Finite Distinguishability and Local Capacity Competition: A Structural Basis for Per-Channel Interaction Dynamics

Keith Taylor
VERSF Theoretical Physics Program

For the General Reader

Physics typically describes the world through smooth, continuous models. Yet quantum theory, thermodynamics, and information theory all point toward a different picture: reality may be built from a finite number of distinguishable states.

If so, physical systems are not infinitely flexible. They operate under genuine limits — only so much information can be stored, only so many configurations can exist, only so many independent events can occur simultaneously.

This paper develops a precise consequence of that idea:

When multiple processes depend on a shared finite constraint, they must compete for it — and that competition is necessarily local.

We show that this leads to a structural selection rule: second-order interaction effects cannot be described as global fluctuations. Instead, they must be expressed as local components competing for limited capacity. This yields a specific mathematical structure — governed by the inverse participation ratio — that determines how capacity is distributed across a system.

We show this principle is structurally consistent with local energy-curvature coupling in General Relativity. In quantum mechanics, it connects formally to the algebraic structure of the theory — in particular superselection rules, which exclude non-attributable operators from the physical observable algebra. It maps directly onto the BCB and TPB principles of the VERSF framework.

Abstract

We propose that **finite distinguishability** imposes a structural constraint on physical systems: shared constraints must act through local capacity allocation, inducing competition between participating degrees of freedom. We formalise two candidate second-order observables — a

global fluctuation term G_2 and a local competition term L_2 — and argue that only L_2 is consistent with the requirement that individual channel contributions remain operationally resolvable.

The distinction between G_2 and L_2 turns on the treatment of cross-channel covariance terms. We show that in a finitely distinguishable system, cross-channel terms correspond to jointly unresolvable contributions and must be excluded from any observable that admits per-channel attribution. This selects L_2 as the physically meaningful second-order observable, with structure:

$$L_2 = (\sum_i w_i^2) \cdot \mathrm{Var}(G)$$

where $\sum_i w_i^2$ is the **inverse participation ratio** (IPR) of the constraint distribution. We connect this structure to locality in General Relativity, to decoherence in quantum mechanics, and to the BCB and TPB principles within VERSF. Within VERSF, this provides a structural derivation — not a fit — for the per-channel second-order correction.

Contents

---

1. Introduction

Modern physical theories typically assume unlimited resolution: fields are continuous, states form infinite-dimensional Hilbert spaces, and probability distributions are unconstrained. Yet several independent lines of evidence point toward a fundamentally discrete or capacity-limited picture:

- **Quantum mechanics** places irreducible limits on joint measurement precision (Heisenberg uncertainty; complementarity).
- **Thermodynamics** bounds the information content of physical systems (Bekenstein bound; Landauer's principle).
- **Gravitational physics** enforces finite entropy for bounded regions (Bekenstein–Hawking entropy; holographic bound).

These observations converge on a single structural principle:

Finite Distinguishability: A physical system can realise only finitely many distinguishable states within any finite region.

From this, a further consequence follows immediately. If the total number of distinguishable states is finite, then constraints that multiple degrees of freedom draw upon simultaneously are **finite resources**: use by one channel genuinely reduces what is available to others.

This paper asks: *what is the second-order structure of an interaction observable in a system governed by finite distinguishability?*

We show that, within the per-channel observable class, the structure is uniquely fixed at this order. The resulting structure is governed by the inverse participation ratio of the constraint distribution, and it excludes global fluctuation models on structural grounds — not by assumption, but by the requirement that individual contributions remain operationally attributable.

2. Formal Definitions

2.1 Finite Distinguishability

A system Σ is **finitely distinguishable** if, for any finite region $\mathscr{R} \subset \Sigma$, the cardinality of the set of physically distinct, operationally separable states within $\mathscr{R}$ is finite:

$|\mathcal{S}(\mathscr{R})| < \infty$

2.2 Shared Constraint

Let G be a scalar constraint shared across n local degrees of freedom $\{X_1, ..., X_n\}$, with allocation weights $\{w_1, ..., w_n\}$ satisfying:

$\sum_i w_i = 1, w_i \geq 0$

The fraction w_i represents the share of G's capacity available to channel X_i. Because $|\mathcal{S}(\mathscr{R})|$ is finite, G itself represents a **bounded resource**: the total capacity available to the system is fixed.

2.3 Operational Resolvability

We define a second-order observable Ω to be **operationally resolvable** if there exists a measurement protocol — even in principle — that can assign a contribution Ω_i to each individual channel X_i, such that $\Omega = \sum_i \Omega_i$ and each Ω_i is independently attributable.

This definition is motivated by finite distinguishability directly: if contributions from distinct channels cannot be separated even in principle, those channels are effectively indistinguishable from each other with respect to the constraint, violating the foundational assumption.

This restriction mirrors a general feature of measurement theory: not all mathematically definable quantities correspond to observables. Only those associated with resolvable measurement outcomes enter the observable algebra. Operational resolvability is the condition that selects physical observables from the broader space of mathematical quantities definable on a system.

---

3. Two Candidate Second-Order Observables

We now construct the two natural candidates for a second-order observable built from the shared constraint G.

3.1 Global Fluctuation Observable

Treating the system as a collective object, one might write:

$G_2 = \mathrm{Var}(\sum_i w_i \cdot G)$

Since all channels share the same underlying variable G:

$G_2 = (\sum_i w_i)^2 \cdot \mathrm{Var}(G) = \mathrm{Var}(G)$

This is independent of the distribution $\{w_i\}$: it collapses all structural information about the allocation into a single number. The weights w_i are invisible in G_2.

3.2 Local Competition Observable

Treating each channel's contribution separately:

$L_2 = \sum_i \mathrm{Var}(w_i \cdot G) = (\sum_i w_i^2) \cdot \mathrm{Var}(G)$

This retains the distribution $\{w_i\}$ through the sum $\sum_i w_i^2$, and assigns to each channel a contribution:

$L_{2i} = w_i^2 \cdot \mathrm{Var}(G)$

which is independently attributable.

3.3 The Cross-Term Structure

The difference between G_2 and L_2 is the set of cross-channel covariance terms. Expanding G_2 explicitly:

$\mathrm{Var}(\sum_i w_i \cdot G) = \sum_i w_i^2 \cdot \mathrm{Var}(G) + 2 \sum_{i<j} w_i w_j \cdot \mathrm{Var}(G)$

So:

$G_2 = L_2 + 2 \cdot \mathrm{Var}(G) \cdot \sum_{i<j} w_i w_j$

The cross terms $2w_i w_j \cdot \mathrm{Var}(G)$ represent **joint fluctuations** between channels i and j. These arise because both channels draw on the same G, and therefore their fluctuations are correlated.

The question of which observable is physically meaningful thus reduces to: **should cross-channel joint fluctuations be included in the second-order physical observable?**

Theorem: Per-Channel Second-Order Selection under Finite Distinguishability

Let a shared constraint G be distributed across degrees of freedom $\{X_i\}$ with weights $\{w_i\}$, $\sum_i w_i = 1$, $w_i \geq 0$. Then, within the class of operationally resolvable ***per-channel*** *second-order observables in a finitely distinguishable system, the unique admissible observable is:*

$L_2 = \sum_i w_i^2 \cdot \mathrm{Var}(G)$

Any observable containing cross-channel covariance terms is inadmissible within this class.

The proof consists of three steps: (1) define the admissible observable class by operational resolvability; (2) show that cross terms fail this criterion because no measurement local to any single channel can recover them; and (3) show that the only permutation-invariant degree-two scalar in the per-channel class is $\sum_i w_i^2$. The proof proceeds via Section 4.

---

4. Exclusion of Cross-Channel Terms Under Finite Distinguishability

4.1 Why Cross Terms Are Present

Channels X_i and X_j are not statistically independent — they draw on the same G. So the cross-covariance $\mathrm{Cov}(w_iG, w_jG) = w_iw_j \cdot \mathrm{Var}(G)$ is non-zero. We do not claim otherwise.

The question is not whether cross-correlations exist in the underlying dynamics. They do. The question is whether a joint fluctuation term $w_iw_j \cdot \mathrm{Var}(G)$ can function as an **operationally resolvable observable**.

4.2 Cross Terms Fail Operational Resolvability

Consider the cross term for a specific pair (i, j):

$C_{ij} = 2w_iw_j \cdot \mathrm{Var}(G)$

This term is not attributable to channel i or channel j individually. It is a joint quantity: it exists only because both channels draw on G simultaneously. Any attempt to assign C_{ij} to channel i leaves channel j's contribution unresolved, and vice versa. The cross term is intrinsically **non-decomposable** into per-channel contributions.

Formally: there is no partition $C_{ij} = C_{ij}^{(i)} + C_{ij}^{(j)}$ that is operationally motivated — any split is arbitrary, not derived from the physics of the constraint.

4.3 Finite Distinguishability Requires Exclusion

In a finitely distinguishable system, any physical observable must correspond to a distinguishable difference in state. Cross terms, being unattributable to any individual channel, represent **jointly indistinguishable** contributions: they cannot be separated into distinct state-changes for X_i versus X_j. No measurement local to channel X_i or X_j alone can recover the cross term; only a joint measurement on the pair can do so — and such a measurement belongs to algebra (ii), not algebra (i).

This is not merely a bookkeeping issue. Under finite distinguishability, an unresolvable joint fluctuation between channels i and j is operationally equivalent to a single collective event — not two separate ones. Including it as part of a multi-channel second-order observable incorrectly attributes resolvability that the system does not possess. Equivalently, a cross term corresponds to a distinguishable fluctuation only of the pair *as a pair*, not of either channel individually; it therefore belongs to a different observable class than the per-channel second-order observable considered here.

Therefore:

Cross-channel terms must be excluded from operationally resolvable second-order observables in finitely distinguishable systems.

4.4 No-Go Statement for Global Observables

We can now state the result in its strongest form:

No-Go Statement: Any second-order observable that assigns weight to jointly unresolvable channel combinations violates finite distinguishability. Therefore, global fluctuation observables are excluded from the physical observable algebra of a finitely distinguishable system.

The global observable G_2 is mathematically well-defined but **physically inadmissible** under finite distinguishability, because it assigns variance to combinations of channels that cannot be distinguished by any operational protocol. It is not merely a less convenient description — it is a description that presupposes distinguishability the system does not possess.

This paper does not claim that cross-channel observables do not exist. It claims only that they do not belong to the algebra of per-channel second-order observables in a finitely distinguishable system. Cross-channel correlations are real; they are simply observables of a different class — pair-level rather than channel-level — and must be treated as such.

Proof sketch: Any observable containing cross terms requires a measurement that jointly resolves contributions from multiple channels simultaneously. Such a measurement collapses channel-level distinguishability — it cannot return a value attributable to channel i or channel j independently. This violates the definition of operational resolvability (Section 2.3). Therefore, observables containing cross terms cannot belong to the per-channel observable algebra. □

4.5 The Two Observable Algebras

We distinguish two observable algebras arising from a shared constraint:

(i) The per-channel observable algebra — consisting of quantities admitting channel-wise attribution; each element can be written as $\sum_i f_i(w_i, G)$ for some per-channel function f_i. L_2 belongs to this algebra.

(ii) The joint observable algebra — consisting of pair- or higher-order observables involving cross terms such as $w_i w_j \cdot \mathrm{Var}(G)$. G_2 belongs to this algebra.

The present paper concerns only algebra (i). Within the per-channel observable sector — the sector relevant to the VERSF application, where the physical quantities of interest are intrinsically channel-level — this restriction follows from finite distinguishability: FD requires that per-channel observables admit per-channel attribution, which excludes cross terms. This is not a claim that algebra (ii) is forbidden by FD; pair-level observables are consistent with FD and physically meaningful. It is a claim that, once one is working within the per-channel sector, FD forces the exclusion of cross terms from that sector's observable algebra. Joint observables in algebra (ii) exist, describe pair-level structure, and are left for future work.

Accordingly, the present result is not a no-go theorem for joint observables as such, but a **selection theorem for the per-channel observable sector**.

4.6 Consistency Check: The Counterfactual

Suppose cross terms were retained in a finitely distinguishable system. Then two distinct configurations — differing only in how contributions are allocated between channels i and j, while holding all per-channel totals fixed — would produce identical values of the observable. These configurations would be operationally indistinguishable despite having different channel-level structure. This contradicts finite distinguishability, which requires that distinct configurations correspond to distinct, separable states.

A referee might object: FD also admits distinguishable *pair-level* states — so why not retain cross terms and work in algebra (ii)? The answer is that algebra (ii) is not being denied. Pair-level observables exist and are physically meaningful. The present paper addresses a specific question: what is the structure of a second-order observable that tracks *individual channel contributions* to a shared constraint? That question is motivated by the VERSF application, in which the physical quantities of interest — update rates, information allocations, per-degree-of-freedom corrections — are intrinsically channel-level quantities, not pair-level ones. Algebra (i) is not the only possible algebra; it is the appropriate algebra for the class of observables the VERSF framework requires. Algebra (ii) remains available for pair-level analyses and is left for future work.

Cross-term retention within algebra (i) is therefore not merely inconvenient — it is inconsistent with the operational requirements of the per-channel observable class.

---

5. Selection of the Local Competition Observable

5.1 L_2 as the Consistent Per-Channel Observable

With cross terms excluded, the physically meaningful second-order observable is:

$L_2 = \sum_i w_i^2 \cdot \mathrm{Var}(G)$

This is operationally resolvable by construction: each term $w_i^2 \cdot \mathrm{Var}(G)$ is attributable to channel i individually, and $L_2 = \sum_i L_{2i}$.

Note that $L_2 \leq G_2$, with equality only when $n = 1$ (a single channel). For $n > 1$, the local competition observable is strictly smaller than the global fluctuation observable, reflecting the genuine reduction in per-channel capacity induced by competition.

5.2 Emergence of the Inverse Participation Ratio

The scalar prefactor in L_2 is:

$\mathrm{IPR} = \sum_i w_i^2$

This is the **inverse participation ratio**, a quantity well-established in condensed matter physics as a measure of distribution concentration. Here it plays a different but structurally motivated role: it measures how concentrated or distributed the constraint capacity is across channels.

At exactly second order — meaning the observable is a polynomial of degree two in the weights $\{w_i\}$ with no additional weight-dependent structure — $\sum_i w_i^2$ is the unique symmetric scalar satisfying all three conditions simultaneously: (a) invariance under permutation of channels, (b) compatibility with per-channel attribution (each term involves only one w_i), and (c) homogeneity of degree two. A candidate such as $\sum_i w_i^2 \cdot f(w_i)$ for a non-constant symmetric function f would be of degree higher than two, violating condition (c) at second order. A candidate such as $(\sum_i w_i)^2 = 1$ satisfies (a) and (c) but fails (b) because it cannot be decomposed into per-channel contributions. No other degree-two symmetric scalar over $\{w_i\}$ exists. The emergence of IPR is therefore not a convenience — it is the unique admissible second-order scalar.

- IPR = 1: all capacity concentrated in one channel ($n = 1$)
- IPR = $1/n$: capacity uniformly distributed across n channels

Thus: $L_2 = \mathrm{IPR} \cdot \mathrm{Var}(G)$.

5.3 Physical Interpretation of Competition

The suppression of L_2 below G_2 has a direct physical interpretation. When channel i uses fraction w_i of constraint capacity, only $1 - w_i$ remains. This reduces the effective variance available to all other channels. The cross terms in G_2 represent the capacity a channel *would* have if the constraint were not shared — they are the spurious contribution of treating channels as independent when they are not. Finite distinguishability forces their exclusion from the observable.

---

6. Connections to Known Physics

6.1 General Relativity

Multi-field stress-energy decomposition. When multiple matter fields φ_i are present, the stress-energy tensor decomposes as:

$T\mu\nu = \sum_i T^{(i)}\mu\nu$

and Einstein's equations give:

$G\mu\nu = (8\pi G/c^4) \cdot \sum_i T^{(i)}\mu\nu$

Now consider the quadratic scalar invariant $G\mu\nu G^\mu\nu$, which appears in quadratic gravity actions and gravitational wave energy expressions. Evaluated via the field equations under full index contraction, this invariant contains both diagonal per-field contributions $T^{(i)}\mu\nu\ T^{(i)}{}^\wedge\mu\nu$ and off-diagonal cross contributions $T^{(i)}\mu\nu\ T^{(j)}{}^\wedge\mu\nu$ $(i \neq j)$. Schematically:

$G\mu\nu\ G^\wedge\mu\nu \propto \sum_i T^{(i)}\mu\nu\ T^{(i)}{}^\wedge\mu\nu + 2\sum_{i<j} T^{(i)}\mu\nu\ T^{(j)}{}^\wedge\mu\nu$

where each term involves the full Lorentzian index contraction. The cross terms $T^{(i)}\mu\nu\ T^{(j)}{}^\wedge\mu\nu$ exist and contribute to invariants such as the Kretschmann scalar. But they are not attributable to field i or field j individually. Per-field attribution of second-order curvature response corresponds precisely to retaining only the diagonal terms $\sum_i T^{(i)}\mu\nu\ T^{(i)}{}^\wedge\mu\nu$; one natural choice of local weighting is given by the ratio of contracted per-field norms to the total contracted norm, $w_i = (T^{(i)}\mu\nu\ T^{(i)}{}^\wedge\mu\nu) / (T\mu\nu\ T^\wedge\mu\nu)$. This is not an analogy. It is the literal instantiation of the paper's G_2 / L_2 distinction in a gravitational context: the global fluctuation observable retains cross terms; the local competition observable does not.

Action locality as structural warrant. The Einstein-Hilbert action:

$S = \int d^4x\ \sqrt{-g}\ [R/(16\pi G) + \mathscr{L}\_matter]$

is manifestly local: the Lagrangian density at each spacetime point depends only on field values and their first derivatives at that point. Locality of the action entails that variation produces local field equations — there is no non-local mixing of contributions from spatially separated regions at the level of the action principle. Cross-field contributions to curvature scalars arise from the nonlinearity of the field equations, not from the action structure; they are second-order collective effects, not per-field observables. The action therefore provides a structural — not merely descriptive — reason why cross-field terms fall outside the per-channel curvature observable.

Bekenstein bound as finite distinguishability. The Bekenstein bound:

$S \leq 2\pi kRE/\hbar c$

bounds the information content — equivalently, the number of distinguishable states — of any finite region with energy E and radius R. This establishes that each spatially bounded region carries a finite state capacity: the GR instantiation of finite distinguishability in the sense of Section 2.1. The bound does not by itself forbid cross-region correlations, but combined with the locality of the action principle, it implies that cross-field contributions to per-region observables cannot be attributed to any individual field within that region — they arise from nonlinear mixing, not from the local capacity allocation. The Bekenstein bound and local field equations therefore jointly imply the finite, locally allocated constraint structure that the present paper formalises: finite capacity per region, allocated per field, with cross-field mixing falling outside the per-channel observable sector.

6.2 Quantum Mechanics

Superselection rules as structural cross-term exclusion. The most direct quantum-mechanical precedent for the paper's argument is not decoherence — which is dynamical — but superselection rules, which are algebraic. In algebraic quantum mechanics, a superselection rule partitions the Hilbert space $\mathscr{H} = \bigoplus_i \mathscr{H}_i$ into sectors between which no physical observable has matrix elements:

$\langle \psi_i | O | \psi_j \rangle = 0$ for $i \neq j$, for all $O \in \mathcal{A}$

This is not dynamic suppression — it is a structural constraint on which operators belong to the physical observable algebra $\mathcal{A}$. Cross terms between sectors exist as mathematical objects in the full operator algebra, but they are excluded from $\mathcal{A}$ on operational grounds: no measurement protocol can resolve a transition between superselection sectors (Wick, Wightman, Wigner). This is structurally identical to the paper's cross-term exclusion: joint covariance terms exist in the underlying mathematics but are inadmissible in the per-channel observable algebra because they fail operational attributability.

Schmidt decomposition and IPR. For any bipartite pure state $|\psi\rangle \in \mathscr{H}_A \otimes \mathscr{H}_B$, the Schmidt decomposition gives:

$|\psi\rangle = \sum_i \sqrt{\lambda_i}\, |a_i\rangle \otimes |b_i\rangle$

The reduced density matrix is $\rho_A = \sum_i \lambda_i |a_i\rangle\langle a_i|$, which is diagonal in the Schmidt basis — cross terms between Schmidt channels are structurally zero, not dynamically suppressed. The purity of ρ_A is:

$\mathrm{Tr}(\rho_A^2) = \sum_i \lambda_i^2 = \mathrm{IPR}(\{\lambda_i\})$

This is the inverse participation ratio of the Schmidt spectrum, appearing here not as a descriptive statistic but as the natural measure of per-channel second-order structure in the operationally reduced state. The connection to L_2 is direct and formal: the per-channel second-order observable in the Schmidt decomposition is exactly the diagonal of ρ_A, with IPR($\{\lambda_i\}$) as its scalar summary.

Algebraic QFT locality. In algebraic quantum field theory (the Haag-Kastler framework), the fundamental object is a net of local algebras $\{\mathcal{A}(\mathcal{O})\}$ indexed by spacetime regions $\mathcal{O}$. The key axiom is local commutativity:

$[A, B] = 0$ for all $A \in \mathcal{A}(\mathcal{O}_1)$, $B \in \mathcal{A}(\mathcal{O}_2)$, $\mathcal{O}_1$, $\mathcal{O}_2$ spacelike separated

This enforces that observables in different regions belong to commuting subalgebras — the physical observable algebra is built from locally attributable pieces. Global observables that mix contributions from spacelike-separated regions are not primitive elements of any local algebra $\mathcal{A}(\mathcal{O})$; they are derived objects outside the operationally local structure. Cross-region second-order terms therefore fall outside the per-channel observable algebra at the foundational level of the theory — this is the AQFT version of the paper's cross-term exclusion, and it is fully structural.

Decoherence as dynamical approach to the algebraic structure. Decoherence — the suppression of off-diagonal density matrix elements $\rho_{ij} = \alpha_i\alpha_j^* \to 0$ $(i \neq j)$ through environmental entanglement — should be understood as the dynamical process by which a system approaches the state already selected by the observable algebra. The algebraic structure (superselection, local algebra commutativity) determines which observables are admissible; decoherence determines on what timescale the dynamics enforces that structure. The paper's argument operates at the algebraic level: cross terms are excluded from the per-channel observable algebra structurally, not because they decay.

6.3 VERSF: BCB and TPB Principles

Within VERSF, two foundational principles enforce finite capacity:

- **Bit Conservation and Balance (BCB):** The total information content of any finite region is conserved and bounded. No new distinguishable states can be created without destroying others.
- **Ticks Per Bit (TPB):** The update rate of any degree of freedom is bounded by its share of the global information budget. Faster evolution by one channel reduces available update rate for others.

Together, BCB and TPB enforce exactly the shared-constraint structure formalised in Section 2.2. The constraint G is the information budget; the weights w_i are TPB-allocated update fractions; and the competition structure follows from BCB's conservation requirement. The present paper provides the structural argument for why this must yield L_2 rather than G_2 as the observable.

6.4 A Common Thread: The Physical Observable Algebra as a Proper Subalgebra

Sections 6.1–6.3 share a structural pattern worth making explicit. In every well-developed physical theory, the physical observable algebra is a **proper subalgebra** of the full mathematical operator algebra:

- In GR, the locality of the Einstein-Hilbert action and the Bekenstein bound restrict which terms enter the physical field equations and which second-order curvature scalars are per-field attributable. The full mathematical tensor algebra contains cross terms; the per-field observable algebra does not.
- In QM, superselection rules and the local algebra structure of AQFT restrict which operators are physical. The full Hilbert space operator algebra contains cross-sector matrix elements; the physical observable algebra $\mathcal{A}$ does not.
- In VERSF, BCB and TPB restrict the information budget and update rate. The full mathematical state space contains global fluctuation terms; the per-channel observable algebra does not.

The operational resolvability requirement of Section 2.3 proposes that **finite distinguishability is the fundamental ground for this restriction** — of which GR locality, quantum superselection, and AQFT commutativity are specific implementations in their respective domains. This paper does not derive GR or QM from finite distinguishability. It argues that the structure they share — the exclusion of non-locally-attributable terms from the physical observable algebra — is a consequence of a single, more primitive principle.

---

7. Application to the VERSF Interface

7.1 The Six-Channel Structure

In the VERSF interface model, a single global constraint is distributed across **six local constraint channels**, corresponding to the minimal symmetry-preserving decomposition of the interface constraint established in the companion paper Taylor (in preparation, AIDA-VERSF-2025-02). The derivation of this six-channel structure is not reproduced here; we take it as an established result of the interface geometry and apply the local competition principle to it.

7.2 Uniform Allocation by Symmetry

Given the six-channel symmetry and the absence of any symmetry-breaking mechanism in the interface structure as established, the natural effective allocation is uniform:

$w_i = 1/6$ for all $i = 1, ..., 6$

7.3 The IPR and Resulting Correction

Substituting into the IPR formula:

$\text{IPR} = \sum_{i=1}^{6} (1/6)^2 = 6 \cdot (1/36) = 1/6$

The second-order correction to an intensive observable R therefore takes the form:

$\delta R \propto -(1/6) \cdot N^{-2}$

where N is the system size parameter of the VERSF interface model. This coefficient 1/6 is a **structural consequence** of the six-channel symmetry of the interface and the local competition principle — it is not fitted to data.

---

8. Relation to Existing Uses of IPR

The inverse participation ratio appears extensively in:

- **Anderson localisation**: IPR measures the spatial concentration of wavefunctions in disordered systems. High IPR signals localisation; low IPR signals delocalisation.
- **Quantum chaos and random matrix theory**: IPR characterises the spread of energy eigenstates across basis states.
- **Many-body localisation**: IPR distinguishes ergodic from non-ergodic phases.

In all these contexts, IPR is a *descriptive* statistic applied to an already-given probability distribution over a basis.

The present use is structurally distinct: here IPR arises as the **selected per-channel second-order scalar** from a derivation based on finite distinguishability and operational resolvability. It is not applied to a wavefunction or eigenstate distribution — it is derived as the natural form consistent with local capacity competition at second order. The connection to known IPR applications is suggestive of deeper structural unity, but the derivation stands independently.

---

9. Discussion

9.1 Scope of the Result

The argument presented here is structural, not dynamical. We have not derived the value of any physical constant, nor specified the particular form of the shared constraint G. What we have shown is:

Whenever a shared constraint governs multiple degrees of freedom in a finitely distinguishable system, the operationally resolvable second-order observable within the per-channel class is uniquely fixed as the local competition observable $L_2 = \text{IPR} \cdot \text{Var}(G)$.

This is a constraint on the form of second-order terms, not on their magnitude.

9.2 What the Argument Does and Does Not Assume

The argument assumes:

- Finite distinguishability (Section 2.1)
- The existence of a shared constraint with well-defined allocation weights (Section 2.2)
- The operational resolvability requirement (Section 2.3)

It does **not** assume:

- Any particular dynamics for G
- Statistical independence of channels (the cross-term exclusion is argued on resolvability grounds, not independence grounds)
- Any specific value of n or the weights w_i

9.3 The Independence Question

A potential objection is that $L_2 = \sum_i \mathrm{Var}(w_i G)$ formally requires treating the channels as independent. We address this directly: **L_2 is not obtained by assuming statistical independence. It is obtained by projecting the full second-order structure onto the subspace of observables that admit per-channel attribution.** The channels are correlated at the level of the shared constraint — that correlation is the source of competition. But the observable that a finitely distinguishable system can register is restricted to per-channel contributions. The cross terms exist in the underlying dynamics; they are excluded from the physical observable algebra precisely because they are not operationally resolvable.

The distinction is not between correlated and uncorrelated variables, but between observable and non-observable components of a correlated system. Correlation is a property of the dynamics; observability is a property of what measurements can resolve. These are independent conditions, and finite distinguishability constrains the latter.

---

10. Conclusion

We have established a structural principle governing second-order interaction terms in finitely distinguishable physical systems:

Shared constraints act through local capacity competition, not global fluctuation.

This principle:

1. Provides a positive physical definition of operational resolvability;

2. Shows that cross-channel covariance terms fail this definition and must be excluded from the per-channel observable algebra;
3. Selects the local competition observable L_2 = IPR · Var(G) as the consistent per-channel second-order observable;
4. Identifies the inverse participation ratio as the governing per-channel second-order scalar;
5. Is structurally consistent with local energy-curvature coupling in General Relativity, and connects formally to superselection-enforced diagonalisation in quantum mechanics — of which decoherence is the dynamical realisation;
6. Provides a structural derivation of the 1/6 coefficient in the VERSF interface correction.

The broader implication is that the **form** of second-order physical interactions — and not merely their magnitude — is constrained by the information-theoretic structure of the systems they act within. In a finite universe, the geometry of capacity distribution is itself a physical observable.