What Meaning Can Copy

Machine-Checked Copy Laws, Braiding Limits and Exact Acceptance for Compositional Quantum Models

Richard Goodman, with the IAOM research workforce
Institute for Applied Ontological Mathematics

Subject areas: quant-ph (primary); cs.CL, cs.LO, math.CT, cs.LG

Abstract

Compositional models of meaning send grammar to diagrams and diagrams to computation, classical or quantum. Every such model copies: a relative pronoun shares its noun, an attention layer fans out its residual stream, a tied weight serves two layers. We ask exactly what copying permits. Over every commutative ring, copying classifies the classical states precisely when the ring's spectrum is connected; each failure is witnessed by a separator, and committing a quantum noun to a classical one costs exactly its logical entropy $1 - \sum p^2$. The same law forces a shared weight to be a sharp basis value. We then ask what braiding adds. In any braided semantics with closed names, closed readings cannot see a crossing label; open frames see it only through the character $\theta^{-\mathrm{writhe}}$, which is blind for Ising and has period five for Fibonacci, and neither theory admits a symmetric fibre functor except Fibonacci at its ramified prime 5. A pre-registered experiment finds no advantage of Ising-braided over symmetric text circuits ($\Delta = -0.0124$, 95% CI $[-0.0253, +0.0004]$). Exact rings decide the physics: the Gauss law tracks Galois transport, not unitarity (refuted at $j \in \{3, 5\}$); Fibonacci lens-space invariants obey a Legendre law for every $L(p, q)$; and an exact parameter-shift rule, a Born head that is a relative norm, a complete Clifford decision and a Dvoretzky–Kiefer–Wolfowitz acceptance rule over the unique order on $\mathbb{Z}[\sqrt{p}]$ close a hybrid loop. Everything is machine-checked in Agda 2.8.0 (`--safe`, postulate-free; 723 planted mutants refused) over a 24-sentence controlled corpus. Acceptance on quantum hardware is pending.

Popular summary

A quantum computer cannot copy an unknown state; a language model copies constantly. A pronoun copies its noun ("Alice, *who* sleeps"), a transformer copies its working memory to every attention head, and a model that reuses one weight in two layers copies that weight. If language models are ever to run partly on quantum hardware, someone has to say exactly which of these copies are allowed, what the forbidden ones cost to repair, and what quantum effects such as braiding (the exchange of particles in two dimensions) could add.

We answer these questions with proofs a computer has checked. Copying works exactly on the "classical" states, and we identify those states over every number system the work uses. Repairing a forbidden copy means measuring first, and the exact price of that measurement is a simple formula in the probabilities. Braiding, it turns out, cannot change the meaning of a finished sentence; it can only mark the unfinished frame around it, and in a pre-registered experiment it gave learned models no advantage. Along the way, exact arithmetic separates the physical versions of two quantum theories from their unphysical twins, and gives an exact gradient rule and an exact statistical test for checking a quantum device against theory. The test on real hardware is the one step still to be done.

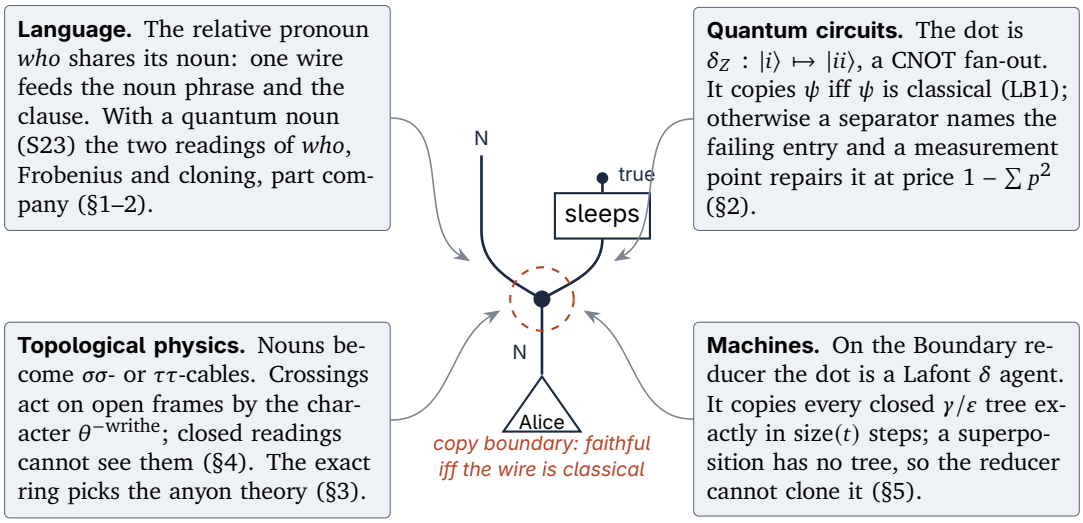


Figure 1. One syntax, many meanings. The diagram of "Alice who sleeps" (S05, centre) is one morphism with four readings. In every reading the black dot is a *copy*. The copy boundary (dashed) marks the copy law of Section 2: the dot is faithful exactly when the wire it copies is classical. The same boundary decides where a sentence needs a measurement (Section 2.2), where a transformer can be quantum (Section 6.1), what braiding can see (Section 4) and what the reducer can duplicate (Section 5.5). Drawn in TikZ from `figures/fig1_one_syntax.tikz`.

What this gives you

- **If you build quantum hardware:** exact gauge-invariant targets for anyon experiments, an acceptance rule whose window and shot budget are decided in advance, and a complete Clifford decision for what your device must actually do.
- **If you build accelerators:** an exact arithmetic lane ($\mathbb{Z}[\frac{1}{2}]$ with canonical forms, proved), and an executable whose 1,440 configurations reproduce byte for byte with the oracles removed.
- **If you design language models:** which parts of an architecture cannot be quantum (every copy of a non-classical value), the exact price of the parts that can, and the proof that a shared weight must be sharp.
- **If you study topological matter:** the pentagon solutions are one gauge orbit, the Gauss law is Galois transport rather than unitarity, and Fibonacci lens-space invariants follow a Legendre law in $(p/5)$.
- **If you engineer regulated systems:** a worked separation of what is proved (the logic) from what is validated (the device), and refusals that name the failing wire.
- **If you work in formal methods:** 44 phases in Agda under `--safe` with no postulates, every claim bound to a declaration, 723 planted mutants refused at the declaration they attack.
- **If you work in category theory or HoTT:** Fox's theorem over rings, Para over a copower, a symmetric-fibre obstruction by Bezout certificates, and open problems with discriminators (Section 9).

Contributions

Each contribution is tagged by kind. *Proved* means machine-checked in Agda 2.8.0 under `--safe` with no postulates, unless the entry names another system. *Refuted* means a kernel-checked counter-model to a statement we or the literature had proposed. *Measured* means pre-registered and reported with its interval. *Proposed* means a hypothesis with a stated discriminator.

C1 *Proved.* The copy law (LB1): over every commutative ring, copying classifies the classical states if and only if every idempotent is 0 or 1 (the spectrum is connected); copiable states are the characters, and normalised ones are the points of $\operatorname{Spec} K$ (Section 2).

C2 *Proved.* The price of committing: δ_Z broadcasts exactly the diagonal densities, and a measurement point costs exactly the logical entropy $1 - \sum p_i^2$ (Section 2.2).

C3 *Proved.* The weight-tying law: tied evaluation is functorial iff the parameter is sharp, and over a connected ring iff it is a basis point; the tied gradient is the sum of the per-copy parameter shifts (Section 2.3).

C4 *Proved and refuted.* Exact rings decide physics: the Gauss-law/unitarity conjecture is refuted at exactly $j \in \{3, 5\}$ and replaced by "the Gauss law holds iff the model is a Galois transport" (local complementation's soundness is shown equivalent in Lean 4); the Fibonacci pentagon locus is one gauge orbit; the Legendre law holds for every lens space (Section 3).

C5 *Proved and refuted.* Braiding's reach: closed readings are routing-blind in every braided semantics with closed names; open frames carry the character $\theta^{-\mathrm{writhe}}$; no symmetric fibre functor exists for Ising or Fibonacci except Fibonacci at 5. The planned closed-sentence separation is refuted (Section 4).

C6 *Measured.* A pre-registered comparison of Ising-braided and symmetric text circuits returns a null result on role binding (Section 4.6).

C7 *Proved.* The exact hybrid loop: an exact parameter-shift rule needing no unitarity, a Born head that is a relative norm, a complete Clifford decision, and DKW acceptance over the unique order on $\mathbb{Z}[\sqrt{p}]$ (Section 5).

C8 *Proved.* From proof to machine: net typing for every accepted term, exact copy of every Lafont γ/ε tree on the reducer's own schedule, and the reducer's actual output proved equal to the model (Section 5.5).

C9 *Refuted.* Three of our own planned statements: a braided lowering of the planar grammar, a flat price bridge under nesting, and the identity between net typing and the term checker.

C10 *Proposed.* Consequences for quantum language models, hardware, accelerators, physics, complex systems and regulated engineering, each with a refuter (Section 6).

Reader map

If you are	Read	The question it answers for you
building quantum hardware or FTQC toolchains	§3, §5, §6.2	Which numbers are observable, what acceptance means at α, what is Clifford.
building accelerators	§5.5, §6.3, App. D	Can exact arithmetic be byte-identical, and at what cost?
designing language models	§1, §2–2.3, §4, §6.1, §8	Where can a model be quantum, at what price, and what braiding adds.
an engineer of regulated or control systems	Popular summary, §1, §6.6, §8	What is proved, what is only validated, and how refusals name the fault.
a topological-matter physicist	§3, §4, §6.4	Gauge orbits, Galois conjugates and lens-space invariants, exactly.
a formal-methods researcher	§5.5, §11, App. A	What the trusted base is and how each claim is bound.
an applied category theorist or HoTT researcher	§2, §2.3, §4, §6.5, §9	Which categorical statements were proved, and what remains open.

1 The pronoun that cannot be copied

Take the noun phrase "Alice who sleeps" (sentence S05 of our corpus). In the compositional distributional model of meaning [12], nouns are vectors in a space N spanned by individuals, sentences are vectors in a space S, and the relative pronoun *who* is built from a Frobenius algebra: it copies the noun wire, hands one copy to the clause, and keeps the other as the meaning of the phrase. Concretely, with individuals $\{\text{alice}, \text{bob}\}$,

$$\text{who}(n \otimes s)_i \;=\; n_i \,\langle \text{true} \,|\, s \rangle,$$

and the copy is the "spider" $\delta_Z : |i\rangle \mapsto |i\rangle \otimes |i\rangle$ (Fig. 1). Alice is a basis vector, so it does not matter whether we read the copy as δ_Z or as a cloner $\psi \mapsto \psi \otimes \psi$: both give $(1, 0)$, "the sleeping Alice".

Now let the noun be a quantum meaning, $\psi = |\text{alice}\rangle + |\text{bob}\rangle$ (sentence S23, which the corpus includes as unsafe by design). The two readings part company. The Frobenius reading copies in the basis and sums the branches,

$$F \;=\; \sum_c \psi_c \,\text{who}\big(e_c \otimes \text{sleeps}(e_c)\big) \;=\; (1, 0),$$

because only Alice sleeps. The cloning reading feeds ψ to both wires, $\psi \,\langle \text{true} \,|\, \text{sleeps}(\psi)\rangle = (1, 1)$, and asserts that Bob is a sleeping person as well. Neither reading is a mistake in arithmetic. The phrase has asked a copy to do something copies cannot do, and the two readings fail differently.

The same object sits inside every transformer. A residual block computes $x + f(x)$ and so uses x twice; attention reads every key and value from every query; a key–value cache stores a vector and reuses it at every later step; weight tying reuses one matrix in two places. In a classical network these copies are free and invisible. In a model whose state is a quantum register, each is a δ, and the question "Alice or Bob?" becomes a design question: which copies are lawful, what the unlawful ones cost to repair, and whether anything quantum survives the repair.

This paper answers those questions exactly, for the compositional fragment where exact answers are possible, and checks every answer in a proof assistant. The answers are, in brief:

- a copy is lawful exactly on classical states, and over every commutative ring we know which states those are (Section 2);
- an unlawful copy is repaired by a measurement point whose exact price is the logical entropy $1 - \sum p_i^2$; for S23 that price is $\frac{1}{2}$, yet the repair leaves the meaning unchanged because the clause selects one referent (Section 2.2);
- a weight shared between layers must be a sharp classical value, and its gradient is the sum over its uses (Section 2.3);
- the exact ring in which amplitudes live decides which anyon theories are physical (Section 3);
- braiding cannot change the meaning of a closed sentence; it can only mark the open frame around one, and in a pre-registered experiment it brought no learning advantage (Section 4);
- the whole hybrid loop, from parameters to accepted hardware samples, can be made exact, and the copy that the interaction-net reducer performs is proved to be the copy of the mathematics (Section 5).

How the claims are checked. The mathematics is formalised in Agda 2.8.0 [42, 49] under `--safe`, with no postulates and no escape pragmas, in four lanes: a core that is `--cubical-compatible`; a cubical lane that re-derives results of our earlier quantum-coherent-topology library; a lane checked `--without-K` for the QNLP compiler; and a lane for an exact golden-ratio state-sum engine. The project ran as 44 phases. Each phase wrote its predictions down before its gate ran, planted mutants into its own statements (723 in all, every one refused by the kernel, with 160 controls accepted), compared its executables against independent oracles, and passed a gate that binds the checked bytes. A self-audit (X4) and a whole-project hostile audit (HA) then raised 31 findings against the work as it stood; all were closed. Several were closed by proving stronger theorems than the paper had been going to claim, and two by refuting the remediation the finding itself had proposed. A controlled corpus of 47 rows fixes the scope: 24 sentences, 3 rings, 5 anyon theories, 5 knots, 5 manifolds, 4 contextuality models and one learning task.

Throughout, four environments mark the status of a claim. A **Theorem** is machine-checked and names its declaration and module. A **Refutation** is a kernel-checked counter-model. A **Measurement** is a pre-registered or recomputed number with its conditions. A **Proposal** is a hypothesis with a *discriminator*, the observation that would decide it. Interpretation stays in the prose. Appendix A maps every claim to its evidence.

2 The copy law

2.1 Copying classifies classical states over a connected spectrum

Work in $\mathrm{Mat}(K)$ for a commutative ring K: objects are natural numbers, morphisms $m \to n$ are $n \times m$ matrices over K, and the tensor is the Kronecker product. The *copy* on n is $\delta_n : K^n \to K^n \otimes K^n$, $e_i \mapsto e_i \otimes e_i$, with counit $\varepsilon_n : e_i \mapsto 1$; this is the Z-spider of the ZX calculus and the comultiplication of the special commutative Frobenius algebra that models relative pronouns.

Definition 1. A state $\psi \in K^n$ is *copiable* if $\delta\psi = \psi \otimes \psi$; equivalently, every coordinate is idempotent ($\psi_i^2 = \psi_i$) and distinct coordinates are orthogonal ($\psi_i\psi_j = 0$ for $i \neq j$). It is *classical* if it is 0 or a basis vector e_k. A *separator* of ψ is an index pair (i, j) with $(\delta\psi)_{ij} \neq \psi_i\psi_j$: a named entry at which the copy equation fails.

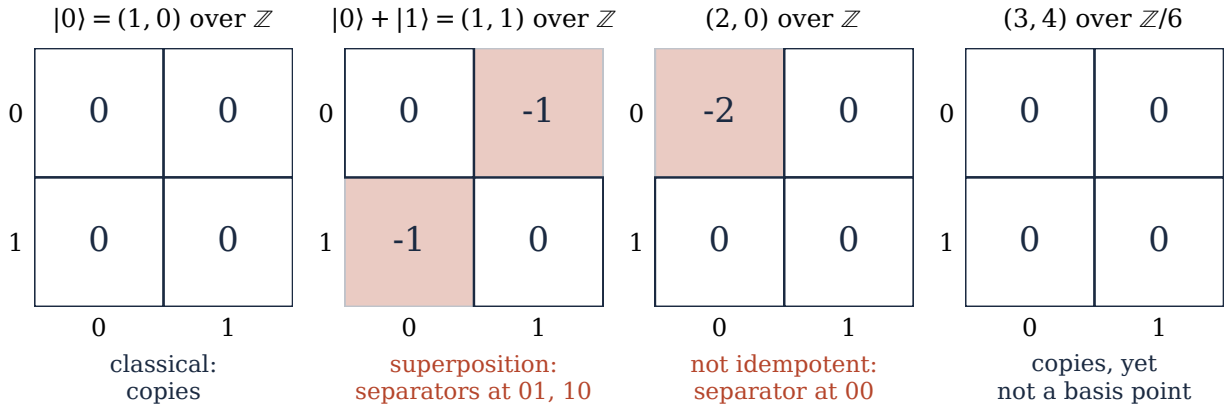


Figure 2. What exactly fails when a copy fails. Each panel shows the copy defect $\delta\psi - \psi\otimes\psi$ of a state $\psi \in K^2$ as a 2×2 array indexed by (i, j); shaded entries are separators. A basis state copies (first panel). A superposition fails off the diagonal (second); a non-idempotent coordinate fails on it (third). Over the disconnected ring $\mathbb{Z}/6$ the state $(3, 4)$ copies although it is not a basis point (fourth): connectedness of $\operatorname{Spec} K$ is the hypothesis LB1 needs. Exact integer arithmetic, `scripts/make_figures.py`; record in `checks/copy_defect.json`.

Theorem 1 (LB1, the copy law). *Over every integral domain K and every n, a state is copiable if and only if it is classical.*
Checked: `copyIffClassical` *in* `MTB.Core.LinearFox`.

The proof is short and constructive, and it shows what the hypothesis is for. Idempotence in a domain forces each coordinate into $\{0, 1\}$, since $e(e-1) = 0$. Orthogonality forces at most one coordinate to be nonzero. A decided scan for the first nonzero coordinate returns the label k. The label is itself the computational content: the proof *is* a decoder from a copiable state to the classical value it carries.

Read backwards, the proof says that a copiable state of K^n is an orthogonal family of idempotents, which is a partition of $\operatorname{Spec} K$ into $n+1$ clopen pieces. So the true hypothesis is not "domain" but "no nontrivial idempotents".

Theorem 2 (sharp form). *Over every commutative ring K: every copiable state of every K^n is classical if and only if every idempotent of K is 0 or 1, i.e.* $\operatorname{Spec} K$ *is connected.*
Checked: `linearFox` *in* `MTB.Core.LinearFox`.

Refutation 1 (the hypothesis cannot be weakened to "reduced"). *Over $\mathbb{Z}\times\mathbb{Z}$ and over $\mathbb{Z}/6$, LB1 fails. In $\mathbb{Z}/6$, which has no nilpotents, $e = 3$ is a nontrivial idempotent and $\psi = (3, 4) \in (\mathbb{Z}/6)^2$ is copiable, normalised ($3 + 4 = 1$) and not classical. Conversely $\mathbb{Z}/4$ is connected but not a domain, and LB1 holds there.*
Checked: `z6RefutesLB1, reducedNotEnough, zzRefutesLB1, z4CopyClassifies` *in* `MTB.Core.LinearFox`.

Figure 2 shows the copy equation failing and succeeding on four two-dimensional states. The fourth panel is the content of Refutation 1: the defect vanishes, so no test that only looks at copying can tell $(3, 4)$ from a classical bit. A disconnected spectrum is a hidden classical resource, a partition of unity that copies perfectly without being a definite value.

Copy is a character; classical states are points. Pair a state with a vector by $\langle\psi, x\rangle = \sum_i \psi_i x_i$. Over every commutative ring, ψ is copiable if and only if $\langle\psi, xy\rangle = \langle\psi, x\rangle\langle\psi, y\rangle$ for all x, y: copiable states are exactly the (non-unital) algebra maps $K^n \to K$ (`copyIsCharacter`, `MTB.Core.CopyCharacter`). This is the characterisation of Coecke, Pavlović and Vicary [13], here proved over a ring rather

than a field. Adding the normalisation $\varepsilon\psi = 1$ excludes the zero state; over a connected nontrivial ring the normalised copiable states are exactly the n basis points, the K-points of $n \cdot \operatorname{Spec} K$ (`pointsOfSpec`). Over a disconnected ring the normalised state $(e, 1-e)$ copies without being classical (`normalisedSplit`). Normalisation does not rescue LB1.

No-cloning is the copy law. Over every commutative ring, a matrix that clones the basis *is* δ_Z, column by column, so it clones exactly the copiable states, and over a domain exactly the classical ones; universal cloning is impossible over every nontrivial ring (`clonerIsDelta`, `clonedIffClassical`, `noUniversalCloner`). The textbook proof's "missing cross term" [54] is precisely "a superposition is not a classical point". Our earlier library's linear no-cloning theorem over $\mathbb{F}_2$ is re-derived from LB1 in its exact original type (`noLinearClone'`, `MTB.Lane.C.NoCloneF2`). Separators make the refusal constructive: over every ring a separator is a failed copy equation, and over a domain one exists if and only if the state is non-classical (`MTB.Core.Separator`).

The other colour, and a coincidence that is not one. Transporting LB1 along a Hadamard H (any h with $h^2 + h^2 = 1$) classifies X-copiable states. It also forces the scalar: $\eta_Z = |0\rangle + |1\rangle$ is X-copied with scalar c if and only if $c = h = 1/\sqrt{2}$ (`xCopyScalarForced`). Our own plan had written $c = \sqrt{2}$; the kernel refuses that mutant. The Hadamard exists in $\mathbb{F}_p$ exactly when $p \equiv \pm 1 \pmod 8$ (measured by the oracle for all $p \le 100$; decided in the kernel for $p = 3, 5, 7$), which is the second supplement to quadratic reciprocity. It is also the condition under which the Galois automorphism σ_k of $\mathbb{Q}(\zeta_{16})$ fixes $\sqrt{2}$ (Section 3). Unitarity of a conjugate theory and the existence of the X-copy scalar are one fact, "$\sqrt{2}$ is in the field", read twice.

One level up. The same shape holds for thin categories: in every monoidal poset whose unit is top, the monoidal product is a categorical product if and only if every object is copiable (`posetFox`, `MTB.Core.PosetFox`, LB2). This is Fox's theorem [20] in the setting where a residuated poset is a Heyting algebra. LB1 is its linear shadow: in $\mathrm{Mat}(K)$ the cartesian core, the maps that commute with copying, is FinSet over a connected ring (`MTB.Core.ClassicalCore`).

A constructor-theoretic reading (interpretation). Constructor theory calls a set of attributes an *information variable* when a constructor can clone it [16]. LB1 identifies the information variables of linear quantum theory over a ring: exactly the bases of clopen pieces of the spectrum. The analogy has a limit. Constructor theory quantifies over all physically possible constructors, while LB1 quantifies over linear maps in $\mathrm{Mat}(K)$. The connection becomes a theorem only after a constructor-theoretic substrate is formalised and shown to send cloning tasks to copy equations; we record this as an open problem (Section 9).

2.2 The price of committing

A non-classical noun cannot be copied, but it can be measured first and the record copied. The question is what this costs and when it changes the meaning. On densities, the copy $\hat{\delta} = \delta_{n^2} \circ \mathrm{mid}$ is obtained by CPM doubling, and the density copy core is again LB1 (`MTB.Core.DensityCopy`).

Theorem 3 (broadcasting and its price). *Over every commutative ring, every dimension, and any ring endomorphism φ standing in for conjugation:*

(a) δ_Z broadcasts a density ρ if and only if ρ is diagonal (`broadcastsIff`);

(b) for a pure state ψ with weights $p_i = \psi_i \varphi(\psi_i)$, $(\sum_i p_i)^2 = \sum_i p_i^2 + \mathrm{sepMass}(\psi)$, where sepMass *is the squared mass of the off-diagonal separator rows (`separatorMass≡logicalEntropy`);*

(c) hence for $\sum_i p_i = 1$ the purity lost by measuring in the copy basis is exactly the logical entropy $h(p) = 1 - \sum_i p_i^2$ (`logicalEntropyNormalised`).

Checked: `broadcastsIff, purityLoss, logicalEntropyNormalised` in `MTB.Core.Broadcast`.

Part (a) is the no-broadcasting theorem [8] at a fixed broadcaster, proved in every dimension over a ring. Parts (b)–(c) say something we did not expect: the separator mass of the copy equation and the logical entropy of the measurement statistics [19] are the same number. The price is *phase-blind*. The state ιe_0, with $\iota^2 = -1$, is fixed by decoherence ($h = 0$) yet fails its copy equation at entry 00, since $\iota \neq \iota^2$ (`phaseBlind`, `MTB.Lane.K.MeasurePoint`). Copying sees phases that measurement does not.

Measurement points in a grammar. A measured term is interpreted as a Kraus family, one pure reading per outcome assignment, and its density is $\sum_\tau |K_\tau\rangle\langle K_\tau|$ (`MTB.Lane.K.MeasurePoint`). The lowering `copySafeLower` inserts a point before each copy of a non-classical noun. We proved:

- it repairs exactly what it can: the lowered term is accepted if and only if the original was structurally well formed (`copySafeLowerAccepted`, `copySafeLowerExact`). The repetition of a spoken word (S22) stays refused, because no measurement repairs a reused occurrence;
- accepted terms are left untouched, and inserting points where nothing is quantum changes no meaning (`copySafeLowerIdle`, `measureHarmless`);
- a measured pronoun decoheres its output, and the measurement is *invisible* if and only if the clause selects at most one referent (`measureIsOutputDecoherence`, `invisibleIffAtMostOne`).

For S23 the lowering inserts one point at price $h = \frac{1}{2}$, and that point is invisible: Bob does not sleep, so $F = (1, 0)$ and the measured density equals the pure Frobenius density (`s23NeededButInvisible`). The phrase "Alice who does not love herself", with the same quantum noun, needs its point and shows it. Three counts are proved ordered, semantic ≤ lowered ≤ static (`semantic≤lowered≤static`).

Measurement 1 (how often a copy needs its point). Over 32,380 generated terms (exhaustive to size 6, random to size 9), the static count of required points totals 2,209 and the semantic count 466. Of the 466 copies that semantically need a point, 441 are invisible. No term with semantic need 0 was visible. The decoherence identity held on all 1,450 lowered terms tested. Source: B2b findings; oracle-reported, not a theorem.

Refutation 2 (prices do not compose under nesting). *On flat terms, in which every relative clause has a relation-free noun and body, the measurement-point prices and the pronoun prices agree point by point (`priceBridge`). In general they do not. S23 nested inside a relative clause about Bob receives two measurement-point prices, one per branch of the outer copy, against one pronoun price.*

Checked: `priceBridge, nestingRefuted, unrestrictedRefuted` in `MTB.Lane.K.PriceBridge`.

2.3 A shared weight must be sharp

Learning adds parameters. Following the categorical account of gradient-based learning [14], a parametrised map $s \to t$ is a pair $(P, f : P \otimes s \to t)$, and composition tensors the parameters: $(Q, g) \circ (P, f) = (Q \otimes P,\ g \circ (Q \otimes f) \circ \alpha)$. In $\mathrm{Mat}(K)$ the input side of $P \otimes -$ is a *copower*, so a parametrised map is a P-indexed family of matrices and composition is fibrewise (`MTB.Core.ParaHybrid`). The category laws hold up to their stated reparametrisation 2-cells, and $P \otimes -$ is an actegory (`act-`

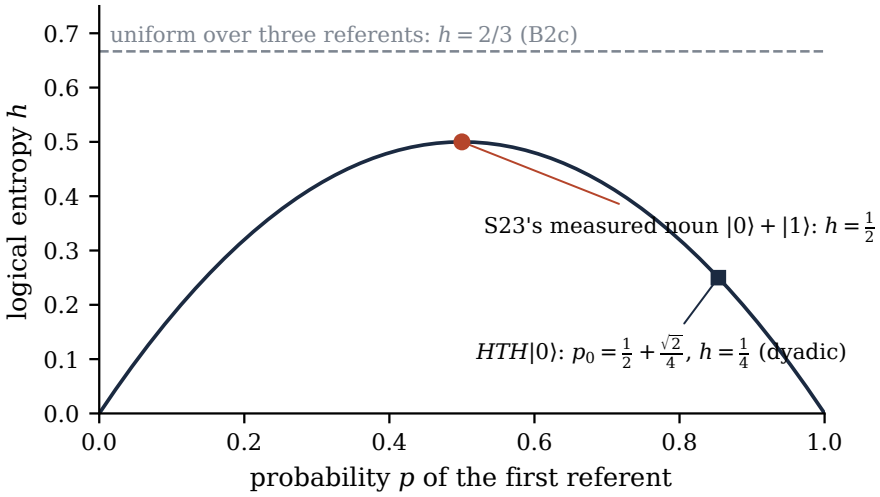


Figure 3. The price of committing. The exact cost of a measurement point is the logical entropy $h = 1 - \sum p_i^2$ (Theorem 3). S23's quantum noun costs $\frac{1}{2}$; uniform uncertainty over three referents costs $\frac{2}{3}$. The one-T-gate state $HTH|0\rangle$ has irrational weights $\frac{1}{2} \pm \frac{\sqrt{2}}{4}$, yet its price is exactly $\frac{1}{4}$: the two weights form one Galois orbit, and a symmetric function of an orbit is rational (`tEntropyDyadic`, §5).

`id`, `act-mult`, `act-unit`). Evaluation at independent parameter states is functorial at every state, quantum ones included.

Weight tying is the one well-typed way to compose without tensoring parameters: reuse P by reparametrising along δ_P. Our plan predicted that tying would break associativity and planted that as a mutant. The kernel refused the mutant because the prediction was false: tying is the coKleisli category of the reader comonad built from the copy, and associativity is coassociativity of δ (`tieAssoc`). What tying breaks is evaluation.

Theorem 4 (weight-tying law). *Over every commutative ring and every parameter shape, tied evaluation at a parameter state p is functorial if and only if p is* sharp*: copiable and normalised. Hence tiable $\iff$ p is a point of the cartesian core, and over a connected nontrivial ring tiable $\iff$ p is a basis point. The zero state copies but cannot be tied, and over $\mathbb{Z}[\frac{1}{2}][\zeta_8]$ neither the mixture $(\frac{1}{2}, \frac{1}{2})$ nor $|+\rangle$ can be tied.*

Checked: `tyingSound↔`, `tyingIffCore`, `tyingIffBasis`, `zeroNotTiable`, `mixNotTiable`, `plusNotTiable` in `MTB.Core.ParaHybrid`.

Over $\mathbb{Z}/6$ the state $(3, 4)$ is sharp, so tying is sound there, but it is not a basis point (`tyingNeedsConnected`). One counter-model refutes both LB1 and the basis form of the tying law without connectedness: the copy law and the tying law are one fact. A random shared weight fails too, because $\mathbb{E}[g_a \circ f_a] \neq \mathbb{E}[g] \circ \mathbb{E}[f]$. A shared weight must be a definite classical value.

The derivative side agrees. Read the parameter circle with coordinate z; for $f(z_1, z_2)$ of degree one in each copy, read at the tied point $z_1 = z_2 = z$,

$$\tfrac{1}{2}\big(f(\iota z, z) - f(-\iota z, z)\big) + \tfrac{1}{2}\big(f(z, \iota z) - f(z, -\iota z)\big)$$

is the dual-number tangent of $f(z, z)$ (`tiedGradient`, `MTB.Lane.K.ShiftSpec`). The gradient of a tied weight is the sum of the per-copy parameter shifts: the chain rule through the copy, as a theorem. Applying the plain shift rule to the tied function is wrong exactly when f has a frequency-two term (`doubleAngleBreaks`).

The hybrid category. Objects are pairs (classical register, density). Quantum gates act by $U \otimes \bar{U}$, classical gates by function matrices, and the parameter action joins the *classical* register undoubled, a choice Theorem 4 forces. Measurement is the transpose of the diagonal embedding, and the Born matrix has entries $U_{ki}\varphi(U_{ki})$, exact and free of square roots (`born`). Measurement commutes with a coherent classical gate if and only if the gate is injective (`measNatIff`): a non-injective gate carries a coherence onto the diagonal, which is interference, while measuring first would have discarded it. A diagonal gate is invisible to measurement exactly when its entries have unit norm (`measPhaseIff`).

3 Exact rings decide which theory is physical

3.1 Rings with canonical forms and Galois maps by universal property

Every amplitude in this work lives in a ring we can compute in exactly:

- $R_8 = \mathbb{Z}[\frac{1}{2}][\zeta_8]$ for Clifford+T and the Born head;
- $R_{16} = \mathbb{Z}[\frac{1}{2}][\zeta_{16}]$ for Ising braiding and parameter shifts;
- $R_{10} = \mathbb{Z}[\zeta_{10}]$ for Fibonacci, with the golden extensions $\mathbb{Z}[\varphi] \subset \mathbb{Z}[\sqrt{\varphi}, \zeta_{10}]$.

One construction serves all of them: $\mathbb{Z}[x]/(f)$ and $\mathbb{Z}[\frac{1}{2}][x]/(f)$ are commutative rings for *every* monic f, with a power basis, a companion-matrix product and a dyadic normal form. The solver's normal form is canonical: two expressions are equal as functions on the ring exactly when they have the same normal form (`canonicity`, `MTB.Core.SolverCanonical`). Decidable equality is therefore a theorem, not a convention, and the same bytes come out on every machine (Section 5.5).

Galois maps are not written down by hand. Each $\sigma_k : \zeta \mapsto \zeta^k$ is obtained from the universal property of $\mathbb{Z}[x]/(f)$ as the ring that corepresents roots of f. It exists, is unique, and the σ_k form a group (`MTB.Core.Galois`, `MTB.Core.GaloisGroup`). The first physical fact follows at once: over R_{16}, σ_k fixes $\sqrt{2}$ if and only if $k \equiv \pm 1 \pmod 8$, for every odd k (`unitaryIff`). The relative norms are Galois-fixed (`normFixed16`, `normFixed10`). The Kauffman bracket is natural under every ring homomorphism, so each Galois map sends the bracket of a theory to the bracket of its conjugate (`bracketGaloisEquivariant`, LB3).

These exact computations corrected our own bindings more than once. The plan bound Ising to $A = \zeta_{16}^3$; the kernel shows this is Kitaev's $\nu = 1$ member [32], while $\mathrm{SU}(2)_2$ is $A = \zeta_{16}$ with loop value $\delta = -\sqrt{2}$. At the Ising point $\sigma^4 = -\iota$, the Hopf link's bracket vanishes, and the Jones polynomial of the torus link $T(2, n)$ at ι is the Jacobi symbol $(2/n)$ (`MTB.Lane.C.IsingBracket`).

3.2 The Gauss law is Galois transport, not unitarity

Our pitch conjectured that, among the Galois conjugates of the ZX scalar model over the dyadic ring $D[\omega]$ ($\omega = \zeta_8$), an internal "Gauss law" detects unitarity, i.e. the sign of $\sqrt{2}$. Galois conjugates of topological phases are a standard way to produce non-unitary theories [23], so the conjecture was natural. It is false.

Refutation 3 (Gauss law ⇏ unitarity). *In the four grid models $\omega \mapsto \omega^j$, $j \in \{1, 3, 5, 7\}$, the Gauss law holds in all four, while σ_j fixes $\sqrt{2}$ only for $j \in \{1, 7\}$ (Table 1). The conjecture fails at exactly $j \in \{3, 5\}$; the table was pre-registered before the check.*

Checked: `gaussLawUnitarityRefuted`, `refutedExactly`, `gridGauss`, `fixesTable` in `MTB.Core.GaussTransport`.

j	Gauss law	$\sigma_j(\sqrt{2}) = \sqrt{2}$	local complementation sound	on the Galois orbit
1	yes	yes	yes	yes
3	yes	**no**	yes	yes
5	yes	**no**	yes	yes
7	yes	yes	yes	yes

Table 1. The pre-registered table, decided. Columns 2–3 are Agda theorems (`gridGauss`, `fixesTable`); columns 4–5 are Lean theorems over the earlier ZX scalar formalisation (`gaussLawTable_eq`, `grid_lc`). The Gauss law follows the last column, not the second.

Theorem 5 (the law behind the table). *For every scalar model over $D[\omega]$ satisfying the phase-group laws, the Gauss law holds if and only if the model is a Galois transport of the standard model, $\exists j.\ \sigma_j\omega = \zeta \wedge \sigma_j\sqrt{2} = root$. Every such model has root $\pm\sqrt{2}$, and every ring endomorphism of $D[\omega]$ is a Galois map.*

Checked: `gaussIffTransport, rootCensus, lawCensus, endoClassify` in `MTB.Core.GaussTransport`.

In Lean 4, over the scalar model of our earlier ZX formalisation, local complementation is exactly sound if and only if the model is a Galois transport (`lc_iff_transport`). The group-law models fall into eight cells, and local complementation holds in exactly the four whose root is $\sigma_j\sqrt{2}$ (`lc_census`, `grid_lc`, `flipped_not_lc`; file `lean/HeytingLean/ZX/GPUFull/Exact/GaussUnitarity.lean`, 59 declarations, axioms $\subseteq$ {`propext`, `Quot.sound`}).

A transported model cannot fail a transported law, and "σ_j fixes $\sqrt{2}$" is not a transported property. The conjecture compared a Galois-invariant property with a Galois-variant one. The other reading survives on a different family: if the character is conjugated but the root is held at $+\sqrt{2}$, then Gauss law, local complementation and "σ_j fixes $\sqrt{2}$" coincide (`fixedRootGauss` for the Gauss part). That statement compares the model with an externally *chosen* $\sqrt{2}$, the same move by which an ordering of a field chooses a real embedding (Section 5). We withdraw the pitch's sentence as a statement about grid models. For the Gauss-law statements the Agda module is the authority, with an empty trust footprint. The local-complementation statements are Lean theorems with the footprint stated above, and they are a reported bridge to the earlier formalisation.

3.3 Fibonacci: one gauge orbit, two braidings, and a Galois twin

For the Fibonacci theory we proved the following, over every golden ring (any commutative ring with $\varphi^2 = \varphi + 1$):

- **One orbit.** The pentagon law is gauge-invariant, and its solution locus is *one* gauge orbit, a torsor. The non-unitary solution F' landed in our earlier library is an explicit gauge transform of the unitary F in $\mathbb{Z}[\sqrt{\varphi}]$, but no gauge transform of it inside $\mathbb{Z}[\zeta_{10}]$ is unitary. The obstruction is found by the unique ring map $\mathbb{Z}[\zeta_{10}] \to \mathbb{F}_5$, at the prime that ramifies over $\sqrt{5}$ (`MTB.Core.FibPentagon`, B3b1). The product formula explains why. The Galois conjugate of φ^{-1} is $-\varphi < 0$, which is not a norm from a CM field, so the obstruction must reappear at a finite place.
- **Hexagons force the pentagon.** Both hexagons hold over every golden ring, and they force the pentagon (`fibHexagons`, LB4). In $\mathbb{Z}[\zeta_{10}]$ there are exactly two braidings, $(R_1, R_\tau) = (\zeta^6, \zeta^3)$ and (ζ^4, ζ^7), whereas the braid relation alone admits exactly 30 pairs, because it sees only the ratio R_1/R_τ (`hexCensus`, `braidCensus`, `MTB.Core.FibHexagonCensus`). A test built on the braid relation alone cannot single out the two genuine braidings among those 30.
- **The Galois twin.** The Galois map σ carries Fibonacci to Yang–Lee, the non-unitary theory with

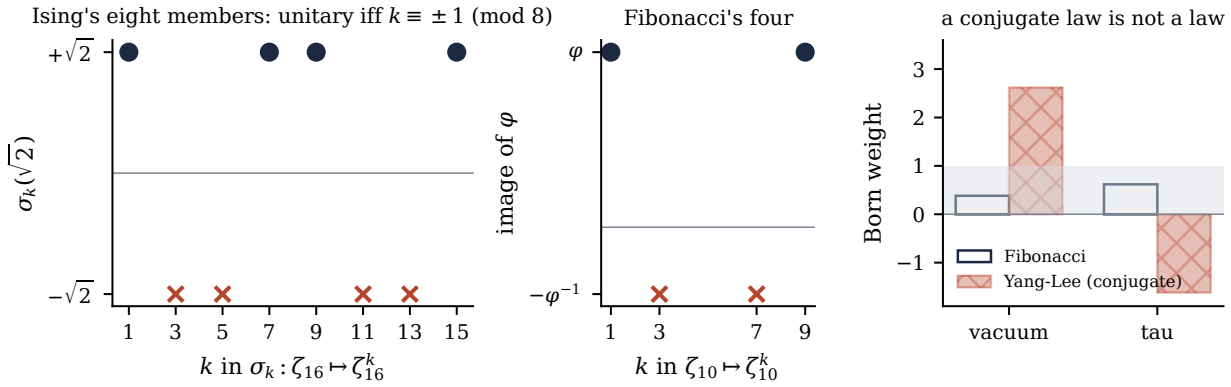


Figure 4. Which theories are physical. Left: the eight Galois members of Ising over $\mathbb{Q}(\zeta_{16})$; σ_k fixes $\sqrt{2}$ exactly when $k \equiv \pm 1 \pmod 8$ (`unitaryIff`). Centre: the four Fibonacci members send φ to φ or to $-\varphi^{-1}$ (Yang–Lee). Right: the Born law of $\tau \otimes \tau \to \{\mathbf{1}, \tau\}$ is $(\varphi^{-2}, \varphi^{-1})$ for Fibonacci; its Galois image $(1+\varphi, -\varphi)$ leaves $[0, 1]$ (shaded) and is not a probability law, so the acceptance rule of Section 5 rejects every histogram against it (`yangLeeRejectsAll`). Values recomputed in `checks/galois.json`.

$d = -\varphi^{-1}$. Within $\mathbb{Z}[\sqrt{\varphi}]$ the unitary members are decided exactly (`fibUnitaryIff`), and the unitary hexagons hold on the landed carrier (`unitaryFibHexagons`).

What is observable (interpretation). Because the pentagon locus is a single orbit, a quantity is a property of the Fibonacci theory exactly when it is constant on that orbit. Gauge-dependent F- and R-entries are not observables, however carefully they are measured. The ratio R_1/R_τ is gauge-invariant, but by the census it does not determine the theory. Together the two facts specify which experimental targets are well posed. Section 6.2 turns this into a proposal.

3.4 Manifold invariants and the Legendre law

The state sums run on certified triangulations. For the toric code ($\mathrm{Vec}_{\mathbb{Z}_2}$) the Turaev–Viro state sum is a count of homomorphisms. For Fibonacci we proved Pachner invariance for every complex, using a symmetric unitary $6j$ that needs no branching structure, together with Reshetikhin–Turaev data from the braiding: S, T, Gauss sums, the Galois law and surgery words. On each of the five corpus manifolds, $\mathrm{TV} = |\mathrm{RT}|^2$ is a kernel theorem (`tvFib-S3` and its siblings, `MTB.Lane.K.TVFib`). The general identity $\mathrm{TV}(C) = \mathrm{RT}(Z(C))$ of Turaev and Walker [31, 44, 50] is *cited*, not proved here.

For lens spaces we go further. Write $p/q = [a_1, \ldots, a_k]$ as a Hirzebruch–Jung continued fraction, and let $W = \tilde{S}T^{a_1}\tilde{S}\cdots T^{a_k}\tilde{S}$ be the corresponding surgery word [28].

Theorem 6 (the Legendre law, for every lens space). *For every coprime (p, q), the Fibonacci invariant of $L(p, q)$ depends only on Legendre symbols mod 5:*

$$\mathrm{TV}_{\mathrm{Fib}}(L(p,q)) = \begin{cases} 1/D^2 & 5 \nmid p,\ (p/5) = +1, \\ 1/\sigma(D^2) & 5 \nmid p,\ (p/5) = -1, \\ \tfrac{1}{2}\big(1 + (q/5)\big) & 5 \mid p, \end{cases} \qquad D^2 = 1 + \varphi^2.$$

The proof tabulates the image of $\mathrm{SL}(2, \mathbb{Z})$ in $\mathrm{SL}(2, \mathbb{Z}/5)$ (120 elements), decides closure under S and T up to a root-of-unity scalar, and reduces every surgery word to its image by induction. No presentation of $\mathrm{SL}(2, \mathbb{Z}/5)$ is needed.

Checked: `fibLegendre` in `MTB.Core.FibLens`.

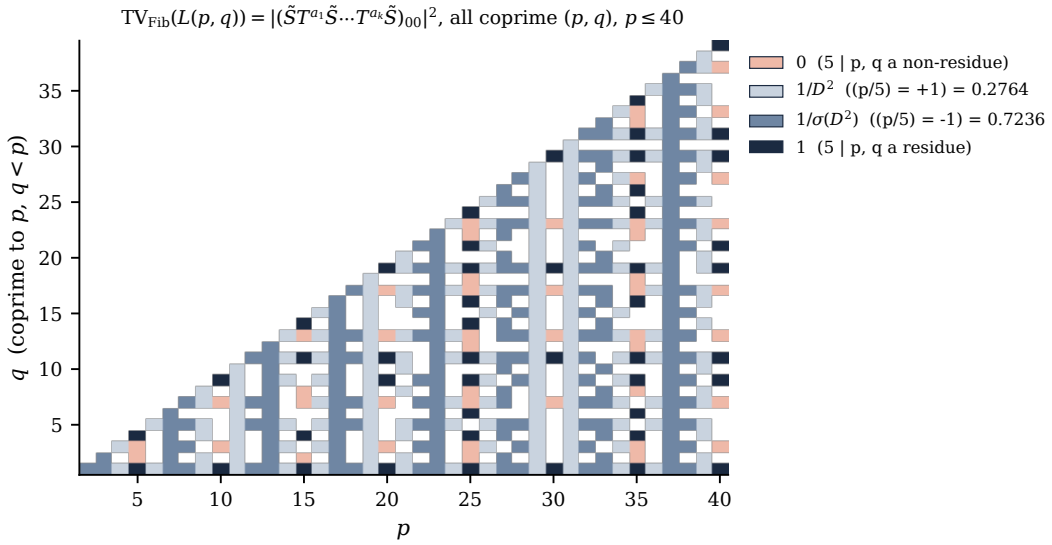


Figure 5. Fibonacci reads a lens space through the prime 5. $\mathrm{TV}_{\mathrm{Fib}}(L(p,q))$ for every coprime pair with $p \le 40$, computed from the modular representation and coloured by value. Only four values occur. Columns with $5 \nmid p$ are constant in q and alternate with the Legendre symbol $(p/5)$; columns with $5 \mid p$ split by the residue of q. The proved law is Theorem 6; this plot is the recomputation of Measurement 2.

This was a conjecture with exhaustive evidence (489 coprime pairs with $p \le 40$) until the hostile audit asked for its authority, because the paper would otherwise have had to call it evidence. Its consequences are concrete:

- Fibonacci separates $L(5,1)$ from $L(5,2)$ by vanishing (values 1 and 0). The two have the same homology and the same toric-code value and are not homotopy equivalent (`fibSeparatesL51L52`).
- It is blind to the homotopy-equivalent but non-homeomorphic pair $L(7,1)$, $L(7,2)$ (`fibBlindL71L72`).
- The toric code reads p mod 2 and the double semion p mod 4 (`MTB.Lane.K.ToricLens`).

Different anyon theories read different primes: each theory sees a lens space through its own Frobenius.

Measurement 2 (independent recomputation in this authoring session). Using floating-point 2×2 matrices for the Fibonacci modular data ($S = \frac{1}{D}\begin{pmatrix} 1 & \varphi \\ \varphi & -1 \end{pmatrix}$, $T = \mathrm{diag}(1, e^{4\pi\imath/5})$), we evaluated $|(ST^{a_1}S\cdots T^{a_k}S)_{00}|^2$ for all 489 coprime pairs $2 \le p \le 40$ and compared each value with Theorem 6. The maximum absolute deviation is 3.1×10^{-15}, and the values for $L(5,1)$, $L(5,2)$, $L(7,1)$ and $L(10,3)$ are those stated above (Fig. 5; `checks/lens_law.json`). This check is numerical and is not part of the kernel authority.

Exactness on an engine. The same amplitudes were also run as ZX diagrams on our exact tensor engine (CPU, no GPU): all 12 certified complexes under both theories, 58 runs over 827 slices, every value equal to the kernel's (reported, B3c4). An independent check by textbook surgery, $|\sum_a d_a^2 \theta_a^p|^2/D^4$ for $L(p,1)$ and $\mathrm{rank}(C)^2$ for T^3, agreed on 60 of 60 configurations (X3).

4 What braiding can and cannot add

Topological quantum computation braids anyons; compositional semantics routes wires. Our pitch asked whether meaning itself can braid: whether a braided monoidal semantics distinguishes readings that a symmetric one conflates, and whether that distinction helps a learner. The answers are mostly negative, and the negative results are exact. We present them as the section's headline.

4.1 The landed grammar is planar

The QNLP compiler we built on is a planar pregroup grammar: its wires never cross. Our plan called for a "forget the braiding" functor from a braided syntax into it. Such a functor cannot exist, because the landed grammar lands in the *symmetric* syntax. Its parse table is exactly the cartesian readback of the braided parses (`MTB.Lane.K.BraidedGrammar`, B4a). We therefore built the braided syntax as a separate grammar, derived the Yang–Baxter equation in it, and proved the landed fragment to be a natural sub-grammar. The writhe sees no corpus sentence; the non-symmetry lives in braid words.

4.2 Points are transparent

Theorem 7 (closed readings are routing-blind). *In every braided semantics in which names are states of the unit, crossing two named nouns is a routing and nothing more: $\sigma_{N,N} \circ (a \otimes b) = (b \otimes a) \circ \sigma_{\mathbb{I},\mathbb{I}} = b \otimes a$. Hence the over- and under-crossed parses of a closed sentence are equal, in every R-matrix model, for every labelling and every continuation.*
Checked: `pointsTransparent, crossBlindOnPoints, closedRoutingBlind` *in* `MTB.Core.BraidedMat`.

Refutation 4 (the planned separation of closed readings). *The plan's statement that "the two parses of S13 ('Alice Bob loves') differ by σ^2 in braided semantics" is false for the closed sentence, under both Ising and Fibonacci (`s13ReadingsEqualIsing, s13ReadingsEqualFib`). What a closed sentence can see is the* routing*, identity versus crossing, i.e. who did what to whom (`s13RoutingVisible`); it cannot see the crossing's label.*

One theorem makes every R-matrix model sound on every wire. If the two atomic crossings cancel, satisfy the Yang–Baxter slide, and let each box slide through an atom, the hexagon-defined crossing is sound for every generator of the braided congruence (`sound`, in the `Sound` submodule of `MTB.Core.BraidedMat`). The proof rests on a four-block interchange for the hexagons and naturality in both arguments derived from one. Two instances follow: the swap model and the Kauffman strand model.

4.3 The open frame carries a character of the braid group

The label lives where the names are not yet filled in. Take the open frame "_ _ loves", with nouns as two-strand cables in the Kauffman model $R = A \cdot 1 + A^{-1}U$.

Theorem 8 (the open frame reads $\theta^{-\mathrm{writhe}}$). *Over every commutative ring with A invertible, the frame of a braid word w satisfies $\delta F_w = \lambda(w)\,\delta N + (1 - \lambda(w))\,C$, where λ is a character of the braid group with $\lambda(over) = A^{-8} = \theta^{-1}$, $\theta = A^8$ the spin-1 twist. The over- and under-readings differ by the double braid's eigenvalue θ^{-2} on the cable's spin-1 channel; the vacuum channel is fixed. Two labels give equal frames if and only if $(A^4 - 1)(A^8 + 1) = 0$, and in an integral domain if and only if $\delta^2 \in \{2, 4\}$.*
Checked: `monodromy, wordLaw, factor, blindIffClassical-⟶` *in* `MTB.Core.BraidWord, MTB.Core.KauffmanFrame`.

- **Ising** ($\theta = -1$) reads only the parity of a word (`isingParity-→`, `isingParity-←`). Its braid representation on the open frame factors through S_2, the same quotient as the parse double cover and as symmetric semantics. The odd (σ) strand still sees the label (`isingStrandApart`), but the nameable sector is symmetric.
- **Fibonacci** has period exactly five. σ^5 reads as the identity routing and $\sigma^1, \ldots, \sigma^4$ do not; σ^4 reads as σ^{-1}, as the oracle predicted (writhe $4 \equiv -1 \bmod 5$).

Conservativity holds for every diagram: the braided semantics agrees with the landed compiler wherever the latter is defined (`braidedConservative`, LB5). The anyon value of a closed corpus sentence is δ^c, where c is the number of separately closed sub-circuits: δ^2 for S11, S18 and S22, and δ for the rest. It depends on the components of the wiring graph and on nothing else. An earlier record stated "δ for every sentence"; the executable caught the error on S18, and the record was superseded rather than rewritten (X1, X2).

4.4 No symmetric fibre functor

Equivariant architectures are representations of groups, and a symmetric monoidal category with a fibre functor is, by Deligne's theorem [15], the category of representations of a (super)group. So whether a braided meaning category can be seen by a group-equivariant model reduces to whether it has a symmetric fibre functor.

Theorem 9 (LB6). *Let F be a braided functor into any K-linear symmetric target with $F(X) \neq 0$. On the hom-algebra* $\mathrm{End}(X \otimes X)$ *the double braiding must map to the identity channelwise. A Bezout certificate $w_c(R_c^2 - 1) = N$, with N a non-zero-divisor on the target, then kills every channel.*

- ***Ising.*** *The deficits are units ($N = 1$), so no symmetric fibre exists in any symmetric target over any base, not even in the super line (`isingNoSymmetricFibre`, `isingBaseChange`, `isingNoSuper`).*
- ***Fibonacci.*** *Over every golden ring and for every hexagon solution, $w_2(t^2-1) = 5$ and $w_4(t^4-1) = 5$, so no symmetric fibre exists in any target where 5 is a non-zero-divisor (`fibNoSymmetricFibre`).*

Checked: `noSymmetricFibre` in `MTB.Lane.C.NoSymmetricFibre`.

Theorem 10 (the Fibonacci dichotomy). *Over every golden integral domain and every braiding the hexagons admit, a faithful symmetric fibre at τ exists in* $\mathrm{Mat}(K^2 \otimes K^2)$ *if and only if* $5 = 0$. *In that case the braiding* is *the swap's signature, and Maschke's decomposition gives the fibre.*
Checked: `fibonacciDichotomy`, `ramifiedGeneric` in `MTB.Core.SymmetricFibreTheories`.

At $p = 5$ this is Ostrik's Ver_5^+ [43]: Fibonacci reduced at its ramified prime is a symmetric fusion category. The certificate cannot be improved, because 5 is the least positive integer in both ideals. Statistics is a separate obstruction: ψ is super-Tannakian and σ is not. The interpretive claim of our pitch, that braided language models lie outside group equivariance, is reduced to what is proved: no symmetric fibre functor exists. The architecture-level statement remains interpretation.

4.5 Ising on cabled nouns is super-symmetric

A noun in the braided Ising semantics is a $\sigma\sigma$-cable, whose fusion channel $\mathbf{1}/\psi$ is a qubit.

Theorem 11 (the Ising crossing of cables is the Koszul swap). *The σ-exchange is diagonal on a pair's channel with the landed R-symbols $R_{\mathbf{1}}^{\sigma\sigma} = -A^{-3}$, $R_{\psi}^{\sigma\sigma} = A$ ($A = \zeta_{16}^3$), and its square is not 1. Nevertheless, the crossing of two $\sigma\sigma$-cables is the Koszul super-swap $|q_0 q_1\rangle \mapsto (-1)^{q_0 q_1} |q_1 q_0\rangle$. For the experiment's wire of two pairs (16 exchanges, 16×16), it is $(\sqrt{2})^{16}$ times the Koszul swap, exactly in $\mathbb{Z}[\zeta_{16}]$.*

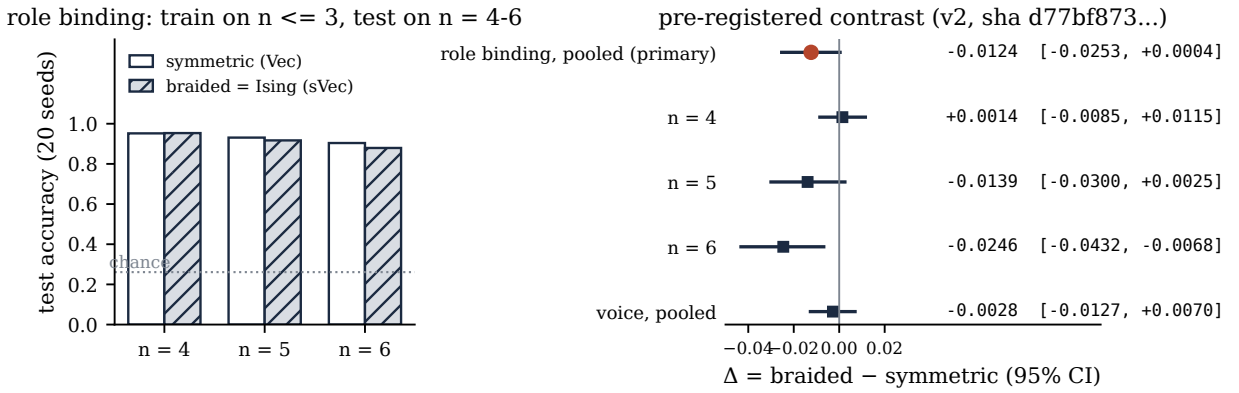


Figure 6. What braiding can add: nothing measurable here. Left: test accuracy per number of entities (training used $n \leq 3$); both semantics generalise at about 3.5× chance. Right: the pre-registered contrast with 95% CIs (hierarchical paired bootstrap). The primary interval contains 0. Data: `data/b4e_v2_results.json`, copied byte for byte from the project; plotted by `scripts/make_figures.py`.

Checked: `exchangeIsIsing`, `cableKoszul`, `wireKoszul`$_2$ in `MTB.Core.IsingCable`, `MTB.Core.IsingCableWire`.

Categorically, $\sigma \otimes \sigma = \mathbf{1} \oplus \psi$ generates sVec inside Ising, and sVec is symmetric. By coherence, every relayout acts on a routed even box by a Koszul permutation, so the Born outputs of an Ising-braided text circuit on cabled nouns cannot depend on layout. This was predicted from theory before any data, and the prediction became the experiment's design.

4.6 The pre-registered experiment

Measurement 3 (braided versus symmetric text circuits). *Question.* On role binding with compositional generalisation in the number of entities, does an Ising-braided text circuit beat the symmetric text circuit of the same architecture and parameter count? *Design.* Stories with n entities from 8 names and $n+1$ events over 3 verbs (active or passive). The question is "who ⟨verb⟩ ⟨object⟩?" with a unique answer. Training uses $n \leq 3$ and testing $n = 4, 5, 6$. Nouns are two cabled σ-pairs (two parity-graded qubits) and verbs are even unitaries. Both models have 1,568 parameters and share data, initialisation and minibatches per seed; they differ only in the crossing (swap versus Koszul swap). There are 20 seeds, 3,000 Adam steps each, and the learning rate is selected on validation. *Decision rule.* Δ = mean over seeds of accuracy(braided) − accuracy(symmetric), with a 95% CI from a hierarchical paired bootstrap ($B = 10^4$). A CI excluding 0 decides the comparison; otherwise the result is null. *Result.* Pooled test accuracy is 0.9289 for symmetric and 0.9166 for braided (chance 0.261), $\Delta = -0.0124$, CI $[-0.0253, +0.0004]$: **null**. Per n: null at 4 and 5; at $n = 6$, $\Delta = -0.0246$, CI $[-0.0432, -0.0068]$, one of three uncorrected secondary intervals, so a trend and not a finding. Both models answered identically under every relayout (1.0000), as predicted. The secondary "voice" task is unidentifiable by theorem (paired accuracy exactly 0.5000 for both models). The CI half-width (≈ 0.013) means a braided advantage of 1.3 points or more would have been detected. *Pre-registration.* v2, sha256 `d77bf873ab2c43c3...`, frozen before its first run. CPU only.

The history of this experiment is part of its result. The first pre-registration (v1, sha256 `de5158fb...`) ran to completion and was null as well, but its braided arm was defective. A box with ordered ports needs its subject on the left, and when the layout put the subject to the right of the object, v1 omitted the Koszul crossing that brings it there, silently using a port-swapped box. Its

own certificate checked the implementation's convention rather than fidelity to the braided category, which is an instance of a gate whose input the author wrote. A first-principles review found the defect: coherence predicts exact relayout invariance, and v1 was not invariant. v2 changed that one sign and nothing else. It was pre-registered after v1's results were known, a fact we disclose rather than hide. Fibonacci was not run, because its cable crossing mixes fusion channels and does not factor through a signed swap.

What the negative result says. On cabled nouns, Ising braiding differs from symmetric meaning only by the Koszul sign rule, which has no bearing on who did what to whom. In this regime braiding is not a source of role-binding advantage. That does not rule out braiding elsewhere: Theorem 8 leaves open frames with a nontrivial character, and Fibonacci has period five there. But it moves the burden. A proposal that braided semantics helps a learner now has to name the open-frame quantity it exploits.

4.7 Contextuality as an exact gate on data

Ambiguous phrases can be contextual in the sheaf-theoretic sense [2, 51], and contextuality is the proved source of a quantum memory advantage in sequence models [6]. We built the exact obstruction and measured what it sees.

Theorem 12 (the Čech obstruction is sound; its converse is false). *For every measurement scenario and every coefficient ring, a nonzero Čech class of the support presheaf certifies contextuality [3]. Certificates are relative 1-cycles, decidable in the kernel, and signalling models are refused, a refusal that carries weight in the proof. The converse fails: a logically contextual model with vanishing class exists, and the false negative is a signed section.*
Checked: `obstructionSound` *in* `MTB.Core.Contextuality`.

Theorem 13 (what the $\mathbb{Z}$-obstruction sees in the (2,2,2) scenario). *The kernel decides the full census: the 2,961 non-signalling supports among* $15^4 = 50{,}625$*; 32,448 sections, of which 2,560 are logically contextual. The $\mathbb{Z}$-obstruction detects 320 of these and misses 2,240. Characterisation: in every non-signalling (2,2,2) model, a logically contextual section is $\mathbb{Z}$-detected exactly when the other three contexts are parity supports (All-versus-Nothing), and its class vanishes otherwise.*
Checked: `coverage, counts, characterisation` *in* `MTB.Core.Contextuality222`.

On the corpus, the Popescu–Rohrlich box is strongly contextual. The first adjective–noun model of Wang et al. [51] is obstructed exactly at one named section. Its second model and the local control have vanishing classes, with global sections exhibited. The census makes the detector's blind spot a number: 87.5% of logically contextual (2,2,2) sections are invisible to the $\mathbb{Z}$-class. Section 6.1 uses this.

5 The exact hybrid loop, and from proof to machine

A hybrid quantum–classical learner has six steps:

(i) it holds parameters;

(ii) it inserts measurement points;

(iii) it reads out through the Born rule;

(iv) it differentiates;

(v) it decides which parts need a quantum device;

(vi) it accepts or rejects what the device returns.

Section 2.3 and Section 2.2 gave (i) and (ii). This section makes (iii)–(vi) exact and then shows that the machine executes the mathematics.

5.1 An exact parameter-shift rule that needs no unitarity

Write $z = e^{\iota\theta}$. Every Pauli rotation with its global phase removed is *affine* in z: $R(z) = \Pi_+ + z\Pi_-$ with $\Pi_\pm = \frac{1}{2}(I \pm P)$. Hence $f(z) = \langle R(z)\psi|O|R(z)\psi\rangle$ is a Laurent polynomial $a_-z^{-1} + a_0 + a_+z$ on every register, for every observable and state. The whole parameter-shift rule [39, 47, 53] is then a single identity of Laurent polynomials.

Theorem 14 (exact grid parameter shift). *Over every commutative ring with $h^2 + h^2 = 1$ and $\iota^2 = -1$: $f(wz) - f(w^{-1}z) = (w - w^{-1})(a_+z - a_-z^{-1})$ and $f' = \iota(a_+z - a_-z^{-1})$, so any c with $c(w - w^{-1}) = \iota$ makes the shift w an exact rule. At the grid $\theta \in (\pi/8)\mathbb{Z}$ every nonzero shift $j\pi/8$, $j = 1, \ldots, 7$, is exact, with weight $c_j = 1/(2\sin(j\pi/8)) \in \frac{1}{2}\mathbb{Z}[\zeta_{16}]$; the half turn has no weight. The hypotheses* not *used are unitarity, $P^2 = I$, normalisation of ψ and Hermiticity of O.*
Checked: `shiftDiff, shiftAny, gridShiftLaurent, weights, noHalfTurnWeight` in `MTB.Core.GridShift`.

One halving suffices for every weight because $1 - \zeta^{2j}$ divides 2 in $\mathbb{Z}[\zeta_{16}]$, where $(2) = (1 - \zeta_{16})^8$ is totally ramified. The localisation that carries $\frac{1}{2}$ and $1/\sqrt{2}$ therefore inverts every grid denominator. The rule is a statement about the frequency support $\{-1, 0, 1\}$, not about physics. It has a converse: two evaluations of the two-term rule, at $z = 1$ and at a second grid point, *decide* whether f has frequency-two content (`ruleNeedsFrequencyOne`, `ruleFromFrequencyOne`). The rotation is also proved to be a unitary one-parameter group on every register whose generator is an involution, with each hypothesis entering exactly where it must (`rotGroup`, `rotUnitary`). Our plan's mutant, "a $\pi/4$ shift fails", turned out to be half a theorem: the $\pi/4$ shift is exact with weight $1/\sqrt{2}$, and what fails is the weight $\frac{1}{2}$. Generators with more than two eigenvalues need one shift per frequency; that multi-term rule is not formalised here.

5.2 The Born head is a relative norm

$\mathbb{Q}(\zeta_8)$ has degree 2 over its real subfield $\mathbb{Q}(\sqrt{2})$, so $|z|^2 = z\,\sigma_7(z) = N_{\mathbb{Q}(\zeta_8)/\mathbb{Q}(\sqrt{2})}(z)$.

Theorem 15 (Born head). *$R_8 \cong \mathbb{Z}[\frac{1}{2}][\sqrt{2}][\iota]$ by maps built only from universal properties, and complex conjugation is the tower's bar. The Born weight $N(\kappa\psi_i) = \mathrm{Re}^2 + \mathrm{Im}^2$ lands in $\mathbb{Z}[\frac{1}{2}][\sqrt{2}]$ by construction. It agrees with the Born matrix of Section 2.3 and with the broadcast Born vector of Section 2.2. It is multiplicative on product states and totally positive (a sum of two squares in both real embeddings). The σ_7-fixed elements of R_8 are exactly the real subring.*
Checked: `bornIsAmpConj, realDescent, bornNormalised, bornTensor, bornTotallyPositive` in `MTB.Core.BornHead`.

Two refusals give this exactness its edge. First, there is no amplitude modulus $R_8 \to \mathbb{Z}[\frac{1}{2}][\sqrt{2}]$: $|1 + \zeta_8|^2 = 2 + \sqrt{2}$ is not a square, because it maps to the non-residue 5 under a checked ring map to $\mathbb{F}_7$ (`noAmplitudeModulus`). Second, a finite fair-coin sampler (a Knuth–Yao tree [33]) has a dyadic, Galois-fixed law. The one-T-gate state $HTH|0\rangle$ has weights $\frac{1}{2} \pm \frac{\sqrt{2}}{4}$, one Galois orbit, so *no* finite fair-coin tree samples it exactly (`noFiniteSampler`, and `noFiniteSamplerGalois` for a second proof). Conversely, every vector of naturals summing to 2^k is the law of an explicit k-coin map (`knuthYao`, `MTB.Core.KnuthYao`). Symmetric statistics of an orbit are rational: $h(HTH|0\rangle) = \frac{1}{4}$ exactly (Fig. 3).

Exact sampling of a Clifford+T Born head needs unboundedly many coins, so acceptance must be statistical.

5.3 A complete Clifford decision

Theorem 16 (routing). *Over any commutative ring, a Clifford witness for U is a record giving, for every Pauli word P, a word Q and a scalar c with $UP = cQU$. This is exactly the Gottesman–Knill tableau [1, 26]. Witnesses compose and tensor, and extend freely from generators to circuits. A refuter is ring-free: if $UP = cQU$, every 2×2 minor of the pair vanishes, so one nonzero minor refutes every scalar at once. Over a domain with ι, the scalar must lie in $\{\pm 1, \pm\iota\}$, and the semantic decision `cliffDec` is complete as well as sound.*

Checked: `routeSound, minorVanishes, cliffDec, check-complete` *in* `MTB.Core.Route`.

Two consequences answer physics questions exactly:

- **Ising.** If $G^2 = -1$ and $2st = 1$, then $s(1 - G)$ is Clifford. A Majorana exchange is of this form, so *every* Ising braid word on four Majoranas routes classical with a composite witness (`isingBraidsClassical`). This is the folklore that Ising braiding is Clifford-only [9, 41], now as a theorem about all words. The noun crossing of Section 4.5 lowers to the word $b_{23}b_{12}b_{34}b_{23}$ (`crossingIsKoszul`).
- **Fibonacci.** Every Fibonacci letter is refuted (`fibBraidsQuantum`), consistent with universality [22].

The syntactic route, which reads only the gate multiset, is sound and provably incomplete: $T;T$ routes "quantum" yet equals S (`routeIncomplete`). The semantic decision resolves it, deciding $T;T$ classical and refuting $H;T$ as a whole circuit. "Quantum" therefore means "a generator is refuted", never "the circuit is not Clifford".

5.4 Acceptance over the unique order on $\mathbb{Z}[\sqrt{p}]$

Comparing a device's empirical distribution with an exact Born law in $\mathbb{Z}[\frac{1}{2}][\sqrt{2}]$ needs an order on that ring, and a bound on $\ln(2/\alpha)$.

Theorem 17 (the order, and the acceptance rule). *On $\mathbb{Z}[\sqrt{p}]$, p prime, the sign of $a + b\sqrt{p}$ is $\operatorname{sgn}(a)\operatorname{sgn}(a^2 - pb^2)$ in the mixed case, and any positive cone containing $\sqrt{p}$ equals this one. An ordering is therefore the choice of a real embedding, and the order inherited from $\mathbb{R}$ is the decided sign. A DKW level carries a decided certificate $\ln(2/\alpha) \le L$ via an exact Taylor sum (for $\alpha = \frac{1}{20}$, $L = \frac{369}{100}$, first certified at $K = 11$). The rule accepts iff every threshold j satisfies $2L_d(C_j - nF_j)^2 \le nL_n$. It is sound and complete with respect to the stated rule, a rejection names its threshold, and the verdict depends only on the histogram.*

Checked: `coneUnique, signIsCone, acceptRuleSound, acceptRuleComplete, rejectRuleWitness` *in* `MTB.Core.QuadOrder, MTB.Core.DKW`.

The probes at $n = 10^4$ are kernel decisions (Fig. 7). A faithful T-state sample $(8536, 1464)$ is accepted. Samples drawn from the 4-coin law $\frac{7}{8}$ and the 2-coin law $\frac{3}{4}$ are rejected. The Galois conjugate swaps the T-state's weights, and the same faithful sample is rejected against the conjugate law: the Born head is Galois-*equivariant*, but acceptance is not Galois-*invariant*, because a sample is a fact in $\mathbb{R}$. For Fibonacci the vacuum weight φ^{-2} is accepted, and *no* histogram of size 10^4 is accepted against Yang–Lee (`yangLeeRejectsAll`, all 10^4+1 decided). A non-unitary conjugate is not a law, and statistics rejects it outright. The DKW inequality [18, 38] is cited in general but *proved* at one instance: at $n = 100$ the exact rejection probability is a finite sum in the law's own ring, and the

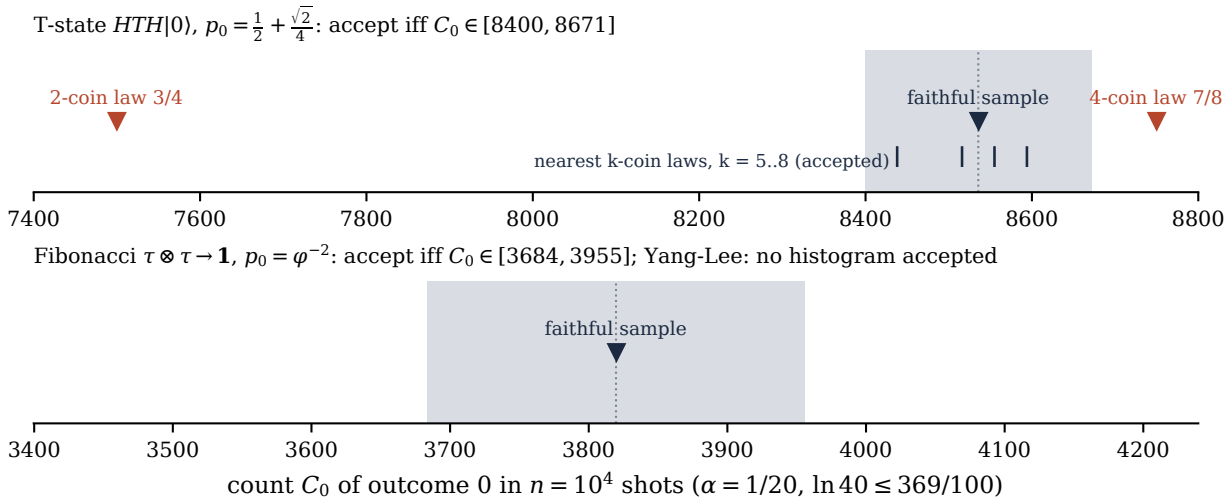


Figure 7. Acceptance windows, decided in advance. At $n = 10^4$ shots and $\alpha = \frac{1}{20}$, the T-state is accepted exactly for $C_0 \in [8400, 8671]$, and both edges are kernel-decided. The 2- and 4-coin dyadic approximations are rejected and those with $k \ge 5$ coins accepted, so at this n the test certifies a sampler to about five bits. The Fibonacci vacuum channel is accepted on $[3684, 3955]$, and no histogram at all is accepted against its Yang–Lee conjugate. The windows were recomputed in this session with 60-digit arithmetic (`checks/dkw.json`); the kernel decisions are B6b's.

kernel decides $10^{-4} \le \text{level} \le 10^{-3}$, hence $20 \cdot \text{level} \le 1$ (`level100`, `MTB.Core.DKWExact`). Twenty runs of an exact integer simulator were all accepted, as pre-registered (reported). **No quantum processor has been run**; the hardware row `B6b_hw` is the project's only open item.

5.5 From proof to machine

An executable that is the mathematics. The Agda deciders compile through MAlonzo to a Haskell executable, `mtb-run`, with argument parsing and JSON in Agda and only a thin Haskell shell around them. Across all 1,440 configurations of the corpus, theories and rings it agreed with independent Python oracles (0 mismatches). An end-to-end acceptance run drew its cases from the NIST randomness beacon [40] after a commitment. It then exhausted the whole choice space: all 1,440 configurations were byte-identical with the oracles removed, and the sweep reproduced the accepted digest `85cf37d7...` (X2, X3).

Copying on an interaction-net reducer. Our Boundary reducer evaluates Lafont's interaction combinators [35]. Its duplicator δ commutes through what it meets, so the copy it performs is *cloning*. By LB7, cloning equals the Frobenius reading exactly when the copied value is classical (`copySafeSound`). Net typing is therefore the condition under which the reducer's copy step means what the grammar means. It has four conjuncts: linearity, sorts, classical kind behind every δ, and single use of each spoken word. Three of them are invisible at the term level.

Theorem 18 (net typing and the reducer's copy).

(a) For every term, acceptance by the grammar's checker implies that its lowered net is net-typed (`acceptedLowers`).

(b) For every closed γ/ε tree t and fuel $\ge \text{size}(t)$, δ against t normalises in exactly $\text{size}(t)$ steps to two exact copies, on the reducer's own replayed schedule (`copyAll`).

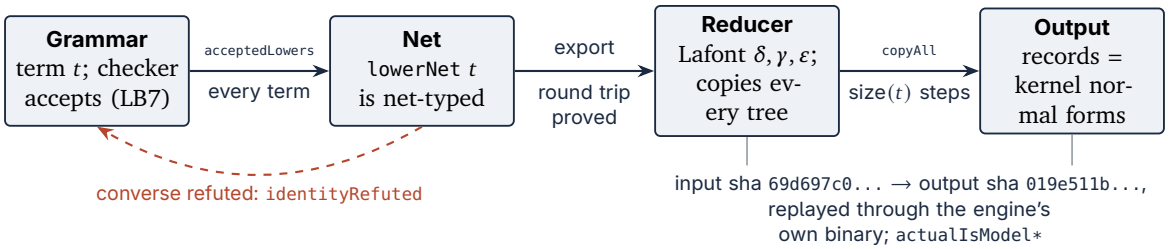


Figure 8. From proof to machine. The grammar's checker (LB7) implies net typing for every term; the reducer copies every closed tree exactly; its actual output on the pre-registered batch is the kernel's normal form (hashes reproduced, records generated into Agda and proved equal). Dashed: the refuted converse (net-typed $\nRightarrow$ accepted).

(c) δ against $\delta(\varepsilon, \varepsilon)$ annihilates instead of copying. Data built from the copier are consumed, not copied.

(d) The reducer's actual output on the pre-registered batch is proved equal to the model's normal forms (`actualIsModel*`).

Checked: `acceptedLowers`, `copyAll`, `deltaNotCopied`, `actualIsModel*` in `MTB.Lane.K.NetTypingGeneral`, `MTB.Core.LafontCopyAll`, `MTB.Core.LafontCopy`, `MTB.Core.LafontActual`.

Refutation 5 (net typing is not the term checker). *The identity "net-typed $\iff$ accepted" is false. A copy used twice at grade 1 and one used once at grade 0 are net-typed but rejected, and a discard is accepted (`identityRefuted`, `gradeBlindG1`, `gradeBlindG0`, `discardAccepted`). The theorem is the implication.*

Part (d) took the most work. The production engine returns only an output hash. We reproduced its input hash `69d697c0...` from the pre-registration, replayed the batch through the engine's own binary to the same output hash `019e511b...`, and generated the replayed records into Agda, where the kernel proves them equal to the model. The reducer has no rule for δ against an opaque generator. A lexical state can be cloned on it only as a closed tree, and a superposition has no tree form. The machine therefore enforces LB1 by construction: it cannot clone what the mathematics says must not be cloned (Fig. 8).

6 Consequences, field by field

Each subsection below can be read on its own. Each follows one pattern:

- the problem as its field states it;
- the lever this work provides, with the result it rests on;
- what can be done now;
- the next investigation, as a **Proposal** with its discriminator.

Everything in this section beyond the cited theorems is prospective.

6.1 Quantum language models

The problem. Proposals for quantum language models fall into four families:

- quantum transformers, such as Quixer's linear combinations of unitaries [30];

- quantum adapters attached to frozen classical models;
- compositional text circuits, such as DisCoCirc and QDisCoCirc [11, 34, 52];
- quantum-inspired tensor networks.

They meet the same walls: classical simulability of the trainable regime, the cost of loading data, the cost of readout in shots, and the absence of unitary nonlinearities. Dequantization results [48] show that apparent speedups often have classical counterparts.

The lever. A transformer is a copying machine, and each place it copies is a δ:

- the residual stream;
- attention fan-out to every query;
- multi-head fan-out;
- the key–value cache;
- tied embeddings.

LB1 (Theorem 1) says such a copy is lawful exactly on classical data. A coherent quantum residual stream is therefore not just hard to build; it is cloning. Theorem 3 prices the repair. Theorem 4 says tied weights must be sharp, and Theorem 18 gives a checker that finds every δ in a graph and refuses the unlawful ones. One loop is copy-safe for free: autoregression. The sampled token is a measurement result, so feeding it back is a classical copy. The architecture that falls out has four parts:

- quantum blocks are linear regions between measurement points;
- a Born head produces the next-token distribution (Theorem 15), whose $|\cdot|^2$ is the model's nonlinearity at the one place a language model needs it most;
- the copy-heavy parts (embedding, residual, cache) stay classical, by theorem;
- exact shifts train it (Theorem 14), a Clifford decision routes it (Theorem 16), and acceptance checks its hardware samples (Theorem 17).

The quantum budget of such an architecture is a number with a proof: the sum of the logical entropies at its measurement points, plus the shots per accepted output.

Where an advantage could live. Contextuality is the proved source of an unconditional memory separation between quantum and classical sequence models [6], and of polynomial expressivity separations in trainable models [5]. Our exact obstruction and census (§4) make contextuality a property of *data* that can be certified before any hardware is built. The census also says what the cheap detector misses (87.5% of logically contextual (2,2,2) sections).

Proposal 1 (copy-typed hybrid language models). A copy-typing pass over published quantum language-model architectures will find that every coherent fan-out either pays for an unlawful copy with a success probability or measurement, or is simulable. The remaining quantum content of each architecture is bounded by its measurement prices plus its non-Clifford layers.

Discriminator: a published architecture with a coherent copy of non-classical data and no compensating measurement, success probability or classical simulation. That would contradict Theorem 1 applied to its lowering.

Proposal 2 (gate quantum sequence models on data contextuality). A quantum sequence model can beat classical memory only on tasks whose data carry a nonzero contextuality certificate. Certify the data first, then build.
Discriminator: a pre-registered separation on a task whose certificate is zero under the full (linear) test, not only the $\mathbb{Z}$-class.

Proposal 3 (replicate quantum adapters against classical orthogonal fine-tuning). Reported gains of small unitary adapters on frozen classical models come from the orthogonal parametrisation, not from quantumness. Decide each adapter circuit's simulability with Theorem 16, emulate it exactly, and compare with classical orthogonal fine-tuning at equal parameter count.
Discriminator: the quantum adapter beats the classical orthogonal baseline beyond the pre-registered interval *and* is not classically simulable.

6.2 Quantum hardware and fault-tolerant toolchains

The problem. Topological-order experiments report F- and R-data, quantum dimensions and braiding phases, but many such quantities depend on a gauge choice. Benchmark-based claims of advantage are hard to verify classically, and recent proposals plant secret structure, for example by Clifford obfuscation [55]. Resource estimates for fault-tolerant algorithms drive roadmaps but come from unverified pipelines.

The lever. The pentagon locus is one gauge orbit (§3), which fixes exactly which quantities are observable. The braid relation alone cannot identify Fibonacci (30 solutions against 2). Every Ising braid is Clifford (Theorem 16). The acceptance rule decides consistency at a stated α with the shot budget known in advance (Theorem 17).

Use now. The acceptance windows of Fig. 7 are pre-computable for any exact Born law in $\mathbb{Z}[\frac{1}{2}][\sqrt{p}]$. A team can state the decision rule before collecting data, which is the hardware counterpart of a pre-registration.

Proposal 4 (a certified table of gauge-invariant targets). For Ising, Fibonacci and the toric code, a finite table of exact gauge-invariant quantities, each with its DKW shot budget, suffices to accept or reject an experimental claim of each topological order.
Discriminator: an entry that changes under a gauge transformation, or a published experiment whose claim the table cannot decide.

Proposal 5 (verify the verifier). The planted-structure equations of a Clifford-obfuscation advantage protocol can be checked exactly in the kernel. Trust then rests on the protocol's hardness assumption alone, which we cannot certify.
Discriminator: an obfuscated instance whose exact verification statistic disagrees with the protocol's claimed value.

6.3 Accelerators and exact arithmetic

The problem. Large-model inference is numerically nondeterministic: accuracy and output length vary with batch size, kernel and precision, because floating-point addition is not associative [56]. One recent remedy restricts INT8 quantisation scales to powers of two and obtains byte-identical generation across GPU kernels [10].

The lever. Arithmetic with power-of-two scales is arithmetic in $\mathbb{Z}[\frac{1}{2}]$. We proved that $\mathbb{Z}[\frac{1}{2}][x]/(f)$ is a commutative ring with canonical forms for every monic f (§3). In exact dyadic arithmetic, reduction order provably cannot change a result. A hardware lane must still bound integer widths. Our executable is byte-identical across all 1,440 configurations with the oracles removed (§5.5).

Proposal 6 (a certified deterministic GEMM lane). A dyadic GEMM with proved width bounds produces the same bytes under every reduction order and on every device, at a measurable overhead relative to BF16.
Discriminator: two reduction orders, or two devices, giving different bytes on the same input within the proved width bound.

6.4 Topological physics and generalised symmetry

The problem. Non-invertible symmetries are described by fusion categories, and the symmetry TFT of a category C is Turaev–Viro on C, equivalently Reshetikhin–Turaev on its Drinfeld centre [21]. The data (F-symbols, centres, condensable algebras) are usually computed by hand or numerically, and depend on a gauge. Separately, claimed signatures of non-Abelian anyons have non-topological explanations, and the field asks for outcomes that only non-Abelian physics can produce.

The lever. §3 gives exact pentagon and hexagon laws over every golden ring, the single gauge orbit, Galois-decided unitarity, kernel state sums with Pachner invariance, and a Legendre law for every lens space. In an Abelian theory every fusion space is one-dimensional, so braids act by commuting phases.

Proposal 7 (certified SymTFT data). The pentagon decider, gauge-orbit classification and Galois unitarity test extend to every multiplicity-free fusion category of small rank. With the Drinfeld centre and its Lagrangian algebras, they yield certified symmetry-TFT data for Ising, Fibonacci, $\mathrm{Rep}(S_3)$ and $\mathrm{Vec}_{\mathbb{Z}_2}$.
Discriminator: a published unitary category that the decider rejects, or a computed gapped-boundary classification that disagrees with the known one.

Proposal 8 (an Abelian-imposter refuter). For a braiding dataset, either the data admit an explicit Abelian model, or a DKW certificate shows that no commuting model reproduces them at α.
Discriminator: a planted Abelian-generated control that the refuter fails to accept; this is run first.

Contextuality is also the resource behind magic-state computation [27]. One cohomology class thus has three readings: linguistic ambiguity, a sequence-model memory separation and magic. Whether those readings coincide quantitatively is open (Section 9).

6.5 Complex systems, semantics and homotopy type theory

The problem. Our Gluing Engine [25] decides whether the parts of a presented system make a whole. It asks four questions:

- congruence: are these the same thing?
- confluence: does it matter who acts first?
- factorisation: what survives noise?
- synchronisation: can local views glue?

It answers them syntactically, with low-dimensional homotopy invariants. It has no account of what the parts mean.

The lever. Meaning is a functor out of a grammar, and a presented system is a grammar. This work supplies a semantics functor into $\mathrm{Mat}(R)$ over exact rings, with three things the engine lacked:

- a copy law (what can be shared);
- a price (what sharing an ambiguous meaning costs);
- a third confluence verdict, "confluent up to a character" (Theorem 8), which is neither commuting nor racing.

Communication between parties is copying a meaning. By LB1 a meaning copies faithfully exactly when it is classical, so a party must commit to a reading before broadcasting an ambiguous one, and it pays $1 - \sum p^2$.

Readings for homotopy type theory (interpretation).

- **Galois loops.** Ring automorphisms are loops in a universe of rings, and every endomorphism of $D[\omega]$ is Galois (Theorem 5). The theorem that replaced our refuted conjecture is literally about transport. Since the order on $\mathbb{Z}[\sqrt{p}]$ is unique (Theorem 17), positivity is the choice of a real embedding, and a Galois loop moves it. Non-unitarity then reads as the holonomy of positivity around a Galois loop.
- **Linear modalities.** Riley's linear homotopy type theory [45] and Sati and Schreiber's quantum monadology [46] separate linear quantum types from classical ones by a modality, with dynamic lifting of measurement results. LB1, the measurement point and its price, and the weight-tying law are concrete semantic witnesses for that split.

Proposal 9 (a validated classical modality). A graded linear univalent fragment whose classical modality ω is interpreted by the copy law satisfies: every ω-typed value copies, and every copiable value over a connected ring is ω-typed.
Discriminator: a model of the fragment in which ω-typed data fail to copy; over $\mathbb{Z}/6$ the converse must fail (Refutation 1).

Proposal 10 (confluent up to a phase). For braided presented systems, confluence has a decidable third verdict whose certificate is a character of the braid group.
Discriminator: a braided system the decider calls confluent up to a phase whose readings differ by more than a scalar.

6.6 Control and regulated engineering

The problem. Natural-language requirements are a known source of defects in safety-critical systems. Current automated auditors detect ambiguity through stochastic variation in a language model's formalisations. Separately, the tools that verify learned components must themselves be qualified before their results count as certification evidence.

The lever. In the controlled fragment, ambiguity is proved: S09's two meanings are two theorems, and a separator names the failing equation. The refusal discipline used throughout this work has two parts:

- a small `--safe` kernel with an explicit trust footprint;

- mutants planted at the claim and refused at its declaration.

This is the shape of the evidence that tool qualification asks for. We keep verification (built right) and validation (right thing) apart. Every result here is verification of logic or mathematics, never a statement about the physical safety of a device.

Proposal 11 (certified requirement disambiguation). On controlled-language requirements, every ambiguity is returned as two proved readings, together with the attribute that separates them as the decided remedy.

Discriminator: a requirement the tool reports unambiguous for which two field engineers, working independently, write inequivalent formalisations within the controlled fragment.

Proposal 12 (a qualification dossier for one checker). One of the checkers used here can be packaged as tool-qualification evidence (a small kernel, its trust footprint and its planted rejections) that a certification reviewer accepts as complete.

Discriminator: a reviewer's list of missing evidence items that the method cannot supply.

7 Positioning

Table 2 compares this work with the nearest lines of research on five axes: exact arithmetic, a machine-checked semantics, kernel refusals (counter-models and planted mutants), runs on quantum hardware, and scale. The rows where we lose are deliberate. We have no hardware result, our corpus is tiny, and our anyon theories are small. What we offer in exchange is that every positive and negative statement in our row is a checked theorem about an exact object.

Line of work	Exact	Checked	Refusals	Hardware	Scale
DisCoCat / lambeq [12, 29, 37]	no	no	n/r	yes	small sentence datasets; toolkit for arbitrary text
DisCoCirc / QDisCoCirc [11, 17, 34, 52]	no	no	n/r	yes [17]	text circuits; complexity results
Quixer [30]	no	no	n/r	simulation	a standard language-modelling benchmark
VyZX [36]	symbolic	yes (Coq)	n/r	no	general ZX diagrams
Categorical deep learning [14, 24]	no	no	n/r	n/a	architectures in general
LHoTT / quantum monadology [45, 46]	n/a	type theory	n/a	no	a programming-language design
This work	**yes**	**yes (Agda)**	**yes**	**pending**	24 sentences; Ising, Fibonacci, toric, semion

Table 2. Positioning. "Exact": amplitudes and verdicts computed without floating point. "Checked": the semantics is formalised in a proof assistant. "Refusals": wrong statements are rejected by kernel counter-models or planted mutants; "n/r" means not reported in the cited work. "Hardware": results on a quantum processor. Cells reflect our reading of the cited papers as of 3 November 2026.

Relative to the compositional QNLP line [12, 34, 37] we change the question. They ask how to compute meanings on quantum hardware; we ask what a meaning must be for a given operation to be lawful, and prove the answer. Relative to verified ZX [36], our scope is narrower in diagrams but reaches braided semantics, anyon theories, statistics and an executing reducer. Relative to categorical deep learning [14, 24], we supply the quantum hybrid case of Para, in which the copy law decides which parameters can be shared. Relative to linear HoTT [45, 46], we provide semantic

evidence for the classical modality rather than a new type theory. Novelty claims are dated: to our knowledge, after searches on 12 October and 3 November 2026, no prior work machine-checks a copy law over rings, a symmetric-fibre obstruction by Bezout certificates, or a Legendre law for Fibonacci lens-space invariants.

8 Limitations

- **Controlled corpus.** The semantics is proved on 24 controlled sentences and finite presentations. The full-grammar meaning square is relative to a Boolean model of the lexicon (e.g. *saw* is extensional). Nothing here is evaluated on open-domain text.
- **Small theories.** Our anyon theories are Ising, Fibonacci, the toric code, the double semion and their Galois conjugates. General fusion categories, multiplicities and 3+1-dimensional theories are not covered.
- **Cited, not proved.** $\mathrm{TV}(C) = \mathrm{RT}(Z(C))$ in general (proved on five manifolds, and for every lens space on the surgery side); the DKW inequality in general (proved at one instance, $n = 100$); $e^x \geq T_K(x)$; the existence of an ordering of $\mathbb{R}$.
- **No hardware.** No quantum processor has been run (`B6b_hw` is open). Every acceptance statement is about the rule and about simulated samples.
- **One experiment, Ising only.** The null result concerns one task family and the Ising theory on cabled nouns. Its v2 pre-registration was frozen after v1's results were known; the reason is disclosed in §4.6.
- **Independence.** The end-to-end acceptance run (X3), the semantic audit (X4) and the hostile audit (HA) were carried out by the implementing session under explicit overrides by the principal researcher. Exhaustion of the choice space and a public-beacon draw replace reviewer choice for X3, but no independent human reviewer has audited this work.
- **Trust base.** The verification of record is the Agda kernel (2.8.0, `--safe`, no postulates), run per phase and re-run for the audits on a network node with a pinned Agda. The repository's server-side acceptance gate has no Agda lane, so it does not typecheck the Agda itself. The Lean bridge modules use only `propext` and `Quot.sound`. The executable's trusted base includes GHC and the MAlonzo backend. The reducer comparison trusts the engine binary that produced the replay.
- **Reported rows.** Several rows are reported observables rather than kernel facts: the exact-engine runs, the experiment, the simulator verdicts, and the text-level bridges to earlier libraries. We label them as such wherever they appear.
- **Oracles.** The independent oracles use separate code paths and representations, but they were written within the same project.
- **Not formalised.** Multi-frequency parameter-shift rules; two-qubit semantic Clifford decisions evaluated in the kernel; a semantic Clifford route for Fibonacci, whose ring contains no ι.

9 Open problems

Each problem below is newly tractable because of a result above. For each we name a first move and a discriminator.

O1 Acceptance on a device. *Question:* does a real processor's output for an exact-shift training step and a Clifford+T readout fall inside the pre-computed DKW window? *Opening:* the windows, the gradient weights and the Clifford routing are all exact and decided in advance (Theorems 14, 16 and 17). *First move:* pre-register one two-qubit circuit with its window, run it, and decide. *Discriminator:* accept or reject at the stated α; a rejection names its threshold. *Payoff:* the project's only open row, and a template for pre-registered hardware claims.

O2 Does any open-frame character help a learner? *Question:* Theorem 8 leaves Fibonacci open frames with a period-five character. Can interference between ambiguous readings of different writhe carry information a symmetric model cannot? *First move:* an exact-engine experiment in which readings with different writhe are held in superposition and resolved by interference, pre-registered against a symmetric control. *Discriminator:* a pre-registered interval excluding zero. *Payoff:* either a first positive role for braiding in meaning, or a sharper negative result.

O3 Certified fusion categories. *Question:* do the pentagon decider, gauge-orbit classification and Galois unitarity test extend to every multiplicity-free fusion category of small rank, and from there to Drinfeld centres and Lagrangian algebras? *Discriminator:* agreement with published tables and gapped-boundary classifications. *Payoff:* certified symmetry-TFT data.

O4 One resource account. *Question:* do the linguistic contextuality class, the sequence-model memory separation and magic-state contextuality [27] coincide quantitatively? Are they related to the measurement price $1 - \sum p^2$? *First move:* compute all three on the (2,2,2) census. *Discriminator:* an empirical model on which they disagree where the theory says they coincide.

O5 Galois holonomy. *Question:* is non-unitarity the holonomy of positivity around a Galois loop, as a theorem in the cubical lane where univalence computes? *Discriminator:* a non-unitary conjugate with trivial holonomy would refute the reading.

O6 No-cloning as a constructor task. *Question:* in a formal constructor-theoretic substrate, does the cloning task reduce to the copy equation, so that LB1 classifies information media? *Discriminator:* a substrate in which a non-classical attribute set admits a cloner.

The most promising next move is O1, because it closes the loop the paper describes and turns every acceptance statement from a rule into an observation. O2 is the most consequential scientific question: our negative result moved the burden of proof, and O2 is the experiment that can meet it.

10 Synthesis: what copying teaches

The paper began with a pronoun that could not be copied and ends with a classification of what can be. Three kinds of content emerge, and the distinction holds in each setting we studied:

- content that can be shared verbatim (classical points; trees on the reducer; sharp weights);
- content that can be shared only after a commitment, at a price of exactly $1 - \sum p^2$ (measured nouns; disambiguated readings);
- content that cannot be shared at all (superpositions; data with no tree form).

Braiding adds less than we hoped: a character on open frames, nothing on closed meanings, and no learning advantage in the one place we measured. Exact rings add more than we expected. They decide which theories are physical, which quantities are observable and what an acceptance test can tell.

The same three kinds describe collaboration among people and SI agents. A claim that is classical, in the sense of having one reading, can be passed along verbatim. An ambiguous claim must be committed before it is broadcast, and someone pays the price of committing, which is a quantity, not a metaphor. Some content, such as an unresolved judgement, cannot be copied without changing it. The work also suggests what kind of trust is worth having in a machine's output. Trust here does not come from fluency. It comes from statements bound to declarations and from refusals that name the failing wire. Whether that discipline scales beyond a controlled corpus is the open question this paper leaves to its readers. We suspect the most useful answers will come from people who will not accept a proof until they have seen what it refuses.

11 Reproducibility and artifact statement

We describe the artifacts against the three criteria of the ACM artifact badges [7]: *available*, *functional* and *reusable*. No artifact-evaluation committee has reviewed them, and we claim no badge.

Available. The formalisation comprises:

- 205 Agda modules under `agda/MTB`;
- the project's scripts, oracles, gates, receipts and findings;
- the claim-graph record.

All of it lives in the Institute's authoritative proof store as project `meaning_that_braids_20261012`. It landed at store generations 5539, 5540, 5544 and 5545, and the hostile-audit remediation at generation 5562. Access is internal at the time of writing; a public release has not been made. The folder that accompanies this paper is self-contained for everything the paper itself computes. It holds:

- the figure and table scripts (`scripts/`);
- byte-for-byte copies of the data they read (`data/`);
- copies of every phase's findings, the summary report and the audit record (`evidence/`);
- the outputs of the checks run for this paper (`checks/`).

Functional. The verification of record runs as follows:

- **Lane aggregators.** `MTB.All`, `MTB.Lane.K.All`, `MTB.Lane.C.All` and `MTB.Lane.G.All` are checked with Agda 2.8.0 under each lane's flags. Each returned rc 0 on the audited bytes, with peak memory 0.87, 1.23 and 1.22 GiB for the first three (Appendix D).
- **Phase gates and the executable.** Each phase gate re-checks its modules, its pins, its planted mutants and its oracle. Across the project, 723 mutants were planted and all 723 refused; 160 controls were planted and all accepted. The executable `mtb-run` is built through MAlonzo, and its oracle sweep covers 1,440 configurations.
- **Claim replay.** The project claim's kernel receipt is replayed by `scripts/mtb/claim_replay.py`: about 640 s warm, and 78 minutes at 2.72 GB peak cold.

For this paper, two scripts reproduce what the paper computes:

- `scripts/make_figures.py` redraws every data figure and re-derives six results independently: the copy defects, the prices, the pre-registered contrast, the Galois images, the 489 lens-space values against the Legendre law, and the DKW windows. All six agree with the kernel records (`checks/summary.json`).
- `scripts/check_decls.py` binds each of the 142 declaration names in this paper to the Agda module that defines it, and each *Checked* line to the modules it names (0 problems; `checks/decl_binding.json`).

Reusable. The modules form layers that can be used independently:

- **exact rings**: `Prelude`, `Ring`, `Solver`, `SolverCanonical`, `Galois`, `Domain`;
- **copy law**: `LinearFox`, `CopyCharacter`, `ClassicalCore`, `Separator`, `Broadcast`;
- **anyon theories**: `TorusBracket`, `GaloisBracket`, `FibPentagon`, `FibHexagon`, `FibLens`, `SymmetricFibre`;
- **the hybrid loop**: `ParaHybrid`, `GridShift`, `BornHead`, `Route`, `QuadOrder`, `DKWRule`;
- **the machine**: `NetTyping`, `Lafont*`;
- **lanes**: K (QNLP, `--without-K`), C (cubical) and G (golden foam).

Every generic module is parametrised by a commutative ring, or by a ring with the named extra structure. A new theory enters by instantiating, not by editing. The import graph and the pinned definitions are recorded, so drift in a reused definition is detected.

Trust footprint.

- **Agda:** `--safe`, no postulates and no escape pragmas, in every lane.
- **Lean bridge modules:** `propext` and `Quot.sound`.
- **The executable:** additionally trusts GHC and MAlonzo.
- **Oracles:** Python, independent of the Agda's route.
- **The reducer comparison:** trusts the engine binary that produced the replay.

What this authoring session re-checked. In preparing this paper we:

- re-derived the six figure results named above;
- confirmed that both pre-registration files hash to their recorded sha256 values;
- bound every declaration name;
- copied every data file byte for byte.

We did *not* re-run the Agda kernel, the executable or the experiment. Those results are reported from their receipts (Appendix A).

Disclosure of super intelligence (SI) assistance

This disclosure follows the 2026 policy of the American Physical Society [4], which asks for the tool's name and version, the nature of its help, how it was directed and how its output was verified. We use "SI" (super intelligence) where other venues write "AI".

- **Tools and versions.** Anthropic's Claude models, operated through the Claude Code command-line agent. This manuscript, its figure scripts and its checks were produced by Claude Opus 5.5 on 3 November 2026. The formalisation, oracles, gates and audits were produced from 12 October to 3 November 2026 in the same agent session lineage (recorded in the project artifacts). Agda checks for the audits ran on a node of the Institute's compute network. The model versions that served each phase are recorded in the session logs, not restated here.
- **Nature of assistance.** The SI wrote the formal statements, proofs, deciders, oracles, gate scripts, figure scripts and this text. It also ran literature searches.
- **How it was directed.** The principal researcher set the research question and the programme, approved the plan and each phase gate, and made every scope decision, including the overrides of independence for X3 and X4. Compute use is under the principal researcher's control: no GPU was used without approval, and no quantum processor was used.
- **How output was verified.**
  - every formal claim by the Agda kernel;
  - every load-bearing statement by mutants planted at it and refused at its declaration;
  - every executable by independent oracles;
  - the project by two audits that raised 31 findings, all closed;
  - this manuscript by scripts that recompute its numbers and bind its declaration names.
- **Figures.** All figures are plots or diagrams drawn by published scripts or TikZ sources from data or from exact computation. None was generated or modified by a generative image model.
- **Authorship.** The SI is not an author. The human author is responsible for the content.

Acknowledgements

This preprint names no external contributor. Contributors of equipment, compute, data or time to the next phase of the programme will be acknowledged here, by name and with their consent.

References

[1] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, 70:052328, 2004.

[2] Samson Abramsky and Adam Brandenburger. The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13:113036, 2011.

[3] Samson Abramsky, Shane Mansfield, and Rui Soares Barbosa. The cohomology of non-locality and contextuality. In *Proceedings of Quantum Physics and Logic 2011*, volume 95 of *EPTCS*, pages 1–14, 2012. arXiv:1111.3620.

[4] American Physical Society. American physical society releases updated AI policy for APS journals. https://www.aps.org/about/news/2026/06/releases-updated-ai-policy-journals, 2026.

[5] Eric R. Anschuetz and Xun Gao. Arbitrary polynomial separations in trainable quantum machine learning, 2024. arXiv:2402.08606.

[6] Eric R. Anschuetz, Hong-Ye Hu, Jin-Long Huang, and Xun Gao. Interpretable quantum advantage in neural sequence learning. *PRX Quantum*, 4:020338, 2023.

[7] Association for Computing Machinery. Artifact review and badging, version 1.1. https://www.acm.org/publications/policies/artifact-review-and-badging-current.

[8] Howard Barnum, Carlton M. Caves, Christopher A. Fuchs, Richard Jozsa, and Benjamin Schumacher. Noncommuting mixed states cannot be broadcast. *Physical Review Letters*, 76: 2818–2821, 1996.

[9] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Physical Review A*, 71:022316, 2005.

[10] Teng-Ruei Chen. Deterministic LLM inference across GPU kernels: Power-of-two INT8 quantization scales and the limits of tolerance-based conformance, 2026. arXiv:2609.00363.

[11] Bob Coecke. The mathematics of text structure, 2019. arXiv:1904.03478.

[12] Bob Coecke, Mehrnoosh Sadrzadeh, and Stephen Clark. Mathematical foundations for a compositional distributional model of meaning. *Linguistic Analysis*, 36:345–384, 2010. arXiv:1003.4394.

[13] Bob Coecke, Duško Pavlović, and Jamie Vicary. A new description of orthogonal bases. *Mathematical Structures in Computer Science*, 23(3):555–567, 2013. arXiv:0810.0812.

[14] Geoffrey S. H. Cruttwell, Bruno Gavranović, Neil Ghani, Paul Wilson, and Fabio Zanasi. Categorical foundations of gradient-based learning. In *Programming Languages and Systems (ESOP 2022)*, 2022. arXiv:2103.01931.

[15] Pierre Deligne. Catégories tensorielles. *Moscow Mathematical Journal*, 2(2):227–248, 2002.

[16] David Deutsch and Chiara Marletto. Constructor theory of information. *Proceedings of the Royal Society A*, 471:20140540, 2015.

[17] Tiffany Duneau, Saskia Bruhn, Gabriel Matos, Tuomas Laakkonen, Katerina Saiti, Anna Pearson, Konstantinos Meichanetzidis, and Bob Coecke. Scalable and interpretable quantum natural language processing: An implementation on trapped ions, 2024. arXiv:2409.08777.

[18] Aryeh Dvoretzky, Jack Kiefer, and Jacob Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *Annals of Mathematical Statistics*, 27(3):642–669, 1956.

[19] David Ellerman. Logical information theory: New logical foundations for information theory. *Logic Journal of the IGPL*, 25(5):806–835, 2017.

[20] Thomas Fox. Coalgebras and cartesian categories. *Communications in Algebra*, 4(7):665–667, 1976.

[21] Daniel S. Freed, Gregory W. Moore, and Constantin Teleman. Topological symmetry in quantum field theory, 2022. arXiv:2209.07471.

[22] Michael H. Freedman, Michael Larsen, and Zhenghan Wang. A modular functor which is universal for quantum computation. *Communications in Mathematical Physics*, 227:605–622, 2002.

[23] Michael H. Freedman, Jan Gukelberger, Matthew B. Hastings, Simon Trebst, Matthias Troyer, and Zhenghan Wang. Galois conjugates of topological phases. *Physical Review B*, 85:045414, 2012. arXiv:1106.3267.

[24] Bruno Gavranović, Paul Lessard, Andrew Dudzik, Tamara von Glehn, João G. M. Araújo, and Petar Veličković. Position: Categorical deep learning is an algebraic theory of all architectures. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024. arXiv:2402.15332.

[25] Richard Goodman and the IAOM research workforce. Where the parts make a whole: A certified instrument for collective sense-making. Institute for Applied Ontological Mathematics, preprint, 2026.

[26] Daniel Gottesman. The Heisenberg representation of quantum computers, 1998. arXiv:quant-ph/9807006.

[27] Mark Howard, Joel Wallman, Victor Veitch, and Joseph Emerson. Contextuality supplies the 'magic' for quantum computation, 2014. Nature 510, 351–355.

[28] Lisa C. Jeffrey. Chern–Simons–Witten invariants of lens spaces and torus bundles, and the semiclassical approximation. *Communications in Mathematical Physics*, 147:563–604, 1992.

[29] Dimitri Kartsaklis, Ian Fan, Richie Yeung, Anna Pearson, Robin Lorenz, Alexis Toumi, Giovanni de Felice, Konstantinos Meichanetzidis, Stephen Clark, and Bob Coecke. lambeq: An efficient high-level Python library for quantum NLP, 2021. arXiv:2110.04236.

[30] Nikhil Khatri, Gabriel Matos, Luuk Coopmans, and Stephen Clark. Quixer: A quantum transformer model, 2024. arXiv:2406.04305.

[31] Alexander Kirillov, Jr. and Benjamin Balsam. Turaev–Viro invariants as an extended TQFT, 2010. arXiv:1004.1533.

[32] Alexei Kitaev. Anyons in an exactly solved model and beyond. *Annals of Physics*, 321:2–111, 2006. arXiv:cond-mat/0506438.

[33] Donald E. Knuth and Andrew C. Yao. The complexity of nonuniform random number generation. In *Algorithms and Complexity: New Directions and Recent Results*, pages 357–428. Academic Press, 1976.

[34] Tuomas Laakkonen, Konstantinos Meichanetzidis, and Bob Coecke. Quantum algorithms for compositional text processing, 2024. QPL 2024; arXiv:2408.06061.

[35] Yves Lafont. Interaction combinators. *Information and Computation*, 137(1):69–101, 1997.

[36] Adrian Lehmann, Ben Caldwell, Bhakti Shah, and Robert Rand. VyZX: Formal verification of a graphical quantum language, 2023. arXiv:2311.11571.

[37] Robin Lorenz, Anna Pearson, Konstantinos Meichanetzidis, Dimitri Kartsaklis, and Bob Coecke. QNLP in practice: Running compositional models of meaning on a quantum computer. *Journal of Artificial Intelligence Research*, 76:1305–1342, 2023. arXiv:2102.12846.

[38] Pascal Massart. The tight constant in the Dvoretzky–Kiefer–Wolfowitz inequality. *Annals of Probability*, 18(3):1269–1283, 1990.

[39] Kosuke Mitarai, Makoto Negoro, Masahiro Kitagawa, and Keisuke Fujii. Quantum circuit learning. *Physical Review A*, 98:032309, 2018.

[40] National Institute of Standards and Technology. NIST randomness beacon. https://beacon.nist.gov. Accessed 2026-11-01.

[41] Chetan Nayak, Steven H. Simon, Ady Stern, Michael Freedman, and Sankar Das Sarma. Non-Abelian anyons and topological quantum computation. *Reviews of Modern Physics*, 80: 1083–1159, 2008.

[42] Ulf Norell. *Towards a Practical Programming Language Based on Dependent Type Theory*. PhD thesis, Chalmers University of Technology, 2007.

[43] Victor Ostrik. On symmetric fusion categories in positive characteristic. *Selecta Mathematica (New Series)*, 26, 2020. arXiv:1503.01492.

[44] Nicolai Reshetikhin and Vladimir G. Turaev. Invariants of 3-manifolds via link polynomials and quantum groups. *Inventiones Mathematicae*, 103:547–597, 1991.

[45] Mitchell Riley. *Linear Homotopy Type Theory*. PhD thesis, Wesleyan University, 2022.

[46] Hisham Sati and Urs Schreiber. The quantum monadology, 2023. arXiv:2310.15735; Quantum Studies: Mathematics and Foundations 12 (2025).

[47] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. Evaluating analytic gradients on quantum hardware. *Physical Review A*, 99:032331, 2019.

[48] Ewin Tang. A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st ACM Symposium on Theory of Computing (STOC)*, pages 217–228, 2019.

[49] The Agda Team. Agda, version 2.8.0. https://github.com/agda/agda, 2025.

[50] Vladimir G. Turaev and Oleg Y. Viro. State sum invariants of 3-manifolds and quantum 6j-symbols. *Topology*, 31(4):865–902, 1992.

[51] Daphne Wang, Mehrnoosh Sadrzadeh, Samson Abramsky, and Victor H. Cervantes. On the quantum-like contextuality of ambiguous phrases. In *Proceedings of the 2021 Workshop on Semantic Spaces at the Intersection of NLP, Physics, and Cognitive Science (SemSpace)*, 2021. arXiv:2107.14589.

[52] Vincent Wang-Maścianica, Jonathon Liu, and Bob Coecke. Distilling text into circuits, 2023. arXiv:2301.10595.

[53] David Wierichs, Josh Izaac, Cody Wang, and Cedric Yen-Yu Lin. General parameter-shift rules for quantum gradients. *Quantum*, 6:677, 2022.

[54] William K. Wootters and Wojciech H. Zurek. A single quantum cannot be cloned. *Nature*, 299: 802–803, 1982.

[55] Bin Yan et al. Classical verification of quantum advantage via Clifford obfuscation, 2026. arXiv:2608.15963.

[56] Jiayi Yuan et al. Understanding and mitigating numerical sources of nondeterminism in LLM inference, 2025. arXiv:2506.09501.

A Claim-to-evidence map

Each row binds a claim to the declarations that state it, the module that defines them, and the phase whose gate accepted it. The receipt column gives the first 12 hex digits of the sha256 of the phase's `receipt.json` (computed on the audited bytes for this paper) and the gate's passed/total checks. "HA" rows are authorised by a hostile-audit check record (Appendix D). The module column is verified by `scripts/check_decls.py`.

Claim	Declarations	Module	Phase / receipt
LB1 copy law; sharp form	`copyIffClassical`, `linearFox`	`MTB.Core.LinearFox`	B1b `b2a01dddd35e` 87/87
Counter-models to LB1	`z6RefutesLB1`, `zzRefutesLB1`, `reducedNotEnough`	`MTB.Core.LinearFox`	B1b
LB2 posetal Fox	`posetFox`	`MTB.Core.PosetFox`	B1a `73ee56242680` 78/78
Copy is a character; points	`copyIsCharacter`, `pointsOfSpec`, `normalisedSplit`	`MTB.Core.CopyCharacter`	B1b
No-cloning = copy law	`clonerIsDelta`, `clonedIffClassical`, `noUniversalCloner`	`MTB.Core.LinearFox`	B1b
QCT no-cloning re-derived	`noLinearClone′`	`MTB.Lane.C.NoCloneF2`	B1b
Broadcasting; price = logical entropy	`broadcastsIff`, `separatorMass≡ logicalEntropy`, `logicalEntropyNormalised`	`MTB.Core.Broadcast`	B2b `0954b90624ba` 61/61
Measurement points	`copySafeLowerAccepted`, `copySafeLowerExact`, `measureHarmless`, `invisibleIffAtMostOne`, `s23NeededButInvisible`, `semantic≤lowered≤static`, `phaseBlind`	`MTB.Lane.K.MeasurePoint`	B2b
LB7 copy-safe grammar	`copySafeSound`	`MTB.Lane.K.GradedUse`	B2a `bb3bf234aabe` 57/57
Price bridge; nesting refuted	`priceBridge`, `nestingRefuted`, `unrestrictedRefuted`	`MTB.Lane.K.PriceBridge`	HA
Para over a copower; tying associative	`tieAssoc, act-mult`	`MTB.Core.ParaHybrid`	B5b `ec8fb5ec97fb` 41/41

Claim	Declarations	Module	Phase / receipt
Weight-tying law	`tyingSound↔`, `tyingIffCore`, `tyingIffBasis`, `tyingNeedsConnected`, `measNatIff`	`MTB.Core.ParaHybrid`	B5b
Tied gradient	`tiedGradient`, `doubleAngleBreaks`	`MTB.Lane.K.ShiftSpec`	B5c `e78cdf221b5b` 48/48
Canonical forms	`canonicity`	`MTB.Core.SolverCanonical`	HA
Galois-fixed norms	`normFixed16, normFixed10`	`MTB.Core.NormFixed`	HA
LB3 Galois-equivariant bracket	`bracketGaloisEquivariant`	`MTB.Core.GaloisBracket`	B3a3 `ef64224db580` 64/64
Gauss law ≠ unitarity (refuted)	`gaussLawUnitarityRefuted`, `refutedExactly`, `gaussIffTransport`, `endoClassify`	`MTB.Core.GaussTransport`	B3a4 `f43c9ac9b839` gate 14/14
LB4 Fibonacci hexagons	`fibHexagons`	`MTB.Core.FibHexagon`	B3b2 `07561d014f14` 86/86
Braid and hexagon censuses	`braidCensus, hexCensus, censusMeetsClassification`	`MTB.Core.FibHexagonCensus`	HA
Unitary Fibonacci hexagons	`unitaryFibHexagons`	`MTB.Core.UnitaryFib`	B3b5 `62a87e0b94f9` 103/103
TV = $\|RT\|^2$ (Fibonacci, S^3)	`tvFib-S3`	`MTB.Lane.K.TVFib`	B3c2 `166ef3788cd8` 72/72
Legendre law, every lens space	`fibLegendre`	`MTB.Core.FibLens`	HA
Points are transparent	`pointsTransparent`, `closedRoutingBlind`	`MTB.Core.BraidedMat`	B4b `c334320765b1` 51/51
LB5 conservativity	`braidedConservative`	`MTB.Lane.K.BraidedMeaning`	B4b
Open-frame character	`wordLaw, factor, blindIffClassical-⟶`	`MTB.Core.KauffmanFrame`	B4c `d907d203a056` 50/50
Fibonacci open frames apart	`readingsApart`	`MTB.Lane.K.Monodromy`	B4c
LB6 no symmetric fibre	`noSymmetricFibre`	`MTB.Lane.C.NoSymmetricFibre`	B5a `ed52624e9108` 45/45
Fibonacci dichotomy	`fibonacciDichotomy`, `ramifiedGeneric`	`MTB.Core.SymmetricFibreTheories`	B5a
Ising cable = Koszul swap	`exchangeIsIsing`, `cableKoszul`	`MTB.Core.IsingCable`	B4e `fbf0de876fe1` 53/53
Contextuality obstruction sound	`obstructionSound`	`MTB.Core.Contextuality`	B4d `e0aab9665ae2` 49/49

Claim	Declarations	Module	Phase / receipt
(2,2,2) census	`coverage, counts, characterisation`	`MTB.Core. Contextuality222`	HA
Exact parameter shift	`shiftDiff, gridShiftLaurent, weights, noHalfTurnWeight`	`MTB.Core.GridShift`	B5c
Born head	`bornIsAmpConj, realDescent, bornTotallyPositive, noFiniteSampler`	`MTB.Core.BornHead`	B5d `24bc81376005` 61/61
Knuth–Yao converse	`knuthYao`	`MTB.Core.KnuthYao`	HA
Clifford decision	`routeSound, minorVanishes, cliffDec, check-complete`	`MTB.Core.Route`	B6a `9f74b704b72d` 69/69
Order on $\mathbb{Z}[\sqrt{p}]$	`coneUnique, signIsCone`	`MTB.Core.QuadOrder`	B6b `1e852edf8fea` 81/81
DKW acceptance	`acceptRuleSound, acceptRuleComplete, rejectRuleWitness`	`MTB.Core.DKW`	B6b
Exact level at $n = 100$	`level100`	`MTB.Core.DKWExact`	B6b
Yang–Lee rejected	`yangLeeRejectsAll`	`MTB.Core.DKWFib`	B6b
Net typing; identity refuted	`acceptedLowers, identityRefuted`	`MTB.Lane.K. NetTypingGeneral`	HA (B6c `b08fcb4846dc` 47/47)
Reducer copies every tree	`copyAll`	`MTB.Core.LafontCopyAll`	HA
Reducer output = model	`actualIsModel*`	`MTB.Core.LafontActual`	HA
LB8 capstone (full grammar)	`lb8Full`	`MTB.CapstoneFull`	X1 `9ac6ff1b6588` 55/55; X4

Reported rows (not kernel facts). The `mtb-run` oracle sweep, 1,440 runs with 0 mismatches (X2 `d68995f9cf38`, 57/57). The end-to-end acceptance run, with 1,440 configurations byte-identical with oracles removed and 60/60 surgery checks (X3 `7d9c5c0617b1`, gate 9/9). The exact-engine runs (B3c4 `bdb28663021e`, 16/16). The experiment (B4e; pre-registrations `de5158fb...`, `d77bf873...`). The simulator verdicts (B6b). The text-level bridges to earlier libraries. **Open:** `B6b_hw`. **Claim graph:** the project claim is *proved*, with a stable fixed-point receipt (two runs) and a planted rejection.

B The controlled corpus

The 24 sentences of the frozen contract corpus, verbatim (`data/contract_corpus.json`). Each is an Agda derivation term. "Copy" names the copy each sentence exercises. The corpus also fixes 3 rings, 5 theories, 5 knots, 5 manifolds, 4 contextuality models and 1 learning task.

Id	Sentence	Class	Copy
S01	Alice loves Bob	control-transitive	none
S02	Bob loves Alice	control-transitive	none

Id	Sentence	Class	Copy
S03	Alice sleeps	control-intransitive	none
S04	Alice does not sleep	control-negation	none
S05	Alice who sleeps	subject-relative	frobenius-noun
S06	Bob who loves Alice	subject-relative	frobenius-noun
S07	Alice whom Bob loves	object-relative	frobenius-noun
S08	Bob whom Alice loves	object-relative	frobenius-noun
S09	Alice who loves Bob who sleeps	nested-relative	frobenius-noun x2
S10	Alice who sleeps and Bob who sleeps	coordination+relative	frobenius-noun x2
S11	Alice sleeps and Bob sleeps	coordination	frobenius-sentence
S12	Alice loves and Bob loves Carol	right-node-raising	frobenius-object
S13	Alice Bob loves	ambiguous-role-order	none
S14	Bob Alice loves	ambiguous-role-order	none
S15	Carol Alice Bob loves	ambiguous-3-noun	none
S16	Alice gave Bob Carol	ditransitive-role-binding	none
S17	Bob gave Alice Carol	ditransitive-role-binding	none
S18	Alice Bob Carol saw help	crossed-dependency	none
S19	Alice who Bob saw sleeps	object-relative-matrix	frobenius-noun
S20	Alice loves Alice	reuse-control (two occurrences, distinct)	none
S21	ill-typed: loves Alice	negative-control-typing	n/a
S22	duplicate-occurrence derivation (Resource.duplicate)	negative-control-linear-use	illegal
S23	Alice who is quantum-meaning (noun space non-classical)	copy-unsafe-by-design	frobenius on non-classical
S24	Bob whom Alice loves and Carol sleeps	mixed	frobenius-noun + frobenius-sentence

C Glossary and terminology crosswalk

We use one primary term throughout. The table gives each community's nearest word once.

Our term	Physics / quantum	Machine learning	Engineering	Category theory / logic	Constructor theory
copy (δ)	fan-out, cloning, broadcast	residual or attention fan-out; weight tying	broadcast, shared signal	comultiplication, diagonal	cloning task
classical (state)	basis or pointer state	a sharp, discrete value	a deterministic value	point of the classical core; copiable	information attribute
measurement point	mid-circuit measurement	sampling step	commit, disambiguate	dynamic lifting; modality unit	—
separator	witness of non-classicality	—	counterexample	failed copy equation	refuter
price ($1 - \sum p^2$)	logical entropy; one minus collision probability	expected disagreement	cost	—	—
machine-checked	—	—	verified (rule set named)	mechanised, kernel-checked	—

Our term	Physics / quantum	Machine learning	Engineering	Category theory / logic	Constructor theory
refutation	counter-model	negative result	failed requirement	counterexample	refuter
acceptance at α	consistency test, DKW band	hypothesis test	acceptance test	—	—
exact	exact cyclotomic arithmetic	bit-exact, deterministic	reproducible, byte-identical	decidable equality	—
anyon theory	unitary modular tensor category	—	—	braided fusion category	—
open frame	—	a template with unfilled slots	—	a morphism with open noun wires	—

Words used only with a qualifier.

- "proved" always names its system.
- "accepted" always names α and n.
- "first" is always dated by a search.

We claim no quantum advantage, and nothing here is a statement about the physical safety of a device.

D Kernel resources of the audit authorities

Each hostile-audit row was authorised by a node run of Agda 2.8.0 on the current bytes, labelled C0. Each planted mutant was re-checked separately. A mutant counts as refused only if the check fails *with* a source position. A check that merely exhausts memory does not count. The table is generated from the copied check records (`data/ha_checks/`) by `scripts/make_appendix_tables.py`. Overall, 47 of 48 mutants are refused with a position. The exception is `LafontActual`'s M6, the orientation of `pairL`: it survives there because the corpus nets are symmetric under it, and it is refused at `LafontProbes` line 468. The audit record discloses this and counts only `LafontActual`'s own refusals.

Module (authority run C0)	rc	peak MiB	mutants refused	Agda
`MTB.All`	0	891	0/0	2.8.0
`MTB.Core.Contextuality222`	0	248	3/3	2.8.0
`MTB.Core.FibHexagonCensus`	0	209	3/3	2.8.0
`MTB.Core.FibLens`	0	353	5/5	2.8.0
`MTB.Core.KnuthYao`	0	150	3/3	2.8.0
`MTB.Core.LafontActual`	0	68	6/7	2.8.0
`MTB.Core.LafontBuild`	0	172	0/0	2.8.0
`MTB.Core.LafontCopyAll`	0	341	5/5	2.8.0
`MTB.Core.LafontCopyStep`	0	801	0/0	2.8.0
`MTB.Core.LafontLists`	0	239	0/0	2.8.0
`MTB.Core.LafontProbes`	0	177	5/5	2.8.0
`MTB.Core.NormFixed`	0	206	3/3	2.8.0
`MTB.Core.SolverCanonical`	0	124	3/3	2.8.0
`MTB.Lane.C.All`	0	1,250	0/0	2.8.0
`MTB.Lane.K.All`	0	1,258	0/0	2.8.0
`MTB.Lane.K.NetTypingGeneral`	0	799	5/5	2.8.0

Module (authority run C0)	rc	peak MiB	mutants refused	Agda
`MTB.Lane.K.PriceBridge`	0	353	4/4	2.8.0
`MTB.Lane.K.QNLPCopyDecide`	0	598	2/2	2.8.0

Other measured costs:

- cold replay of the capstone claim: 78 min at 2.72 GB;
- the golden-foam lane's cold check: about 25 min per gate module;
- one concrete closed-sentence request of the braided executable: about 10 s and 4.8 GB.

The heaviest kernel traps, each measured and removed, are recorded in the phase findings (`evidence/findings/`).

E The pre-registration of the braiding experiment

The v2 pre-registration (`data/b4e_v2_prereg.json`, sha256 `d77bf873ab2c43c359011a330aba2344ba63ef2384a8124da23ba6fe33bed0f7`) fixes the following, verbatim in substance:

- **Question.** On role binding with compositional generalisation in the number of entities, does an Ising-braided text circuit beat the symmetric text circuit of the same architecture and parameter count?
- **Theory before data.**
  - Both models' Born outputs are relayout-invariant (T1′, exact).
  - The voice task is layout-only, so both models have paired voice accuracy exactly $\frac{1}{2}$ (T3′).
  - The role-binding prediction is open, with no direction stated.
- **Data.** Train $n \in \{2, 3\}$ (2,000 stories per n); validate on 250 and test on 500 per n, for $n \in \{4, 5, 6\}$. Splits are seeded and disjoint.
- **Training.**
  - 3,000 Adam steps, batch 32 per n, learning rate from $\{0.003, 0.01, 0.03\}$ selected on validation;
  - per seed, both models share data, initialisation and minibatches;
  - CPU only.
- **Decision.** Δ is the mean over 20 seeds of the paired accuracy difference, with a 95% CI from a hierarchical paired bootstrap ($B = 10^4$, seed 20261022). A CI excluding 0 decides; otherwise the result is null. A model is informative only if its validation accuracy is at least chance +0.10.
- **Stopping.** All $2 \times 2 \times 20$ runs; none repeated or dropped; nothing changed after any result was seen.
- **Disclosure.**
  - v2 supersedes v1 (`de5158fb...`) after a model-fidelity defect was found by review. v1's braided arm omitted the Koszul crossing for subject-right events: the v1 sign is wrong on 192 even units at $n = 4$ and fails coherence 576 times.
  - v2 changes only that sign.

– v2 was registered after v1's results were known.