hexagon:2611.01032v1

Time is the order in which the universe makes its decisions

Aiden Smith

November 2026

Abstract

Suppose that definite outcomes are decisions taken by a hidden layer at the moment which-branch information leaks irreversibly, all in one global order, and that this order is what time is at the level of facts. We show that Born's rule is exactly the condition that hides that order. Born's rule is equivalent to the invariance of outcome statistics under the layer's choice of stopping rule, and the order of spacelike decisions leaves no trace in any marginal statistics if and only if every detector follows Born's rule. A layer driven by smooth noise cannot decide fairly: some stopping rule sees a violation of order $\sqrt{\kappa}$, where κ is the decision rate times the memory of the noise, while the final odds move only at order κ. For two-state noise the decision rule has an exact closed form that runs from Born's rule to majority rule, with zones of certainty above $\kappa = \frac{1}{2}$; for any number of outcomes driven by exchangeable bounded noise such zones never appear below $\kappa = \frac{1}{2}$, and for Gaussian noise the first-order distortion favours the largest single branches. The same argument shows that exact colored-noise collapse models can assign Born probabilities only at a final time fixed in advance. Where Born's rule fails, the order becomes visible as a one-directional signal between entangled parties, with an optimal strength of $(3\sqrt{2} - 4)k$. Public data from loophole-free Bell tests bound $\kappa < 0.017$ for photodetection decisions, whichever station decides first; a direct order-swap test on a superconducting quantum processor, where the order is set by the circuit, gives $\kappa < 0.006$ for readout decisions; and Z-boson decays bound κ below 2.4×10^{-3} to 1.1×10^{-2} for decay decisions, depending on which channels count as outcomes.

1 Introduction

Quantum mechanics supplies probabilities; experiments deliver single outcomes. In any picture in which outcomes are produced by an underlying process, a hidden layer that turns weights into facts, three experimental facts constrain how that process is ordered. Bell correlations require the layer to act on the joint configuration of distant systems. Experiments with moving beam splitters exclude models in which each detector decides first in its own rest frame [31], so the decisions share a single order. And hidden influences of any finite speed would permit signalling [5], so the shared order is laid along instantaneous slices. We take this order to be what time is at the level of facts: the sequence in which definite outcomes come into existence.

Relativity treats all inertial frames alike, and no experiment has revealed a preferred simultaneity. Our central result explains the silence. The order of spacelike decisions leaves no trace in any statistics exactly when every detector follows Born's rule (Theorem 2). Born's rule is the condition that hides time's order.

We assume a single bridge principle: decisions are stopping times of a hidden process, taken when which-branch information leaks irreversibly, in one global order. Everything else is derived. Sections 3 and 4 treat a layer whose decisions are fair. Sections 5–9 treat a layer driven by noise with a finite memory τ, where fairness fails; one dimensionless number, $\kappa = c^2\tau$, the decision rate times

the memory, controls every deviation. Section 10 shows how the order becomes visible, Section 11 confronts the predictions with data, and Section 12 draws two consequences for the global structure of time. The main results are:

1. Born's rule is equivalent to invariance of outcome statistics under the choice of stopping rule; threshold devices suffice, and every continuous fair decision is Brownian motion run on a hidden clock (Theorem 1).
2. The decision order is statistically invisible if and only if every detector follows Born's rule, with no symmetry assumption (Theorem 2). This is the time-ordered form of the no-signalling derivations of Born's rule [18, 32].
3. A layer with smooth paths cannot decide fairly, and some stopping rule sees a violation of order $\sqrt{\kappa}$ (Theorem 3).
4. For two-state noise the decision rule is exactly solvable, with a sharp transition at $\kappa = \frac{1}{2}$ (Theorem 4); for many outcomes the first-order law favours the largest single branches (Proposition 3), and for exchangeable bounded noise certainty zones need $\kappa \geq \frac{1}{2}$ (Proposition 4).
5. Exact colored-noise collapse models assign Born probabilities only at a final time fixed in advance, not at the state-dependent moments at which detectors click (Theorem 5).
6. The signalling channel opened by a deformed rule has the exact optimum $(3\sqrt{2}-4)k$ in its natural family, maximally entangled states are immune to it, and public Bell-test data and a direct test on a superconducting processor find no effect, bounding κ at the level of 10^{-2} (Sections 10 and 11).

2 Setting and axioms

The layer carries each measurement's outcome weight as a random process and stops it at hidden moments.

Axiom 1 (weights)**.** For a two-outcome measurement, the layer carries a process $W_t \in [0, 1]$, adapted to the layer's filtration $(\mathcal{F}_t)$, with $W_0 = |c_+|^2$, the quantum weight of the outcome "+".

Axiom 2 (decisions are stopping times)**.** A decision happens at a stopping time T of $(\mathcal{F}_t)$. The outcome is "+" with conditional probability W_T; when the weight runs to 0 or 1, the outcome is the value it reaches.

Axiom 3 (decisions at leakage)**.** No decision is taken before which-branch information has leaked irreversibly into an environment. Interference experiments enforce this: an earlier decision would destroy interference that is observed.

Axiom 4 (one decision order)**.** Decided events carry a strict total order $\prec$. If the setup of decision B depends on the outcome of decision A, then $A \prec B$. The order is laid along the slices of a foliation, on which the layer acts instantaneously.

Axiom 5 (branch update)**.** When a decision selects an outcome, every weight not yet decided is replaced by its value in the state conditioned on that outcome. A decision's outcome probabilities depend on the state only through the weights of its own outcomes.

Definition 1. When the dynamics are Markov in the weight, the *decision rule* is $F(x) = \Pr(+ \mid W_0 = x)$; Born's rule is $F(x) = x$. The *decision rate* c^2 is the rate at which the weight is driven

during a decision, roughly the inverse of the time for one bit of which-branch information to leak. The *memory* τ is the correlation time of the layer's noise. Their product $\kappa = c^2\tau$ is the dimensionless number that controls everything below.

Thesis. At the level of facts, time is the order $\prec$. Clocks, light cones and proper time remain properties of averages. The rest of the paper asks when this order can be seen.

3 Fairness: Born's rule as invariance under stopping

Different devices, environments and delays stop the layer at different hidden moments. Born's rule for a device with stopping time T means $\mathbb{E}[W_T] = W_0$.

Theorem 1 (fairness)**.**

(1) $\mathbb{E}[W_T] = W_0$ *for every bounded stopping time* T *if and only if* W *is a martingale.*

(2) Let W *be a regular diffusion on* $(0,1)$*,* $dW = \mu(W)\,dt + \sigma(W)\,dB$ *with* $\sigma > 0$*. If every* threshold device*, which stops when* W *first leaves an interval* $(a,b) \ni W_0$*, obeys Born's rule, then* $\mu \equiv 0$*.*

(3) If W *is a continuous martingale, it is a Brownian motion run on the clock* $\langle W\rangle$ *and absorbed at* 0 *and* 1*:* $W_t = W_0 + B(\langle W\rangle_t)$*.*

(4) If W *is a continuous martingale whose clock runs at a rate bounded below on every interval* $[\delta, 1-\delta]$*, then* W *ends at* 0 *or* 1 *almost surely.*

Proof. (1) If W is a martingale, optional stopping gives $\mathbb{E}[W_T] = W_0$. Conversely, take $s < t$ and $A \in \mathcal{F}_s$. The rule "stop at s off A, at t on A" is a bounded stopping time; comparing it with "stop at s" gives $\mathbb{E}[W_t \mathbf{1}_A] = \mathbb{E}[W_s \mathbf{1}_A]$, the martingale property. (2) The device ends at b with probability $(s(x) - s(a))/(s(b) - s(a))$, where s is the scale function. Born's rule requires $(x-a)/(b-a)$ for all $a < x < b$, so s is affine, and since $s'' = -(2\mu/\sigma^2)s'$, μ vanishes. (3) This is the Dambis–Dubins–Schwarz theorem [28]. (4) The clock's final value satisfies $\mathbb{E}\langle W\rangle_\infty = \mathbb{E}[W_\infty^2] - W_0^2 \le 1$, so it is finite almost surely. A bounded martingale converges, and a limit inside $(0,1)$ would keep the clock running forever. □

Born's rule is thus a symmetry statement: outcome statistics do not depend on when the hidden layer stops, and the conserved quantity is the weight. Threshold devices are ordinary sequential measurements ("stop once the confidence reaches a or b"), so part (2) needs nothing exotic. Part (3) says that every continuous fair decision is one object, Brownian gambler's ruin, whose only freedom is the clock $\langle W\rangle$. That clock is the hidden timing of the decision, and it cannot be observed, because fairness makes every statistic independent of it. Pearle's gambler's-ruin collapse [25] is the best-known instance.

4 Born's rule is what hides the order

Alice and Bob share a two-qubit state and each makes a two-outcome decision, Alice with rule F_A and Bob with rule F_B. Let w and $1-w$ be the weights of Alice's outcomes, b_1 and b_2 Bob's weights conditional on each of them, and $B_0 = wb_1 + (1-w)b_2$ his unconditional weight. If Alice decides first in $\prec$, Axiom 5 puts Bob's qubit in one of the two conditional states; if Bob decides first, his decision acts on the entangled state:

$$P_B(\text{Alice first}) = F_A(w)\,F_B(b_1) + \big(1 - F_A(w)\big)F_B(b_2), \qquad P_B(\text{Bob first}) = F_B(B_0). \tag{1}$$

Theorem 2 (order invisibility)**.** *Let F_A and F_B be continuous with $F(0) = 0$ and $F(1) = 1$. Bob's marginal is the same whichever party decides first, for every two-qubit state and every pair of measurements, if and only if $F_A(x) = F_B(x) = x$.*

Proof. If both rules are the identity, both expressions equal B_0. Conversely, three families of configurations suffice; by the Hughston–Jozsa–Wootters theorem [19], a projective measurement on Alice's qubit steers Bob's qubit to any two-state ensemble of his reduced state, so each family is realisable. (i) With the state $\sqrt{w}\,|00\rangle + \sqrt{1-w}\,|11\rangle$ and both parties measuring in that basis, $b_1 = 1$ and $b_2 = 0$, and agreement gives $F_A(w) = F_B(w)$; call the common rule F. (ii) With $b_1 = b$ and $b_2 = 0$, agreement gives $F(w)F(b) = F(wb)$. (iii) With $b_1 = 1$ and $b_2 = b$, it gives $1 - F(w + (1-w)b) = (1 - F(w))(1 - F(b))$. So F and $G(x) = 1 - F(1-x)$ are both multiplicative on $[0, 1]$. A continuous multiplicative map of $[0, 1]$ with $F(0) = 0$ and $F(1) = 1$ is a power: F cannot vanish inside $(0, 1]$, since $F(x_0) = 0$ would give $F(x_0^{1/n}) = 0$ for every n and hence $F(1) = 0$, and then $-\ln F(e^{-s})$ is continuous and additive in s. So $F(x) = x^p$ and $G(x) = x^q$ with $p, q > 0$, and $x^p = 1 - (1-x)^q$ for all x. Comparing both sides as $x \to 0$ gives $p = q = 1$. □

In a world with one decision order, the operational equivalence of inertial frames, applied to outcome statistics, is therefore not a separate principle: it is equivalent to Born's rule, for every detector at once. Time's order can exist and stay invisible only if every decision is fair.

Without a decision order, the same equivalence appears as a no-signalling derivation of Born's rule: if outcome probabilities are any function of the quantum weight other than the identity, a party who steers can signal to the other [18, 32, 33]. The decision order adds a direction. Only the party who decides later sees its statistics shift, so the signal runs one way and reverses when the order does (Section 10). Both arguments use rules that depend on an outcome's own weight alone, as in Axiom 5. Outside that class, Galley and Masanes give a modified measurement postulate that does not signal, and show that every modification violates purification and local tomography [15].

Corollary 1 (remote-decision detector)**.** *Write $F(x) = x + k\,g(x)$ for both parties. When Alice's detector is fair, Bob's own detector still reveals whether Alice has decided, through the Jensen gap*

$$\Delta_B = k\big[w\,g(b_1) + (1-w)\,g(b_2) - g(B_0)\big] + O(k^2). \tag{2}$$

For $g(x) = x(1-x)(2x-1)$, the distortion of a Gaussian smooth layer (Section 5),

$$\frac{\Delta_B}{k} = w(1-w)(b_1 - b_2)^2\big[3 - 2(1+w)\,b_1 - 2(2-w)\,b_2\big] + O(k), \tag{3}$$

and its largest magnitude over $[0, 1]^3$ is exactly $\frac{1}{8}$, reached at $w = \frac{1}{3}$, $b_1 = \frac{1}{4}$, $b_2 = 1$ and its mirror images.

Proof. In terms of B_0 and $d = b_1 - b_2$ the bracket in (3) equals $3 - 6B_0 - 2d(1-2w)$, so for fixed w and d the gap falls linearly as B_0 grows, and its minimum puts the larger of b_1, b_2 at 1. With $b_2 = 1$ and $b_1 = 1 - e$ it equals $-w(1-w)e^2[3 - 2e(1+w)]$, which is smallest at $e = 1/(1+w)$, with value $-w(1-w)/(1+w)^2$. That is smallest at $w = \frac{1}{3}$, where it equals $-\frac{1}{8}$. The case $b_1 = 1$ is the same with w and $1-w$ exchanged, and $g(1-x) = -g(x)$ turns the minimum into the maximum $+\frac{1}{8}$. Every point of the cube is realisable by [19]. □

Alice need not read her outcome: letting her qubit decohere is enough.

5 Smooth layers cannot decide fairly

A weight with smooth paths cannot be fair: a smooth path has a direction at every instant, and a fair game cannot have one.

Theorem 3 (no smooth fair decisions)**.**

(1) A continuous martingale whose paths have finite variation is constant.

(2) Suppose W has C^1 paths near $t = 0$ with $W'(0) \neq 0$ almost surely. For the threshold device with $a = x - \delta$ and $b = x + k\delta$ $(k > 0)$, the probability of ending at b tends to $\Pr(W'(0) > 0)$ as $\delta \to 0$, for every k, whereas Born's rule requires $1/(1 + k)$. So no weight with C^1 paths is fair for all threshold devices.

(3) (A floor.) *Let W have continuous paths that are monotone up to a stopping time $R > 0$, the first reversal of direction, and let the initial direction be known at time $0+$. Then*

$$\sup_T |\mathbb{E}[W_T] - W_0| \;\geq\; \max\big(\mathbb{E}(W_R - W_0)^+,\; \mathbb{E}(W_R - W_0)^-\big) \;\geq\; \tfrac{1}{2}\,\mathbb{E}|W_R - W_0|. \tag{4}$$

Proof. (1) Its quadratic variation vanishes, so $\mathbb{E}[(W_t - W_0)^2] = \mathbb{E}\langle W\rangle_t = 0$. (2) Where $W'(0) > 0$ the path increases on some interval $[0, t_0(\omega)]$, so for small enough δ it reaches $x + k\delta$ before $x - \delta$; where $W'(0) < 0$ the reverse holds. Dominated convergence gives the limit $\Pr(W'(0) > 0)$, whatever k is. (3) Stop at R on the paths that start upward and at once on the others. This is a stopping time, because the direction is fixed at $0+$ (decide at a small time ε by the sign of $W_\varepsilon - W_0$ and let $\varepsilon \to 0$), and it gives $\mathbb{E}[W_T] - W_0 = \mathbb{E}(W_R - W_0)^+$. The mirror rule gives $-\mathbb{E}(W_R - W_0)^-$. □

Brownian paths escape only because they cross their starting value infinitely often in every interval after the start.

The smooth-noise model. With $y = \operatorname{artanh} z$ and $x = (1 + z)/2$ the outcome weight, let

$$\dot{y} = c\,\eta(t) + c^2 g(y), \qquad W = f(y) = \tfrac{1}{2}(1 + \tanh y), \tag{5}$$

where η is stationary noise with mean zero, unit intensity $\int_{-\infty}^{\infty} \mathbb{E}[\eta(t)\eta(0)]\,dt = 1$, variance $s^2 = 1/(2\tau)$ and correlation time τ. The drift $g = \tanh$ makes W a martingale in the white-noise limit $\tau \to 0$, where the model reduces to the standard collapse equation.

Size of the floor. In the model (5) the floor (4) is of order $\sqrt{\kappa}$. With two-state noise (Section 6) and $\kappa < \frac{1}{2}$, the noise outruns any drift of size at most one, so the velocity reverses exactly when the noise flips; R is then exponential with rate $1/(2\kappa)$ (units $c = 1$) and $\mathbb{E}(W_R - x)^\pm = x(1-x)\sqrt{2\kappa}\,(1 + O(\sqrt{\kappa}))$. With Ornstein–Uhlenbeck noise, $d\eta = -\eta\,dt/\tau + dB_t/\tau$ gives $\int_0^R \eta\,dt = \tau(\eta_0 - \eta_R) + B_R$, so the area under the noise until its first zero has mean $\tau\eta_0$, and $\mathbb{E}(W_R - x)^\pm = x(1 - x)\sqrt{\kappa/\pi}\,(1 + O(\sqrt{\kappa}))$. The leading terms do not depend on the drift as long as it is bounded. Figure 1 compares exact quadrature (two-state noise) and simulation (Gaussian noise, 200 000 paths, with a bridge correction for zero crossings inside a time step) with these laws. At $\kappa = 10^{-3}$ the floor is about 4κ for Gaussian noise and 9κ for two-state noise, while the final odds move by about 0.1κ.

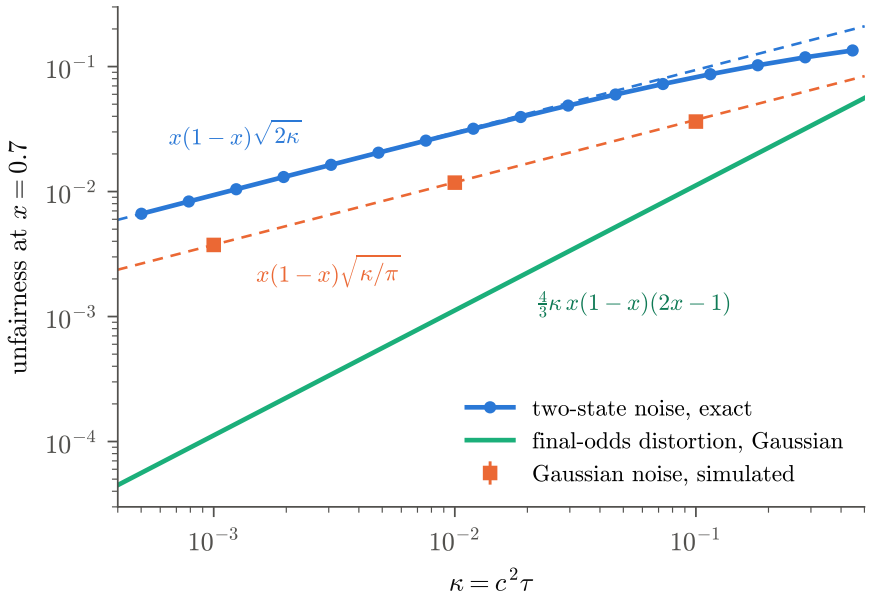

Figure 1: The unfairness floor of Theorem 3(3) at $x = 0.7$: the mean of $\mathbb{E}(W_R - x)^+$ and $\mathbb{E}(W_R - x)^-$ for two-state noise (exact) and Gaussian noise (simulation), against the leading-order laws (dashed). The distortion of the final odds, $\frac{4}{3}\kappa\, x(1-x)(2x-1)$, is parametrically smaller.

Interruption at fixed times. Suppose a device reads the weight through any smooth increasing function u ($u = f$ for Born's rule) at a fixed time t. Its statistics are unchanged only if $\mathbb{E}\, u(y_t) = u(y_0)$ for all t. Smooth paths carry no Itô term, so for $t \ll \tau$

$$\mathbb{E}\, u(y_t) - u(y_0) = c^2 u'(y_0)\, g(y_0)\, t + \tfrac{1}{2} c^2 s^2 u''(y_0)\, t^2 + O(t^3). \tag{6}$$

The first term forces $g \equiv 0$; then the second forces $u'' \equiv 0$, and an affine function of y that stays between 0 and 1 on the whole line is constant. No choice of drift or readout makes a smooth layer fair under interruption. Simulation with Ornstein–Uhlenbeck noise at $\kappa = 0.1$ and $x = 0.7$ (200 000 paths) gives a shift of $\mathbb{E}f$ of $+0.00160 \pm 0.00002$ at $t = \tau/10$ (the formula gives $+0.00160$) and $+0.0103 \pm 0.0002$ at $t = \tau$; a drift tuned to first order in κ, $g = \tanh y\,(1 - 2c^2\tau\, \mathrm{sech}^2 y)$, gives $+0.00132$ and $+0.0076$, and white noise gives -0.00005 ± 0.00009 and -0.0004 ± 0.0003.

Final odds. With $g = \tanh$ and Gaussian noise the final odds are sharpened,

$$F_\kappa(x) = x + \tfrac{4}{3}\kappa\, x(1-x)(2x-1) + O(\kappa^2), \tag{7}$$

the two-outcome case of Proposition 3. Simulation at $\kappa = 0.1$ (400 000 paths) gives $F - x = +0.0114 \pm 0.0007$, -0.0124 ± 0.0006 and $+0.0124 \pm 0.0006$ at $x = 0.7$, 0.2 and 0.8, against $+0.0112$, -0.0128 and $+0.0128$. The tuned drift removes the first-order term (residual -0.0046 ± 0.0010 at $\kappa = 0.1$) but, by the floor, not the unfairness.

What a smooth layer can still do. It can keep the weight frozen and decide by a sudden fair jump, as flash models do. It can predetermine outcomes, as pilot-wave theory does: given the whole layer, the weight is 0 or 1 throughout. Or the weight can be a rough function of a smooth but genuinely random layer. For a stationary Gaussian layer the last option fails when paths are

analytic in time, because the Kolmogorov–Krein criterion then makes the past determine the future [12]. A smooth layer that keeps producing new randomness must be infinitely differentiable but not analytic.

6 Exact crossover from Born's rule to majority rule

For a layer driven by two-state noise the decision rule can be found in closed form. The noise is a telegraph signal $s(t) = \pm\sigma$ that flips at rate γ in each direction, and

$$\dot{y} = c\, s(t) + c^2 \tanh y, \qquad \sigma^2 = \gamma, \qquad \tau = \frac{1}{2\gamma}, \qquad \kappa = c^2\tau. \tag{8}$$

The choice $\sigma^2 = \gamma$ gives the noise unit intensity. The layer starts in its stationary state, $s = \pm\sigma$ with probability $\frac{1}{2}$ each. Write $z = \tanh y$ and $x = (1+z)/2$.

Theorem 4 (exact crossover)**.** *With $\sigma^2 = 1/(2\kappa)$, the decision rule is the distribution function $F_\kappa(x) = \int_{\text{lo}}^{2x-1} \rho_\kappa / \int \rho_\kappa$, where*

$$\rho_\kappa(z) \propto \begin{cases} (1-z^2)^{\,b-1}(\sigma^2 - z^2)^{-b-1}, & b = \dfrac{1}{1-2\kappa}, \quad \kappa < \tfrac{1}{2}, \\ (1-z^2)^{-2}\, e^{-1/(1-z^2)}, & \kappa = \tfrac{1}{2}, \\ (\sigma^2-z^2)^{\,a-1}(1-z^2)^{-a-1} \quad \text{on } |z|<\sigma, & a = \dfrac{1}{2\kappa-1}, \quad \kappa > \tfrac{1}{2}, \end{cases} \tag{9}$$

and the lower limit is the bottom of the support. For $\kappa > \frac{1}{2}$, $F_\kappa(x) = 0$ when $2x - 1 \le -\sigma$ and $F_\kappa(x) = 1$ when $2x - 1 \ge \sigma$.

Proof. Set $c = 1$ and let $u_\pm(z)$ be the probability of "+" starting from z with the noise in state $\pm\sigma$. The backward equations are

$$(z+\sigma)(1-z^2)\, u_+' = \gamma\,(u_+ - u_-), \qquad (z-\sigma)(1-z^2)\, u_-' = -\gamma\,(u_+ - u_-). \tag{10}$$

The difference $D = u_+ - u_-$ solves a separable equation,

$$\frac{D'}{D} = \frac{2\gamma z}{(1-z^2)(z^2-\sigma^2)} \quad \Longrightarrow \quad D = C\left|\frac{\sigma^2 - z^2}{1-z^2}\right|^{\gamma/(1-\sigma^2)}, \tag{11}$$

and the exponent $\gamma/(1-\sigma^2)$ equals $-b$ below $\kappa = \frac{1}{2}$ and a above it. The averaged probability has derivative $\frac{1}{2}(u_+' + u_-') = \gamma\sigma D/((1-z^2)(\sigma^2 - z^2))$, which is ρ_κ. Below $\kappa = \frac{1}{2}$ ($\sigma > 1$), D vanishes at $z = \pm 1$ and the boundary values $u_\pm(-1) = 0$, $u_\pm(1) = 1$ fix the constants. Above it ($\sigma < 1$), both noise states push z the same way once $|z| > \sigma$, so the outcome there is certain and D vanishes at $\pm\sigma$. At $\kappa = \frac{1}{2}$ both forms converge to the middle line, because $\big(1-(\sigma^2-1)/(1-z^2)\big)^{\sigma^2/(\sigma^2-1)} \to e^{-1/(1-z^2)}$. For the verification, the functions $u_\pm$ are bounded and solve the backward equations, so $u_{s(t)}(z_t)$ is a bounded martingale. Below $\kappa = \frac{1}{2}$, $y_t \to \pm\infty$ almost surely, because $\tanh y$ pushes outward while the noise has mean zero, so the martingale converges to the indicator of "+"; above it, the process enters $|z| \ge \sigma$ in finite time, where the outcome is fixed. Optional stopping identifies $u_\pm$ with the outcome probabilities. □

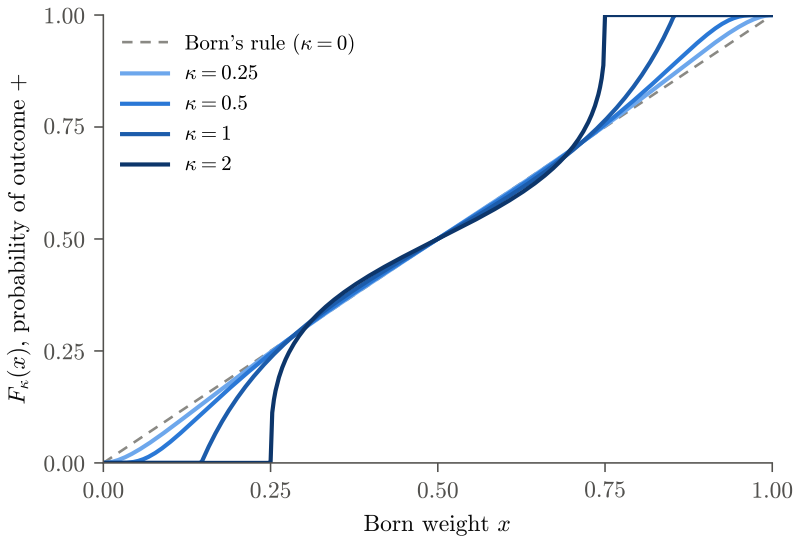

Figure 2: The exact decision rule of Theorem 4 for a layer driven by two-state noise. Above $\kappa = \frac{1}{2}$, states with $|2x - 1| \geq 1/\sqrt{2\kappa}$ are decided with certainty for the likelier outcome.

Table 1: Theorem 4 against direct simulation of the telegraph-driven model (60 000 paths per entry, standard error about 0.002): simulated / exact. At $\kappa = 0.05$ and $x = 0.9$, simulation with a fine time step gives 0.9040 ± 0.0008 against 0.9040.

κ	$x = 0.6$	$x = 0.7$	$x = 0.9$
0.25	0.5942 / 0.5978	0.6989 / 0.6995	0.9241 / 0.9225
0.5	0.5934 / 0.5953	0.6994 / 0.6985	0.9532 / 0.9531
1	0.5901 / 0.5895	0.6959 / 0.6960	1.0000 / 1.0000
2	0.5799 / 0.5797	0.7033 / 0.7017	1.0000 / 1.0000

Consequences. Born's rule is the limit $\kappa \to 0$: there $b \to 1$, $\sigma \to \infty$, ρ_κ becomes flat and $F_\kappa(x) = x$. Expanding $\rho_\kappa = 1 + 2\kappa[\ln(1 - z^2) + 2z^2] + O(\kappa^2)$ inside the integrals, where the logarithm is integrable so the remainder is uniform, gives

$$F_\kappa(x) = x + \kappa\, h(x) + O(\kappa^2), \qquad h(x) = \Phi(2x - 1) - (4 \ln 2 - \tfrac{8}{3})x, \tag{12}$$

with $\Phi(z) = (1 + z)\ln(1 + z) - (1 - z)\ln(1 - z) - 2z + 2\ln 2 - 2 + \frac{2}{3}(z^3 + 1)$. This is not a pure sharpening: $h(0.6) = -0.0080$, $h(0.7) = -0.0010$, $h(0.9) = +0.0789$, whereas Gaussian noise with the same correlation function gives the pure sharpening (7). The shape of the distortion fingerprints the law of the noise, not only its memory. At $\kappa = \frac{1}{2}$ the rule changes character: above it, every state with $|2x - 1| \geq 1/\sqrt{2\kappa}$ is decided with certainty for the likelier outcome, and this zone widens as κ grows. As $\kappa \to \infty$ the undecided zone shrinks to $x = \frac{1}{2}$ and F_κ tends to a step: majority rule (Figure 2). Table 1 compares the theorem with direct simulation.

7 Pulsed decisions and slow layers

Decisions that are fast compared with the layer's memory follow majority rule, whether the speed comes from pulsing the decision or from a slow layer.

Table 2: Pulsed decisions at $\kappa = 0.1$ and $x = 0.7$ (100 000 paths).

t/τ	λ	simulated F	$I_x(\lambda, \lambda)$
0.1	20.67	0.9962 ± 0.0002	0.9962
0.3	7.35	0.9420 ± 0.0007	0.9423
1	2.72	0.8251 ± 0.0012	0.8240
3	1.46	0.7475 ± 0.0014	0.7447
10	1.11	0.7164 ± 0.0014	0.7119

Proposition 1 (pulsed decisions)**.** *In the model (5) with Ornstein–Uhlenbeck noise, let the decision run only during pulses of length t, separated by gaps long compared with τ. In the limit of many weak pulses, $c^2 t \to 0$ with t/τ fixed,*

$$F(x) \to I_x(\lambda, \lambda), \qquad \lambda = \frac{t}{t - \tau(1 - e^{-t/\tau})} \;\approx\; \frac{2\tau}{t} \quad (t \ll \tau), \tag{13}$$

where I_x is the regularized incomplete beta function.

Derivation. During one pulse the drift moves y by $c^2 \tanh(y)\, t$, while the noise adds an increment of variance $c^2[t - \tau(1 - e^{-t/\tau})]$, independent from pulse to pulse because the gaps exceed τ; the change of $\tanh y$ within a pulse corrects these moments by relative amounts of order $c^2 t$. In the limit, the pulse-to-pulse chain converges to the diffusion $dy = c^2 \tanh y\, dT + (c/\sqrt{\lambda})\, dB$ in pulse time T. Its backward equation $(1/2\lambda)u'' + \tanh y\, u' = 0$ gives $u' \propto \cosh^{-2\lambda} y$, that is $du/dx \propto [x(1-x)]^{\lambda - 1}$, so $u = I_x(\lambda, \lambda)$. □

Since λ depends only on t/τ, pulses shorter than the memory give order-one departures from Born's rule even as $\kappa \to 0$. Table 2 compares the limit with simulation at $\kappa = 0.1$; the longest pulses there have $c^2 t = 1$ and sit outside the limit.

Proposition 2 (slow layers)**.** *In the model (5) with Ornstein–Uhlenbeck noise: (a) for $x \neq \frac{1}{2}$, $F_\kappa(x) \to \mathbf{1}[x > \frac{1}{2}]$ as $\kappa \to \infty$; (b) in the window $2x - 1 = \zeta/\sqrt{2\kappa}$, $F_\kappa \to \Phi_{\mathcal{N}}(\zeta)$, the standard normal distribution function.*

Argument. In units $c = 1$ the noise has variance $1/(2\kappa)$ and memory κ, so while a decision runs it is nearly frozen at $\eta(0)$. The equation $\dot{y} = \eta(0) + \tanh y$ has an unstable fixed point at $\tanh y = -\eta(0)$, and the outcome is "+" exactly when $\eta(0) > -z_0$, which gives $\Phi_{\mathcal{N}}(z_0\sqrt{2\kappa})$ and hence (b). Part (a) follows because a path that has escaped returns only through a large deviation of the noise. □

Simulation at $\zeta = 0.5$ gives 0.7152, 0.6990, 0.6950 and 0.6931 at $\kappa = 3$, 10, 30 and 100, against $\Phi_{\mathcal{N}}(0.5) = 0.6915$; the gap shrinks roughly as $(\ln \kappa)/\kappa$. Born's rule is the slow-decision limit of a smooth layer, and decisions fast relative to its memory are majority votes.

8 Many outcomes

Give each of N outcomes a log-weight y_i, with weights $p = \mathrm{softmax}(y)$, and let

$$\dot{y}_i = \sqrt{2}\, c\, \eta_i(t) + 2c^2 p_i, \tag{14}$$

with independent unit-intensity noises η_i. In the white-noise limit this is the standard N-outcome collapse $dp_i = \sqrt{2}\, c\, p_i(dB_i - \sum_j p_j\, dB_j)$, a martingale, and the rate at which any two branches

Table 3: The many-outcome law for three outcomes: Feynman–Kac evaluation of V (40 000–60 000 paths of the fair collapse, errors about ± 0.001) against the replicator part $\frac{4}{3}p(p-S)$. Direct simulation of the colored-noise model at $\kappa = 0.05$ and weights $(0.6, 0.3, 0.1)$ gives $(P-p)/\kappa = (0.110,\ -0.058,\ -0.052) \pm (0.008,\ 0.007,\ 0.005)$. For two outcomes the same code returns 0.1129 ± 0.0011 and 0.0959 ± 0.0008 at $x = 0.7$ and 0.9, against 0.1120 and 0.0960.

weights	V, Feynman–Kac	$\frac{4}{3}p(p-S)$
(0.6, 0.3, 0.1)	(0.116, −0.064, −0.052)	(0.112, −0.064, −0.048)
(0.5, 0.25, 0.25)	(0.088, −0.044, −0.044)	(0.083, −0.042, −0.042)
(0.4, 0.4, 0.2)	(0.023, 0.023, −0.046)	(0.021, 0.021, −0.043)

separate does not depend on N. For two outcomes it is the model (5) with the same c, so $\kappa = c^2\tau$ keeps its meaning. Write $S = \sum_j p_j^2$ and $S_3 = \sum_j p_j^3$, and let

$$\mathcal{M}_0 = \sum_{i,j} p_i p_j (\delta_{ij} - p_i - p_j + S)\, \frac{\partial^2}{\partial p_i \partial p_j} \tag{15}$$

be the generator of the fair collapse, with time in units of $1/c^2$.

Proposition 3 (first-order law, Gaussian layer)**.** *With independent Ornstein–Uhlenbeck noises the probability that outcome k wins is $P_k = p_k + \kappa V_k(p) + O(\kappa^2)$, where V_k vanishes at the vertices of the simplex and solves*

$$\mathcal{M}_0 V_k = 4p_k\Big[-2S^2 + 2p_k S - \tfrac{5}{2}p_k^2 + \tfrac{5}{2}S_3 - \tfrac{1}{2}S + \tfrac{1}{2}p_k\Big]. \tag{16}$$

Equivalently, $V_k = \frac{4}{3}p_k(p_k - S) + \delta V_k$ with $\mathcal{M}_0\, \delta V_k = -\frac{2}{3}p_k[p_k(1-p_k) - (S - S_3)]$, so that δV_k is two thirds of the expected time integral of $p_k[p_k(1-p_k) - (S-S_3)]$ along the fair collapse. Moreover:

(a) for two outcomes $\delta V = 0$ and $V = \frac{4}{3}x(1-x)(2x-1)$;

(b) for three or more outcomes $\delta V \neq 0$, so the law is not of replicator form;

(c) for one outcome of weight x against a field split into many branches of negligible weight, $V = \frac{1}{3}x(1-x)(4x+1)$.

The derivation, a multiple-scale expansion in $\sqrt{\tau}$, is given in Appendix A; the identities were checked in exact rational arithmetic for $N = 2$ to 6. Table 3 compares the Feynman–Kac evaluation of V with the replicator part $\frac{4}{3}p(p-S)$ and with direct simulation of the colored-noise model. The correction δV is small but real: the largest outcome gains, and the smallest loses, slightly more than the replicator form predicts.

Splitting matters. The distortion is not a function of an event's total weight. Split an event A of weight 0.7 into M equal branches and its complement into M equal branches, each branch with its own noise. Then V for A falls slowly with M (Table 4). Splitting only the rival has the opposite effect: by (c), an outcome of weight $\frac{1}{2}$ against a finely split field has $V = \frac{1}{4}$, against 0 for two outcomes, and with 32 rival branches Feynman–Kac gives 0.200 ± 0.002. The distortion favours whichever single branches are largest. Its predictions for coarse events therefore depend on which branches count as the outcomes of one decision.

Proposition 4 (certainty needs $\kappa \geq \frac{1}{2}$ for any number of outcomes)**.** *Let the noises of the outcomes be exchangeable and bounded, able to dwell near any of their joint values for any length of time with*

Table 4: First-order distortion V of an event of weight 0.7 when both it and its complement are split into M equal, independently driven branches (Feynman–Kac, errors about ± 0.0015).

M	1	2	4	8	16	32
V	0.113	0.064	0.040	0.027	0.020	0.015

positive probability, as telegraph noise can, and such that the largest push of one branch against another, $\eta_j - \eta_i$, can be exerted against all rivals at once. Let D be that largest push divided by the root-mean-square value of $\eta_j - \eta_i$, and let τ be the integrated correlation time of $\eta_j - \eta_i$. Zones of weights from which one outcome wins with probability one exist if $\kappa > D^2/2$ and do not exist if $\kappa < D^2/2$. Since $D \geq 1$, they never exist below $\kappa = \frac{1}{2}$, whatever N is.

Proof. Measure time in units of $1/(2c^2)$, so that $\dot{y}_i = s_i + p_i$, where the rescaled pushes have root-mean-square difference $1/\sqrt{2\kappa}$ and largest difference $\Delta = D/\sqrt{2\kappa}$. If $\Delta < 1$, take $\Delta < \Delta' < 1$ and a point where $p_k - p_j \geq \Delta'$ for every $j \neq k$; such points exist near the vertex. There $\frac{d}{dt}(y_k - y_j) \geq \Delta' - \Delta > 0$ whatever the noise does, so every ratio p_j/p_k falls, p_k rises, and $p_k - p_j = p_k(1 - p_j/p_k)$ rises: the region is invariant, every $y_k - y_j$ grows without bound, and outcome k wins for every noise history. If $\Delta > 1$, take $0 < \delta < \Delta - 1$, any interior point, any k and any $j \neq k$. With positive probability the noise stays for a time T within δ of a value that pushes j against all rivals at once; meanwhile every $y_j - y_i$ grows at a rate at least $\Delta - \delta - 1 > 0$, so at the end each exceeds any level L we choose. From there, as long as every $y_j - y_i$ stays above $L/2$, each grows with drift $p_j - p_i \geq 1 - 2Ne^{-L/2}$ plus $\int (s_j - s_i)\, dt$, which has mean zero and finite memory; so with probability tending to one as L grows these differences never fall below $L/2$, and j wins. Outcome k therefore loses with positive probability from every interior point. Finally, exchangeability makes $\eta_j - \eta_i$ symmetric about zero, so its largest value is its largest magnitude, which is at least its root-mean-square value: $D \geq 1$. □

Two-state noise for two outcomes has $D = 1$, which recovers the threshold $\kappa = \frac{1}{2}$ of Theorem 4. Independent two-state noises for N outcomes have $D = \sqrt{2}$, so the threshold is $\kappa = 1$ for every N; noise that pushes one randomly chosen outcome at a time has $D = \sqrt{N/2}$ and threshold $N/4$; Gaussian noise is unbounded, $D = \infty$, and the second half of the proof applies at every κ, so it has no certainty zones. The threshold is set by the shape of the noise, not by the number of outcomes. A large field of small rivals behaves differently, because its total weight grows deterministically: for independent two-state noises, a single large outcome becomes certain, with probability tending to one as the field grows, once $1/(2\sqrt{\kappa}) + (\sqrt{1 + 4\kappa} - 1)/(4\kappa) < 1$, that is above $\kappa \approx 0.596$. The first term is the outcome's worst push and the second the field's growth rate, in units of $1/(2c^2)$.

9 Colored-noise collapse models

The same reasoning applies to the exact collapse models built on colored noise. In these models a linear stochastic Schrödinger equation

$$\frac{d}{dt}|\phi_t\rangle = \big[-iH + c\,\eta(t)A\big]|\phi_t\rangle + K_t|\phi_t\rangle \tag{17}$$

is driven by a Gaussian noise η under a reference measure $\mathbb{Q}$, where the term K_t, local in time for white noise and a functional of the noise history otherwise, is chosen so that $\mathbb{E}_{\mathbb{Q}}\|\phi_T\|^2 = 1$ for every fixed T. The physical probability of a noise history is the norm-weighted measure $d\mathbb{P}_T = \|\phi_T\|^2 d\mathbb{Q}$ at a chosen final time T, which reproduces Born's rule and excludes signalling at that time [3].

Table 5: Largest shifts of Bob's statistics, in units of k.

protocol	largest $\lvert\Delta P_B\rvert/k$	where
Bob's detector alone, Alice fair	$\frac{1}{8}$ (Corollary 1)	$w=\frac{1}{3}$, $b_1=\frac{1}{4}$, $b_2=1$
Alice decides now or later	0.1266 (numerical)	$\cos^2\chi\approx 0.18$
Alice chooses between two directions	$3\sqrt{2}-4\approx 0.2426$	$\cos^2\chi\approx 0.178$, Bob on the equator

Theorem 5 (collapse at state-dependent times)**.** *Suppose the noise in (17) has continuous paths, so that $|\phi_t\rangle$ is continuously differentiable in t, and that the norm $N_t=\|\phi_t\|^2$ is not almost surely constant. Then there is a bounded stopping time σ of the noise filtration with $\mathbb{E}_{\mathbb{Q}}N_\sigma\neq 1$. The norm-weighted rule therefore cannot assign normalized probabilities to outcomes registered at state-dependent times, and renormalizing makes those probabilities depend on the state beyond its weights.*

Proof. N_t is adapted and has C^1, hence finite-variation, paths. If $\mathbb{E}_{\mathbb{Q}}N_\sigma=N_0=1$ for every bounded stopping time σ, then N is a martingale by Theorem 1(1), and a continuous martingale of finite variation is constant by Theorem 3(1), contrary to assumption. □

For white noise the norm is a genuine $\mathbb{Q}$-martingale, which is why the white-noise models are consistent at every stopping time. The size of the inconsistency follows from the floor (4). For a two-outcome measurement with $A=\frac{1}{2}\sigma_z$ and weight x, $N_t-1\approx c(2x-1)\int_0^t\eta\,ds$ at early times, which is monotone until η changes sign; so $\sup_\sigma|\mathbb{E}_{\mathbb{Q}}N_\sigma-1|$ is of order $|2x-1|\sqrt{\kappa}$.

Detectors register outcomes at state-dependent times: a click is the first crossing of a threshold, not a moment fixed in advance. Theorem 5 says that exact colored-noise collapse models assign Born probabilities only to the state at a final time chosen by hand, and not to the moments at which outcomes become definite. This sharpens the observation that non-Markovian quantum trajectories admit no continuous-measurement interpretation [38]. Approximate, causal versions of these models, such as equations truncated at low order in the noise strength, carry no such guarantee outside the weak-noise regime in which they were derived; the exact construction obeys Born's rule, but only at its final time.

10 Making the order visible

Where the decision rule departs from Born's rule, the order becomes measurable. Take $F(x)=x+k\,g(x)$ with $g(x)=x(1-x)(2x-1)$ and $k=\frac{4}{3}\kappa$, the Gaussian layer, and the state $\cos\chi|00\rangle+\sin\chi|11\rangle$, with Alice deciding along a Bloch direction $\mathbf{n}$ and Bob along $\mathbf{m}$. To first order, the shift in Bob's probability caused by Alice deciding first is

$$\Delta P_B=k\Big[g(w)\,(b_1-b_2)+w\,g(b_1)+(1-w)\,g(b_2)-g(B_0)\Big]+O(k^2). \tag{18}$$

The first term comes from Alice's detector and the rest from Bob's (Corollary 1). Torres Alegre proposed testing a deformed rule by comparing Bob's statistics between two steering choices of Alice chosen to maximise this Jensen gap, without predicting its size [33]. At the two-direction optimum of Table 5, Alice's own term is exactly as large as the gap, so it supplies half the signal. Exact and first-order shifts agree, for example 3.2355×10^{-4} against 3.2352×10^{-4} at $\kappa=10^{-3}$. Table 5 lists the optima.

Closed form for the two-direction protocol. Let Bob measure along a direction on the equator of his Bloch sphere, so that $B_0 = \frac{1}{2}$ and $g(B_0) = 0$. Write $C = \cos 2\chi$, $S = \sin 2\chi$, let θ be the polar angle of Alice's direction, and choose its azimuth to maximise the effect. With $p = C\cos\theta$ and $q = S\sin\theta$,

$$\frac{\Delta P_B}{k} = \frac{pq}{4}\Big[1 + \frac{2q^2}{(1-p^2)^2}\Big]. \tag{19}$$

The pairs (p, q) fill the triangle $|p| + q \le 1$, $q \ge 0$, and for $p > 0$ the shift grows with q, so its maximum lies on the edge $q = 1 - p$, where it equals $\frac{1}{4}p(1-p)[1 + 2/(1+p)^2]$. The derivative vanishes at $p = \sqrt{2} - 1$, where $(1+p)^2 = 2$ and the shift is $(3\sqrt{2} - 4)/2$. Alice's mirror direction $\theta \to \pi - \theta$ reverses the sign, so the largest difference between two of her choices is $3\sqrt{2} - 4$, reached at $\theta = 2\chi$ with $\cos^2 2\chi = \sqrt{2} - 1$, that is $\cos^2\chi = (1 - (\sqrt{2}-1)^{1/2})/2 \approx 0.178$. A numerical search over all states and directions finds nothing larger.

Proposition 5 (maximal entanglement is immune)**.** *For $\chi = \pi/4$ the shift vanishes to all orders, for every rule with $F(1-x) = 1 - F(x)$.*

Proof. Alice's weights are $\frac{1}{2}$ and Bob's conditional weights are b and $1 - b$, so $P_B(\text{Alice first}) = \frac{1}{2}[F(b) + F(1-b)] = \frac{1}{2} = F(\frac{1}{2}) = P_B(\text{Bob first})$. □

Bell tests with maximally entangled states are therefore blind to this effect; it needs unbalanced entanglement.

When the order shows. The shift exists only when Alice's decision precedes Bob's in $\prec$. If $\prec$ is laid along the rest frame of the cosmic microwave background, through which the Sun moves at 369.82 km/s toward galactic $(l, b) = (264.021°, 48.253°)$ [27] (right ascension 11.20 h, declination $-6.94°$), the laboratory timing offset at which the shift switches on is

$$\delta t^* = -\frac{\mathbf{v}\cdot(\mathbf{x}_A - \mathbf{x}_B)}{c^2}, \qquad |\delta t^*| \le \beta\frac{L}{c} \approx 41 \text{ ns per 10 km of baseline.} \tag{20}$$

For an east–west baseline the projection turns with sidereal time at 99% of full amplitude, because the dipole lies near the celestial equator, so the switch-on offset traces a sinusoid with a period of one sidereal day and a phase fixed by the dipole's right ascension. When the two detections are separated in the laboratory by more than $\beta L/c$, the same station decides first at every sidereal time, and the prediction becomes a fixed, one-directional signal: the later station's statistics depend on the earlier station's setting, never the reverse. Every shift runs forward in $\prec$, so no causal loop can form (Section 12).

11 Confronting data

All upper limits on κ below are at 95% confidence and use the constraint $\kappa \ge 0$.

11.1 Loophole-free Bell tests

The photonic loophole-free tests use unbalanced states to tolerate detector inefficiency [13]: $(|VH\rangle - 2.9|HV\rangle)/\sqrt{9.41}$ in Vienna [16], $0.961|HH\rangle + 0.276|VV\rangle$ at NIST in 2015 [30], and $0.982|HH\rangle + 0.191|VV\rangle$ in the NIST experiment of 2017 [10]. On each pulse a pair is present with small probability μ, so the state reaching the detectors is mostly vacuum, and each station's decision is click or no click. Under Axiom 3, nothing records whether a pair was present before the photons reach the

Table 6: Predicted relative change of the later station's click rate between the earlier station's two settings, per unit κ, for the Gaussian layer at the published states, angles and efficiencies.

experiment	later station	at its first setting	at its second setting
Vienna 2015	Bob ($-6.5°$, $25.5°$)	0.362	0.073
Vienna 2015	Alice ($94.4°$, $62.4°$)	0.323	0.069
NIST 2015	Bob ($-4.2°$, $25.9°$)	0.469	0.047
NIST 2015	Alice ($4.2°$, $-25.9°$)	0.478	0.047
NIST 2017	Bob ($3.7°$, $-23.6°$)	0.65	0.025
NIST 2017	Alice ($-3.7°$, $23.6°$)	0.65	0.025

detectors, so that branch is decided there too. Applying (18) to these weights (Appendix B) gives, to leading order in μ, the relative change of the later station's click rate between the earlier station's settings x_1 and x_2:

$$r(y) = \kappa \, \frac{\eta_A}{\eta_B \, P_B(y)} \Big[P_A(x_1) \, \phi\big(\eta_B q(x_1, y)\big) - P_A(x_2) \, \phi\big(\eta_B q(x_2, y)\big) \Big], \qquad \phi = \tfrac{4}{3} g, \tag{21}$$

where P_A and P_B are single-photon pass probabilities, $q(x, y)$ is the probability that Bob's photon passes given that Alice's did, and η_A, η_B are the system efficiencies. Alice's own distortion cancels against Bob's unheralded term at this order, because g is linear near zero; what remains is Bob's detector acting on heralded weights, the term of Corollary 1. The effect is large where Bob's heralded click probability lies on opposite sides of $\frac{1}{2}$ for Alice's two settings. Table 6 lists the predicted relative shifts.

Vienna 2015. The per-setting counts tabulated in a reanalysis of the public data [7] agree across the other station's settings to within 1.3 standard deviations, at 0.2–0.3% precision. Fitting the predicted pattern gives $\kappa = 0.010 \pm 0.008$ if Alice decides first and -0.003 ± 0.009 if Bob does, so $\kappa < 0.03$ for either order, and also for any mixture of the two.

NIST 2015 and 2017. The raw data of both experiments are public [23, 24]. Tables 8 and 9 in Appendix B list the per-setting singles for each run, and Table 7 and Figure 3 give the fits. The published geometry does not fix which station's detection comes first in the laboratory, so both orders are fitted: if Alice's station decides first, Bob's click rate depends on Alice's setting, and if Bob's does, the reverse. Both fits are null. Combining the two experiments gives $\kappa = -0.0006 \pm 0.0045$ if Alice's station decides first and $\kappa = +0.0065 \pm 0.0045$ if Bob's does, so $\kappa < 0.008$ or $\kappa < 0.014$ for the Gaussian layer. If the order varied during the runs, the two fits measure complementary fractions of κ, and their sum, 0.0059 ± 0.0063, gives $\kappa < 0.017$ for any mixture of the two orders. With the telegraph distortion (12) the bounds are $\kappa < 0.002$ for either order and $\kappa < 0.0024$ for any mixture. If the pair-or-vacuum branch were decided before detection rather than at it, the larger predicted shift would move to the other setting and change sign; the fits remain null, with $\kappa < 0.006$, 0.013 and 0.013 in the three cases.

11.2 The fastest decisions

Every effect above is set by $\kappa = c^2\tau$, so fast decisions give the sharpest limits on the memory. The Z boson has width 2495.2 ± 2.3 MeV and lifetime 2.6×10^{-25} s, and its decay channel leaks irreversibly at hadronization, within about 10^{-24} to 10^{-23} s. LEP measured the ratio of hadronic to leptonic widths as $R_\ell = 20.767 \pm 0.025$, a precision of 0.12% [4]. Evaluating Proposition 3 by

Table 7: Fits of the predicted pattern (Gaussian layer, Table 6, efficiency 0.75 at both stations) to the public NIST data, for each order of the two stations. "Alice first" fits Bob's click rate against Alice's setting; "Bob first" fits Alice's rate against Bob's setting. Two of the nine 2015 runs had the Pockels-cell window shifted by ±200 ns.

data	κ, Alice first	κ, Bob first
NIST 2015, seven standard runs	$+0.0064 \pm 0.0091$	$+0.0109 \pm 0.0091$
NIST 2015, all nine runs	$+0.0079 \pm 0.0079$	$+0.0150 \pm 0.0078$
NIST 2017, five data sets	-0.0047 ± 0.0054	$+0.0024 \pm 0.0055$
combined (2015 all runs and 2017)	-0.0006 ± 0.0045	$+0.0065 \pm 0.0045$

Feynman–Kac, a decision among five classes (hadrons, neutrinos and the three charged-lepton flavours) moves R_ℓ by 1.00κ in relative terms, and a decision among the eleven quark and lepton channels by 0.22κ. Against twice the measured error this gives $\kappa < 2.4 \times 10^{-3}$ or $\kappa < 1.1 \times 10^{-2}$, about twice that if part of a shift is absorbed into the fitted strong coupling. With decisions at hadronization ($t_{\mathrm{leak}} \lesssim 10^{-23}$ s), τ is then below about 2×10^{-26} s for decay classes and 10^{-25} s for channels. Splitting channels further into independently decided final states changes these numbers (Table 4). Valentini and Varma have proposed testing Born's rule, in the quantum-nonequilibrium framework, for particles with lifetimes near 10^{-25} s, among them τ leptons from Z decays [36].

11.3 Decision speed and memory

If a decision is made when the first bit of which-branch information leaks irreversibly, then $c^2 \approx 1/t_{\mathrm{leak}}$ and $\kappa \approx \tau/t_{\mathrm{leak}}$. Photodetection leaks in roughly 10^{-13} to 10^{-12} s, so the Bell-test bounds give $\tau \lesssim 10^{-15}$ to 2×10^{-14} s. A memory as short as the Planck time, 5.4×10^{-44} s, gives κ below 10^{-30} in a photodetector, which no experiment will reach. If the memory is the same for every kind of decision, the tightest bound applies to all of them. With decay decisions at hadronization, $\tau < 10^{-25}$ s gives $\kappa < 10^{-12}$ in a photodetector and $\kappa < 10^{-17}$ for superconducting readout, where which-state information leaks over 10^{-8} to 10^{-7} s; the photodetection bound alone gives $\kappa \lesssim 10^{-6}$ there. A table-top signal therefore needs a memory specific to the kind of decision, or decays that are not decided at hadronization. The direct test of Section 11.5 bounds readout decisions without either assumption. Pulsed decisions (Proposition 1) give order-one effects once a pulse of irreversible leakage is shorter than τ, which for the memories still allowed requires attosecond switching.

11.4 The cost of computing with a deformed rule

A sharpened Born rule is computationally powerful: a rule $|\psi|^p$ with $p \neq 2$ lets a quantum computer solve PP-complete problems in polynomial time [1], and nonlinear quantum gates solve NP-complete problems [2]. Here the power has a price that grows exponentially.

Proposition 6 (energy cost of reading a bias)**.** *Suppose a bias ε in a weight $\frac{1}{2}+\varepsilon$ is to be estimated from n decisions on independently prepared copies, the i-th a pulsed decision of length t_i, followed by classical processing. Each outcome carries Fisher information about ε below $\frac{16}{\pi}(1 + 2\tau/t_i)$. If each decision must also produce orthogonal records within its pulse, the Mandelstam–Tamm bound [20] gives $t_i \geq \pi\hbar/(2\Delta E_i)$, and any locally unbiased estimator with standard deviation ε needs*

$$n + \frac{4\tau}{\pi\hbar} \sum_i \Delta E_i \;\geq\; \frac{\pi}{16\,\varepsilon^2}, \tag{22}$$

so either $n \geq \pi/(32\varepsilon^2)$ decisions or a total energy spread $\sum_i \Delta E_i \geq \pi^2\hbar/(128\,\tau\varepsilon^2)$.

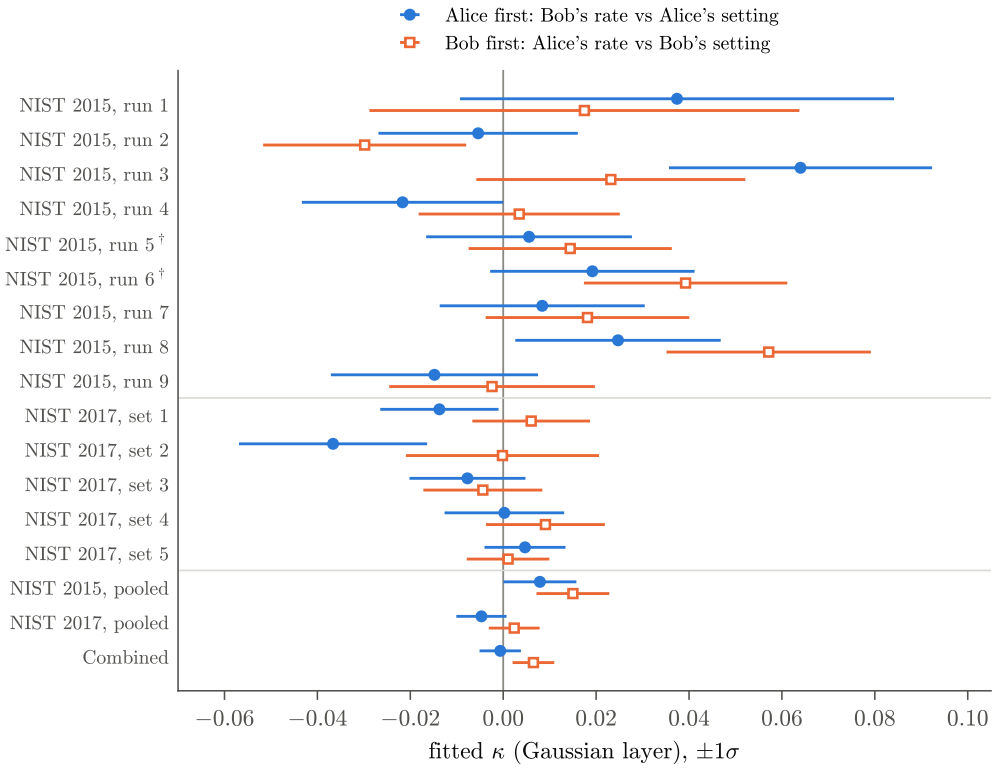

Figure 3: Fitted κ for each run of the NIST loophole-free tests of 2015 and 2017, for each order of the two stations. Runs marked † had the Pockels-cell window shifted by ±200 ns.

Proof. The derivative of $I_x(\lambda,\lambda)$ at $x=\frac{1}{2}$ is $A(\lambda)=(2/\sqrt{\pi})\,\Gamma(\lambda+\frac{1}{2})/\Gamma(\lambda)$, which is below $2\sqrt{\lambda/\pi}$ by Wendel's inequality [37], so one outcome's Fisher information at $\varepsilon=0$ is $4A^2<16\lambda/\pi$. Also $\lambda(u)=u/(u-1+e^{-u})\le 1+2/u$ for $u=t/\tau$, because $u-2+(u+2)e^{-u}$ vanishes at $u=0$ and increases. Fisher information adds over independent outcomes, and the Cramér–Rao bound requires the total to reach $1/\varepsilon^2$. □

For $\varepsilon=2^{-m}$ both alternatives grow as 4^m. With $\tau=10^{-25}$ s, a single decision that resolves such a bias must be a pulse shorter than $(32/\pi)\,\tau\,4^{-m}\approx 10\,\tau\,4^{-m}$, which is below the Planck time for $m\ge 33$. Pulsing gives no exponential advantage over plain sampling, so the computational power of the deformed rule stays out of reach, provided decisions are the only nonlinear element and act on independent copies.

11.5 A decisive test

The cleanest test does not need spacelike separation. Quantum mechanics forbids setting-dependent marginals at any timing, while in the model the earlier of two timelike-separated decisions always comes first. Prepare $\cos\chi|00\rangle+\sin\chi|11\rangle$ with $\cos^2\chi=0.178$, measure one qubit along a direction in the xz-plane at polar angle 130.1° or 49.9°, and afterwards measure the other along x. The probability that the second qubit gives "+" then differs between the two settings of the first by $+0.3235\,\kappa$, times the readout fidelity of the second qubit, while Born's rule predicts zero. Three features separate the effect from ordinary crosstalk: it must move to the other qubit when the measurement order is reversed; it must vanish for a product state; and it must vanish for a maximally

entangled state (Proposition 5). The test can be run with two photodetectors at different path lengths, or with mid-circuit measurements on trapped-ion or superconducting qubits, where it probes a decision of a different physical kind. Its sensitivity is about $\kappa \sim 10^{-3}$ with 10^8 trials and 10^{-4} with 10^{10}, which bright photon-pair sources collect in a day.

A table-top version. With polarization-entangled photons the test needs only standard parts: a Sagnac source of $\cos\chi|HH\rangle + \sin\chi|VV\rangle$, the sender's polarizer at 65° or 25° from horizontal (Bloch polar angles 130.1° and 49.9°), switched at random, and the receiver's polarizing beam splitter at 45°, all rotated in the same sense. Placing the receiver's beam splitter before a fibre delay of about 50 ns in each output fixes the order in every frame, because the delay exceeds the light time between the stations; the fibre then carries which-path rather than polarization information, so polarization drift in it does not matter. Moving the delay to the sender's side reverses the order with nothing else changed. The quantity to record is the receiver's raw click ratio, sorted by the sender's setting, rather than coincidences. Born's rule fixes that ratio whatever happens at the sender, whereas a 5% efficiency mismatch between the sender's two detectors shifts a coincidence-based estimate by the equivalent of $\kappa \approx 4.5 \times 10^{-4}$. The predicted difference is $0.3235\,\eta\kappa$, where η is the fraction of the receiver's photons whose partners reach the sender's detectors, so $\sigma(\kappa) = 1/(0.3235\,\eta\sqrt{N})$ for N receiver counts. At 10^6 counts per second, $\sigma(\kappa) = 10^{-4}$ takes about three hours with $\eta = 0.3$ and half an hour with $\eta = 0.75$. Beyond the controls above, a signal must reverse sign when the state is mirrored to $\cos^2\chi = 0.822$ and grow in proportion to η.

A run on a quantum processor. The test has been run on the IBM Heron r2 processor `ibm_marrakesh`, with mid-circuit dispersive readout of qubits 5 and 6, the measurement order set by the circuit, and both orders used: 1 760 000 shots in eleven jobs. Dividing the difference in the later qubit's probability of "+" between the earlier qubit's two settings by 0.3235 times the later qubit's measured readout visibility, about 0.99, and combining the two orders gives $\kappa = -0.0014 \pm 0.0029$, so $\kappa < 0.0048$. The product-state and maximally entangled controls are consistent with zero (largest $|z| = 1.03$), and the scatter between jobs and between repeated copies of each circuit matches shot noise. The later qubit's conditional state keeps about 95% fidelity with the ideal one across the readout window. Any loss of its polarization lowers the predicted signal, because the two terms of (18) scale as its first and third powers; allowing for a loss of up to 10%, the bound is $\kappa < 0.006$ for this kind of decision.

12 Order structure

Because decided events form one strict order, no recorded fact can travel around a closed timelike curve, and inside a black hole's universal horizon every fact comes after every fact outside.

Proposition 7 (no facts on time loops)**.** *Under Axiom 4, no decision can lie on a closed causal chain that carries its own record.*

Proof. A closed timelike curve that carries the record of decision A back to A's apparatus makes A's setup depend on A's outcome, so $A \prec A$. A longer cycle gives the same by transitivity, and a strict order has no cycles. □

Anything that leaves records, including every object that gives off heat, therefore cannot traverse a closed timelike curve; only systems that make no decisions could. For record-bearing matter, chronology protection becomes a consequence of the decision order.

Proposition 8 (universal horizons)**.** *In gravity theories with a preferred foliation, stationary black holes have a universal horizon inside the event horizon, a leaf of the foliation that traps signals of any speed [6, 8]. If $\prec$ is laid along that foliation in an eternal stationary black hole, every decision inside the universal horizon follows every decision outside.*

Proof. The Killing flow maps leaves to leaves and shifts the preferred time T by a constant, so $\chi \cdot \nabla T = 1$ in the exterior. The universal horizon is the leaf to which χ is tangent, where $\chi \cdot \nabla T$ would vanish, so it carries no finite value of T: $T \to +\infty$ on approach from outside [9]. Every real value of T is taken by some exterior leaf, because the flow shifts T by any amount. T increases continuously along every future causal curve, so a curve from an exterior point reaches the universal horizon only after crossing every exterior leaf with larger T, and causal curves cross it inward only. An order laid along the leaves that respects causal dependence (Axiom 4) must therefore put every leaf inside after every leaf outside. □

The history of facts then has an order type longer than the integers: the outside history is an unbounded increasing sequence, and the inside facts all come after it.

13 Discussion

If the bridge holds and the layer is fair, time's order exists but can never be seen; if the layer has memory, the order leaks out exactly where decisions are fast, and it does so as a one-directional signal fixed by the order of detections.

Relation to prior work. Born's rule from a martingale in collapse models is Pearle's gambler's-ruin mechanism [25]. That nonlinearity or Born-rule deviations allow signalling is due to Gisin [17] and Valentini [34], whose signal-locality theorem covers deterministic hidden variables; Valentini also tied such signalling to absolute simultaneity [35]. Han and Choi derived Born's rule from relativistic causality [18], Torres Alegre derived it from no-signalling under steering for rules $P = \Phi(\tau)$ of the transition probability, in generalized probabilistic theories and in infinite dimensions [32, 33], and Galley and Masanes showed that every alternative measurement postulate violates purification and local tomography [15]. Bancal and coauthors showed that hidden influences of any finite speed permit signalling [5], and Salart and coauthors bounded their speed in any preferred frame [29]. Wiseman and Gambetta showed that non-Markovian quantum trajectories lack a continuous-measurement interpretation [38]. Adler and Bassi built collapse models with non-white noise and the norm-weighted probability rule [3], and Mukherjee and van Wezel obtained Born's rule in colored-noise state reduction from a fluctuation–dissipation relation in the white-noise limit [21]. Nelson's stochastic mechanics keeps Born's rule exact with a drift given by Einstein's relation [22]. Aaronson and Abrams–Lloyd showed what deformed Born rules and nonlinear gates would compute [1, 2], and Childs and Young studied state discrimination and search under the Gross–Pitaevskii nonlinearity [11]. Peres argued that nonlinear variants of quantum mechanics violate the second law [26]; Weinberg and, recently, Fedida and Kent dispute the entropy that argument assumes [14]. Bednorz reanalysed the public data of the 2015 loophole-free tests for setting dependence of the marginals [7].

New here are the reading of Born's rule as invariance under the choice of stopping rule; the time-ordered form of the no-signalling equivalence, in which only the later party's statistics shift, with the remote-decision detector and its exact maximum; the quantitative floor for smooth layers; the exact crossover for a two-state layer; the pulsed-decision law and its energy cost; the many-outcome first-order law, its dependence on splitting, and the certainty threshold for any number of outcomes;

the stopping-time inconsistency of exact colored-noise collapse models; the closed-form signalling optimum and the immunity of maximal entanglement; and bounds on κ from public Bell-test data, from a direct order-swap test on a superconducting processor, and from Z-boson decays.

Limitations. Everything rests on the bridge, and while the layer is fair the bridge cannot be tested: by Theorem 2, a fair layer gives exactly the statistics of quantum mechanics. The smooth-layer results come from models driven by Ornstein–Uhlenbeck or telegraph noise, with one log-weight per outcome and independent noise for each branch, and the coefficients depend on the noise law. Which branches count as the outcomes of one decision is not fixed by the framework, and the distortion of coarse outcomes depends on it. Identifying the decision rate with the leakage rate is an assumption, and when a real detector decides is not pinned down. Propositions 1 and 3 rest on limit arguments whose remainders are not controlled here, and Proposition 2 is argued rather than proved. The universal-horizon result assumes that $\prec$ follows the preferred foliation of those gravity theories.

Outlook. The likeliest world has a fair layer, or a memory too short ever to matter. There the lasting content is structural: Born's rule is the price of a hidden time. If the memory is long enough, the places to look are fast decisions, where pulsed which-path leakage and the fastest decays test it, and unbalanced entanglement, where the order-reversal test of Section 11.5 would show a one-directional signal that ordinary physics cannot produce.

A Derivation of the many-outcome law

Measure time in units of $1/(2c^2)$, so that $\dot y_i = s_i + p_i$ with independent Ornstein–Uhlenbeck pushes of correlation time $\tau' = 2\kappa$. Write $s_i = \zeta_i/\sqrt{\tau'}$, where ζ_i is a unit-rate Ornstein–Uhlenbeck process of stationary variance $\frac12$ with generator $L_{\mathrm{OU}} = -\zeta\partial_\zeta + \frac12\partial_\zeta^2$ in time units of τ'. The probability $u(y,\zeta)$ that outcome k wins solves $\mathcal{L}u = 0$ with

$$\mathcal{L} = \varepsilon^{-2}\sum_i L_{\mathrm{OU}}^{(i)} + \varepsilon^{-1}\sum_i \zeta_i\partial_i + \sum_i p_i\partial_i, \qquad \varepsilon = \sqrt{\tau'}, \tag{23}$$

where $\partial_i = \partial/\partial y_i$. Expanding $u = u_0 + \varepsilon u_1 + \varepsilon^2 u_2 + \dots$, the orders ε^{-2} and ε^{-1} give $u_0 = u_0(y)$ and $u_1 = \sum_i \zeta_i\partial_i u_0$. Solvability at order ε^0 gives $\frac12\Delta u_0 + \sum_i p_i\partial_i u_0 = 0$, solved with the right boundary values by $u_0 = p_k$, because $\partial_i p_k = p_k(\delta_{ik} - p_i)$ makes p_k a martingale of the fair collapse. Then $u_2 = \frac12\sum_{ij}(\zeta_i\zeta_j - \frac12\delta_{ij})\partial_i\partial_j u_0 + v$, and solvability at order ε^2 gives

$$\tfrac12\Delta v + \sum_m p_m\partial_m v = -\tfrac14\Delta^2 p_k - \tfrac12\sum_j \partial_j\Big(\sum_m p_m\partial_j\partial_m p_k\Big). \tag{24}$$

Evaluating the right side with $\partial_i p_k = p_k(\delta_{ik} - p_i)$ and $\partial_i S = 2p_i(p_i - S)$ gives $p_k\big[-2S^2 + 2p_kS - \frac52 p_k^2 + \frac52 S_3 - \frac12 S + \frac12 p_k\big]$. Averaged over the stationary noise, $P_k = p_k + \tau' v + O(\tau'^2)$, so $V = 2v$; and in time units of $1/c^2$ the generator is $\mathcal{M}_0 = 2\big(\frac12\Delta + p\cdot\nabla\big)$, which in the weights is the operator of Section 8. This gives (16), and Dynkin's formula gives the time-integral form. Subtracting $\mathcal{M}_0[\frac43 p_k(p_k - S)]$ leaves the source of δV_k.

For two outcomes, $\mathcal{M}_0 = \frac12(1-z^2)^2\,d^2/dz^2$ and the source is $-z(1-z^2)^2$, so $V'' = -2z$ and $V = z(1-z^2)/3 = \frac43 x(1-x)(2x-1)$. For a single outcome of weight x against a field of negligible branches, the field's total weight evolves deterministically, $S \to x^2$ and $S_3 \to x^3$, $\mathcal{M}_0 \to x^2(1-x)^2 d^2/dx^2$, and the source becomes $-4x^2(1-x)^2(2x-\frac12)$; integrating with $V(0) = V(1) = 0$ gives $V = \frac13 x(1-x)(4x+1)$.

B Click-or-no-click decisions and the Bell-test data

On a pulse that carries a pair with probability μ, Alice's click weight is $w = \mu\eta_A P_A(x)$, Bob's weight given her click is $b_1 = \eta_B q(x, y)$, his unconditional weight is $B_0 = \mu\eta_B P_B(y)$, and his weight given no click is $b_2 = (B_0 - w b_1)/(1 - w)$. Inserting these into (18) with $k\,g = \kappa\phi$ and expanding in μ, the terms $\phi(w)(b_1 - b_2)$ and $(1-w)\phi(b_2) - \phi(B_0)$ cancel at order μ because $\phi(x) = -\frac{4}{3}x + O(x^2)$, leaving $\Delta P_B = \kappa\, w\, \phi(b_1) + O(\kappa\mu^2)$. Dividing by B_0 and taking the difference between Alice's settings gives (21). For the telegraph distortion h, which is not linear near zero, the full expression is evaluated numerically at the per-pulse pair probability.

Tables 8 and 9 list the per-setting trials N and the numbers of Alice's and Bob's clicks used in Table 7: for 2015, pulse slots 1 to 9 of the released processed data [23]; for 2017, all fourteen pulse slots of the five released data sets [24]. The settings are $a = 4.2°$, $a' = -25.9°$, $b = -4.2°$, $b' = 25.9°$ in 2015 and $a = -3.7°$, $a' = 23.6°$, $b = 3.7°$, $b' = -23.6°$ in 2017.

Table 8: NIST 2015: trials N / Alice's clicks / Bob's clicks for each setting pair. Runs 5 and 6 had the Pockels-cell window shifted by +200 and −200 ns.

run	(a, b)	(a, b')	(a', b)	(a', b')
1	10 411 791 / 4 089 / 4 032	10 416 900 / 4 026 / 14 010	10 229 839 / 12 247 / 3 896	10 231 061 / 12 558 / 13 696
2	50 928 270 / 18 036 / 18 842	50 924 566 / 18 329 / 60 608	50 912 956 / 57 416 / 18 851	50 915 670 / 57 113 / 60 952
3	26 789 179 / 10 427 / 11 004	26 775 145 / 10 305 / 34 040	26 768 995 / 32 016 / 10 655	26 776 277 / 32 021 / 34 074
4	45 674 302 / 18 592 / 18 415	45 681 090 / 18 557 / 57 771	45 641 427 / 52 823 / 18 562	45 663 036 / 52 916 / 58 102
5	44 679 233 / 18 251 / 17 687	44 686 255 / 18 190 / 58 020	44 678 267 / 50 990 / 17 688	44 659 286 / 50 321 / 57 484
6	44 767 955 / 18 349 / 17 990	44 754 487 / 18 012 / 57 429	44 739 350 / 51 887 / 17 866	44 739 769 / 51 733 / 56 850
7	44 699 512 / 18 241 / 17 864	44 712 421 / 18 033 / 57 615	44 694 021 / 51 471 / 17 772	44 695 377 / 52 015 / 57 764
8	44 469 638 / 18 243 / 17 871	44 473 907 / 17 762 / 57 066	44 463 141 / 51 259 / 17 646	44 459 075 / 51 081 / 57 133
9	45 550 482 / 17 718 / 17 422	45 544 623 / 17 788 / 55 598	45 533 203 / 51 789 / 17 517	45 532 907 / 51 265 / 55 832

Table 9: NIST 2017: trials N / Alice's clicks / Bob's clicks for each setting pair.

set	(a, b)	(a, b')	(a', b)	(a', b')
1	7 465 764 / 29 643 / 29 167	7 464 290 / 29 521 / 113 616	7 466 617 / 111 420 / 29 446	7 468 649 / 111 455 / 113 331
2	7 451 054 / 10 981 / 11 455	7 457 940 / 10 989 / 38 349	7 451 477 / 41 042 / 11 724	7 449 499 / 41 117 / 38 507
3	7 455 075 / 28 928 / 30 394	7 451 379 / 28 972 / 102 312	7 453 601 / 108 145 / 30 572	7 458 904 / 108 850 / 101 599
4	7 465 798 / 29 075 / 28 653	7 460 670 / 28 860 / 110 445	7 461 505 / 106 914 / 28 645	7 458 849 / 107 491 / 110 078
5	15 030 166 / 60 191 / 62 597	15 025 581 / 60 086 / 206 973	15 028 277 / 212 958 / 62 402	15 026 186 / 214 113 / 206 890

References

[1] S. Aaronson, Quantum computing, postselection, and probabilistic polynomial-time, *Proc. R. Soc. A* **461**, 3473 (2005), arXiv:quant-ph/0412187.

[2] D. S. Abrams and S. Lloyd, Nonlinear quantum mechanics implies polynomial-time solution for NP-complete and #P problems, *Phys. Rev. Lett.* **81**, 3992 (1998), arXiv:quant-ph/9801041.

[3] S. L. Adler and A. Bassi, Collapse models with non-white noises, *J. Phys. A* **40**, 15083 (2007), arXiv:0708.3624.

[4] ALEPH, DELPHI, L3, OPAL and SLD Collaborations, LEP Electroweak Working Group, and SLD Electroweak and Heavy Flavour Groups, Precision electroweak measurements on the Z resonance, *Phys. Rept.* **427**, 257 (2006), arXiv:hep-ex/0509008.

[5] J.-D. Bancal, S. Pironio, A. Acín, Y.-C. Liang, V. Scarani and N. Gisin, Quantum non-locality based on finite-speed causal influences leads to superluminal signalling, *Nat. Phys.* **8**, 867 (2012), arXiv:1110.3795.

[6] E. Barausse, T. Jacobson and T. P. Sotiriou, Black holes in Einstein-aether and Hořava–Lifshitz gravity, *Phys. Rev. D* **83**, 124043 (2011), arXiv:1104.2889.

[7] A. Bednorz, Analysis of assumptions of recent tests of local realism, *Phys. Rev. A* **95**, 042118 (2017), arXiv:1511.03509.

[8] J. Bhattacharyya, M. Colombo and T. P. Sotiriou, Causality and black holes in spacetimes with a preferred foliation, *Class. Quantum Grav.* **33**, 235003 (2016), arXiv:1509.01558.

[9] J. Bhattacharyya, A. Coates, M. Colombo and T. P. Sotiriou, Evolution and spherical collapse in Einstein-æther theory and Hořava gravity, *Phys. Rev. D* **93**, 064056 (2016), arXiv:1512.04899.

[10] P. Bierhorst *et al.*, Experimentally generated randomness certified by the impossibility of superluminal signals, *Nature* **556**, 223 (2018), arXiv:1702.05178.

[11] A. M. Childs and J. Young, Optimal state discrimination and unstructured search in nonlinear quantum mechanics, *Phys. Rev. A* **93**, 022314 (2016), arXiv:1507.06334.

[12] H. Dym and H. P. McKean, *Gaussian Processes, Function Theory, and the Inverse Spectral Problem* (Academic Press, 1976).

[13] P. H. Eberhard, Background level and counter efficiencies required for a loophole-free Einstein–Podolsky–Rosen experiment, *Phys. Rev. A* **47**, R747 (1993).

[14] S. Fedida and A. Kent, The thermodynamics of readout devices and semiclassical gravity (2025), doi:10.1103/bpwh-1pkf, arXiv:2509.07096.

[15] T. D. Galley and L. Masanes, Any modification of the Born rule leads to a violation of the purification and local tomography principles, *Quantum* **2**, 104 (2018), arXiv:1801.06414.

[16] M. Giustina *et al.*, Significant-loophole-free test of Bell's theorem with entangled photons, *Phys. Rev. Lett.* **115**, 250401 (2015), arXiv:1511.03190.

[17] N. Gisin, Weinberg's non-linear quantum mechanics and supraluminal communications, *Phys. Lett. A* **143**, 1 (1990).

[18] Y. D. Han and T. Choi, Quantum probability assignment limited by relativistic causality, *Sci. Rep.* **6**, 22986 (2016), arXiv:1307.2026.

[19] L. P. Hughston, R. Jozsa and W. K. Wootters, A complete classification of quantum ensembles having a given density matrix, *Phys. Lett. A* **183**, 14 (1993).

[20] L. Mandelstam and I. Tamm, The uncertainty relation between energy and time in non-relativistic quantum mechanics, *J. Phys. USSR* **9**, 249 (1945).

[21] A. Mukherjee and J. van Wezel, Colored-noise-driven unitarity violation causing dynamical quantum state reduction, *Phys. Rev. A* **109**, 032214 (2024), arXiv:2306.05849.

[22] E. Nelson, Derivation of the Schrödinger equation from Newtonian mechanics, *Phys. Rev.* **150**, 1079 (1966).

[23] L. K. Shalm, M. Stevens and S. W. Nam, Bell test research software and data, NIST (2015), `https://www.nist.gov/node/773781`.

[24] L. K. Shalm, Experimentally generated randomness certified by the impossibility of superluminal signals (data), NIST (2017), `https://data.nist.gov/od/id/63B8D44FD72FCC6BE0531A57068170351888`, doi:10.18434/T4/1425490.

[25] P. Pearle, How stands collapse I, *J. Phys. A* **40**, 3189 (2007), arXiv:quant-ph/0611211.

[26] A. Peres, Nonlinear variants of Schrödinger's equation violate the second law of thermodynamics, *Phys. Rev. Lett.* **63**, 1114 (1989); S. Weinberg, reply, *ibid.* **63**, 1115 (1989).

[27] Planck Collaboration, Planck 2018 results. I. Overview and the cosmological legacy of Planck, *Astron. Astrophys.* **641**, A1 (2020).

[28] D. Revuz and M. Yor, *Continuous Martingales and Brownian Motion*, 3rd ed. (Springer, 1999).

[29] D. Salart, A. Baas, C. Branciard, N. Gisin and H. Zbinden, Testing the speed of 'spooky action at a distance', *Nature* **454**, 861 (2008), arXiv:0808.3316.

[30] L. K. Shalm *et al.*, Strong loophole-free test of local realism, *Phys. Rev. Lett.* **115**, 250402 (2015), arXiv:1511.03189.

[31] A. Stefanov, H. Zbinden, N. Gisin and A. Suarez, Quantum correlations with spacelike separated beam splitters in motion: experimental test of multisimultaneity, *Phys. Rev. Lett.* **88**, 120404 (2002).

[32] E. O. Torres Alegre, Causal consistency selects the Born rule: a derivation from steering in generalized probabilistic theories (2025), arXiv:2512.12636.

[33] E. O. Torres Alegre, Causal rigidity of Born-type probability rules in infinite-dimensional operational theories (2026), arXiv:2602.09056.

[34] A. Valentini, Signal-locality in hidden-variables theories, *Phys. Lett. A* **297**, 273 (2002), arXiv:quant-ph/0106098.

[35] A. Valentini, Hidden variables and the large-scale structure of space-time, arXiv:quant-ph/0504011.

[36] A. Valentini and M. Varma, Towards a test of the Born rule in high-energy collisions (2025), arXiv:2505.07510.

[37] J. G. Wendel, Note on the gamma function, *Amer. Math. Monthly* **55**, 563 (1948).

[38] H. M. Wiseman and J. M. Gambetta, Pure-state quantum trajectories for general non-Markovian systems do not exist, *Phys. Rev. Lett.* **101**, 140401 (2008), arXiv:0806.3101.

AI assistance. The author developed this work with substantial assistance from AI systems, including Anthropic's Claude, which carried out derivations, simulations and drafting. The results have not yet been independently reviewed.