A Statistical Quantum-Correlation Inference Framework

large Statistical Remote-Action Inference, Entanglement Controls, Causal Validation, and a Conditional Human-Level Communication Horizon

Yuval Fradkin

August 2026
Submission-candidate consistency edition

Abstract

We formulate a falsifiable mathematical and experimental framework for testing whether a remote local choice in one wing of a pre-established entangled system is associated with any statistically detectable change in the complete measurement record available at the other wing. The framework does not assume that such an effect exists and does not claim a loophole in the no-signaling theorem. Under standard quantum mechanics, local trace-preserving operations predict identical reduced states and zero locally accessible signaling information. We define randomized encoding, full-distribution comparison, matched controls, likelihood decoding, information-theoretic capacity estimates, and spacelike-separated validation.

1 Introduction

The technological motivation is statistical inference rather than single-shot state reading. Alice and Bob may prepare and synchronize an entangled architecture in advance. After separation Alice later selects a local operation; Bob records a large local data set and asks whether its probability law depends on Alice's choice. Standard quantum mechanics predicts that it does not. The framework therefore treats no-signaling as the null hypothesis and defines the strongest controlled test of any residual signature.

2 Inference as the Mechanism; Communication as an Emergent Application

The central conceptual distinction of this framework is between *physical message transmission* and *statistical inference*. The proposed mechanism is not formulated as a classical signal, symbol, or material carrier propagating from Alice to Bob after separation. Instead, Bob acquires a local record R_B and applies a pre-agreed statistical decision rule to infer which hypothesis concerning Alice's later operation is better supported by his data.

Operationally, the proposed chain is

$$\text{Alice's local choice} \longrightarrow \text{Bob's local statistical record} \longrightarrow \text{large-}N\text{ inference} \longrightarrow \text{decoded conclusion.} \tag{1}$$

If repeated validated inferences can be associated through a pre-shared codebook with symbols, words, or commands, then two human users obtain an application-level communication capability.

In this sense,

$$\boxed{\text{statistical inference is the proposed mechanism, while human communication is a conditional emergent use.}} \tag{2}$$

This distinction is operationally important but does not weaken the information-theoretic test. If Bob can infer Alice's later free choice above chance from his local record, then that record contains information about the choice:

$$P(\hat{X} = X) > \frac{1}{2} \quad \Longrightarrow \quad I(X; R_B) > 0 \tag{3}$$

for a genuinely informative binary decision problem. Thus the terminology "inference" is not used to evade no-signaling. Rather, it specifies the architecture to be tested while retaining the standard no-signaling prediction as the null hypothesis.

3 Relation to Prior Work

The framework sits at the intersection of several established lines of research. First, the no-communication/no-signaling structure of quantum mechanics provides the null prediction: local trace-preserving operations do not change the remote reduced state [14, 7]. Second, Bell and CHSH theory, together with modern loophole-free experiments, establish nonclassical correlations under stringent locality and random-setting conditions while preserving the distinction between nonlocal correlation and controllable signaling [3, 4, 5, 6, 15, 16]. Third, quantum state and channel discrimination provide the mathematical language for distinguishability, optimal decision rules, adaptive protocols, and error bounds [8, 10, 17, 18]. Fourth, weak and continuous measurement theory motivates the use of complete time-resolved local records rather than only single-shot means [11, 12]. Fifth, classical information theory and finite-sample mutual-information analysis are required before a measured statistical dependence can be converted into a defensible capacity claim [13, 20, 21].

Recent high-statistics tests on public quantum processors are particularly relevant experimentally. Rybotycki et al. reported statistically significant no-signaling deviations on IBM Quantum devices, including idle configurations, while explicitly noting that their space and time scales do not exclude ordinary subluminal communication and that technical imperfections remain a central concern [2]. Their result therefore motivates stronger controls and strict spacelike validation; it is not taken here as evidence that entanglement-only inference or communication has already been demonstrated. Modern statistical analyses of Bell experiments further show the importance of likelihood-based testing and explicit treatment of deviations from no-signaling equalities [19].

The present contribution is not a new derivation of Bell nonlocality, the no-signaling theorem, or quantum channel discrimination. Its proposed novelty is the integration of these ingredients into a technology-oriented validation hierarchy: full-distribution remote-action testing, entanglement-differential controls, blind inference, conservative finite-sample information bounds, causal isolation, and a conditional transition from a validated physical signature to human-readable communication.

4 Physical Baseline: No-Signaling

Let ρ_{AB} be a bipartite state and Λ_A a CPTP map applied locally by Alice. Bob's reduced state is

$$\rho'_B = \mathrm{Tr}_A[(\Lambda_A \otimes I_B)\rho_{AB}] = \mathrm{Tr}_A(\rho_{AB}) = \rho_B.$$

For every POVM $\{E_y\}$ available to Bob, $P(y|x) = \mathrm{Tr}(E_y \rho_B^{(x)})$ is independent of Alice's local choice x. Hence

$$P(Y|X=0) = P(Y|X=1), \qquad I(X;Y) = 0.$$

The same applies to Bob's complete register B^N, so collective measurements, weak-measurement records, higher moments, and classifiers remain subject to the same null prediction.

5 Central No-Signaling Equivalence Theorem

Theorem. Let $X \in \{0,1\}$ denote Alice's freely selected local CPTP operation. If Bob's complete reduced states satisfy

$$\rho^{(0)}_{B^N} = \rho^{(1)}_{B^N},$$

then every Bob-only record has identical conditional laws,

$$P(R_B|X=0) = P(R_B|X=1),$$

and consequently

$$D_{\mathrm{TV}} = D_{\mathrm{KL}} = C_{\mathrm{Ch}} = I(X;R_B) = C = 0.$$

This follows because equal density operators generate equal outcome distributions for every POVM. Thus large ensembles can amplify an existing nonzero distinguishability, but cannot create distinguishability when the underlying conditional states are identical.

6 Operational Architecture

Before separation Alice and Bob may establish entangled resources, synchronized clocks, trial boundaries, codebooks, measurement bases, calibration rules, and pre-registered analyses. These resources must not encode Alice's later free choice. After separation Alice selects $X_i \in \{0,1\}$, for example I_A or a pre-agreed local operation U_A. Bob records $R_{B,i}$. The empirical question is

$$P(R_B|X=0) \stackrel{?}{=} P(R_B|X=1).$$

7 Primary Statistical Model

For a binary Bob outcome define

$$p_0 = P(Y=+1|X=0), \quad p_1 = P(Y=+1|X=1), \quad S = p_1 - p_0.$$

The null is $H_0 : S = 0$. With N_0, N_1 trials, $\hat{S} = \hat{p}_1 - \hat{p}_0$ and

$$\widehat{\mathrm{SE}}(\hat{S}) = \sqrt{\frac{\hat{p}_1(1-\hat{p}_1)}{N_1} + \frac{\hat{p}_0(1-\hat{p}_0)}{N_0}}.$$

Large N reduces uncertainty; it does not create nonzero S from identical distributions.

8 Full-Distribution Detection

Let $f_0(r) = P(R_B = r|X=0)$ and $f_1(r) = P(R_B = r|X=1)$. Define

$$T = D_{\mathrm{TV}}(f_0, f_1) = \frac{1}{2}\int |f_1(r) - f_0(r)|\, dr.$$

Under exact no-signaling, $T = 0$. A second measure is

$$D_{\mathrm{KL}}(f_1 \| f_0) = \int f_1(r) \log \frac{f_1(r)}{f_0(r)}\, dr.$$

This formulation tests variance, multimodality, higher moments, autocorrelation, spectral structure, and multivariate temporal patterns without assuming a mean shift.

9 Weak and Continuous Measurements

For continuous measurement let $R_B = \{Y(t) : 0 \leq t \leq T\}$. The relevant object is the complete probability functional $\mathcal{P}[R_B|X]$. Weak measurements can expose time-resolved structure with limited disturbance, but remain local measurements and inherit the no-signaling null prediction.

10 Entanglement-Specific Differential Controls

Define four modes: E (entangled), U (unentangled pulse-matched), I (idle), and N (hardware-null). For each mode m, define $S_m = p_{1,m} - p_{0,m}$. The elementary differential statistic is

$$D_E = S_E - S_U.$$

For complete records define

$$\Gamma = D_{\mathrm{TV}}(P_E(R_B|0), P_E(R_B|1)) - D_{\mathrm{TV}}(P_U(R_B|0), P_U(R_B|1)).$$

A candidate entanglement-specific anomaly requires more than $S_E \neq 0$: it must exceed matched controls and replicate independently.

11 Pre-Registered Primary Endpoint and Uncertainty

To avoid an implicit multiple-testing strategy, the confirmatory experiment must designate one primary endpoint before data acquisition. For high-dimensional local records, we define the primary endpoint as the entanglement-differential full-record distinguishability

$$\Gamma = T_E - T_U, \qquad T_m = D_{\mathrm{TV}}(P_m(R_B|0), P_m(R_B|1)), \tag{4}$$

with $m \in \{E, U\}$. The confirmatory alternative is not merely $\hat{\Gamma} > 0$, but

$$\Gamma_{\mathrm{lower}}^{95\%} > 0, \tag{5}$$

where the lower bound is constructed by a pre-specified resampling, permutation, or finite-sample procedure whose assumptions match the record model.

The binary mean-shift statistics S_E and D_E remain interpretable secondary endpoints. Additional spectral, higher-moment, and machine-learning features are exploratory unless explicitly pre-registered and multiplicity-adjusted.

For continuous or high-dimensional records, naive plug-in estimation of total-variation distance can be strongly biased. The analysis must therefore pre-specify either (i) a finite feature map and density/distribution estimator validated by simulation under the null, or (ii) a classifier-based two-sample statistic with nested training and untouched test data. In either case, the reported quantity must include uncertainty calibration under the complete null pipeline.

12 Randomization and Blinding

Alice's choices should be unpredictable before each trial, ideally $X_i \sim \mathrm{Bernoulli}(1/2)$. Experimental mode should also be randomized where possible. Labels remain concealed during acquisition and primary preprocessing. Exclusion rules, stopping rules, feature sets, significance thresholds, and analysis code are fixed before unblinding. Multiple-comparison correction is mandatory.

13 Power and Sample-Size Planning

The confirmatory sample size must be fixed from a pre-specified minimum scientifically relevant effect rather than chosen after inspection of the data. For the binary endpoint, let

$$\delta_{\text{min}} = |p_1 - p_0| \tag{6}$$

be the smallest conditional probability difference considered worth detecting. For approximately balanced groups and probabilities near p, a normal-approximation planning rule is

$$N_{\text{per group}} \approx \frac{2p(1-p)\left(z_{1-\alpha/2} + z_{1-\beta}\right)^2}{\delta_{\text{min}}^2}, \tag{7}$$

with a larger pre-registered value used when discreteness, dependence, overdispersion, or multiplicity requires it. The protocol should state α, desired power $1-\beta$, δ_{min}, expected loss/exclusion rate, and the resulting total number of trials before acquisition.

For the primary full-record endpoint Γ, analytic power may be unavailable. In that case sample-size planning must be performed by simulation from a frozen generative/null model and a family of pre-specified alternatives, running the complete estimator and confidence-bound pipeline end to end. The simulation code and assumptions should be archived with the preregistration.

14 Likelihood Decoder

If independent calibration establishes distinct conditional laws, define

$$\Lambda(R_B^{(n)}) = \log \frac{P(R_B^{(n)}|X=1)}{P(R_B^{(n)}|X=0)}.$$

With equal priors, decide $\hat{X} = 1$ if $\Lambda > 0$. For independent binary outcomes with k positive outcomes,

$$\Lambda(k) = k \log \frac{p_1}{p_0} + (n-k) \log \frac{1-p_1}{1-p_0}.$$

The decoder must be evaluated on untouched blind data.

15 Ensemble Scaling and Error Exponents

Suppose $p_1 = p + \delta$ and $p_0 = p$, $|\delta| \ll 1$. Then

$$D_{\text{KL}}(P_1 \| P_0) \approx \frac{\delta^2}{2p(1-p)}.$$

Chernoff information is

$$C_{\text{Ch}} = -\min_{0 \le s \le 1} \log \sum_y P_0(y)^s P_1(y)^{1-s}.$$

If $P_0 \neq P_1$, decoding error may decrease approximately as $e^{-nC_{\text{Ch}}}$. If $P_0 = P_1$, then $C_{\text{Ch}} = 0$ and no ensemble size improves decoding beyond chance.

16 Information-Theoretic Criterion

For a stationary memoryless empirical inference channel with one local record R_B per input choice X, the single-use capacity is

$$C_1 = \max_{P(X)} I(X; R_B). \tag{8}$$

The present framework, however, also permits continuous records, autocorrelation, temporal dependence, and memory across trials. For a general block process, define

$$C_n = \frac{1}{n} \max_{P(X^n)} I(X^n; R_B^n), \tag{9}$$

and the asymptotic capacity rate

$$\boxed{C = \limsup_{n\to\infty} \frac{1}{n} \max_{P(X^n)} I(X^n; R_B^n)}. \tag{10}$$

When the empirically validated process is stationary and memoryless, this reduces to $C = C_1$. Under exact no-signaling, Bob's complete conditional law is independent of Alice's input sequence, so

$$P(R_B^n|X^n) = P(R_B^n) \quad \Longrightarrow \quad I(X^n; R_B^n) = 0 \tag{11}$$

for every n, and therefore $C = 0$.

A reproducible conditional difference in an apparatus is not by itself a technological claim. The relevant inference process must support a strictly positive conservative lower bound on the appropriate information rate. For experimental comparison, define

$$\Delta C = C_E - C_U, \tag{12}$$

as a differential diagnostic metric, not a new standard quantum-capacity definition.

17 Machine-Learning Secondary Decoder

A classifier $g_\theta(R_B) \mapsto \hat{X}$ may search for nonlinear signatures in high-dimensional records. Data must be separated into training, validation, and untouched test sets. The decisive criterion is blind-test performance above chance with pre-specified confidence bounds. Machine learning is a detector, not evidence by itself.

18 Causal Isolation and Spacelike Validation

Let L be spatial separation and Δt the interval from Alice's completed random setting choice to Bob's completed measurement event. A spacelike arrangement requires

$$c|\Delta t| < L.$$

With timing uncertainty σ_t and safety factor $\eta < 1$, use

$$c(|\Delta t| + k\sigma_t) < \eta L.$$

Verified spacelike geometry excludes direct causal influence propagating at or below c from Alice's setting event to Bob's measurement event. It does not by itself exclude common causes, pre-existing correlations, synchronization leakage, shared electronics, or instrumental coupling; those possibilities must be addressed separately by randomization, physical isolation, telemetry, matched controls, and independent replication.

19 Three-Gate Validation Hierarchy

A candidate effect advances only through the following pre-specified gates:

1. **Gate 1 — blind primary statistical validation.** The frozen primary full-record analysis must pass its pre-registered confirmatory criterion on untouched blind data, with the estimator, uncertainty procedure, sample size, and multiplicity plan fixed before acquisition. The binary quantities S_E and D_E remain secondary endpoints unless explicitly designated otherwise in advance.

2. **Gate 2 — entanglement-specific differential validation.** The conservative lower confidence bound for the primary entanglement-differential endpoint must satisfy

$$\boxed{\Gamma^{95\%}_{\text{lower}} > 0}, \tag{13}$$

relative to the pulse-matched unentangled control, while idle and hardware-null modes remain compatible with their pre-registered null ranges.

3. **Gate 3 — causal isolation and independent replication.** The effect must persist under verified spacelike separation and independent replication. Spacelike geometry constrains direct Alice-to-Bob influence propagating at or below c; common causes, pre-existing correlations, synchronization leakage, shared electronics, and other instrumental alternatives must be separately constrained by the control architecture.

Failure at any gate prevents interpretation as an entanglement-specific remote-action inference effect. Only after all three gates are passed may information-rate estimation and technological decoding be treated as confirmatory rather than exploratory.

20 From Statistical Inference to Human-Level Message Reconstruction

If, contrary to the standard null prediction, calibration establishes a stable conditional channel, Alice may encode symbols in blocks and Bob may decode from local records. Operational success is

$$P(\hat{X} = X) > \frac{1}{2} + \epsilon$$

on blind data, followed by error-correcting coding. At the human interface this would appear as message reconstruction from statistical evidence. Such a result would not evade no-signaling by terminology: if Bob's local data contain information about Alice's later free choice, then $I(X; R_B) > 0$. A successful spacelike result would therefore be a fundamental physics result requiring extraordinary validation.

21 Pre-Shared Keys and Synchronization

A pre-shared key K can define codewords, timing, bases, and interpretation rules, but cannot by itself create information about a later independent choice X. If

$$P(R_B|X, K) = P(R_B|K),$$

then $I(X; R_B|K) = 0$. The key becomes useful for decoding only if the physical record already contains a conditional difference or if post-choice side information is supplied.

22 Adaptive Bell-State and Noise-Aware Layer

Several technological components remain meaningful regardless of the signaling test. The preparation system may select among Bell states or other resources according to measured noise, decoherence, readout fidelity, or hardware asymmetry. If q indexes candidate resources and ν a noise profile, define

$$q^*(\nu) = \arg\max_q F(q, \nu),$$

where F is a pre-defined fidelity or task-performance functional.

23 Experimental Protocol

A rigorous implementation should: (1) prepare entangled and matched controls under randomized labels; (2) generate Alice's setting unpredictably after preparation; (3) record Bob's raw stream without Alice labels; (4) preserve timestamps, pulse logs, calibration and environmental telemetry; (5) apply pre-registered binary and full-distribution tests; (6) repeat in unentangled, idle, and null modes; (7) use permutation and time-shift controls; (8) replicate across devices and laboratories; (9) only after reproducible differential evidence implement a blind likelihood decoder; and (10) finally repeat under verified spacelike separation.

24 Interpretation of Outcomes

Four outcomes are distinguished. First, $S_E = 0$ and controls are null: standard no-signaling is confirmed. Second, $S_E \neq 0$ but controls show a similar effect: hardware coupling, drift, or leakage is favored. Third, an entanglement-specific differential effect appears without spacelike separation: the result is interesting but ordinary causal influence remains open. Fourth, an entanglement-specific effect survives spacelike separation and independent replication: this is an anomalous result requiring exhaustive scrutiny and potentially a new physical explanation.

25 Relation to the Earlier Statistical-Inference Proposal

The framework preserves the original intuition that a very small reproducible effect, if physically present, can be extracted by large-ensemble statistics. It changes the logical order. Rather than assuming a remote statistical signature and optimizing inference, the new framework first tests whether the signature exists, then whether it is entanglement-specific, and finally whether it survives causal isolation. Data rate and message decoding are derived only after positive empirical channel capacity is established.

26 Limitations

The framework does not prove that entanglement-only message reconstruction is possible. Standard quantum mechanics predicts the null result. A positive result on a non-spacelike quantum processor is not evidence of superluminal signaling because crosstalk, common-mode electronics, control leakage, timing correlations, and other causal mechanisms remain possible. Conversely, a null result does not invalidate the statistical architecture; it constrains the class of mechanisms available for exploitation.

27 Terminological and Physical Interpretation

Throughout this work, "remote-action inference" denotes a statistical hypothesis test performed locally by Bob. "Communication" denotes only the higher-level human use that would become possible if a sequence of such inferences were reproducibly informative and could be mapped to symbols.

No claim rests on a semantic distinction between inference and communication. If Bob's record remains independent of Alice's later choice, then no amount of inference, coding, machine learning, or ensemble averaging creates information. Conversely, if Bob's local record becomes reproducibly informative about Alice's later choice, the resulting positive mutual information is physically significant regardless of whether the user-facing process is described as inference, reconstruction, or communication.

Accordingly, the framework is conservative in its physics and precise in its engineering interpretation: the mechanism under test is statistical inference; a communication capability is only a conditional application of a validated positive-capacity inference process.

28 Conditional Communication and Technology Horizon

The framework deliberately separates a test of physics from a technological claim. Under standard quantum mechanics, the null prediction remains zero locally accessible information about Alice's later free choice. Consequently, if the null prediction is confirmed, entanglement-only message reconstruction is unavailable within the tested architecture.

If, however, a reproducible conditional difference is observed, technological interpretation is activated only after the full validation hierarchy is passed. Define

$$\Gamma^{95\%}_{\text{lower}} > 0 \Longrightarrow \text{independent spacelike replication} \Longrightarrow C^{95\%}_{\text{lower}} > 0 \Longrightarrow \text{reliable blind decoding.} \quad (14)$$

Here $C^{95\%}_{\text{lower}}$ denotes a conservative 95% lower confidence bound on the empirically estimated channel-capacity rate appropriate to the validated record process. Because mutual-information estimation from finite data can be biased or statistically fragile, the estimator and confidence construction must be pre-specified and validated for the relevant alphabet or record model; distribution-free or finite-alphabet bounds should be preferred when their assumptions apply [20, 21]. A minimum technology-activation criterion is

$$C^{95\%}_{\text{lower}} > 0 \quad (15)$$

together with post-error-correction blind decoding satisfying, for example,

$$P_{\text{blind}}(\hat{X} = X) \geq 0.99. \quad (16)$$

The value 0.99 is an illustrative engineering target, not a fundamental physical threshold. A different application-specific target may be used if it is justified and pre-registered before the final blind test.

If these conditions are reproducibly satisfied, the statistical layer provides an operational route

$$\text{weak conditional signature} \rightarrow \text{large-}N\text{ accumulation} \rightarrow \text{likelihood decoding} \rightarrow \text{error correction} \rightarrow \text{symbols} \rightarrow \text{hu} \quad (17)$$

At the human interface, Alice and Bob would then possess an operational communication system in which the conventional physical synchronization and resource-distribution stage occurred before separation, while no conventional communication channel would be active during the later inference stage.

This is a conditional technological statement, not a prior assertion that standard no-signaling quantum mechanics already permits such communication. A positive entanglement-specific effect surviving verified spacelike separation would itself be a fundamental empirical result. Before attributing it to a new quantum mechanism, experiments must exclude leakage, ordinary causal coupling, common-mode electronics, timing correlations, selection effects, calibration drift, and statistical artifacts, and the result must be independently replicated.

Thus the assumptions of the framework remain consistent with established mathematics and physics. The empirical outcome is intentionally left open. If the standard null is observed, no entanglement-only channel is obtained. If a positive-capacity, entanglement-specific, spacelike-replicated effect is observed, the mathematics developed here specifies how that measured effect could be converted into reliable communication while the physical origin of the effect becomes a separate foundational question.

29 Reference Scope and Evidentiary Use

The references in this paper serve distinct evidentiary roles. Foundational sources support the standard quantum-mechanical null model and Bell/nonlocality context; state- and channel-discrimination sources support the statistical decision framework; weak-measurement sources support the measurement architecture; and information-theoretic sources support the transition from empirical distinguishability to finite-sample information claims. None of the cited Bell, weak-measurement, or channel-discrimination results is treated as prior evidence that Bob can infer Alice's later free choice from a spacelike-separated local record. The Rybotycki et al. result is used specifically as motivation for high-statistics controls and causal isolation, not as confirmation of the proposed effect.

30 Conclusion

We have formulated a falsifiable mathematical architecture for testing statistical remote-action inference in pre-established entangled systems without assuming a violation of known physics. The central theorem connects equality of Bob's conditional reduced states to equality of all Bob-only measurement laws and therefore to vanishing total-variation distance, KL divergence, Chernoff information, mutual information, and channel capacity.

The experimental architecture then asks whether an actual physical implementation departs reproducibly from this null prediction. Any candidate effect must first be statistically established on blind data, then shown to be specific to the entangled condition relative to matched controls, and finally survive verified spacelike separation and independent replication.

Only after these gates are passed does the framework permit a technological interpretation. A strictly positive conservative lower bound on channel capacity, combined with reliable blind decoding after error correction, would provide the mathematical bridge from a weak measured signature to human-readable communication without a conventional post-separation communication channel. Such an outcome would not retroactively make it a consequence of the standard no-signaling model; it would constitute a new empirical fact whose physical mechanism would require independent determination.

The framework therefore has two scientifically well-defined outcomes. A null result reinforces the standard no-signaling prediction and constrains the proposed architecture. A validated positive result supplies both a foundational physics anomaly and, through the statistical and information-theoretic layer developed here, a conditional route toward a new communication technology.

References

[1] Y. Fradkin, *Quantum Correlation Networking / Information Extraction from a Quantum System for Indirect Information Transfer*, 2025, prior manuscript underlying the present framework. No external persistent identifier is supplied here.

[2] T. Rybotycki, T. Bialecki, J. Batle, and A. Bednorz, "Violation of No-Signaling on a Public Quantum Computer," *Advanced Quantum Technologies* **8**, 2400661 (2025). doi:10.1002/qute.202400661.

[3] J. S. Bell, "On the Einstein Podolsky Rosen Paradox," *Physics Physique Fizika* **1**, 195–200 (1964).

[4] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, "Proposed Experiment to Test Local Hidden-Variable Theories," *Physical Review Letters* **23**, 880–884 (1969).

[5] B. Hensen et al., "Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres," *Nature* **526**, 682–686 (2015). doi:10.1038/nature15759.

[6] B. Hensen et al., "Loophole-free Bell test using electron spins in diamond: second experiment and additional analysis," *Scientific Reports* **6**, 30289 (2016). doi:10.1038/srep30289.

[7] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th Anniversary ed., Cambridge University Press (2010; print publication 2011).

[8] C. W. Helstrom, *Quantum Detection and Estimation Theory*, Academic Press (1976).

[9] A. S. Holevo, *Probabilistic and Statistical Aspects of Quantum Theory*, North-Holland (1982).

[10] K. M. R. Audenaert et al., "Discriminating States: The Quantum Chernoff Bound," *Physical Review Letters* **98**, 160501 (2007).

[11] Y. Aharonov, D. Z. Albert, and L. Vaidman, "How the Result of a Measurement of a Component of the Spin of a Spin-1/2 Particle Can Turn Out to be 100," *Physical Review Letters* **60**, 1351–1354 (1988).

[12] H. M. Wiseman and G. J. Milburn, *Quantum Measurement and Control*, Cambridge University Press (2010).

[13] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed., Wiley (2006).

[14] G. C. Ghirardi, A. Rimini, and T. Weber, "A general argument against superluminal transmission through the quantum mechanical measurement process," *Lettere al Nuovo Cimento* **27**, 293–298 (1980). doi:10.1007/BF02817189.

[15] L. K. Shalm et al., "Strong Loophole-Free Test of Local Realism," *Physical Review Letters* **115**, 250402 (2015). doi:10.1103/PhysRevLett.115.250402.

[16] M. Giustina et al., "Significant-Loophole-Free Test of Bell's Theorem with Entangled Photons," *Physical Review Letters* **115**, 250401 (2015). doi:10.1103/PhysRevLett.115.250401.

[17] M. Piani and J. Watrous, "All Entangled States are Useful for Channel Discrimination," *Physical Review Letters* **102**, 250501 (2009). doi:10.1103/PhysRevLett.102.250501.

[18] S. Pirandola, R. Laurenza, C. Lupo, and J. L. Pereira, "Fundamental limits to quantum channel discrimination," *npj Quantum Information* **5**, 50 (2019). doi:10.1038/s41534-019-0162-y.

[19] R. D. Gill, "Optimal Statistical Analyses of Bell Experiments," *AppliedMath* **3**, 446–460 (2023). doi:10.3390/appliedmath3020023.

[20] A. G. Stefani, J. B. Huber, C. Jardin, and H. Sticht, "Confidence Intervals for the Mutual Information," *International Journal of Machine Intelligence and Sensory Signal Processing* **1**(3), 202–212 (2014). doi:10.1504/IJMISSP.2014.066430.

[21] D. McAllester and K. Stratos, "Formal Limitations on the Measurement of Mutual Information," in *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, PMLR **108**, 875–884 (2020).

A Core Validation Quantities

For the binary and full-record tests,

$$S = p_1 - p_0, \qquad T = D_{\mathrm{TV}}(P(R_B|0), P(R_B|1)), \tag{18}$$

$$D_E = S_E - S_U, \qquad \Gamma = T_E - T_U. \tag{19}$$

For a general process with possible memory, the information-rate criterion is

$$\boxed{C = \limsup_{n\to\infty} \frac{1}{n} \max_{P(X^n)} I(X^n; R_B^n)}. \tag{20}$$

For a stationary memoryless process this reduces to

$$C = \max_{P(X)} I(X; R_B). \tag{21}$$

The standard no-signaling prediction for an ideal local-operation experiment with no post-choice physical channel from Alice to Bob is

$$S = T = D_E = \Gamma = C = 0. \tag{22}$$

B Minimum Decision Rule

A technological decoding claim should not be made unless all of the following conditions are met:

1. the primary endpoint, estimator, uncertainty procedure, sample size, and multiplicity plan are pre-registered;
2. a blind primary effect is independently replicated;
3. the entanglement-differential lower bound satisfies $\Gamma_{\text{lower}}^{95\%} > 0$ and the result exceeds matched idle and hardware-null controls;
4. a decoder fixed before final unblinding predicts Alice's setting on untouched data;
5. the effect persists across devices and independent laboratories;
6. verified spacelike geometry excludes direct Alice-to-Bob influence at or below c, while common-cause and instrumental alternatives are separately constrained;
7. the conservative capacity bound satisfies
$$C_{\text{lower}}^{95\%} > 0;$$
8. post-error-correction blind decoding reaches the pre-registered technological target; for the illustrative target used in the main text,
$$P_{\text{blind}}(\hat{X} = X) \geq 0.99,$$
with the understanding that another justified application-specific threshold may be substituted if fixed before the final blind test.

C Inference–Communication Distinction: Compact Statement

For the purposes of this framework, define:

$$\mathcal{M}_{\text{phys}} : \text{the physical preparation and measurement process,} \tag{23}$$
$$\mathcal{I} : R_B \mapsto \hat{X}, \quad \text{the local statistical inference map,} \tag{24}$$
$$K : \text{the pre-shared codebook,} \tag{25}$$
$$\mathcal{H} : (\hat{X}_1, \ldots, \hat{X}_m, K) \mapsto \text{human-readable message.} \tag{26}$$

The proposed research target is the empirical validation of $\mathcal{I}$, not the assumption of a post-separation message carrier. If

$$I(X; R_B) = 0, \tag{27}$$

then $\mathcal{I}$ cannot outperform chance and $\mathcal{H}$ cannot create a message about Alice's later choices. If instead a rigorously validated experiment yields

$$I(X; R_B) > 0, \tag{28}$$

then repeated applications of $\mathcal{I}$ combined with K and error correction may support $\mathcal{H}$. The latter is the sense in which a human-level communication capability is an emergent application of a statistically informative inference mechanism.