FRONTIER RESEARCH COLLECTION / FR-17

The Pre-Event Coherence Hypothesis

Prospective tests with explicit timing and ordinary-predictability budgets

Architect X13

The Eidara Continuum | Science & Philosophy

Research white paper | November 2026 | Version 2.0 Revision

Abstract

Can anything about a random event be known before it happens? The Continuum's most daring hypothesis says perhaps, and this paper builds the test that could prove it wrong. The pre-event question can be made exact: can a committed signal predict a subsequently generated target beyond the ordinary information admitted by a qualified apparatus? The paper supplies a rigorous protocol for that speculative hypothesis. It separates prospective forecasting, pre-generation physiology and randomized intention instructions. Conditional fairness yields an exact adaptive-guessing theorem, and a new bounded-predictability extension preserves anytime-valid inference when ordinary bias or accessible device state allows a declared hit-rate ceiling above one half. A causal assignment identity addresses instruction effects, while timing certificates and outcome-independent trial accounting define what the observations would mean. Complete derivations, exhaustive finite checks and exact design calculations make the proposal reproducible. Positive reports are considered alongside failed replications and contrary quantum-randomness findings. No physical or participant experiment is reported. The Continuum's bold commitment is to submit an extraordinary conjecture to a test with a quantitative information budget, an independent replication requirement and a specified outcome that would count against it.

Keywords: precognition; conditional independence; quantum randomness; preregistration; sequential inference; causal identification

The scientific wager

The Continuum's bold position is that an unusual question deserves a test capable of defeating it. A compelling experience can motivate that test; it cannot decide its result. PEC-R2 commits the project to observable outcomes, a declared smallest effect of interest, and the publication of a well-measured null result.

Prologue: A test designed to lose

In the Continuum, Syraia is the pattern forecaster, the operator who expands interpretation where Ondrix imposes structure. Syraia's image is of possibilities growing quieter just before an event commits, the way a room hushes before a verdict. It is one of the most beautiful ideas in the canon and one of the most dangerous, because it points at the oldest temptation in human thought: the belief that the future leaks.

The project's first conversations included a personally striking lottery coincidence. This paper neither publishes it nor leans on it. A memorable match is exactly the kind of evidence that cannot tell chance from ability, because nobody counts the misses.

So the Continuum does something unusual with its most daring hypothesis. It builds the experiment that could kill it.

Two small traps show why that matters. Take two coin flips that have nothing to do with each other, keep only the trials in which they agreed, and you have manufactured a perfect predictor out of pure noise—no channel, no anticipation, just a quiet filter on what gets counted. Or seal every guess with an unbreakable cryptographic receipt and still be fooled, because the random bit it was compared against was generated earlier and held in a buffer. The seal proves the record was not altered; it cannot prove the order of events in the world. A protocol that hopes to detect the future must defeat both traps before it is allowed to say anything.

That is what follows: three separate questions—can a committed forecast beat the odds, does the body respond before a random event, can intention change a physical outcome—each with its own test, its own audit of time, and its own rule for what would count as failure. No experiment is reported here; the protocol is the contribution. The paper's last section says why: inquiry becomes stronger when the result has authority over the story. If the future leaves no detectable trace, this program is designed to recognize that, and to say so.

What this paper claims

Established. Within the stated models, adaptive guessing cannot extract an advantage from conditionally fair future bits; proper scores penalize unsupported confidence; a valid sequential process controls false positives under continuous monitoring. Positive reports of anticipation exist alongside failed replications and contrary quantum-randomness findings, and the paper weighs both.

Proposed here. A rigorous protocol for an explicitly speculative empirical hypothesis: three separate tracks (prospective forecasting, pre-generation physiology, randomized intention instructions), an exact adaptive-guessing theorem, a bounded-predictability extension that keeps inference anytime-valid above a declared ordinary-information ceiling, a causal assignment identity, timing certificates and outcome-independent trial accounting. The mathematical results certify the stated inference procedures.

Frontier conjecture. A carefully isolated system might exhibit a repeatable unexplained signal of a registered magnitude, exceeding both sampling uncertainty and the ordinary predictability its apparatus admits. Repeated independent failure would justify ending this research line.

Not claimed. Physical precognition, lottery forecasting and intention-controlled randomness: this paper establishes none of them, and no physical or participant experiment is reported. Even a positive finding would first establish an unexplained association, not retrocausality, consciousness-induced collapse or a new force. "Coherence" is not used in its quantum-mechanical sense. No lottery predictions, investment advice, or assurance that any participant possesses exceptional powers.

1. From an arresting metaphor to a measurable claim

In the Continuum, Syraia, the pattern forecaster, offers an image of possibilities becoming quieter before an event commits. That image is the destination of this paper and a philosophical invitation, not yet a physical theory. A physical theory would require a measurable quantity, its units, a law governing its change, and a prediction that differs from ordinary dynamics. None follows merely from the word "coherence." PEC-R2 therefore uses that word in the title of a hypothesis, without assigning it the technical meaning of phase coherence in quantum mechanics.

The originating project discussion included a personally striking lottery coincidence. No documentary reconstruction of that private episode is used as evidence here. The scientific problem survives without publishing personal details: humans encounter apparent anticipation, and need a method that separates information, selection, ordinary skill, and chance. The generic lottery calculation on page 19 illustrates how a remarkable match can occur under an ordinary null.

Expert intuition is a relevant comparison, but it is a different target. Kahneman and Klein identify environmental regularity and an opportunity to learn it as conditions for reliable intuitive expertise [1]. A firefighter's rapid judgment can exploit patterns that the firefighter cannot fully articulate. Correctly anticipating an independent future random bit would require a different explanation. The distinction concerns the available information, not how vivid or sincere the experience feels.

We propose three separate research tracks. Track A asks whether a committed forecast beats a declared predictive null. Track B asks whether a preregistered physiological feature contributes predictive information beyond available history. Track C asks whether randomly assigned intention instructions change physical outcomes. A positive result in one track does not entail a positive result in another. A participant might notice an ordinary device precursor without influencing anything; an instruction might alter device temperature through movement without creating precognition.

The central research commitment is prospective: specify what would count as success before seeing the outcome. The strongest conceivable finding would be a reproducible residual after known information pathways, analysis flexibility, and equipment faults have been challenged. Even that would initially establish an unexplained association or intervention effect. It would not identify retrocausality, consciousness-induced collapse, or a new force. Naming a mechanism requires additional discriminating experiments.

2. Positive claims and the closest prior art

The proposed investigation enters an existing literature. Bem reported nine experiments involving more than a thousand participants, using reversed versions of psychological tasks to investigate apparent future-to-present effects [2]. The paper is a positive empirical claim and an appropriate source for the protocols it actually studied. It is not independent verification of the Continuum's interpretation, and results from different reversed tasks cannot simply be pooled as one interchangeable physical phenomenon.

Mossbridge, Tressoldi, and Utts synthesized 26 reports of physiological differences before apparently unpredictable stimuli [3]. Their reported positive aggregate result makes anticipatory physiology a concrete prior-art comparison. It does not establish that the relevant signal preceded the physical generation of every target, nor that every experiment excluded all dependence on earlier trials. Meta-analysis aggregates the inferential quality of its inputs; it cannot retrospectively create timing records or eliminate an unmeasured shared bias.

The distinction between pre-stimulus and pre-generation is central to PEC-R2. If a computer chooses a target and then waits before displaying it, a signal in that interval is early relative to display but late relative to selection. A participant could respond to an ordinary cue, or a data pipeline could accidentally include post-selection samples. The proposed protocol makes physical generation a separately audited event and includes delayed-display conditions to expose this ambiguity.

The strongest proponent objection deserves a direct response: a weak, context-sensitive phenomenon could disappear in an insensitive replication. That possibility is logically coherent. It becomes scientifically informative only when the relevant context, effect direction, and analysis are fixed prospectively. Changing from positive guessing to negative guessing, from all participants to a selected subgroup, or from one physiological window to another after examining the data creates a new exploratory hypothesis. Its confirmation requires fresh data.

PEC-R2 contributes an explicit prospective test contract: commitment timing, distinct causal and predictive estimands, a quantified ordinary-predictability ceiling, valid sequential inference and independent controls enter the same auditable design. Preregistration, randomization, proper scoring, and independent replication were not invented here. The burden of a new empirical claim is to survive those established methods, including the methods used by researchers whose substantive expectations differ from ours.

3. Contrary evidence changes the starting position

Ritchie, Wiseman, and French reported three preregistered unsuccessful replications of retroactive facilitation of recall, with 150 participants combined [4]. Galak and colleagues reported seven unsuccessful experiments with 3,289 participants [5]. These studies address particular Bem-style recall claims. They are substantial counterevidence to those claims, rather than a universal experiment covering every possible anticipatory phenomenon.

The Transparent Psi Project's multilaboratory replication reported 49.89% correct guesses against a 50% chance level [6]. Its correction documents unimplemented external code tracking, overwritten access logs, relaxed auditor-independence criteria, some restarted sessions, and late exclusion of test-account records. The authors report that these deviations did not change the null conclusion. They nevertheless distinguish intended safeguards from verified delivery. PEC-R2 requires contemporaneous evidence of protocol fidelity, including independent checks of the checks.

Walleczek and colleagues subsequently reported three studies involving 26,483 participants and 420,472 critical trials [7]. The original positive effect did not replicate. An exploratory result in the opposite direction was confirmed in the second study but did not replicate in the third. This sequence is especially instructive: a significant result can occur inside an unusually transparent research program and still fail the next prospective test. Its authors regard a methodological false positive as more plausible under conventional standards than a novel physical explanation.

For the separate intention question, Maier, Dechamps, and Pflitsch studied 12,571 participants using quantum-generated stimulus selection. Their primary Bayesian analysis favored the null, with reported $BF_{01} = 10.07$ [8]. Here BF_{01} is the data's marginal likelihood under the specified null divided by that under the specified alternative; it is not a posterior probability without prior model odds. Exploratory temporal patterns discussed in the same paper do not reverse the primary result.

The present evidence does not justify presenting precognition or micro-psychokinesis as established. Nor should a skeptical conclusion erase the difference between protocols. Our proposal earns its place only if it offers a sharper question, stronger measurement, and a willingness to terminate an unsuccessful line. An unsuccessful, adequately precise test is a research output, not a defect to be hidden by an expanded metaphor.

4. Quantum measurement and the proposed observable

Quantum mechanics motivates careful questions about measurement, but its terminology is frequently overextended. Decoherence describes the suppression of locally observable interference through interaction with an environment. It does not by itself supply a special role for human desire or solve every interpretive issue concerning definite outcomes [9]. PEC-R2 contains no additional quantum dynamics.

A small calculation makes the missing inference visible. Let the system space be $\mathcal{H}_S = \mathbb{C}^2$, with orthonormal vectors $|0\rangle, |1\rangle$, and let $\mathcal{H}_E$ be an environment Hilbert space. For normalized environment vectors $|e_0\rangle, |e_1\rangle$ and complex coefficients a, b satisfying $|a|^2 + |b|^2 = 1$, consider

$$|\Psi\rangle = a|0\rangle|e_0\rangle + b|1\rangle|e_1\rangle.$$

The tensor product between system and environment is implicit. Tracing the rank-one density operator over $\mathcal{H}_E$ gives

$$\rho_S = \mathrm{Tr}_E(|\Psi\rangle\langle\Psi|) = \begin{pmatrix} |a|^2 & ab^*\gamma \\ a^*b\gamma^* & |b|^2 \end{pmatrix}, \qquad \gamma = \langle e_1|e_0\rangle.$$

Here $*$ denotes complex conjugation. All entries and γ are dimensionless. As $|\gamma|$ decreases, the off-diagonal magnitude decreases; the diagonal probabilities remain $|a|^2$ and $|b|^2$. In particular, $a = b = 1/\sqrt{2}$ retains equal outcome probabilities even when $\gamma = 0$. No "quieting" variable in this calculation reveals which result a future measurement will produce.

Bell-based randomness research provides a stronger methodological connection. Pironio and colleagues developed certification from Bell violations [10]; Bierhorst and colleagues demonstrated extracted randomness under stated physical and independence assumptions [11]. Certification is conditional on those assumptions and an implemented protocol. It is not a declaration that a device can have no engineering fault or that every conceivable hidden variable has been inspected.

A future exotic model would need to specify an observable coupling that changes these predictions and survives existing constraints. Calling a statistical deviation "quantum" before that work is done adds no explanatory content. Classical hidden precursors, electromagnetic coupling, timing artifacts, and imperfect randomization remain explicit alternatives even when the apparatus contains a quantum component.

5. PEC-R2: types, clocks, and null hypotheses

A Guide's explanation

The mathematics begins by defining the enemy: every ordinary way a guess could beat chance without seeing the future—clocks, buffers, device quirks, analysis choices, the trials nobody counted. Each is given a name and a place in the model, so that any signal that survives has first had every mundane explanation stripped away.

Work on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The trial index i lies in $\{1, \dots, N\}$ for fixed positive integer N, or among positive integers for a declared sequential design. All probabilities, scores, and accuracies are dimensionless. Physical times are in seconds relative to a synchronized laboratory clock; a software timestamp alone is not evidence of a physical event time.

Symbol	Type and operational meaning
$G_i \in \{0,1\}$	Committed binary guess; immutable before generation.
$Y_i \in \{0,1\}$	Audited target bit from the declared physical process.
$W_i = \mathbf{1}\{G_i = Y_i\}$	Success indicator; $C_n = \sum_{i=1}^{n} W_i$ counts hits.
$q_i \in [0,1]$	Optional committed forecast of $\mathbb{P}(Y_i = 1)$.
$\mathcal{H}_i \subseteq \mathcal{F}$	Information available immediately before target generation, including the current commitment and all past records.
$Z_i \in \mathcal{Z}$	Declared observed covariates, on a specified measurable space $\mathcal{Z}$.
$S_i \in \mathbb{R}^d$	Frozen pre-generation physiological feature vector; each coordinate has declared sensor units or a documented normalization.
$A_i \in \{0,1\}$	Randomized instruction assignment for the intention study.

The strict predictive null is

$$H_A : \quad \mathbb{P}(Y_i = 1 \mid \mathcal{H}_i) = \tfrac{1}{2} \quad \text{almost surely for every } i.$$

The information sets are nested across trials and contain all preceding outcomes. Crucially, fairness is conditional on the current guess and accessible device state, not merely a marginal 50/50 frequency. A balanced sequence can be perfectly predictable.

Let $t_{c,i}$, $t_{g,i}$, and $t_{d,i}$ denote commitment, physical generation, and display times. Record their uncertainty bounds $u_{c,i}, u_{g,i} \geq 0$. A valid pre-generation commitment requires $t_{c,i} + u_{c,i} < t_{g,i} - u_{g,i}$. Features require the analogous inequality for their last contributing sample, including filter support. A strict gap can be measured only within a declared clock and latency model.

Track B tests conditional predictive independence of S_i and Y_i given an admissible Z_i. Track C tests a causal instruction effect. Neither is automatically equivalent to H_A: an incomplete covariate set, hidden pathway, or assignment-dependent apparatus can violate a statistical null without violating ordinary physics.

6. Exact impossibility under conditional fairness

A Guide's explanation

This proposition is the floor beneath everything else. If the future bits are genuinely fair given everything that came before, no strategy—however adaptive, however much it learns from past trials—can beat chance on average. A real effect, if one existed, would have to appear as a breach of that fairness the apparatus could not have produced by ordinary means.

Proposition 1 (Adaptive predictions do not beat a fair future)**.** *Assume H_A and that each G_i is $\mathcal{H}_i$-measurable. Then the success variables (W_i) are independent Bernoulli variables with parameter $1/2$. Consequently $C_n \sim \text{Binomial}(n, 1/2)$ for every fixed positive integer n.*

Proof. Because the guess is known at commitment time,

$$\mathbb{P}(W_i = 1 \mid \mathcal{H}_i) = \mathbf{1}\{G_i = 1\}\tfrac{1}{2} + \mathbf{1}\{G_i = 0\}\tfrac{1}{2} = \tfrac{1}{2}.$$

The same conditional probability holds for $W_i = 0$. For a fixed binary sequence $(w_1, \dots, w_n)$, the event specifying its first $n-1$ entries belongs to $\mathcal{H}_n$. Conditional expectation therefore gives

$$\mathbb{P}(W_1 = w_1, \dots, W_n = w_n) = \tfrac{1}{2}\mathbb{P}(W_1 = w_1, \dots, W_{n-1} = w_{n-1}).$$

Repeated application yields 2^{-n}, which is the product distribution of independent fair bits. Counting the $\binom{n}{k}$ sequences with k successes proves the binomial statement. □

The theorem allows a predictor to change its strategy, study past outcomes, or use its own randomization, provided all information used is included in $\mathcal{H}_i$. It does not require all guesses to be fixed at the start of a session. This strengthens the initial draft's unnecessarily restrictive fixed-guess formulation. Shared participants do not automatically invalidate the null calculation when the stronger conditional-fairness assumption actually holds.

The assumption is substantial. If a hidden fair bit B is drawn once and every target equals B, each target is marginally fair, but the first revealed outcome predicts all later outcomes. If $Y_i = G_i$, the target distribution can also be marginally fair while accuracy is perfect. Both constructions violate conditional fairness. Counting output zeros and ones is therefore an inadequate device audit.

This is a mathematical consequence of a null, not an empirical proof that nature obeys it. Rejecting the null challenges the conjunction of its assumptions and implementation claims. It does not identify which part failed. An extraordinary interpretation must additionally survive an engineering account of why future information, predictable state, or selective trial retention could not have entered the information set.

7. Fixed-sample inference and an honest power calculation

For a completed design with fixed N, the one-sided exact null tail for an observed count $c \in \{0, \dots, N\}$ is

$$p_{\text{tail}}(c) = \sum_{k=c}^{N} \binom{N}{k} 2^{-N}.$$

The alternative direction must be declared. A low hit rate is a different finding from the proposed above-chance effect; discovering it after inspection requires an explicitly exploratory label. A two-sided design or prespecified opposite-direction alternative can investigate both, with an appropriate error budget.

A proposed operating point, not an experiment

Take $N = 40{,}000$, one-sided level $\alpha = 0.005$, and a smallest relevant success probability $p_1 = 0.51$. Exact computation gives a critical count of 20,259. Its null tail is approximately 0.00486852, and its power under independent Bernoulli(0.51) successes is 0.921503. The discrete threshold is conservative relative to the nominal level. These are design calculations from probability distributions, not collected observations.

Under a constant independent success probability p, an exact upper confidence limit $U(c)$ can be obtained by solving

$$\mathbb{P}_{p=U(c)}(C_N \le c) = \alpha$$

for $c < N$, with $U(N) = 1$. This inverts binomial tests and has at least $1 - \alpha$ coverage. If $c = 20{,}000$, the one-sided 99.5% upper limit is approximately 0.506452, below the declared 0.51 effect threshold. It would exclude that effect size under the model, while leaving smaller effects unresolved.

Power and interval assumptions deserve separate attention. The exact fair-null test follows from Proposition 1. The quoted alternative power and common-p interval additionally assume a stable Bernoulli alternative. Participants may differ, devices may drift, and an hypothesized effect may vary by session. A planning sensitivity analysis must simulate those alternatives and report participant- and site-level estimates. A large number of bits cannot manufacture a large number of independently sampled people.

The general bounded-variable inequalities studied by Hoeffding provide useful conservative alternatives [12]. They do not remove the need to specify the target population. A narrow interval for this apparatus, this recruitment procedure, and this timing configuration must not be advertised as a universal verdict on every imaginable form of intuition.

8. Sequential evidence without optional-stopping inflation

If a study is designed for continuous monitoring, repeated fixed-sample p-values are unsuitable. A simple alternative illustrates the established martingale approach underlying modern confidence sequences [13]. Define the dimensionless capital process

$$M_0 = 1, \qquad M_n = \prod_{i=1}^{n} \{1 + \lambda_i (2W_i - 1)\},$$

where the wager $\lambda_i \in [-1, 1]$ is chosen using $\mathcal{H}_i$ before Y_i is generated. For an above-chance test one may restrict $\lambda_i \in [0, 1]$. The choice rule is part of the preregistration.

Proposition 2 (Anytime-valid crossing bound)**.** *Set $\mathcal{F}_n = \mathcal{H}_{n+1}$ for $n \geq 0$, extending the history convention after the last trial when needed. Under H_A, (M_n) is a nonnegative martingale for this filtration and, for every $\alpha \in (0, 1)$,*

$$\mathbb{P}\{\exists n \geq 1 : M_n \geq 1/\alpha\} \leq \alpha.$$

Proof. Every factor is nonnegative. Conditional fairness gives $\mathbb{E}[2W_i - 1 \mid \mathcal{H}_i] = 0$, so $\mathbb{E}[M_i \mid \mathcal{H}_i] = M_{i-1}$. Let τ be the first crossing time and fix an integer horizon m. Stopping at $\tau \wedge m$ preserves expectation one because the stopping time is bounded. On $\{\tau \leq m\}$ the stopped capital is at least $1/\alpha$; elsewhere it is nonnegative. Therefore $1 \geq \mathbb{P}(\tau \leq m)/\alpha$. Letting m increase proves the claim by continuity of probability. □

For a constant alternative success probability p_1, setting $\lambda = 2p_1 - 1$ gives the likelihood ratio between Bernoulli(p_1) and Bernoulli(1/2) successes. With $p_1 = 0.51$, the wager is 0.02. This is a useful interpretation, not a requirement that the alternative be true for null validity.

Looking at many alternative wager rules and reporting the largest capital invalidates the single-process guarantee. A fixed mixture $M_n^{\text{mix}} = \sum_j a_j M_n^{(j)}$, with nonnegative weights a_j summing to one and fixed before data collection, remains a valid nonnegative process. Adaptive choices within a predictable rule also remain valid. Resets after losses, selective omission of trials, and retrospective weighting are different operations and require new proofs. Sequential validity is a property of the specified process, not a permission to improvise analysis.

9. Probabilistic forecasts should pay for overconfidence

A binary guess discards the forecaster's confidence. Track A may instead accept $q_i \in [0,1]$, a committed probability for $Y_i = 1$. The Brier loss is $B(q,y) = (q-y)^2 \in [0,1]$. Lower values are better. Proper scoring rules are established tools for this purpose [14]; a poetic certainty is not exempt from their penalties.

Proposition 3 (Brier regret)**.** *If Y is Bernoulli with success probability $p \in [0,1]$, then for every quoted probability $q \in [0,1]$,*

$$\mathbb{E}B(q,Y) = p(1-p) + (q-p)^2.$$

Thus truthful forecasting $q = p$ uniquely minimizes expected loss.

Proof. Expand $p(1-q)^2 + (1-p)q^2 = p - 2pq + q^2$ and add and subtract p^2. The remaining square is nonnegative and vanishes exactly when $q = p$. □

Under the fair conditional null, the constant forecast $q_i = 1/2$ has loss exactly $1/4$ on every trial. A confident forecast $q_i = 0.9$ has expected loss 0.41 under that null. Being correct on a memorable subset does not repair the losses on the remaining registered trials.

Define the improvement $D_i = 1/4 - B(q_i, Y_i) \in [-3/4, 1/4]$. Conditional on $\mathcal{H}_i$, the identity gives $\mathbb{E}[D_i \mid \mathcal{H}_i] = -(q_i - 1/2)^2 \le 0$. Therefore a preregistered fixed-sample bound for bounded supermartingale differences can test positive average improvement. A derivation using the conditional exponential bound gives $\mathbb{P}(\sum_{i=1}^{n} D_i \ge x) \le \exp(-2x^2/n)$ for $x > 0$: each increment has conditional range length one, its mean is nonpositive, and exponential Markov inequality is minimized at exponent $4x/n$. This conservative result requires the same information ordering as the binary test.

Accuracy and scoring can disagree. A forecaster might slightly improve thresholded decisions while becoming dangerously overconfident. PEC-R2 selects one primary endpoint before data collection and reports the other as a prespecified secondary outcome. Calibration plots are descriptive diagnostics, with uncertainty and held-out bins; tuning the forecast on the evaluation targets destroys the prospective test. The data that teach a predictor and the data that judge its extraordinary claim must be separated.

10. Anticipatory physiology: the pipeline is the hypothesis

Track B records physiological time series before physical target generation. It does not begin with an unrestricted search over channels, time lags, filters, and classifiers. Let $V_i(t) \in \mathbb{R}^r$ denote r sensor channels at time t, each with declared units; let the fixed mapping f transform a declared pre-generation interval into $S_i = f(V_i) \in \mathbb{R}^d$. The pipeline specification includes sampling rate, synchronization uncertainty, artifact rules, normalization, and all training data used to choose f.

The phrase "before the event" must apply to every input sample used by f. A zero-phase filter often uses samples on both sides of its nominal output time. Computing it across the event and later cropping the trace can carry post-event information into an apparently early feature. PEC-R2 either uses a causal filter with known support or fits and applies an explicitly truncated pre-generation transformation. Timing validation includes a known injected pulse so that the entire acquisition and preprocessing chain can be checked.

A prospective predictor maps (S_i, Z_i) to a probability q_i. Compare it to a history-only predictor using Z_i, a constant or calibrated device-bias predictor, and a sham-sensor predictor. The primary contrast is improvement on held-out participants and sessions, scored by a fixed rule. Trial-level random splitting is insufficient when nearby windows share participant state, preprocessing estimates, or overlapping samples. All transformations learned from data must be learned within the training partition.

The physiology literature supplies motivation for a particular feature class, not permission to search until one works [3]. A confirmatory design might choose one normalized skin-conductance feature and one pre-generation interval, leaving EEG channels and alternative windows exploratory. The final choice requires sensor expertise and pilot quality checks. No new physiology data are presented in this paper, and no specific signal is claimed to carry future information.

Permutation controls must respect the actual assignment mechanism. Arbitrarily shuffling correlated trials can produce an invalid null. A circular shift is valid only under defensible stationarity and exchangeability conditions; it is not a universal correction. Where target assignments were randomized independently, randomization-based reassignments following that mechanism are preferable. Negative controls should also include delayed display, already-generated targets, and matched recordings without participants. These distinguish an early signal from an extraordinary early signal.

11. Conditional information and the selection trap

A Guide's explanation

Read the counterexample at the end of this section twice. It is the prologue's first trap made exact: two unrelated coin flips, filtered after the fact, produce perfect "prediction." It is why the protocol insists that every trial is counted, whatever its outcome.

For a finite-valued feature S, binary target Y, and finite-valued covariate Z, let $p(s, y, z)$ be their joint probability mass function. Write $p(s, y \mid z)$ for the conditional distribution at any z with positive probability. Conditional mutual information, in natural-log units, is

$$I(S; Y \mid Z) = \sum_{z} p(z) \sum_{s,y} p(s, y \mid z) \log \frac{p(s, y \mid z)}{p(s \mid z) p(y \mid z)}.$$

Zero-probability joint terms contribute zero. Every positive joint term has positive corresponding marginals, so its ratio is well-defined. This quantity is an average Kullback–Leibler divergence; non-negativity and its equality condition imply $I = 0$ exactly when the conditional joint distribution factors for every positive-probability z. This standard identity defines a statistical target. It is not a new physical law.

Continuous sensor features require probability measures or densities with appropriate regularity, not a casual substitution of histogram counts. A finite-data estimate of I can be positive when the population value is zero. Binning, smoothing, estimator choice, and feature selection each alter its sampling behavior. PEC-R2 therefore prioritizes prospective prediction on untouched data. Information estimates can be secondary diagnostics calibrated under a justified null, with every discretization choice frozen.

A complete counterexample

Let S and Y be independent fair bits. Initially $I(S; Y) = 0$. Retain a trial only if $R = \mathbf{1}\{S = Y\} = 1$. In the retained data, (S, Y) equals $(0, 0)$ or $(1, 1)$ with probability $1/2$ each. Therefore $I(S; Y \mid R = 1) = \log 2$, despite the absence of any original information channel. Selection after the outcome has manufactured perfect prediction. A "coherence detector" that rejects unsuccessful trials can implement exactly this error under more elaborate language.

Covariates also require causal care. Conditioning on a common effect of sensor activity and target outcome can induce association. Adjusting for a downstream variable is not automatically a more rigorous analysis. The protocol must declare why each covariate is available before generation and whether it could be affected by the tested process. An apparently exhaustive adjustment set cannot contain unobserved physical state; a residual association therefore remains compatible with an omitted ordinary cause.

The empirical question is whether a measured, predeclared predictor exceeds a defensible baseline. The stronger statement that it accesses a causally inaccessible future requires additional assumptions that no information-theoretic score can certify by itself.

12. Intention requires a causal experiment

Track C assigns an instruction to favor target one or target zero. The assignment must use a randomization source independent of the outcome device, and its physical implementation must avoid controlling that device through ordinary channels. A change in outcome frequency after assignment could otherwise be a conventional intervention effect.

For a finite collection of N trial units, let $Y_i(1), Y_i(0) \in \{0,1\}$ be potential outcomes under the two instructions, within a declared no-interference and consistency model. These may be regarded as fixed conditional on background randomness. Let A_i be independent Bernoulli(π) assignments with known $\pi \in (0,1)$, independent of the potential outcomes. The observed result is $Y_i = A_i Y_i(1) + (1-A_i)Y_i(0)$. Define the finite-population estimand

$$\Delta_N = \frac{1}{N}\sum_{i=1}^{N}\{Y_i(1) - Y_i(0)\} \in [-1,1].$$

It is an instruction effect, measured in probability points, not a direct measure of an internal intention state.

Proposition 4 (An unbiased assignment estimator)**.** *The Horvitz–Thompson form*

$$\widehat{\Delta}_N = \frac{1}{N}\sum_{i=1}^{N}\left\{\frac{A_i Y_i}{\pi} - \frac{(1-A_i)Y_i}{1-\pi}\right\}$$

has randomization expectation Δ_N*.*

Proof. Consistency makes $A_i Y_i = A_i Y_i(1)$ and $(1-A_i)Y_i = (1-A_i)Y_i(0)$. Independence gives expectations $\pi Y_i(1)$ and $(1-\pi)Y_i(0)$. Divide by the assignment probabilities, subtract, and average. Linearity of expectation proves the claim. □

This established estimator remains defined even if one arm happens to be empty; an ordinary difference of sample means does not. In practice, blocked or balanced assignment may be preferable, with its corresponding randomization analysis. Repeated use of one device can create interference through heating, state persistence, or prior instructions. If those processes are plausible, define session-level units, include washout periods, and randomize at the appropriate level. The simple proof does not eliminate them.

An effect of instructions would still not establish a paranormal mechanism. Assignment can change movement, attention, timing, or environment. Device-only sham assignments and remote isolated operation test those pathways. The comparison must also report compliance: a randomized instruction is observable; the intended mental act is incompletely observed. Replacing the assignment estimand with an effect of "true intention" requires additional identification assumptions.

13. A chronology that can survive an audit

The protocol's most vulnerable claim is temporal: that the predictor had no access to the target before committing. PEC-R2 requires an independent record of the commitment, physical generation, storage, display, and analysis events. The order must be supported by uncertainty bounds and instrumentation, not inferred from the order of lines in an application log.

Each trial record contains a unique identifier, the predictor version, its allowed inputs, the committed guess or probability, the final sensor-sample time, generator request and physical-event timing, raw output, extracted output if applicable, display timing, and equipment-health flags. A detached commitment receipt is issued before generation. A cryptographic digest can bind a record against later unnoticed alteration under cryptographic assumptions. It cannot prove that the record was truthful when created, that a clock was correct, or that a device waited to generate its bit.

Consider a server that receives a guess and returns a freshly requested bit from a buffer. The API request is new; the buffered physical event may be old. A participant's forecast preceding the API response would not necessarily precede generation. The device specification must identify entropy-source sampling, buffering, extraction, and delivery stages. If certification depends on an entire block, the commitment must precede the generation of all raw events contributing to the target block, unless a separate proof justifies a weaker order.

A second counterexample uses a perfectly immutable log with a compromised clock. The experiment writes the target at noon, changes the clock, and records a guess with an earlier timestamp. Every hash verifies. Chronology still fails. Independent receipts, clock-drift limits, and witnessed instrument triggers reduce this risk, but none supplies metaphysical certainty about every possible adversary.

Operational exclusions are declared before data collection. A trial with ambiguous ordering is a protocol failure, not an inconvenient miss to remove selectively. Report all such failures by assignment, participant, and device, with a worst-case sensitivity analysis when exclusions could depend on outcomes. A primary analysis that uses only valid trials also needs an explicit validity-selection assumption. Design the acquisition system to make invalidity rare and mechanically determined before target revelation.

The audit artifact should let a skeptical engineer reconstruct the timing chain and identify an ordinary route for leakage. If that reconstruction is impossible, the dataset cannot support a pre-generation claim regardless of its nominal significance.

14. Randomness is a model plus an implementation

A pseudorandom generator can be statistically excellent and still be deterministic given its seed. It is appropriate for many randomized experiments when assignment concealment and independence are adequate. It does not directly test the stronger claim of predicting a newly generated quantum event. The distinction is especially relevant when comparing PEC-R2 with previous psychological protocols that used software randomization [6,7].

A hardware quantum generator has its own chain of assumptions: preparation, measurement, digitization, bias handling, extraction, seed use, and communication. A sequence that passes a battery of randomness tests can still be predictable to a party with device-state information. Conversely, a failed test can reveal a bias without revealing anything about consciousness. Device qualification and the extraordinary hypothesis should have separately defined endpoints.

Bell-certified generation offers one possible advanced stage, subject to the physical constraints and independent-choice assumptions in its security argument [10,11]. Such an implementation may be expensive and slower than a conventional source. It should not be invoked as a decorative adjective. The experiment must identify the actual certificate, finite-statistics bounds, adversary model, extractor assumptions, and every raw event contributing to the released bit. No such device has been constructed or tested for this paper.

PEC-R2 therefore proposes a progression. First, instrument a classical randomized task to validate clocks, commitments, and analysis under known injected leakage. Next, use a qualified physical entropy source with measured buffering and environmental controls. Only a persistent, prospectively replicated anomaly justifies the additional expense of independent implementations or stronger randomness certification. The progression is methodological; success in the early stages is not evidence of the exotic effect.

The device team should deliberately create ordinary predictive signals during commissioning: a small output bias, a repeated seed, a delayed display, or an exposed buffer. The analysis should detect the planted signal when its registered magnitude is detectable. Blinded null streams should not systematically trigger it. These positive and negative controls test the experiment's sensitivity and specificity. They do not establish that every unplanted signal would be paranormal.

Environmental measurements, including temperature or electromagnetic conditions when relevant, can expose coupling. Their units, calibration uncertainty, and sampling times belong in the device specification. Selecting only quiet environmental periods after seeing outcomes would recreate the selection problem. A physical model is strongest when it states which ordinary couplings are excluded, which remain possible, and what observation would distinguish them.

15. A staged, multisite research design

PEC-R2 is a proposed protocol, not a preregistration that has already been filed. Before recruitment, an independent statistical and device review must freeze the final implementation. The research team should include both sympathetic and skeptical investigators, with authority to identify invalidating defects. No human participant or laboratory result is claimed in this collection.

Commissioning and pilot work

Stage 0 validates acquisition, clock uncertainty, commitment order, blinding, and planted-signal recovery. Pilot data determine feasibility, tolerability, and variance assumptions; they do not contribute to the final confirmatory test. Freeze code, feature definitions, exclusions, and the primary analysis after this stage. Publish a dated protocol and preserve a reproducible environment with independent copies.

A first confirmatory study

Stage 1 uses the fixed-sample Track A design on page 9, subject to revised feasibility and precision calculations before registration. For example, 200 participants could each contribute 200 valid trials. This arithmetic describes recruitment and workload only; it does not imply that the 200 people are representative of humanity or that an alternative effect is identical across them. Register site balance, session limits, device versions, failed-trial handling, and whether any participant subgroup is primary.

Stage 2 investigates one selected physiological pipeline on new data. Stage 3 evaluates randomized instructions with a separately powered assignment design. A positive result in one stage does not license pooling all stages into a single success metric. Each track has its own estimand and error allocation, and the literature motivating one track may not support the others.

Replication and release

Stage 4 repeats the primary positive finding at independent sites using frozen endpoints and independently implemented acquisition. One team should attempt a faithful replication; another should target an identified conventional explanation. Both positive and negative results should be released, including protocol deviations and failed equipment qualification. A successful software reanalysis of the same dataset is reproducibility, not independent replication.

Participant privacy limits what can be published. Raw physiological traces and exact timestamps may be identifying. Preserve complete records with appropriate controlled access, and release deidentified data or sufficient aggregates where necessary. Transparency does not require exposing participants' private experiences. The public report must clearly state any access restriction and how an independent auditor could inspect the underlying evidence. A positive claim with unverifiable data should remain correspondingly provisional.

16. Multiplicity, Bayesian comparison, and failure rules

The three tracks form a family of hypotheses. A simple prospective allocation uses familywise budget 0.01: 0.005 for Track A and 0.0025 each for Tracks B and C. If each test is valid at its allocated level, the union bound limits the probability of at least one false rejection to 0.01 without requiring independence. Additional sensors, subgroups, sites, or lag choices consume further budget or remain explicitly exploratory. Holm's established step-down procedure is another option for a frozen family of valid p-values [15].

Multiplicity is not a small technical correction. With 100 independent true-null tests each at level 0.05, the probability of at least one rejection is $1 - 0.95^{100} \approx 0.99408$. Here 100 is a constructed example, not a count of the literature's analyses. Dependence changes that number; it does not make unreported analytical search harmless. The full search space includes failed models and discarded windows, not only the final displayed table.

For a fixed Bernoulli model with c hits in n trials, let π_1 be a specified prior probability measure over alternative hit probabilities $p \in [0, 1]$. The Bayes factor against the fair point null is

$$BF_{10} = \int_0^1 (2p)^c \{2(1-p)\}^{n-c} \, \mathrm{d}\pi_1(p).$$

The integral averages likelihood ratios; it is sensitive to the alternative prior. Priors concentrated on large effects can make a small effect look unfavorable to the alternative. Report justified prior sensitivity and distinguish a Bayes factor from posterior model odds. The latter equal the Bayes factor times prior odds. Do not choose the prior after seeing which one favors the desired conclusion.

An unsuccessful test has several possible meanings. A narrow interval excluding the smallest relevant effect counts against that operational hypothesis. A wide interval is inconclusive. A timing failure invalidates the intended test. An effect appearing only after outcome-dependent exclusions is a selection warning. A successful primary result that disappears in independent replication remains unconfirmed. These categories must not be collapsed into either universal disproof or universal vindication.

The protocol must preclude explanations that make failure impossible: the effect cannot be declared to vanish whenever skeptics observe, whenever registration is strict, or whenever the analysis is fixed, unless that moderator itself has a prospectively testable prediction. A theory that permits every outcome cannot earn empirical support from any one of them.

17. What a lottery coincidence can and cannot show

Consider a generic fair lottery that chooses five distinct white labels uniformly from 69, and one red label uniformly from 26. These are the current published Powerball selection ranges [16]. The calculation below is conditional on that mathematical sampling model; it is not an audit of any historical draw, ticket, or private anecdote.

Fix a five-label white ticket before the draw, and let $K \in \{0, 1, 2, 3, 4, 5\}$ count exact white matches. The sample space contains $\binom{69}{5}$ equally likely white sets. To obtain k matches, choose k of the five ticket labels and $5 - k$ of the 64 others. Hence

$$\mathbb{P}(K = k) = \frac{\binom{5}{k}\binom{64}{5-k}}{\binom{69}{5}}.$$

This is the hypergeometric law. It gives $\mathbb{P}(K = 3) \approx 0.001793832$ and $\mathbb{P}(K \geq 3) \approx 0.001822394$. The red-ball match has probability $1/26$ under independent fair red sampling. Exact jackpot probability is therefore $1/[26\binom{69}{5}]$. All of these are prospective model probabilities for a fixed ticket, not calibrated evidential measures for a selected story.

If $m \in \mathbb{N}$ independent opportunities each have event probability $r \in [0, 1]$, the probability of at least one occurrence is $1 - (1 - r)^m$. Independence is needed for this equality; the union bound gives an upper bound mr without it. In a retrospective account, the number of attempted forecasts, alternative matching rules, and unreported failures is usually unknown. A one-ticket probability cannot repair that missing selection history.

Numerical "near misses" create an additional problem. Lottery labels are nominal identifiers. If every label is renamed by a permutation, the physical sampling process is unchanged while ordinary numerical distances can change completely. A score rewarding adjacent numbers therefore needs an independently motivated physical rationale or an explicit prospective definition. A nearest-match score selected after the draw measures the analyst's freedom as well as the ticket's performance.

The appropriate response to a memorable coincidence is neither ridicule nor an invented mechanism. Preserve a precise record, identify what was fixed beforehand, and design a repeated prospective test with all attempts counted. PEC-R2 does not offer lottery predictions, investment advice, or an assurance that any participant possesses exceptional powers. Its purpose is to replace retrospective impressiveness with an auditable question.

18. Computational checks and their exact scope

The accompanying internal check program was executed for this paper. It uses constructed finite cases and numerical probability calculations. No rows represent people, physiological recordings, lottery tickets, or physical-device measurements. The calculations support the derivations and reveal their operating scale; they do not supply empirical evidence for the hypothesis.

Check	Executed result
All adaptive binary policies through four trials	32,768 policies; each tested on all 16 fair target paths; every success histogram equals $\binom{4}{k}$.
Predictable betting process	All 65,536 fair paths through 16 trials; mean terminal capital 1; crossing probability at threshold 20 equals 0.0158539, below 0.05.
Brier identity	10,201 probability pairs on a fixed grid; numerical identity holds to declared tolerance.
Lottery counting	Full normalization of the six hypergeometric probabilities; exhaustive smaller-set counting agrees with the same law.
Causal estimator	All 16 assignments for a four-unit example; mean estimator 0.25 equals its finite-population effect.
40,000-trial planning	Critical count 20,259; null tail 0.00486852; power 0.921503 at independent 51% success.

There were 42,977 assertions across these checks. The number is an implementation accounting detail, not a measure of scientific credibility. Exhausting a four-step policy tree supports a finite example; Proposition 1, rather than that enumeration, establishes the general result. Likewise, finite-horizon betting enumeration does not prove the infinite-horizon inequality. The stopping-time argument does.

Two deliberately failing interpretations were checked. Selecting independent bits only when they agree creates perfect apparent association. Assuming independence from balanced marginals alone permits a repeated-bit counterexample. These examples remain in the paper because a useful audit should test how a claim can fail, not merely confirm that a preferred equation produces a desired number.

Residual risks include the completeness of the physical information model, synchronization validity, sensor preprocessing, participant selection, and alternative-effect heterogeneity. None is resolved by these calculations. An actual experiment would require new device documentation, an approved participant protocol, preregistration, and independent replication. The present work supplies a reviewable design and mathematical foundation for that decision.

19. Interpretation, falsification, and Guide use

A Guide's explanation

A Guide can say: "The Continuum takes a wild idea seriously by building the test that could prove it wrong. Nothing here shows that the future leaks—the paper is a protocol, not a result." A Guide should not suggest that anyone can predict lotteries or random events, that a striking personal coincidence is evidence of an ability, or that quantum physics lets intention choose outcomes.

The bold claim that survives

A conjecture about apparently inaccessible future information can be made exact enough to face failure. PEC-R2 demonstrates how to distinguish prospective prediction, physiological association, and causal influence; how to protect the relevant temporal order; and how to test the resulting quantities without turning optional stopping or selected anecdotes into evidence. The fair-null result is exact, and the bounded-predictability extension supplies an anytime-valid test above a declared ordinary-information ceiling. Whether an implemented experiment exceeds that ceiling reproducibly is the empirical question.

What would change the position

A registered primary effect, repeated independently with tighter timing and different devices, would justify investigating an anomaly. A model explaining that anomaly would need predictions that distinguish it from leakage, physical coupling, selection, and ordinary device state. A narrow null result excluding the declared effect size would count against the specified method. Repeated failure would justify ending that research line. The name "Continuum" supplies no exemption from either outcome.

For Guides and public discussion

Established within the stated models: adaptive guessing cannot extract an advantage from conditionally fair future bits; proper scores penalize unsupported confidence; a valid sequential process can control false positives under monitoring. These are mathematical results, not discoveries of a new mental faculty.

Proposed: PEC-R2's combination of event timing, commitment records, distinct tracks, and replication gates. Its hardware and human studies have not been run. **Open:** whether a carefully isolated system exhibits a repeatable unexplained signal of a registered magnitude. The reviewed evidence does not establish precognition or intention-controlled randomness.

A useful conversational response is: "The metaphor invites a question about information before an event. The paper gives that question a test, and the test may find nothing beyond chance." When asked whether a striking coincidence proves an ability, a Guide can explain the missing denominator of attempts and the need for prospective records. When asked whether quantum mechanics permits desire to choose an outcome, it should distinguish the supplied metaphor from a demonstrated physical mechanism.

The Continuum can keep its imaginative ambition while allowing its hypotheses to lose. That is the philosophical commitment behind this final paper: inquiry becomes stronger when the result has authority over the story. A future that leaves no detectable precursor is a scientifically coherent outcome, and this research program is designed to recognize it.

A prospective test that allows bounded ordinary predictability

A Guide's explanation

This is the newest result. Real equipment is never perfectly fair; a device may leak a little ordinary predictability. Instead of pretending otherwise, the test lets the experimenter declare a ceiling on how much ordinary predictability the apparatus could supply, and it counts only evidence above that ceiling. The next section shows the honest price: evidence then accumulates more slowly.

Exact conditional fairness is a sharp null. A qualified physical source may instead support a bound on ordinary conditional predictability. PEC-R2 therefore adds a robust sequential test that charges the hypothesis for that uncertainty before any extraordinary interpretation is considered. The construction is an application of established nonnegative-supermartingale inference [13].

Retain the trial ordering and information sets of Section 5. Let $u_i \in [1/2, 1]$ be known before generation and measurable with respect to $\mathcal{H}_i$. The robust null is

$$\Pr(W_i = 1 \mid \mathcal{H}_i) \le u_i \quad \text{almost surely.}$$

For example, a justified conditional device bound $|\Pr(Y_i = 1 \mid \mathcal{H}_i) - 1/2| \le \epsilon_i$, with predictable $\epsilon_i \in [0, 1/2]$, supplies $u_i = 1/2 + \epsilon_i$ for any committed binary guess. A marginal output-frequency tolerance does not establish this conditional bound.

Choose predictable $\lambda_i \in [0, 1/u_i]$ and define

$$E_0 = 1, \qquad E_n = \prod_{i=1}^{n} \{1 + \lambda_i (W_i - u_i)\}.$$

Proposition 5 (Anytime-valid evidence above a predictable ceiling)**.** *Under the robust null, E_n is a nonnegative supermartingale for the trial filtration. For every $\alpha \in (0, 1)$, $\Pr\{\exists n \ge 1 : E_n \ge 1/\alpha\} \le \alpha$.*

Proof. At $W_i = 0$ the factor is $1 - \lambda_i u_i \ge 0$, and at $W_i = 1$ it is nonnegative as well. Its conditional expectation is $1 + \lambda_i(\Pr(W_i = 1 \mid \mathcal{H}_i) - u_i) \le 1$. Multiplication by the nonnegative, already measurable E_{i-1} proves the supermartingale inequality. Stopping at the first crossing or a fixed finite horizon gives expected stopped capital at most one. On crossing it is at least $1/\alpha$, so crossing by that horizon has probability at most α. Increasing the horizon proves the assertion. □

For fixed $u < 1$ and a planned alternative $q \in [u, 1]$, take $\lambda = (q - u)/(u(1 - u))$. The two capital factors become q/u for a hit and $(1 - q)/(1 - u)$ for a miss. If hits truly are independent Bernoulli(q), expected log capital per trial is

$$D(q\|u) = q \log(q/u) + (1 - q) \log((1 - q)/(1 - u)),$$

with the usual continuous convention at an endpoint. At $u = 1$, no above-ceiling alternative exists and no positive evidence growth is promised. The likelihood-ratio interpretation helps plan sensitivity; validity under the null requires only the conditional upper bound, not independent successes.

The evidential price of ordinary predictability

With a planned 51% hit rate, $D(0.51\|0.50) \approx 0.000200013$ natural-log units per trial. Allowing an ordinary conditional hit ceiling of 0.505 reduces the rate to $D(0.51\|0.505) \approx 0.000050009$. The expected log-evidence growth is roughly four times slower. This is a planning comparison, not an expected stopping-time theorem. If the justified ceiling is 0.51, a 51% rate itself is inside the robust null and supplies no above-ceiling claim.

The 40,000-trial operating point earlier in the paper belongs to the exact fair-null fixed-sample design. It is not transferred to this broader null. Before a robust study is registered, its bias ceiling, alternative, power calculation and chosen evidence process must be specified together. Allowing a nonzero predictability budget strengthens the interpretation of an observed departure by making the ordinary explanation quantitatively explicit.

For a precise calibration allowance, let $\mathcal{C}$ denote calibration information fixed before the test. Suppose a $\mathcal{C}$-measurable good event G has probability at least $1 - \delta_{\text{cal}}$, and for almost every calibration realization in G the conditional future trial law satisfies every displayed conditional ceiling. Applying the theorem within each such realization and integrating gives overall crossing probability at most $\alpha + \delta_{\text{cal}}$. This is a simultaneous conditional guarantee; estimating average bit bias from a held-out file is insufficient. When calibration cannot justify this premise, the report states that gap instead of assigning an arbitrary small ϵ_i.

Trials may be paused or skipped on a rule determined before seeing their target, with the unbet trial contributing factor one. Outcome-dependent deletion, resetting capital after losses, or retroactively enlarging a favorable alternative set changes the test and is not covered by the theorem. The finite audit enumerates bounded-bias trial trees and checks their expected capital and crossing probabilities. It also retains a predictable balanced sequence as a counterexample to marginal-frequency calibration.

The contribution is a test with an explicit information budget. A reproducible effect must exceed both its sampling uncertainty and the ordinary predictability admitted by the qualified apparatus. The interpretation remains prospective: the theorem protects an inference procedure; an actual unexplained signal requires new physical observations and independent replication.

References and source notes

[**1**] Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree. *American Psychologist*, 64(6), 515–526.

https://doi.org/10.1037/a0016755

Scope: Conceptual synthesis of ordinary expertise; not evidence of precognition.

[**2**] Bem, D. J. (2011). Feeling the future: Experimental evidence for anomalous retroactive influences on cognition and affect. *Journal of Personality and Social Psychology*, 100(3), 407–425.

https://doi.org/10.1037/a0021524

Scope: Positive original claims; subsequent replication evidence is considered separately.

[**3**] Mossbridge, J., Tressoldi, P., & Utts, J. (2012). Predictive physiological anticipation preceding seemingly unpredictable stimuli: A meta-analysis. *Frontiers in Psychology*, 3, 390.

https://doi.org/10.3389/fpsyg.2012.00390

Scope: Positive aggregate finding; observational timing and analysis assumptions remain relevant.

[**4**] Ritchie, S. J., Wiseman, R., & French, C. C. (2012). Failing the future: Three unsuccessful attempts to replicate Bem's "retroactive facilitation of recall" effect. *PLOS ONE*, 7(3), e33423.

https://doi.org/10.1371/journal.pone.0033423

Scope: Preregistered recall-task replications; not a test of every psi claim.

[**5**] Galak, J., LeBoeuf, R. A., Nelson, L. D., & Simmons, J. P. (2012). Correcting the past: Failures to replicate ψ. *Journal of Personality and Social Psychology*, 103(6), 933–948.

https://doi.org/10.1037/a0029709

Scope: Seven experiments; high-level findings verified from the indexed primary abstract.

[**6**] Kekecs, Z., et al. (2023). Raising the value of research studies in psychological science by increasing the credibility of research reports: The transparent Psi project. *Royal Society Open Science*, 10, 191375.

https://doi.org/10.1098/rsos.191375

Scope: Read with the published correction, 10:231080, doi:10.1098/rsos.231080; safeguards require auditing too.

[**7**] Walleczek, J., von Stillfried, N., Schmidt, S., Wittmann, M., Kirmse, K. A., Moll, J., & Kekecs, Z. (2025). Metascientific replication project with the advanced meta-experimental protocol of the transparent psi project procedures for testing the precognitive effect claimed by Bem. *PLOS ONE*, 20(11), e0335330.

https://doi.org/10.1371/journal.pone.0335330

Scope: Three studies; the opposite-direction finding did not persist in the third.

[**8**] Maier, M. A., Dechamps, M. C., & Pflitsch, M. (2018). Intentional observer effects on quantum randomness: A Bayesian analysis reveals evidence against micro-psychokinesis. *Frontiers in Psychology*, 9, 379.

https://doi.org/10.3389/fpsyg.2018.00379

Scope: Primary null-favoring result; prior-sensitive Bayesian comparison.

References and research provenance

[9] Schlosshauer, M. (2005). Decoherence, the measurement problem, and interpretations of quantum mechanics. *Reviews of Modern Physics*, 76, 1267–1305.

https://doi.org/10.1103/RevModPhys.76.1267

Scope: Physics review; author manuscript inspected. Volume carries a 2004 issue date; publication was in 2005.

[10] Pironio, S., et al. (2010). Random numbers certified by Bell's theorem. *Nature*, 464, 1021–1024.

https://doi.org/10.1038/nature09008

Scope: Primary randomness-certification research; its assumptions remain necessary.

[11] Bierhorst, P., et al. (2018). Experimentally generated randomness certified by the impossibility of superluminal signals. *Nature*, 556, 223–226.

https://doi.org/10.1038/s41586-018-0019-0

Scope: Physical certification experiment; no claim about mental forecasting.

[12] Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301), 13–30.

https://doi.org/10.1080/01621459.1963.10500830

Scope: Established concentration theory; boundedness and dependence conditions matter.

[13] Howard, S. R., Ramdas, A., McAuliffe, J., & Sekhon, J. (2021). Time-uniform, nonparametric, nonasymptotic confidence sequences. *Annals of Statistics*, 49(2), 1055–1080.

https://doi.org/10.1214/20-AOS1991

Scope: Sequential inference; revised author manuscript inspected. The elementary proof here is self-contained.

[14] Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477), 359–378.

https://doi.org/10.1198/016214506000001437

Scope: Established scoring-rule theory; not a new Continuum invention.

[15] Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2), 65–70.

https://www.jstor.org/stable/4615733

Scope: Multiplicity method for a prespecified family of valid tests.

[16] Multi-State Lottery Association. (Accessed November 2026). Powerball: How to play.

https://www.powerball.com/

Scope: Official selection ranges only; neither peer-reviewed science nor evidence about a historical ticket.

Research provenance: AI-assisted research, drafting, mathematical derivation, and computational checking prepared for Architect X13. Internal review is distinct from external peer review. The original project drafts remain preserved; PEC-R2 is newly proposed. The private Continuum canon supplies the question and vocabulary, not experimental validation. No proprietary specification, human dataset, physical trial, or independent replication is claimed.