Bath or Ledger

The Mixing Question as a Fact About Pre-Factual Conserved Weight — Relocating Q-MIX to the Structure of Bit Balance

Keith Taylor VERSF Theoretical Physics Programme versf-eos.com

---

For the General Reader

Two companion papers reduced the puzzle of why quantum probability is the *square* of an amplitude to a small number of substrate questions. The sharpest survivor is this: does the substrate's pre-commitment dynamics include continuous motions that *mix* two distinct possible outcomes — coherently moving "possibility-weight" from one alternative into another — or only motions that leave each alternative's weight where it is and merely rotate its phase? Call this the mixing question. If mixing is present, the square follows (given the other ingredients); if absent, other exponents are allowed and reality is, at the level of reversible dynamics, classical.

It is tempting to settle this by appeal to meaning: surely, one wants to say, if the alternatives are *genuine* possibilities of one shared process — not yet committed facts — then the process must be able to turn one into another, so mixing is forced. This paper resists that temptation, because it is a trick of words. A perfectly ordinary classical die, before it lands, has six jointly-available alternatives sharing one total probability; landing selects one. That is a real possibility space, and nothing in it turns one face continuously into another. So "genuine possibility" cannot, by itself, force mixing — classical probability is a counterexample, and pretending otherwise just hides the assumption inside the word "genuine."

What this paper does instead is find the *structural* fact that actually decides the question, and show it is checkable rather than a matter of definition. The fact is this: before commitment, is the conserved possibility-weight kept as **one shared pool** that can be freely reallocated among the alternatives, or as **many separate accounts**, one per alternative, each conserved on its own? We call these the bath reading and the ledger reading. And we prove, cleanly, that mixing exists *exactly* in the bath case. If the weight is one shared pool, the dynamics can move it between alternatives — that motion *is* mixing. If it is separate accounts, the dynamics can only rephase each account in place — that is the classical, no-mixing case. So the whole mixing question turns on one sharp question about the substrate's conservation law: **one pool, or many accounts?**

This is honest progress rather than a triumph — and it goes one step further than "find the structural fact." When you chase the question down to the framework's deepest starting point — the rule that fixes how much can be distinguished inside a region of reality — something striking happens. That rule is written about a *region*, not about a list of pre-labelled outcomes: the capacity to tell things apart belongs to the unresolved whole, and the separate outcomes only come into being when a fact is committed. The "one pool or many accounts" question then becomes: is the way that regional capacity divides up *fixed* before commitment, or can the dividing lines *move* until commitment freezes one? And that turns out to be a question the programme had already run into and named, in a completely different context — the "refinement" question about whether every legitimate way of sub-dividing is one the basic picture can already see. The pool-or-accounts question is not a new question. It is that same refinement question, in disguise.

So the three companion papers, taken together, point toward something clean. The square would follow from just two things: one mild bookkeeping principle (that physics cannot depend on the arbitrary way you describe a configuration), and one already-known open question (the refinement question) — which, it now appears, is secretly responsible for *four* things at once: the refinement structure itself, the linearity of the dynamics, the mixing of outcomes, and how the basic capacity divides. Two honest cautions keep this from being more than it is. First, the claim that the mixing question and the refinement question are really *the same* question is, at the deepest level, an identification we have argued for but not yet fully proven — they concern different objects (the flow of weight on one hand, the sub-dividing of measurements on the other), and the bridge between those objects is a lemma still to be supplied, not a settled fact. Second, even granting that bridge, answering the one remaining question the "natural" way — that the dividing lines are alive and only settle at commitment, as the framework's whole picture of fact-formation suggests — is a *choice of reading*, not something the starting rule forces. The classical, rigid alternative remains logically available, and a universe built on it would simply be classical rather than quantum.

So the honest headline is not "the square is now derived, modulo one natural choice." It is: the many scattered-looking conditions for the square have been shown to be, very probably, *one* condition plus a mild principle — and that one condition has a coherent classical answer the universe might instead have taken. What the three papers achieve is not the square itself but a dramatic narrowing: the entire remaining mystery of why probability is a square now rests on a single, sharply stated question about how reality divides its possibilities before it commits to one — together with the lemma that would confirm the scattered hinges really are that one question.

Abstract

The Squaring Residue reduced the Born exponent to the existence of continuous off-diagonal reversible mixing of distinct outcomes (Q-MIX); the linearity companion showed the surrounding linearity reduces to a mild decomposition principle plus the refinement obstruction. This paper attacks Q-MIX. It declines the natural ontological argument — "genuine pre-factual possibilities must be interconvertible, hence mixing is forced" — on the ground that it settles the question by the meaning of "genuine," and that classical probability (jointly-available alternatives, no interconversion) is a standing counterexample to the inference. The decisive structure is not the meaning of possibility but the structure of the conserved pre-factual weight.

We introduce a sharp dichotomy on the conservation law. With per-outcome weight $w_i := h(|c_i|)$ and conserved normalization $N = \Sigma_i w_i$ (separable, inherited from the prior papers' BCB-additivity N1):

- **Ledger reading (L):** each w_i is separately conserved by pre-factual reversible dynamics — d independent accounts;
- **Bath reading (B):** only the total N is conserved — one shared, fungible pool, with individual w_i reallocable subject to fixed total.

The results are:

1. **(Ledger ⟹ Torus — proven.)** Under L, with h strictly monotone, every pre-factual reversible transformation preserves each $|c_i|$, hence is diagonal: the connected dynamics is the phase torus $\mathbb{T}$_d, with at most discrete outcome relabelling. No off-diagonal mixing. This is classical reversible probability.
2. **(Bath ⟹ Mixing — proven, modulo the continuity upgrade.)** Under B, pre-factual dynamics contains transformations that change individual w_i while fixing the total — weight reallocation between outcomes. Granted that pre-factual reversibility is continuous (the temporal-extensibility / holonomy-continuity upgrade the programme already flags), such reallocation generates a continuous off-diagonal generator: genuine mixing.
3. **(The Possibility-Connectivity Theorem, fiat-free.)** Combining: **off-diagonal mixing exists iff the bath reading holds.** Q-MIX is therefore equivalent to a structural fact about the conservation law — one shared bath versus many separate ledgers — not to a verdict about what counts as "genuine" possibility. The ontological intuition (that genuine connected possibility forces mixing) is correct *for the bath reading* and false for the ledger reading, and the two readings are distinguished by the conservation structure, not by definition.
4. **(Connectivity for free.)** If the bath is shared across *all* d outcomes (a single pool, not block-diagonal sub-pools), weight is reallocable between any pair, so the mixing graph is complete — discharging the connectivity/symmetry leg of the linearity companion's conjunctive hinge without extra assumption.
5. **(Relocation, then collapse.)** The bath/ledger question is the unmined surplus of Bit Conservation and Balance the prior papers set aside; additivity (N1) fixes that N is separable but is silent on whether the individual w_i are separately conserved. Pursued one level deeper to the packing primitive FP1 (which bounds the capacity of a *region*, not pre-assigned per-outcome accounts), the fork becomes whether a region's capacity is *rigidly* partitioned before commitment (ledger) or *dynamically* reallocable until commitment localizes one continuation (bath) — which we identify with the *Packing* paper's refinement obstruction (Obstruction B / FP3): **bath ⟺ dynamical partition** (near-definitional), and **dynamical partition ⟺ Obstruction B's reverse inclusion** (a *cross-object* identification — flows of weight versus decompositions of measurements — resting on a shared structural gloss, with the formal translating lemma outstanding). FP1 as written is silent between the rigid and dynamical readings — which is exactly why Obstruction B was left open.

The contribution is a relocation and a *unification*, not a closure, and conditional on that lemma: we do **not** assert the bath/dynamical reading (that would be the fiat in structural clothing, and FP1 does not force it — the region-sourced classical ledger is consistent with the axiom). We prove the bath reading equivalent to mixing, argue that pursued to FP1 it is the refinement obstruction (modulo the cross-object lemma), and exhibit the classical ledger as the coherent alternative. The upshot is a collapse of the decision tree — modulo that lemma — bringing the Born exponent, assembled across the three companion papers, to **$\ell^2 \Leftarrow$ PC ∧ Obstruction B**: one mild new decomposition principle (PC), and the single standing refinement obstruction, which on the strength of the identifications governs four things at once (ODG's refinement structure, transport linearity, the bath reading / mixing, and the dynamical-partition reading of FP1; the single-bath step supplied by one-region-one-pool). Affirming Obstruction B — reading admissible decomposition as the dynamical refinement of a single shared capacity, as the commitment ontology favours — cascades through all four to the square. Both the identification lemma and the affirmative reading are flagged, not smuggled: the honest collapse is "two conditions, one conditionally identified with the other," and a substrate taking the rigid reading would be classical.

Contents

1. Introduction: the mixing question and the temptation of fiat
2. Inherited setup
3. Two readings of pre-factual weight: bath and ledger
4. The Possibility-Connectivity Theorem
5. Why this is not the fiat: the ledger reading is classical probability
6. The bath reading delivers connectivity
7. Relocation to BCB-balance, and the collapse to the refinement obstruction
8. Assembly: the consolidated state of the Born exponent
9. Limitations and Open Problems
10. Conclusion

Scope and Conditional Status

This paper is the third companion to *The Squaring Residue*, alongside the linearity paper, and inherits their setup and conventions. It attacks the mixing question Q-MIX, relocates it to the structure of pre-factual conserved weight (bath versus ledger), and then — pursuing the fork to the packing primitive — identifies it with the *Packing* paper's refinement obstruction (Obstruction B). It does not *close* that obstruction; it shows the mixing question, the linearity question, and the partition structure of the primitive are one question, collapsing the Born exponent's residual dependence to a mild new principle (PC) plus the single standing refinement obstruction.

Epistemic labelling is maintained: **proven**, **conditional** (on a named, separately-open hypothesis or the continuity upgrade), **conjectural**. Results 1 and 3–4 are **proven** (result 2, and the "⟸" of result 3, **conditional** on the continuity upgrade); result 5's collapse is **conditional on Claim 7.1.1** (the identification of the bath/ledger fork with Obstruction B, asserted as an identification of content, formal proof outstanding — §9.5). The resolution of the obstruction itself — rigid (ledger, classical) versus dynamical (bath, quantum) partition — is **open** and handed to the BCB/TPB analysis. Time is emergent throughout; "pre-factual" denotes the substrate configuration *before* an irreversible commitment event, and "reversible dynamics" the substrate maps on that pre-commitment configuration, never motion through a temporal continuum.

A methodological note, because it is the point of the paper. There is a fast argument that mixing is forced: pre-factual alternatives are not yet separate facts, so a genuine shared possibility space must let them transform into one another, hence mixing. We reject this argument as stated. It settles the question through the word "genuine," and classical probability — jointly-available alternatives with no interconversion — satisfies every uncontroversial sense of "shared possibility space" while having no mixing. The fast argument therefore smuggles its conclusion into a definition, in exactly the way the programme has refused throughout (cf. the *Squaring* paper's refusal of "reversible = isometry," and the linearity paper's refusal to let "respects superposition" stand undischarged). This paper replaces the word-verdict with a structural one — bath versus ledger — that the substrate decides.

---

1. Introduction: the mixing question and the temptation of fiat

The companion papers leave the Born exponent resting on a short list of substrate conditions, of which the sharpest is **Q-MIX**: does substrate-admissible reversible pre-factual dynamics contain a continuous off-diagonal generator — a motion that coherently redistributes weight between two distinct possible outcomes — or is it exhausted by per-outcome phase rotation and discrete relabelling? The *Squaring Residue* paper showed Q-MIX (with linearity and connectivity) decides the squared norm; the linearity paper supplied the linearity leg. Q-MIX itself was left open and handed to the substrate.

There is an appealing way to settle Q-MIX from the programme's own ontology. The substrate's defining move is the commitment event: a fact is the irreversible selection of one outcome from a set of pre-factual alternatives. Before commitment, the argument runs, the alternatives are *not* separate facts — they are possibilities of one shared process — and a dynamics that could only rephase each alternative in place, never turning one into another, would be treating them as already-separate classical labels rather than as joint possibilities of a single pre-factual structure. So, the argument concludes, a genuine possibility-to-fact substrate must support continuous redistribution among alternatives, which is exactly off-diagonal mixing.

The conclusion may well be true. But the argument, as stated, does not establish it, and we must say why clearly, because the failure is instructive. The argument turns on the step: *if the dynamics only rephases and permutes, the alternatives are "merely a list of classical labels," not a "genuine shared possibility space."* This is a verdict delivered by the words "merely" and "genuine," not a derived fact. The decisive counterexample is ordinary classical probability. A fair die before it lands has six jointly-available pre-factual alternatives; they share one total weight (the probabilities sum to one); the landing is an irreversible commitment selecting one. By every uncontroversial criterion this is a shared possibility space with a commitment event — and its pre-commitment "dynamics" (reweighting, relabelling) contains no continuous interconversion of faces whatsoever. So "shared possibility space with commitment" does *not* entail mixing; classical probability is a standing witness. To get mixing out, the fast argument must define "genuine" possibility so that connectedness-as-interconvertibility is built in — and then it has assumed what it set out to prove.

This is the same failure mode the programme has policed repeatedly: a contested conclusion riding into a proof on the back of a word. The remedy is the one used before — find the *structural* condition the word was gesturing at, state it sharply, and let the substrate decide it. The structural condition here is not the meaning of "possibility." It is the **conservation structure of pre-factual weight**: whether the conserved quantity is one shared pool (reallocable among alternatives) or many separate accounts (each fixed). We show mixing is equivalent to the former. That equivalence is a theorem; which structure the substrate has is a separate, open, checkable question.

---

2. Inherited setup

We use the companion papers' framework.

- **The carrier (H0).** Pre-factual state $\psi = (c_1, \ldots, c_d) \in \mathbb{C}^d$, a path-sum amplitude vector over d operationally distinguishable outcomes.
- **The conserved normalization.** $N(\psi) = \Sigma_i h(|c_i|)$, separable and phase-blind (from BCB-additivity N1 and the diagonal phase C2), faithful and continuous (N4). We add one mild regularity property, flagged:
    - **(H-mono)** h is strictly monotone increasing on $[0, \infty)$: more distinguishability content in an outcome means strictly more weight. This is natural (weight tracks amplitude modulus) and is used only to invert h. Where it is needed we say so.
- **Per-outcome weight.** Define $w_i(\psi) := h(|c_i|)$, so $N(\psi) = \Sigma_i w_i(\psi)$. The w_i are the pre-factual "possibility-weights" of the individual outcomes; N is their conserved total.
- **Pre-factual reversible dynamics.** R, the substrate-admissible reversible maps on the pre-commitment configuration, bijections preserving N (conservation), with $R \supseteq \mathbb{T}_d$ (the inherited diagonal phases, C2). Q-MIX asks whether R^0 exceeds $\mathbb{T}_d$ by an off-diagonal generator.

The single new question this paper poses: *which* quantities does R conserve — the total N only, or each w_i separately?

---

3. Two readings of pre-factual weight: bath and ledger

That $N = \Sigma_i w_i$ is conserved does not say whether the *summands* are individually conserved. Two structurally distinct conservation laws are compatible with "N conserved," and they are the fork.

Definition 3.1 (Ledger reading, L). Pre-factual reversible dynamics conserves each per-outcome weight separately: for every $U \in R$ and every i,

$$w_i(U\psi) = w_i(\psi).$$

(For unlabelled outcomes this is relaxed to conservation of the multiset $\{w_i\}$, permitting discrete relabelling; the continuous content is unchanged.) Under L, the pre-factual weights are **d independent accounts**, each balanced on its own.

Definition 3.2 (Bath reading, B). Pre-factual reversible dynamics conserves only the total: for every $U \in R$, $\Sigma_i w_i(U\psi) = \Sigma_i w_i(\psi)$, but R contains transformations with $w_k(U\psi) \neq w_k(\psi)$ for some k — weight is **reallocated** between outcomes subject to the fixed total. Under B, the pre-factual weight is **one shared, fungible pool**.

These are exhaustive of the relevant possibilities (either every w_i is separately conserved, or some U moves weight while fixing the total) and mutually exclusive. Note what neither reading touches: both are compatible with the *separable form* $N = \Sigma h(|c_i|)$ (inherited from N1). N1 fixed the *shape* of the conserved functional; the bath/ledger fork concerns the *conservation structure of the dynamics on that functional* — a different and, as the prior papers noted, as-yet-unmined question.

Remark 3.2.1 (The fork is BCB-balance, not BCB-additivity). This is the place to locate the distinction within the programme's primitive. BCB-additivity (N1) — distinguishable outcomes carry additive bit-content — was extracted and used; it gives separability and nothing about bath-versus-ledger. The remaining, unextracted content of "Bit Conservation **and Balance**" is precisely the *balance* structure: whether bit-content is balanced as one shared reservoir across alternatives (B) or balanced account-by-account (L). The prior papers flagged this surplus as a co-decider (Squaring §9.7); here it acquires a definite job — it decides the bath/ledger fork, hence (by §4) Q-MIX.

---

4. The Possibility-Connectivity Theorem

Theorem 4.1 (Ledger ⟹ Torus — proven, given H-mono). Under the ledger reading L, the connected pre-factual reversible dynamics is the diagonal phase torus: $R^0 = \mathbb{T}\_d$. No off-diagonal mixing exists.

Proof. By L, every $U \in R$ preserves $w_i(\psi) = h(|c_i|)$ for each i and all ψ. By H-mono, h is injective, so preservation of $h(|c_i|)$ gives preservation of $|c_i|$ for each i: $|(U\psi)_i| = |c_i|$. Consider a one-parameter subgroup $\exp(tX) \subseteq R^0$ with generator X. Differentiating $|(\exp(tX)\psi)_i|^2 = |c_i|^2$ at $t = 0$ gives $\mathrm{Re}(\bar{c}_i (X\psi)_i) = 0$ for all i and all ψ. Writing $(X\psi)_i = \Sigma_j X\_\{ij\} c_j$ and varying the relative phases of the c_j, the off-diagonal terms ($j \neq i$) average to zero only if $X\_\{ij\} = 0$ for $j \neq i$; the diagonal term forces $\mathrm{Re}(X\_\{ii\}) = 0$. Hence X is diagonal and skew — $X \in t\_d$. So every generator lies in t_d, $R^0 = \mathbb{T}\_d$, and the only further admissible maps are discrete (outcome relabellings preserving the multiset of weights). No continuous off-diagonal generator exists. ■

Theorem 4.2 (Bath ⟹ Mixing — proven modulo the continuity upgrade). Under the bath reading B, pre-factual reversible dynamics contains a transformation reallocating weight between two outcomes. If pre-factual reversibility is continuous (the temporal-extensibility / holonomy-continuity upgrade), this reallocation is generated by a continuous off-diagonal generator: genuine mixing exists.

Proof. By B, there is $U \in R$ and an outcome k with $w_k(U\psi) \neq w_k(\psi)$ for some ψ, while $N(U\psi) = N(\psi)$. By H-mono this means $|(U\psi)_k| \neq |c\_k|$: U changes outcome k's modulus, and since the total is fixed, it does so by moving weight to (at least) one other outcome ℓ. Thus U is not modulus-preserving on coordinate k, hence $U \notin \mathbb{T}\_d$ (the torus preserves every modulus) and U is not a mere relabelling of the configuration with k's modulus intact: U has genuine off-diagonal content. Granted continuity of pre-factual reversibility, U lies on a continuous one-parameter subgroup $\exp(tX) \subseteq R^0$; its generator X must have a nonzero off-diagonal entry coupling k and ℓ (a purely diagonal X would, by Theorem 4.1's computation, preserve every modulus, contradicting that the flow changes $|c\_k|$). So $X \notin t\_d$: a continuous off-diagonal generator exists. **Status: conditional on the continuity upgrade** (the existence of weight-reallocating maps is unconditional under B; their assembly into a continuous generator uses continuity). ■

Theorem 4.3 (Possibility-Connectivity Theorem — proven modulo the continuity upgrade). Continuous off-diagonal mixing of distinct outcomes exists in pre-factual reversible dynamics **iff** the bath reading holds:

Q-MIX ⟺ B (one shared fungible bath, not d separate ledgers).

Proof. (⟸) Theorem 4.2. (⟹) Contrapositive of Theorem 4.1: if not B then L (the readings are exhaustive), and L gives $R^0 = \mathbb{T}\_d$, no mixing. ■

The fiat-free reading of the ontological intuition. Theorem 4.3 vindicates the *content* of the possibility argument while removing its fiat. The intuition was: genuine connected possibility forces mixing. The precise version: *the bath reading* — weight as one shared fungible pool —

forces mixing, and *the ledger reading* — separate accounts — forbids it. "Connected possibility space," used carefully, names the bath reading; "list of alternatives" names the ledger reading. The argument is correct once "connected/genuine" is cashed out as the bath structure of the conservation law — a structural fact — rather than left as an evaluative word. The work the fast argument tried to do by definition is now done by the conservation law, and the conservation law is something the substrate fixes, not something we legislate.

---

5. Why this is not the fiat: the ledger reading is classical probability

It is essential that the ledger reading is not a strawman or a degenerate case to be dismissed. It is exactly classical reversible probability, and recognizing this is what keeps the paper honest.

Under L, the conserved data is the tuple (w_1, …, w_d) of separately-conserved per-outcome weights, with Σw_i fixed; the admissible dynamics rephases each outcome and (for unlabelled outcomes) permutes them. Strip the operationally-inert phases and this is a probability vector evolving by permutations — classical reversible probability, with the w_i/N as the probabilities and commitment as the sampling. The classical die is the canonical instance: d = 6, weights the face-probabilities, commitment the landing, no interconversion of faces.

So the ledger reading is a fully coherent substrate, not a failure of "genuine possibility." A reality whose BCB-balance instantiates L would be classical-at-the-level-of-reversible-dynamics, would carry a non-ℓ^2 (indeed any separable) conserved N, and would settle the Born exponent negatively. This is the same honest stance the *Squaring* paper held about its Torus branch: the classical option is live, and the programme's job is to determine whether the substrate excludes it, not to define it away. Theorem 4.3 locates that determination precisely — it is the bath/ledger fork — and §7 hands it to BCB-balance.

The methodological gain over the fast argument is exactly this refusal to dismiss L. The fast argument had to call L "merely a list," a non-genuine possibility space, to force mixing — and that dismissal is the fiat. Theorem 4.3 needs no such dismissal: it states the biconditional and lets L stand as the coherent classical alternative it is. Mixing is forced *if* the substrate is a bath; it is genuinely absent *if* the substrate is a ledger; and which holds is not ours to stipulate.

---

6. The bath reading delivers connectivity

A dividend. The linearity companion's conjunctive hinge required, beyond mixing, a *connectivity/symmetry* condition — that the off-diagonal generators connect all outcome pairs, so that "ℓ^2 on each connected pair" propagates to global isotropic ℓ^2. Under the bath reading this comes for free, provided the bath is genuinely shared across all outcomes.

Proposition 6.1 (Connectivity from a single shared bath — proven modulo continuity). If the bath reading holds with a *single* pool shared across all d outcomes (not a block-diagonal decomposition into separate sub-baths over disjoint outcome groups), then weight is reallocable between every pair of outcomes, so the mixing graph is complete (connected), and the connectivity leg of the linearity companion's hinge is satisfied.

Proof. A single shared pool means that for any pair (k, ℓ) there is an admissible reallocation moving weight from k to ℓ (otherwise the pool would split along the k–ℓ boundary into separate conserved sub-pools, contradicting "single bath"). By Theorem 4.2 each such reallocation generates an off-diagonal generator coupling k and ℓ. Ranging over all pairs, the generators connect all outcomes; the mixing graph is complete. ■

Remark 6.1.1 (The block-diagonal caveat is itself meaningful). The hypothesis "single pool, not block-diagonal sub-baths" is not free, and its failure is interpretable. If the bath splits into sub-pools — say outcomes {1,2,3} share one reservoir and {4,5,6} another, with no reallocation across the split — then mixing exists within each block but not across, the mixing graph has those two components, and the conserved N is ℓ^2 *within* each block but may carry independent weights across blocks (a block-diagonal positive form, not global isotropic ℓ^2). This is the same partial-ℓ^2 outcome the linearity companion flagged when connectivity fails, now given a substrate meaning: super-selection-like sectors are sub-baths between which pre-factual weight does not flow. Whether the VERSF substrate has a single bath or sub-baths is a refinement of the bath/ledger question, handed forward with it.

---

7. Relocation to BCB-balance, and the collapse to the refinement obstruction

The companion papers repeatedly noted that BCB was only partially mined — additivity (N1) extracted, the "Balance" content set aside — and flagged the surplus as a co-decider of the Born exponent (Squaring §9.7). This paper gives the surplus a definite identity.

The claim. The bath/ledger fork *is* the BCB-balance surplus. Additivity (N1) says distinguishable outcomes carry additive bit-content — it fixes $N = \Sigma h(|c_i|)$ and is silent on whether the summands are separately conserved. The balance content of "Bit Conservation and **Balance**" is exactly the statement of *how* bit-content is balanced across the commitment process: as one shared reservoir reallocable among pre-factual alternatives (bath, B), or as separate per-outcome balances (ledger, L). So:

Q-MIX (mixing) ⟺ B (bath) ⟺ the BCB-balance surplus asserts a single shared reservoir.

This is the relocation the paper delivers. The mixing question — previously "is there a continuous off-diagonal generator?", a question about the dynamical group with no obvious substrate handle — becomes "does BCB-balance balance bit-content as one shared bath or as separate ledgers?", a question about the conservation primitive itself, which the BCB papers either already answer or can be made to answer by a precise statement. That is a far more tractable target, and it is substrate-native rather than imported (contrast the Hardy/Masanes–Müller continuous-reversibility axiom, which simply *postulates* the bath structure under the name "continuous reversibility"; here it is referred to the conservation law for adjudication).

What we do not do. We do not assert B. To assert it would be the fiat in structural clothing — "the substrate is a bath because a bath is what a genuine substrate is." Whether BCB-balance is a bath or a ledger is a fact about the unstated balance content, to be settled by stating it, not by preference. We have proven the equivalence and located the question; the answer is the TPB/BCB paper's to give.

7.1 Chasing the fork to the packing primitive

The bath/ledger fork can be pursued one level deeper, to the finite-packing primitive FP1 itself, and doing so reveals that it is not a *new* question at all. FP1 is stated as a bound on the distinguishability capacity of an operational *region*:

|Σ(M)| ≤ Vol_op(M) / Δ_op^{d_op}.

The primitive object is Vol_op(M) — the capacity of the unresolved region M — not a tuple (C_1, …, C_d) of pre-assigned per-outcome capacities. This is a textual fact about how the primitive is written, and it leans bath-like: capacity is borne by the region, and in the commitment ontology the determinate outcomes do not exist as separate facts before commitment localizes one continuation. The ledger reading, by contrast, requires per-outcome capacities C_i that *exist before commitment* — which presupposes exactly the pre-committed outcome-structure the commitment ontology denies.

That is suggestive, but it does not by itself force the bath, and we must be exact about why — because the rigid reading has a comeback, and it is the same die in new clothing. FP1's region capacity *decomposes*: the *Packing* paper's distinguishability partitions (its §5) give additive sub-capacities Vol_op(M) = Σ_i Vol_op(M_i). A skeptic grants that capacity lives on the region and then says: *but the region's capacity partitions into sub-region capacities, and those Vol_op(M_i) are my per-outcome accounts — I am a ledger sourced from a region, and FP1 plus your own partition structure permits me.* So "capacity is borne by the region" (true, textual) does **not** give "capacity is non-partitioned" (the bath claim). The region-sourced ledger is consistent with the bare axiom.

What separates the readings is therefore not whether the capacity decomposes, but whether the decomposition is **rigid or dynamical**:

- **Rigid partition (ledger):** the sub-capacities $Vol\_op(M_i)$ are fixed before commitment; commitment merely reveals which sub-region was realized. Nothing reallocates. This is the region-sourced classical die.
- **Dynamical partition (bath):** the partition boundaries can move and capacity reallocates across them pre-commitment; the decomposition is frozen only at commitment, which localizes one continuation from the shared unresolved capacity. This is the bath.

Claim 7.1.1 (The bath/ledger fork ≡ Obstruction B — identification of content). The rigid-versus-dynamical-partition question is precisely the *Packing* paper's refinement obstruction (Obstruction B / FP3). FP3's load-bearing reverse inclusion asks whether every admissible decomposition is a resolution-stable refinement — i.e. whether admissible decompositions are rigidly fixed or are the (movable) resolution-stable refinements of a shared capacity. A rigid partition is the negation of the reverse inclusion (decompositions fixed independently of refinement); a dynamical partition is its affirmation (decompositions are exactly the resolution-stable refinements of the unresolved region). Hence: **dynamical partition ⇔ Obstruction B's reverse inclusion holds ⇔ bath reading.** The bath/ledger fork, chased to FP1, is the refinement obstruction.

We assert this as an identification of content, with the same status the linearity companion gave to RS ≡ Obstruction B. We are careful, in light of §9.5, about exactly what is and is not definitional here: "bath ⇔ dynamical partition" we grant as two vocabularies for one statement, but "dynamical partition ⇔ FP3's reverse inclusion" is a *cross-object* identification (flows of weight versus decompositions of measurements) that rests, so far, on a shared gloss rather than a translation between the objects, and the formal lemma supplying that translation is outstanding (§9.5). Where the full biconditional is not yet earned we fall back on the weaker sufficient statement — "the fork is no harder than Obstruction B" — which carries the consolidation of §8 without asserting more than is proven.

Remark 7.1.2 (A dynamical partition of one region is a single bath). The connectivity dividend of §6 (Proposition 6.1) required not merely the bath reading but a *single* shared bath across all outcomes, as opposed to block-diagonal sub-baths. Under the dynamical reading this is not a further independent assumption: a dynamical partition is, by construction, the movable refinement of *one* unresolved region M, whose capacity Vol_op(M) is a single pool. If the capacity instead split into sub-pools between which no reallocation occurred, those sub-pools would be separately conserved — which is the *rigid* (ledger) reading applied at the level of the blocks, not the dynamical reading. So "the partition of a single unresolved region is dynamical" entails "the bath is single": one region, one pool. Sub-baths arise only if the substrate presents the pre-factual configuration as *several* unresolved regions with no cross-region reallocation — a structural fact distinct from the rigid/dynamical fork, addressed in §9.4. Granting that commitment resolves *one* unresolved region (the commitment ontology's picture), the dynamical reading delivers a single bath, and the §8 cascade to *isotropic* ℓ^2 needs no third input beyond the affirmative reading of Obstruction B.

The crucial honesty. FP1 *as written is silent* between the rigid and dynamical readings — which is exactly why the *Packing* paper left Obstruction B open. So we do not claim FP1 forces the bath. We claim the bath/ledger fork *is* Obstruction B (modulo the outstanding cross-object lemma, §9.5), and that reading the partition as dynamical — the reading the programme's commitment ontology favours, on which capacity belongs to the unresolved region and is localized only at commitment — is the affirmative resolution of Obstruction B, not a derivation from the bare axiom. The rigid (classical, region-sourced ledger) reading remains consistent with FP1 and is the coherent alternative the substrate could instead instantiate. Closing the fork by adopting the dynamical reading is permitted; closing it by stealth — harvesting the bath while pretending FP1 forced it — is the over-claim we refuse.

---

8. Assembly: the consolidated state of the Born exponent

With three companion papers in hand, the Born exponent's dependence can be stated in consolidated form.

From *The Squaring Residue* (corrected by the linearity companion): the conserved normalization is the squared (ℓ^2) norm — the squaring residue closes, PAMV Open Problem 10 resolves positively — **iff**

(linear reversible transport) ∧ (off-diagonal mixing) ∧ (connectivity/symmetry).

The linearity companion reduced the first leg: **linear transport ⟸ PC (mild decomposition-independence) + Obstruction B (the refinement obstruction)**. The present paper reduces the second and third legs: **mixing ⟸ bath reading B**, and **connectivity ⟸ single shared bath** (Proposition 6.1), with the single bath itself supplied by the dynamical reading of one unresolved region (Remark 7.1.2: one region, one pool). And §7.1 identifies the bath reading itself with the refinement obstruction: **bath ⟺ dynamical partition, and dynamical partition ⟺ Obstruction B's reverse inclusion** — the second equivalence being a cross-object identification whose formal lemma is outstanding (§9.5). Substituting through, *modulo that lemma*:

ℓ^2 ⟸ PC ∧ Obstruction B.

The three apparently-separate conditions collapse toward two, because the third (bath-balance) is — modulo the outstanding cross-object translation — not independent of the second: it is Obstruction B, viewed as the reallocability of pre-factual capacity rather than as the refinement of measurements. The honest statement of the collapse is therefore "**two conditions, one of which is conditionally identified with the other**," the conditionality being precisely Claim 7.1.1's cross-object lemma. So the surviving substrate conditions are:

1. **PC** — reversible transport may not depend on bookkeeping decomposition. *Mild, the one genuinely new commitment across all three companion papers, to be stated and checked.*
2. **Obstruction B** — the refinement obstruction, which on the strength of the identifications (the linearity companion's RS ≡ Obstruction B, and this paper's Claim 7.1.1) governs four

things at once: ODG's refinement structure (the *Packing* paper), the linearity of reversible transport (the linearity companion), the bath reading / mixing (this paper, §4–5), and the dynamical-partition reading of the packing primitive (§7.1). Affirming it — reading admissible decomposition as the dynamical refinement of a single shared unresolved capacity — cascades through all four to isotropic ℓ^2, the single-bath step supplied by Remark 7.1.2.

This is the consolidated honest state. It is a genuine collapse of the decision tree — **conditional on the cross-object lemma of §9.5** — bringing the Born exponent's residual dependence to **one mild new principle (PC) and one standing refinement question (Obstruction B)**, the deep substrate fact the programme already owed. Settle Obstruction B affirmatively (and supply the lemma) and PC is the only survivor; the squaring (Diagonal-Torus) then follows by the proven machinery.

Remark 8.1 (The dynamical reading, and its status). Two distinct things are conditional here and should not be conflated. *First*, the identification of the bath reading with Obstruction B rests on the cross-object lemma (§9.5), which is outstanding; until it is supplied the collapse is "two conditions, one conditionally identified with the other," and we use the weaker sufficient form ("the fork is no harder than Obstruction B") where the full biconditional is not earned. *Second*, even granting the identification, the cascade to ℓ^2 is reached by *affirming* Obstruction B's reverse inclusion — reading the partition of pre-factual capacity as dynamical (§7.1) — which is a reading of the primitive, favoured by the commitment ontology but **not** forced by the bare FP1 axiom (silent between rigid and dynamical). The rigid reading — region-sourced classical ledger — remains consistent with FP1: a substrate instantiating it would be classical-at-the-level-of-reversible-dynamics, carry a non-ℓ^2 conserved N, and settle the Born exponent negatively. So what this paper establishes is not closure but *unification* — that linearity, mixing, connectivity, and the partition structure of FP1 are (modulo the lemma) one question — leaving the entire residual content of the Born exponent, beyond the mild PC, as the single affirmative reading of Obstruction B. Removing three branches by showing them one is the result; both the identification lemma and the reading that resolves the surviving branch are flagged, not smuggled.

---

9. Limitations and Open Problems

9.1 Bath or ledger? — the question, not the answer (Q-MIX relocated)

The paper proves Q-MIX ⇔ bath reading and, in §7.1, identifies the bath/ledger fork with the refinement obstruction (Obstruction B); it does **not** determine which way that obstruction resolves. That determination requires stating the balance content of BCB and the partition structure of FP1 precisely — reading off whether admissible decomposition is rigid (ledger, classical) or dynamical (bath, quantum). We record no prior here beyond the programme's commitment ontology — which *favours* the dynamical reading (capacity belongs to the unresolved region, localized only at commitment) but, as §1, §5, and §7.1 insist, does not *force* it from the bare axiom (the region-sourced classical ledger is consistent with FP1). The honest status: **open**, now in its sharpest form — it is Obstruction B — and handed to the BCB/TPB paper.

9.2 The continuity upgrade

Theorem 4.2 (bath ⇒ mixing) is conditional on pre-factual reversibility being continuous — the same temporal-extensibility / holonomy-continuity upgrade flagged in the Squaring paper (its §11.3 / Physical Necessity §C.8). Under the bath reading the *existence* of weight-reallocating maps is unconditional; assembling them into a continuous one-parameter generator uses continuity. If reversibility were only discrete, the bath reading would give discrete weight-reallocation (a permutation-like shuffling of weight) without a continuous generator, and the Diagonal-Torus selection would not engage. So the continuity upgrade is load-bearing on the mixing side exactly as on the phase side. **Conditional**, with the named roadmap inherited.

9.3 H-mono

Theorems 4.1–4.2 use strict monotonicity of h to invert weight-preservation into modulus-preservation. H-mono is mild (weight strictly increases with amplitude modulus) but it is an assumption; a non-monotone h (weight not tracking modulus) would break the clean equivalence and is not obviously excluded by the inherited inputs. We judge H-mono natural and flag it **conditional**, to be confirmed from the BCB definition of bit-content.

9.4 Single bath versus sub-baths — one region or several?

Proposition 6.1's connectivity dividend (global *isotropic* ℓ^2) requires a single shared bath, not block-diagonal sub-baths. Remark 7.1.2 argues this is not a hidden third input *given the dynamical reading*: a dynamical partition is the movable refinement of one unresolved region, whose capacity is one pool, so one-region-one-pool delivers the single bath. The residual question is therefore not "single bath versus sub-baths" as an independent fork, but whether the substrate presents the pre-factual configuration as **one** unresolved region or **several** with no cross-region reallocation. The latter would give super-selection-like sectors (Remark 6.1.1) and block-diagonal partial-ℓ^2. This is a structural fact about how the substrate packages pre-factual configurations — distinct from the rigid/dynamical fork, and unresolved here. We note it bears on whether the cascade reaches *global isotropic* ℓ^2 or only block-wise ℓ^2; under the commitment ontology's picture of commitment resolving a single unresolved region, the single-region reading is the natural one, but it is a separate reading and we do not assume it silently.

9.5 The outstanding lemma is a cross-object translation

§7.1 (Claim 7.1.1) and §8 rest the collapse on the identification *dynamical partition ⟺ Obstruction B's reverse inclusion ⟺ bath*. The two halves do not have the same status, and honesty requires separating them.

- **Bath ⟺ dynamical partition** is nearly definitional and we grant it: "weight reallocates across movable partition boundaries before commitment" and "the partition is dynamical" are one statement in two vocabularies.
- **Dynamical partition ⟺ FP3's reverse inclusion** is the half doing the real work, and it is **not** definitional. FP3's reverse inclusion is a statement about *measurement decompositions* — every admissible decomposition is a resolution-stable refinement. The bath statement is about *reallocability of conserved weight under reversible dynamics*. These are claims about different objects — decompositions of measurements versus flows of weight — and our argument that they coincide has so far been an appeal to a shared gloss ("admissible structure introduces nothing beyond resolution-stable refinement of a shared capacity"), not a translation between the objects.

We flag this squarely, because two propositions sharing an English paraphrase is exactly the kind of word-bridge §1 and §5 refused elsewhere: the fast argument was disarmed by not letting "genuine" carry the inference, and Claim 7.1.1 must not be allowed to let "resolution-stable refinement of a shared capacity" carry an inference of the same logical type one register up. **The outstanding lemma is therefore specifically the cross-object translation — weight-reallocability under reversible dynamics ⟺ FP3 reverse inclusion on measurement decompositions — not a tidying-up of something already morally settled.** Until it is supplied, the §8 collapse is properly read as "**two conditions, one of which is conditionally identified with the other**," not flatly "two conditions." We fall back, where the weaker statement suffices, on "the fork is no harder than Obstruction B" — which does carry the consolidation's force (any progress on the obstruction transfers) without asserting the full biconditional. The parallel to the linearity companion's RS ≡ Obstruction B identification is real and the same lemma would likely settle both cross-object translations at once; that single lemma is the clean capstone, and it is genuinely outstanding, not cosmetic.

9.6 Dependence on inherited inputs

As with the companions, the argument rests on H0 (the complex carrier), N1 (separability), and the correctness of the inherited Diagonal-Torus selection (used in §8's assembly). The new load-bearing inheritance is the claim that the BCB-balance surplus is well-posed as the bath/ledger fork — a referee should confirm that BCB, fully stated, *has* a determinate answer to "shared reservoir or separate ledgers," since the relocation is only as sharp as that question is well-defined.

---

10. Conclusion

The mixing question — does pre-factual reversible dynamics turn one possible outcome continuously into another, or only rephase each in place? — is the sharpest surviving hinge of the Born exponent. There is a fast way to settle it from the substrate's commitment ontology: genuine possibilities of one process must be interconvertible, so mixing is forced. We have refused that argument, because it settles the question by the meaning of "genuine," and classical probability — jointly-available alternatives, no interconversion — is a standing counterexample to the inference. A possibility space need not have mixing; the classical die does not.

What decides mixing is not the meaning of possibility but the structure of the conserved pre-factual weight. We proved a clean dichotomy: weight is either one shared fungible bath, reallocable among alternatives, or many separate ledgers, each conserved alone; and **mixing exists iff the bath reading holds** (Theorem 4.3, modulo the continuity upgrade). The ledger reading is precisely classical reversible probability — a coherent substrate, not a non-genuine one — and the bath reading is the quantum one. The ontological intuition was right in content and wrong in method: "connected possibility space" correctly names the bath structure, but the connectedness must be read off the conservation law, not asserted of the word.

This relocates the mixing question to a substrate-native, checkable form — and then, pursued one level deeper to the packing primitive (§7.1), it goes further. The bath/ledger fork, chased to FP1, becomes the question of whether a region's distinguishability capacity is *rigidly* partitioned before commitment (ledger) or *dynamically* reallocable until commitment localizes one continuation (bath) — and that question we identify with the refinement obstruction (Obstruction B) the *Packing* paper already isolated. The identification has two halves of unequal status: that the bath reading *is* a dynamical partition we grant as near-definitional; that a dynamical partition *is* Obstruction B's reverse inclusion is a cross-object claim (flows of weight versus decompositions of measurements) resting, so far, on a shared structural gloss rather than a proven translation — the lemma supplying that translation is outstanding (§9.5). Modulo that lemma, the bath reading is not a *third* substrate condition beside the refinement obstruction; it is that obstruction, seen as the reallocability of pre-factual capacity. FP1 as written is silent between the rigid and dynamical readings, which is exactly why the obstruction was left open.

Assembled across the three companion papers, the Born exponent therefore collapses — modulo the cross-object lemma — toward **two** conditions, not three: a mild decomposition-independence principle (PC), and the refinement obstruction (Obstruction B), which on the strength of the two identifications (RS ≡ Obstruction B in the linearity companion, and Claim 7.1.1 here) governs four things at once: ODG's refinement structure, the linearity of reversible transport, the bath reading and hence mixing, and the dynamical-partition reading of the packing primitive (with the single-bath step supplied by one-region-one-pool, Remark 7.1.2). PC is the single genuinely new commitment across all three papers; the rest is one standing question. We do not assert the dynamical reading that resolves it — that would be the fiat in structural clothing, and FP1 does not force it (the rigid, region-sourced classical ledger is consistent with the axiom). What we have shown is the *unification* — four apparently-separate hinges are, modulo one outstanding lemma, one — so that the entire residual content of the Born exponent, beyond PC, is the single affirmative reading of Obstruction B. Removing three branches by showing them one is the result; both the identification lemma and the reading that would close the surviving branch are flagged, not smuggled.

The square is no longer waiting on a vague "is there a beam-splitter?", nor even on a separate "bath or ledger?". It waits — modulo the cross-object lemma that ties the bath reading to the refinement obstruction — on one fact the programme has owed since the *Packing* paper: is admissible decomposition rigid or dynamical; does pre-factual distinguishability capacity belong to the unresolved region and localize at commitment, or sit pre-partitioned in separate accounts? Read it dynamically — as the commitment ontology favours — and, with the mild PC, the square follows. That reading, honestly labelled and not assumed, together with the lemma that would make the unification a theorem, is now the whole of what stands between finite distinguishability and the Born rule.