2) = f(N_1) + f(N_2)$ on positive integers is of the form $f(N) = C \log N$ for some $C > 0$ and any base. Since Condition (iii) gives $f(1) = 0$, which is satisfied automatically ($C \log 1 = 0$), and monotonicity requires $C > 0$, we obtain $f(N) = C \ln N$. In closure units, set $C = 1$.

Therefore $\tilde{S}(N) = \ln N$. ■

Remark. No extension to real arguments is required --- the integer domain is sufficient. The logarithmic form is not assumed: it is the unique monotone solution to the multiplicative Cauchy equation on positive integers. Any alternative ($N - 1$, $\sqrt{N}$, etc.) violates additivity or monotonicity.

3.2 Lemma B --- Binary Minimality

Lemma B (Binary Minimality). *Under finite distinguishability (VERSF axiom A1), the smallest nontrivial irreversible refinement step increases partition cardinality by exactly 1. In particular, starting from a single indistinguishable class ($N = 1$), the smallest nontrivial refinement produces $N = 2$.*

Proof. A partition refinement that produces a new committed record must separate at least one previously indistinguishable pair into distinct classes --- otherwise the partition is unchanged and no new distinguishable fact is produced (the event is trivial). Therefore any nontrivial refinement increases cardinality by at least 1. The minimum increase of exactly 1 is achievable (split any one indistinguishable pair), so the minimum is 1. For the primitive case starting from $N = 1$ (a single fully indistinguishable class), the minimum nontrivial refinement increases cardinality to $N = 2$, which is the binary partition $\{s_1\} \mid \{s_2\}$. ■

3.3 Corollary C --- Primitive Entropy Quantum in Closure Units

Corollary C (Primitive Entropy Quantum). *The entropy change associated with the smallest nontrivial irreversible closure event is*

$\tilde{\Theta}_0 = \ln 2$ (in closure units).

Proof. By Lemma B, the smallest nontrivial refinement changes cardinality from N = 1 to N = 2. By Theorem A with C = 1:

$\tilde{\Theta}_0 = \tilde{S}(2) - \tilde{S}(1) = \ln 2 - \ln 1 = \ln 2$. ■

This result is unconditional within VERSF: it requires only Theorem A and Lemma B, with no reference to kB or thermodynamics.

4. Primitive Commitment: Admissible Class and Spectral Theorem

4.1 The Admissible Primitive Class

Define the **admissible primitive class** $\mathcal{A}$prim as the set of closure configurations satisfying two conditions:

$\mathcal{A}$prim := { $\Psi \neq 0 : \Psi \perp$ ker(Lcell), Ψ irreducible under decomposition into cheaper admissible excitations }.

Condition 1 --- Non-triviality. Ψ is orthogonal to the kernel of Lcell (the constant sector) and is nonzero.

Condition 2 --- Irreducibility. Ψ cannot be written as $\Psi = \Psi_1 + \Psi_2$ where both Ψ_1 and Ψ_2 are individually nonzero, orthogonal to ker(Lcell), and each achieves a Rayleigh quotient strictly less than that of Ψ:

$\langle\Psi_i, \text{Lcell } \Psi_i\rangle / \langle\Psi_i, \Psi_i\rangle < \langle\Psi, \text{Lcell } \Psi\rangle / \langle\Psi, \Psi\rangle$, i = 1, 2.

The class $\mathcal{A}$prim makes no reference to any specific eigenvalue. The variational minimum is defined as

λ^* := min { $\langle\Psi, \text{Lcell } \Psi\rangle / \langle\Psi, \Psi\rangle : \Psi \in \mathcal{A}$prim }.

4.2 Theorem D --- Spectral Minimum for Primitive Commitment

Theorem D (Spectral Minimum). *For the K = 7 closure cell with Laplacian Lcell and spectrum {0, 2, 2, 4, 4, 5, 7}, λ* = 2. The minimum is achieved by configurations in the two-dimensional eigenspace at eigenvalue 2.*

Proof.

Step 1. By Condition 1, all $\Psi \in \mathcal{A}$prim are nonzero and orthogonal to ker(Lcell). By Courant--Fischer--Weyl, the minimum of $\langle\Psi, \text{Lcell } \Psi\rangle/\langle\Psi, \Psi\rangle$ over ker(Lcell)⊥ is the first nonzero eigenvalue μ_1. Since any eigenvector at μ_1 satisfies Condition 2 --- as Step 2 shows --- the constrained minimum over $\mathcal{A}$prim equals μ_1.

Step 2. Let Ψ^* be an eigenvector at μ_1. Suppose for contradiction $\Psi^* = \Psi_1 + \Psi_2$ with $\Psi_i \neq 0$, $\Psi_i \perp \ker(L_{cell})$, and Rayleigh quotient of each Ψ_i strictly less than μ_1. But each Ψ_i lies in $\ker(L_{cell})^\perp$, so the Courant--Fischer minimum theorem gives $R(\Psi_i) \geq \mu_1$ for any nonzero $\Psi_i \perp \ker(L_{cell})$. This contradicts the assumption $R(\Psi_i) < \mu_1$. Therefore no such decomposition exists, Ψ^* is irreducible, and $\Psi^* \in \mathcal{A}_{prim}$.

Step 3. By explicit computation of the K = 7 cell spectrum, $\mu_1 = 2$. ■

Remark. The twofold degeneracy at $\lambda^* = 2$ is not treated here as a mere coincidence: given binary primitive commitment, it is the minimal multiplicity compatible with the required relational support structure, as Bridge Lemma F (Section 5) establishes.

4.3 Spectral Lemma --- Primitive Closure Cost

Spectral Lemma (Primitive Spectral Normalization). *The dimensionless closure action cost of the primitive commitment mode is*

$$\Delta\tilde{S}_{prim} := (r/2) \cdot \lambda^* = (r/2) \cdot 2 = r.$$

The Wilson prefactor 1/2 and the cell eigenvalue λ = 2 cancel exactly; no residual numerical factor remains.*

Proof. Direct substitution of $\lambda^* = 2$ from Theorem D. ■

Both inputs are independently fixed: the 1/2 from the quadratic entropy penalty structure, the 2 from the K = 7 combinatorics. Their product is a non-trivial property of the K = 7 cell. Theorem E below proves this cancellation is unique to K = 7.

5. Bridge Lemma: Connecting Spectral Structure and Binary Refinement

The previous sections establish two independent minimality results:

Entropy minimality. Lemma B shows that the smallest irreversible refinement of a distinguishability partition is binary: $1 \to 2$.

Spectral minimality. Theorem D shows that the smallest nontrivial admissible spectral sector of the closure cell has eigenvalue $\lambda^* = 2$ with multiplicity two.

These results arise from different structures --- combinatorial distinguishability and graph spectral theory. To complete the derivation we must show that they refer to the same primitive commitment event. The following lemma establishes that connection.

Bridge Lemma F (Minimal Spectral Realization of Binary Refinement). *Let a primitive commitment event be defined as the smallest irreversible binary refinement of a distinguishability partition (Lemma B). Then the λ* = 2 eigenspace of the K = 7 closure cell is the unique lowest-roughness spectral sector capable of realizing such a refinement.*

Proof.

Step 1 --- Physical motivation: primitive commitment requires at least two distinguishable directions.

This step establishes the physical picture motivating the lemma; a full proof would require deriving the relational representation theory from the VERSF axioms, which is beyond the scope of the present paper.

A binary partition refinement produces two mutually distinguishable alternatives. In relational frameworks of distinguishability, alternatives are encoded as directions in a representation space relative to the constant (indistinguishable) sector. Two mutually independent alternatives require two orthogonal directions: one separating the alternatives from the constant sector, and one distinguishing the alternatives from each other. A one-dimensional support space can only encode a total ordering --- it cannot represent two independent alternatives that are each non-trivially distinct from both the constant sector and from each other. This is consistent with the broader VERSF relational-geometry results that one-dimensional structures collapse relational distinctions, whereas at least two independent directions are required for nontrivial distinguishability structure.

Accepted as physical motivation: any primitive commitment configuration occupies a support space of dimension at least two. The subsequent steps are rigorous given this dimension bound.

Step 2 --- Minimal admissible support sector.

By construction of the admissible primitive class $\mathcal{A}$prim, any primitive configuration must lie outside ker(Lcell). By the Courant--Fischer theorem, the smallest admissible roughness outside the kernel occurs in the first nonzero eigenspace. For the $K = 7$ closure cell this eigenvalue is $\lambda^* = 2$, and the corresponding eigenspace has multiplicity two. The $\lambda^* = 2$ sector is therefore the lowest-roughness spectral sector providing the minimum dimensionality --- at least two --- required for binary distinguishability.

Step 3 --- Exclusion of higher spectral sectors.

Any configuration with support in a higher eigenspace ($\lambda > 2$) has strictly larger closure action cost

$\Delta\tilde{S} = (r/2)\lambda > r$.

Such configurations violate the minimality requirement of primitive commitment and cannot represent the primitive refinement. ■

Conclusion. The $\lambda^* = 2$ eigenspace is the unique spectral sector simultaneously satisfying: non-triviality (orthogonality to the kernel); minimal roughness (first nonzero eigenvalue); and sufficient dimensionality (multiplicity two) to represent binary alternatives.

Interpretation. The bridge lemma shows that the two minimality results derived earlier are not coincidental:

- the combinatorial minimum $\tilde{\Theta}_0 = \ln 2$, and
- the spectral minimum $\lambda^* = 2$

both describe the same primitive closure event. The eigenvalue identifies the lowest-cost admissible closure excitation; the eigenspace multiplicity supplies the minimal relational structure required to encode the two distinguishable outcomes of the binary refinement. The bridge is not inferred from eigenspace dimension alone, but from the conjunction of binary minimality, relational support requirements, and roughness minimality. The direction of implication runs from binary distinguishability to the spectral sector --- not the reverse. Binary commitment requires at least two relational directions; the $\lambda^* = 2$ eigenspace is the minimal spectral sector with that property. Thus the entropy and spectral analyses converge on a single primitive commitment structure.

6. Wheel Uniqueness

Theorem E (Uniqueness of Primitive Spectral Normalization). *Let WK denote the wheel graph on K vertices total: one central hub vertex connected to a cycle of K − 1 perimeter vertices (so the cycle has K − 1 nodes and K − 1 edges). Let* λ(K) be the smallest nonzero eigenvalue of the unweighted combinatorial Laplacian of WK. Then*

$(1/2) \cdot \lambda^*(K) = 1$ if and only if $K = 7$.

Proof. The Laplacian of WK block-decomposes into a hub sector and a perimeter (cycle) sector (see Appendix for derivation). The perimeter modes satisfy

$\lambda m(K) = 3 - 2\cos(2\pi m/(K-1)),\ m = 1, ..., K-2,$

and the hub mode has eigenvalue K. For $K \geq 4$ the hub eigenvalue $K > 3 \geq \lambda m(K)$ for small m, so the smallest nonzero eigenvalue is a perimeter mode at $m = 1$:

$\lambda^*(K) = 3 - 2\cos(2\pi/(K-1)).$

Verification for K = 7: $K - 1 = 6$, so the perimeter modes are $\lambda m = 3 - 2\cos(\pi m/3)$ for $m = 1, ..., 5$. At $m = 1, 5$: $\lambda = 3 - 2\cos(\pi/3) = 3 - 1 = 2$ ✓. At $m = 2, 4$: $\lambda = 3 - 2\cos(2\pi/3) = 3 + 1 = 4$ ✓. At $m = 3$: $\lambda = 3 - 2\cos(\pi) = 5$ ✓. Hub mode: $K = 7$ ✓. This recovers the spectrum $\{0, 2, 2, 4, 4, 5, 7\}$.

The condition $(1/2)\cdot\lambda^*(K) = 1$ becomes $\cos(2\pi/(K-1)) = 1/2$. The general solution is $2\pi/(K-1) = \pm\pi/3 + 2\pi n$. For $K \geq 4$ (cycle has at least 3 vertices), the unique integer solution is $2\pi/(K-1) = \pi/3$, giving $K - 1 = 6$, hence $K = 7$. ■

For the full wheel-cell family:

$(1/2) \cdot \lambda^*(K) > 1$ for $K < 7$ (over-committed), $(1/2) \cdot \lambda^*(K) = 1$ for $K = 7$ (exact normalization), $(1/2) \cdot \lambda^*(K) < 1$ for $K > 7$ (under-committed).

$K = 7$ is the unique crossing point. The spectral cancellation $\Delta\tilde{S}prim = r$ is a structural property of $K = 7$ alone.

The significance of the $K = 7$ closure cell is therefore twofold. Spectrally, it is the unique wheel graph whose primitive eigenvalue satisfies the normalization condition $(1/2)\lambda^* = 1$. Informationally, its two-dimensional primitive eigenspace is the minimal relational structure capable of supporting the binary partition refinement required for irreversible commitment. The coincidence of these two minimalities --- spectral and informational --- is what allows the primitive closure cost $\Delta\tilde{S}prim = r$ to be derived rather than postulated.

7. The Barrier Formula and Physical Normalization

7.1 Proposition --- Barrier Formula (Unconditional)

Define the entropy quantum Θ_0 as the physical entropy cost of one primitive commitment event, and let $\eta := \Theta_0/(kB \ln 2)$. The physical entropy cost of the primitive mode is then

$$\Delta Sc = (\Theta_0 / \tilde{\Theta}_0) \cdot \Delta\tilde{S}prim,$$

where $\tilde{\Theta}_0 = \ln 2$ (Corollary C) is the closure-unit primitive entropy quantum. Since the primitive entropy quantum in closure units is $\tilde{\Theta}_0 = \ln 2$, the ratio $\Theta_0/\tilde{\Theta}_0$ converts closure entropy units into physical entropy units. Since $\Delta\tilde{S}prim = r$ (Spectral Lemma), this gives $\Delta Sc = (\Theta_0 / \ln 2) \cdot r$.

Proposition (Barrier Formula).

$\Phi c = \eta\, r$. [unconditional]

Proof. From the definitions above, $\Delta Sc = (\Theta_0/\tilde{\Theta}_0)\cdot\Delta\tilde{S}prim = (\Theta_0/\ln 2)\cdot r$. The commitment barrier is

$$\Phi c = \Delta Sc / (kB \ln 2) = \Theta_0\, r / (kB \ln 2) = \eta r. \ ■$$

7.2 Conditional Theorem G --- Physical Entropy Identification

Conditional Theorem G (Physical Entropy Identification). *If VERSF closure entropy is identified with the physical thermodynamic entropy of irreversible record formation --- that is, if VERSF distinguishability classes coincide with the microstate classes of statistical mechanics --- then*

$$S = kB\, \tilde{S},$$

and therefore for primitive binary commitment

$\Theta_0 = k_B \ln 2$, $\eta = 1$.

Proof. Theorem A establishes $\tilde{S}(N) = \ln N$ unconditionally. The identification of VERSF classes with Boltzmann microstates means both $\tilde{S}(N) = \ln N$ and $S_B(N) = k_B \ln N$ are computed over the same N states. By the Matching Lemma (Section 8.1), since both are the unique admissible entropy on that shared structure up to a positive constant, proportionality forces $\alpha = k_B$ and hence $S = k_B \tilde{S}$. Applying to Corollary C: $\Theta_0 = k_B \cdot \tilde{\Theta}_0 = k_B \ln 2$. Then $\eta = \Theta_0/(k_B \ln 2) = 1$. ■

Support chain. The physical entropy identification is supported by four distinct lines of argument already developed in the VERSF programme:

1. **Binary irreducibility / minimal record.** Lemma B establishes that the minimal irreversible record corresponds to a binary partition refinement. The smallest irreversible fact is one bit. This is not an assumption --- it follows from the admissibility axioms.
2. **Structural counting gives ln 2.** Corollary C derives $\tilde{\Theta}_0 = \ln 2$ in closure units unconditionally from Theorem A and Lemma B. The dimensionless value is structurally forced, not postulated.
3. **Thermodynamic commitment cost gives kB ln 2.** The Matching Lemma shows that once VERSF closure states are identified with thermodynamic microstates, the conversion factor k_B is uniquely determined. The physical cost $\Theta_0 = k_B \ln 2$ then follows from the structural value $\tilde{\Theta}_0 = \ln 2$, not from an independent postulate.
4. **Bekenstein threshold route.** Theorem I derives the invariant $\chi(L)$ from the action budget, and the coherence scale satisfies $\chi(\xi) = 1$ --- marking the threshold at which one primitive commitment event first becomes admissible. This independently requires one-bit threshold behaviour at the scale ξ, consistent with $\Theta_0 = k_B \ln 2$ as the per-event entropy cost.

Together, these four lines converge on the same threshold. The identification is not a single isolated assumption --- it is the joint conclusion of multiple independent structural arguments within the programme.

What this does and does not establish. Proportionality $S = \alpha \tilde{S}$ is structurally forced --- it is not a choice. The value $\alpha = k_B$ follows from the identification of state classes, which is a physical interpretation claim rather than an empirical calibration. What the theorem does *not* do is derive k_B from pure VERSF combinatorics without any external input: k_B is a measured constant and its value enters through the Boltzmann formula.

Remark on Landauer. $\Theta_0 = k_B \ln 2$ exactly saturates Landauer's bound $\Delta S \geq k_B \ln 2$ for bit erasure. The structural counting gives $\tilde{\Theta}_0 = \ln 2$ unconditionally; Theorem G converts this to $k_B \ln 2$, which thermodynamics separately requires as a minimum. The derivation and Landauer run in opposite directions and meet at the same value.

7.3 Conditional Corollary H

Conditional Corollary H *(conditional on Theorem G).*

$\Phi c = r$.

Proof. By Theorem G, $\eta = 1$. Substituting into the Proposition: $\Phi c = 1 \cdot r = r$. ■

Since the closure-stiffness paper establishes $r = \mathcal{O}(1)$, this gives $\Phi c \sim \mathcal{O}(1)$ conditionally. The unconditional floor is the spectral statement $\Delta\tilde{S}prim = r$.

7.4 Falsification Paragraph

The conditional results depend on three structural properties. If any of the following fails, the derivation weakens to $\Phi c = \eta r$ with $\eta \neq 1$:

- **If closure entropy is not logarithmic** (Theorem A fails --- one of conditions (i)--(iii) is violated): the primitive entropy quantum $\tilde{\Theta}_0$ changes, and η acquires a correction factor from the non-logarithmic entropy measure.
- **If primitive commitment is not binary-minimal** (Lemma B fails --- primitive events involve 3 or more alternatives simultaneously): $\tilde{\Theta}_0 = \ln N$ for some $N > 2$, and $\eta = \ln N / \ln 2 > 1$.
- **If condition (iv) of Theorem G fails** --- that is, if VERSF distinguishability classes do not coincide with Boltzmann microstate classes --- then the Matching Lemma cannot be applied, $\alpha \neq kB$, and $\eta \neq 1$. In this case $\Phi c = \alpha r/kB$, with the numerical value depending on the ratio of the two entropy scales.

Each failure mode is distinguishable in principle: they predict specific deviations $\eta \neq 1$ that would show up as $\Phi c \neq r$ in experimental measurements of both Ec and Φc.

8. The Commitment-Capacity Chain

8.1 The Physical Normalization of Closure Entropy

The gap between the unconditional and conditional tiers is precisely: what fixes the constant α in $S = \alpha \tilde{S}$? This section shows that proportionality is structurally forced, that α is fixed by a structural identification rather than a calibration, and states the residual assumption precisely.

Step 1: Proportionality is forced.

Both $\tilde{S}$ and S represent the same ordering of distinguishability refinements. $\tilde{S}$ satisfies additivity for independent subsystems, monotonicity under refinement, and the null-singleton condition. Any physical entropy S describing the same distinguishability structure must satisfy the same three conditions. Since both are logarithmic in N by Theorem A, and both encode the same ordering, they must be related by a constant scale factor:

$$S = \alpha \tilde{S} = \alpha \ln N$$

for some universal constant $\alpha > 0$. This is structurally forced --- not a choice.

Step 2: The Matching Lemma fixes $\alpha = kB$.

Matching Lemma. *Suppose the distinguishability classes of the VERSF framework coincide with the microstate classes of statistical mechanics --- that is, two states are VERSF-distinguishable if and only if they belong to different Boltzmann microstates. Then for any distinguishability class of size N,*

$\tilde{S}(N) = \ln N$ (VERSF closure entropy, Theorem A), $SB(N) = kB \ln N$ (Boltzmann entropy),

and both entropies are computed over the same set of N states. Since $\tilde{S}$ and SB are both logarithmic entropies on the same shared distinguishability structure, uniqueness up to a positive multiplicative constant implies proportionality; comparing with Boltzmann's formula fixes that constant to kB.

$SB = kB\,\tilde{S}$, hence $\alpha = kB$. ■

The conditionality of Theorem G therefore rests on a single structural identification: that VERSF distinguishability classes and Boltzmann microstate classes are the same objects. This is a physical interpretation claim about what VERSF states *are* --- it is more precise than simply asserting empirical agreement, and its locus is exactly the question of whether the VERSF zero-entropy substrate recovers statistical mechanics.

What this does not prove. The Matching Lemma does not derive kB from pure VERSF combinatorics without external input. The Boltzmann entropy $SB = kB \ln N$ is an empirically established formula, and kB is a measured constant. The lemma shows that *if* VERSF distinguishability classes coincide with microstates, then $\alpha = kB$ follows structurally --- but the identification of the two class systems is the conditional step. The real achievement remains: deriving the *form* $\tilde{S} = \ln N$ from axioms (Theorem A), with the *scale* fixed once the physical interpretation is granted.

This identification links closure entropy to thermodynamic entropy at the level of commitment events, not general dynamical evolution.

8.2 Theorem I --- Derivation of the Quartic Commitment-Capacity Invariant

Theorem I (Quartic Commitment-Capacity Invariant). *Let a bounded causal region (in the sense of a light-crossing domain) of size L contain energy density ρ. Under the VERSF axioms of finite distinguishability and finite localisation capacity, and using the quantum action scale ħ, the maximum number of irreversible commitment events the region can support per causal crossing scales as $\rho L^4/(\hbar c)$. The natural dimensionless commitment-capacity invariant is, up to order-one convention-dependent factors,*

$\chi(L) = \rho\, L^4 / (\hbar c)$.

χ(L) should be understood specifically as the capacity for irreversible commitment (fact-formation) events, not as a general measure of causal influence or computational capacity. Nor does it quantify the persistence or stability of records once formed, which is governed by separate dynamical conditions (persistence capacity).

Definition (Capacity Stratification). Within VERSF, three distinct capacities govern physical systems:

- **Coherence capacity:** reversible distinguishability in the pre-commitment layer.
- **Commitment capacity $\chi(L)$:** the ability to produce irreversible facts; quantified by $\chi(L) = \rho L^4/(\hbar c)$.
- **Persistence capacity:** the stability of committed records over time, governed by separate dynamical conditions.

These capacities operate at different stages of physical reality and should not be conflated.

The quartic power L^4 is forced by the action budget.

Proof.

Step 1: Entropy budget of a causal region. By Theorem A, a region capable of resolving N distinguishable commitment alternatives carries closure entropy $\tilde{S} = \ln N$. The total energy contained in a region of size L with energy density ρ is

$E_L \sim \rho L^3$.

Step 2: Action budget. The causal crossing time --- the time for information to traverse the region --- is $\tau_L = L/c$. The available action during one causal crossing is therefore

$\mathcal{A}_L = E_L \cdot \tau_L = \rho L^3 \cdot (L/c) = \rho L^4 / c$.

Step 3: Convert action to an upper bound on distinguishability capacity. The quantum of action is $\hbar$. The Margolus--Levitin theorem provides an upper bound on the rate of distinguishable state transitions: a system of energy E can produce at most $E/(\pi\hbar/2)$ orthogonal state transitions per unit time. This gives an upper bound on the maximum number of distinct primitive commitment events proportional to $\mathcal{A}_L/\hbar$:

$N_{max} \lesssim \mathcal{A}_L / \hbar = \rho L^4 / (\hbar c)$, (up to order-one factors from the Margolus--Levitin bound).

The commitment-capacity invariant $\chi(L)$ measures this upper bound on commitment capacity (fact-formation capacity) per causal crossing. The coherence scale ξ defined by $\chi(\xi) = 1$ is therefore the scale at which the capacity *bound* equals one primitive commitment event --- not necessarily where capacity equals one exactly. Within the framework this is the natural definition of the primitive commitment scale, since it marks the smallest region for which the bound permits one irreversible fact.

Step 4: The quartic scaling follows from the action budget, not merely energy. The key point is that commitment capacity --- the capacity for irreversible fact formation --- is determined by the *action budget* $\mathcal{A}$L, not by the energy EL alone. Energy determines how many state transitions per unit time are possible, but it is the product of energy and available time --- the action --- that determines how many total transitions occur within the causal cell. Doubling the energy while halving the causal time leaves the action (and hence the commitment capacity) unchanged. This is why the relevant scaling is

$N_{max} \sim \mathcal{A}L/\hbar = (\rho L^3)(L/c)/\hbar = \rho L^4/(\hbar c)$,

with L^4 rather than L^3. The table makes the forced scaling explicit:

Quantity Scaling

Energy in region L^3

Causal crossing time L

Action budget = energy × time L^4

Commitment capacity ∝ action/$\hbar$ L^4

No other exponent is consistent with commitment capacity being proportional to the action budget.

Define the natural dimensionless commitment-capacity invariant (up to order-one convention-dependent factors from the Margolus--Levitin prefactor):

$\chi(L) := \rho L^4 / (\hbar c)$. ■

Remark on conditionality. Theorem I is unconditional within VERSF: it uses only the energy content of the causal region, the causal crossing time, and the quantum action scale $\hbar$. It does not require the entropy identification $S = k_B \tilde{S}$. The invariant $\chi(L)$ and the coherence scale ξ derived from it are therefore both fully established.

8.3 Corollary --- The Coherence Scale

Corollary (Primitive Commitment Defines the Coherence Scale). *A primitive commitment event is the smallest bounded causal region whose action budget first permits a primitive commitment event. Therefore $\chi(\xi) = 1$, giving*

$\xi = (\hbar c / \rho)^{(1/4)}$.

Proof. Set $\chi(\xi) = \rho \xi^4/(\hbar c) = 1$ and solve for ξ. ■

With vacuum-density estimate $\xi \approx 8 \times 10^{-5}$ m, the Wilson sector gives

$Ec \sim r \cdot \hbar c/\xi = r \cdot (\hbar c)$^(3/4) $\cdot\ \rho$^(1/4) $\approx r \times 2.5$ meV.

This energy scale is **unconditional in r**. The specific value 2.5 meV requires r = 1, which is conditional.

8.4 Theorem J --- Minimal Admissible Wheel Cell

The wheel-family uniqueness established by Theorem E --- $(1/2)\cdot\lambda^*(K) = 1 \Leftrightarrow K = 7$ --- raises the question of whether K = 7 can be derived as the minimal admissible closure cell more broadly, or whether the result is specific to the wheel family. This section states what is established, what remains open, and what the next step requires.

What is established

Theorem E already proves that within the wheel family WK, K = 7 is the unique cell satisfying the primitive spectral normalization condition. The five spectral conditions that characterise an admissible closure cell are:

1. **One-dimensional kernel:** dim ker(L) = 1 (single constant sector).
2. **Primitive eigenvalue exactly 2:** $\lambda_1 = 2$ (first nonzero eigenvalue).
3. **Twofold degeneracy:** mult(λ_1) = 2 (two independent binary commitment modes).
4. **Spectral gap:** $\lambda_2 > \lambda_1$ (primitive sector separated from higher roughness).
5. **Hub-compatible transport structure:** the graph admits a distinguished hub vertex mediating reversible transport across the cell.

For the wheel family, conditions 1--5 uniquely select K = 7. The explicit spectrum {0, 2, 2, 4, 4, 5, 7} satisfies all five, and no other WK does.

Theorem J (Wheel Minimality). *Within the wheel-graph family WK, K = 7 is the unique cell satisfying all five admissibility conditions above. In particular, it is the minimal admissible wheel cell.*

Proof. Conditions 1 and 5 are satisfied by all WK by construction (wheel graphs have one-dimensional kernels and a central hub). Condition 2 requires $\lambda^*(K) = 2$, which by the formula $\lambda^*(K) = 3 - 2\cos(2\pi/(K-1))$ holds if and only if K = 7 (Theorem E proof). Conditions 3 and 4 are verified by inspection of the K = 7 spectrum: mult($\lambda^* = 2$) = 2, and the next eigenvalue is 4 > 2. No WK with K < 7 satisfies condition 2; no WK with K > 7 satisfies condition 2 either. Therefore K = 7 is the unique, and hence minimal, admissible wheel cell. ■

What is not yet established

Theorem J does not prove that W_7 is the minimal admissible closure cell among *all* connected graphs. Three questions remain open:

1. **Must the closure cell belong to the wheel family?** The hub-plus-cycle structure is motivated by the closure-flow derivation, but the requirement that the cell be a wheel graph has not been derived from the VERSF axioms alone.
2. **Are there non-wheel graphs with fewer than 7 vertices satisfying all five conditions?** This requires a graph-spectral exclusion argument: systematically checking connected graphs on $n < 7$ vertices for conditions 1--5. Such a check is finite but has not been performed.
3. **Is the spectrum {0, 2, 2, 4, 4, 5, 7} uniquely realized by W_7 among admissible graphs?** Graph isomorphism and cospectral graph theory would be needed to answer this.

The honest current position

The strongest defensible statement is:

W_7 is the unique wheel cell satisfying the primitive spectral normalization condition $(1/2)\cdot\lambda^* = 1$. Within the natural hub-plus-cycle closure family, it is therefore the minimal admissible cell. Whether this uniqueness extends to all connected closure graphs is an open graph-spectral problem whose resolution requires an exclusion theorem for non-wheel candidates on fewer than 7 vertices.

The next step

The natural programme is to define the exclusion conditions formally:

$\dim \ker(L) = 1$, $\lambda_1 = 2$, $\mathrm{mult}(\lambda_1) = 2$, $\lambda_2 > 2$

and check all connected graphs on $n = 4, 5, 6$ vertices. Since the number of connected graphs on $n \leq 6$ vertices is finite, this exclusion problem is in principle a finite spectral classification problem --- tractable by direct enumeration or catalogue. If no such graph exists on $n < 7$ vertices, Theorem J would be elevated to a full uniqueness theorem without the wheel-family restriction.

Consistency Theorem. *Suppose:*

(i) The primitive spectral normalization gives $\Delta\tilde{S}prim = r$ (Spectral Lemma, unconditional).

(ii) The primitive entropy quantum in closure units is $\tilde{\Theta}_0 = \ln 2$ (Corollary C, unconditional).

(iii) The fact-formation dimensionless driving variable is normalized as $\Phi := \Delta\tilde{S}/\tilde{\Theta}_0$ --- measured in units of the primitive entropy quantum.

(iv) The entropy identification $S = kB\,\tilde{S}$ holds (Theorem G).

Then the closure-stiffness coefficient satisfies $r = 1$.

Proof. Condition (iii) fixes the closure-unit barrier to $\tilde{\Phi}c = \tilde{\Theta}_0/\tilde{\Theta}_0 = 1$ by definition of the normalization. Condition (iv) promotes this to $\Phi c = 1$ in physical units. Corollary H gives $\Phi c = r$. Therefore $r = 1$. ■

Remark on the logical structure. Condition (iii) is a normalization choice --- it sets the scale of the driving variable. It is natural (measuring commitment strength in units of the primitive commitment event itself), but it is a choice, not an independent constraint. Given that choice, $r = 1$ follows from condition (iv) alone: the entropy identification (iv) is the load-bearing conditional step, and (i) and (ii) supply the unconditional spectral and counting results that make the physical result meaningful. The theorem is best read as: *under the entropy identification and the natural primitive-unit normalization, the spectral, counting, and threshold sectors are mutually consistent at exactly r = 1.* The result $r = 1$ is therefore a primitive consistency value --- not a prediction derived simultaneously from three independent constraints.

9. Experimental Implications

9.1 Predictions at Two Strengths

Unconditional. The spectral identity $\Delta\tilde{S}prim = r$ and the energy scale

$$Ec \sim r \cdot \hbar c/\xi$$

are fully established. A measurement of Ec constrains r independently of the entropy identification.

Conditional (on the physical entropy identification of Theorem G). Theorem G gives $\eta = 1$, Corollary H gives $\Phi c = r$, and the Consistency Theorem selects $r = 1$ as the primitive consistency value. The framework therefore predicts $Ec \sim r \cdot \hbar c/\xi$. If the primitive consistency value $r = 1$ holds, this corresponds to

$$Ec \approx 2.5 \text{ meV}, \quad \Phi c \approx 1.$$

These are conditional predictions, not unconditional values.

9.2 The Two-Observable Test

Observable 1: $Ec \sim r \cdot \hbar c/\xi$ [unconditional; ≈ 2.5 meV if $r = 1$] Observable 2: $\Phi c = r$ [conditional on Theorem G]

Observable 1 is unconditional in its form; the specific value requires $r = 1$. Observable 2 is conditional throughout. A joint measurement tests the physical entropy identification and the primitive consistency value simultaneously. Any deviation $\Phi c \neq r$ from independent measurements of Ec and Φc would identify which of the conditions in the Consistency Theorem requires revision.

Most robust handle. Measuring Ec gives $r = Ec \cdot \xi/(\hbar c)$ without any use of the entropy identification. This is the experiment that can be designed and interpreted entirely within the unconditional tier.

Accessible scales. The millielectronvolt range corresponds to thermal energies at ∼30 K, overlapping with low-temperature quantum systems, Josephson junction physics, and precision atom interferometry.

9.3 Physical Intuition for K = 7

The spectral proof of Theorem J establishes K = 7 mathematically. It is also helpful to see why K = 7 is physically plausible as the minimal admissible closure cell. The spectrum {0, 2, 2, 4, 4, 5, 7} corresponds to four distinct spectral roles, each with a physical interpretation:

Eigenvalue Multiplicity Role

---

0 1 Constant sector --- the null/indistinguishable state

2 2 Primitive transport modes --- the two binary commitment outcomes

4, 5 3 Stabilizing modes --- buffer between primitive and UV sectors

7 1 UV closure mode --- suppressed by the Wilson penalty

Counting: 1 + 2 + 2 + 2 = 7. The four roles require a minimum of seven independent spectral degrees of freedom. This is not a proof --- the spectral exclusion argument of Theorem J does the rigorous work --- but it makes the result intuitive: K = 7 is the smallest cell that can simultaneously hold a null state, two binary commitment channels, a stabilizing buffer, and a UV closure mode without collapsing any of these sectors together.

9.4 The Unified Derivation Chain

Before the formal summary, it is useful to see the entire argument as a single vertical chain, with the conditional gate clearly marked:

VERSF axioms: finite distinguishability, irreversible commitment

↓

Theorem A: $\tilde{S}(N) = \ln N$ [logarithmic form forced]

↓

Lemma B: $1 \to 2$ binary minimality

↓

Corollary C: $\tilde{\Theta}_0 = \ln 2$ [primitive entropy quantum, closure units]

↓

Theorem I: $\chi(L) = \rho L^4/(\hbar c)$ [action budget]

↓

Corollary: $\chi(\xi) = 1 \Longrightarrow \xi = (\hbar c/\rho)^{(1/4)}$

↓

Theorem J: $K = 7$ minimal admissible wheel cell

↓

Theorem D: $\lambda^* = 2$ [Rayleigh--Ritz over $\mathcal{A}$prim]

↓

Spectral Lemma: $\Delta\tilde{S}$prim = r [exact cancellation]

↓

Theorem E: $(1/2)\cdot\lambda^*(K) = 1 \Leftrightarrow K = 7$ [structural]

↓

Bridge Lemma F: $\lambda^* = 2$ mode $\leftrightarrow$ binary partition refinement

↓

Proposition: $\Phi c = \eta\, r$ [unconditional barrier formula]

↓

── CONDITIONAL GATE: physical entropy identification

Theorem G: $S = kB\,\tilde{S}$ [VERSF classes = Boltzmann microstates;

supported by binary irreducibility, structural counting, Matching Lemma, Bekenstein threshold]

↓

Corollary H: $\Phi c = r$

↓

Consistency: $r = 1$ [primitive consistency value]

↓

Final: $\Phi c = r = 1$, $Ec \sim \hbar c/\xi \approx 2.5$ meV [conditional]

Everything above the gate is unconditional. The single conditional step is the identification of VERSF distinguishability classes with the microstate classes of irreversible record formation.

10. Summary: Logical Structure

10.1 The Full Theorem Chain

── UNCONDITIONAL

Theorem A: $\tilde{S}(N) = \ln N$ [monotonicity + additivity + null singleton]

Lemma B: Smallest nontrivial irreversible refinement: $1 \to 2$

Corollary C: $\tilde{\Theta}_0 = \ln 2$ [primitive entropy quantum in closure units]

Theorem I: $\chi(L) = \rho L^4/(\hbar c)$ [action budget; L^4 forced]

Corollary: $\chi(\xi) = 1 \Longrightarrow \xi = (\hbar c/\rho)^{(1/4)}$ [coherence scale derived]

Theorem J: $K = 7$ minimal admissible wheel cell

[five spectral conditions; unique in wheel family;

broader exclusion open]

Theorem D: $\lambda^* = 2$ [Rayleigh--Ritz over $\mathcal{A}$prim; K=7 spectrum]

Spectral Lemma: $\Delta\tilde{S}prim = (r/2)\cdot\lambda^* = r$ [exact cancellation]

Theorem E: $(1/2)\cdot\lambda^*(K) = 1 \Leftrightarrow K = 7$ [cancellation structural]

Bridge Lemma F: binary commitment requires ≥2 relational directions

→ $\lambda^* = 2$ eigenspace is the unique minimal spectral realization

Proposition: $\Phi c = \eta\, r$ [unconditional barrier formula]

── CONDITIONAL ON PHYSICAL ENTROPY IDENTIFICATION

Support: (1) binary irreducibility → 1-bit minimal record [Lemma B]

(2) structural counting → $\tilde{\Theta}_0 = \ln 2$ [Corollary C]

(3) Matching Lemma → $\alpha = kB$ [Section 8.1]

(4) Bekenstein threshold → 1-bit threshold at ξ [Theorem I]

Thm G: $S = kB\, \tilde{S}$, $\Theta_0 = kB \ln 2$, $\eta = 1$

Cond. Corollary H: $\Phi c = r$

Consistency Thm: $r = 1$ [primitive consistency value under entropy identification

and natural normalization; (iv) is the sole load-bearing condition]

── INDEPENDENT PHYSICAL PREDICTION

$Ec \approx r \times 2.5$ meV [unconditional in r; $r = 1$ conditional]

10.2 Status of Each Result

Result Status

$\tilde{S}(N) = \ln N$ --- forced logarithmic form **Unconditional** --- Theorem A

$1 \to 2$ binary minimality **Unconditional** --- Lemma B

$\tilde{\Theta}_0$ = ln 2 in closure units **Unconditional** --- Corollary C

$\chi(L) = \rho L^4/(\hbar c)$ --- L^4 forced by **Unconditional** --- the action budget Theorem I

$\xi = (\hbar c/\rho)^{(1/4)}$ from $\chi(\xi) = 1$ **Unconditional** --- Theorem I Corollary

K = 7 minimal admissible wheel cell **Unconditional** --- Theorem J (wheel family); broader uniqueness open

$\mathcal{A}$prim defined without reference to λ* **Unconditional** --- Section 4.1

λ* = 2 by Rayleigh--Ritz over $\mathcal{A}$prim **Unconditional** --- Theorem D

$\Delta\tilde{S}$prim = r --- no residual factor **Unconditional** --- Spectral Lemma

$(1/2)\cdot\lambda^*(K) = 1 \Leftrightarrow K = 7$ **Unconditional** --- Theorem E

Bridge Lemma F: binary commitment requires **Unconditional** --- ≥2 relational directions; λ* = 2 is the Bridge Lemma F unique minimal spectral realization

$E_c \sim r \cdot \hbar c/\xi$ (form) **Unconditional** in r --- Wilson sector + Theorem I

$\Phi_c = \eta r$ **Unconditional** --- Proposition

If two entropy measures share the same **Unconditional** --- distinguishability structure, they must be structural result, Section 8.1 proportional

VERSF distinguishability classes coincide **Conditional** --- with Boltzmann microstate classes physical interpretation claim (Theorem G)

$\alpha = k_B$ (from Boltzmann matching given **Conditional consequence** shared classes) --- Matching Lemma

$\Theta_0 = k_B \ln 2$, $\eta = 1$ **Conditional** --- Theorem G (Physical Entropy Identification)

$\Phi_c = r$ **Conditional** --- Corollary H

r = 1 --- primitive consistency value under **Conditional** --- entropy identification and natural Consistency Theorem normalization

Ec ≈ 2.5 meV (requires r = 1) Conditional --- physical entropy identification

10.3 Unified Statement

Unconditional:

$\Delta\tilde{S}prim = r$, $Ec \sim r \cdot \hbar c/\xi$

Conditional (on physical entropy identification, Theorem G):

$\Phi c = r = 1$, $Ec \approx 2.5$ meV [if r = 1 holds]

Throughout, $\chi(L)$ is used strictly in the sense of commitment capacity as defined in the Capacity Stratification above, distinct from both coherence capacity and persistence capacity.

The conditional tier rests on the Physical Entropy Identification of Theorem G: the claim that VERSF distinguishability classes are the microstate classes of irreversible record formation. This identification is supported by four distinct lines within the programme --- binary irreducibility, structural counting, the Matching Lemma, and the Bekenstein threshold route --- rather than resting on a single assumption. Proportionality $S = \alpha\,\tilde{S}$ is structurally forced; the Matching Lemma then fixes $\alpha = kB$ once the class identification is granted. Every other step is a combinatorial theorem, a functional-equation argument, an action-budget calculation, or a direct consequence of the VERSF axioms. Full uniqueness of K = 7 beyond the wheel family remains an open graph-spectral exclusion problem.

Appendix: Spectral Rigidity of the K = 7 Closure Cell

A.1 Derivation of the Wheel Graph Laplacian Spectrum

The wheel graph WK has K vertices: one hub vertex h connected to every vertex of a (K−1)-cycle C_{K−1}. The Laplacian L of WK can be written in block form relative to the hub and perimeter sectors.

Hub equation. The hub h has degree K − 1 (connected to all perimeter vertices). Its Laplacian equation is

$$(K-1)\, fh - \Sigma_i f_i = \lambda\, fh,$$

where f_i are the perimeter vertex values.

Perimeter equations. Each perimeter vertex i has degree 3 (two cycle neighbours and the hub):

$$3\, f_i - f_{i-1} - f_{i+1} - fh = \lambda\, f_i.$$

Decoupling by Fourier modes on the cycle. The (K−1)-cycle C_{K−1} has discrete Fourier modes

$$f_i\hat{}(m) = e\hat{}(2\pi i\cdot m\cdot i/(K-1)),\ m = 0, 1, ..., K-2.$$

For $m \neq 0$ (non-constant modes), the mode satisfies $\Sigma_i f_i\hat{}(m) = 0$, so the hub equation gives $fh = 0$. The perimeter equation for mode m then reduces to

$(3 - 2\cos(2\pi m/(K-1)))\ f_i\hat{}(m) = \lambda\ f_i\hat{}(m)$,

giving perimeter eigenvalues

$\lambda m = 3 - 2\cos(2\pi m/(K-1))$, $m = 1, ..., K-2$.

The m = 0 (constant) sector. For the $m = 0$ Fourier mode all perimeter values are equal: $f_i = a$ for all i. Two eigenvectors span this sector.

Zero mode. Setting $fh = a$ and $f_i = a$ gives the uniform vector. The hub equation: $(K-1)a - (K-1)a = 0$, and the perimeter equation: $3a - a - a - a = 0$. Both give $\lambda = 0$. ✓

Hub-breathing mode. Set $fh = (K-1)b$ and $f_i = -b$ (hub moves opposite to perimeter). Hub equation: $(K-1)(K-1)b - (K-1)(-b) = \lambda(K-1)b$, giving $(K-1)b + b = \lambda b$, so $\lambda = K$. Perimeter equation: $3(-b) - (-b) - (-b) - (K-1)b = \lambda(-b)$, giving $-b - (K-1)b = \lambda(-b)$, so $\lambda = K$. ✓

Therefore the $m = 0$ sector contributes eigenvalues 0 and K.

Summary for WK:

$\mathrm{Spec}(WK) = \{0\} \cup \{3 - 2\cos(2\pi m/(K-1)) : m = 1, ..., K-2\} \cup \{K\}$.

For $K = 7$ (cycle of length 6):

m	$\lambda m = 3 - 2\cos(\pi m/3)$
1, 5	$3 - 2\cdot(1/2) = 2$
2, 4	$3 - 2\cdot(-1/2) = 4$
3	$3 - 2\cdot(-1) = 5$

hub K = 7

This gives $\mathrm{Spec}(W_7) = \{0, 2, 2, 4, 4, 5, 7\}$, confirming the spectrum used throughout the paper.

A.2 Spectral Rigidity

The $K = 7$ closure cell Laplacian has spectrum $\{0, 2, 2, 4, 4, 5, 7\}$ fixed by the combinatorial structure of the unweighted $K = 7$ graph. For this unweighted model, the spectrum is a graph isomorphism invariant and is spectrally rigid in that sense.

This rigidity should not be overstated. If edge weights are allowed to vary, or a weighted Laplacian is used, eigenvalues move continuously. The present derivation uses the unweighted combinatorial Lcell throughout. Whether the framework is stable under controlled weighting perturbations is a question for the closure-cell construction paper.

Within the unweighted model, the spectral gap $\Delta\lambda = 2$ means that $\Phi c = \eta r$ cannot be continuously tuned by perturbations preserving the graph isomorphism class --- any shift in λ^* requires a change in combinatorial structure (e.g. changing K or the graph topology).

Derivation status summary. Theorem E establishes that $(1/2)\cdot\lambda^*(K) = 1$ is unique to $K = 7$ within the wheel family. Theorem J elevates this to a wheel-minimality theorem: W_7 is the unique wheel cell satisfying all five spectral admissibility conditions. Bridge Lemma F connects the $\lambda^* = 2$ mode to binary minimality. Theorem A forces the logarithmic entropy form. Theorem I derives $\chi(L) = \rho L^4/(\hbar c)$ from the action budget. Together the unconditional tier delivers $\Delta\tilde{S}prim = r$, $\tilde{\Theta}_0 = \ln 2$, and $Ec \sim r \cdot \hbar c/\xi$. The conditional tier, gated by $S = kB\ \tilde{S}$, delivers $\Phi c = r$ and $r = 1$. Full uniqueness of $K = 7$ beyond the wheel family --- requiring exclusion of non-wheel graphs on fewer than 7 vertices --- remains open.