

What Does a Bell Test Preserve?

Recursive Structure, Information Sufficiency, and a Falsifiable Research Program for Scientific Representation

Jeffery Scott Allbright | AIIA Technologies, Inc.

Abstract

Bell tests establish powerful constraints on local hidden-variable models through experimentally estimated relationships among measurement settings and outcomes. Those conclusions do not require that the underlying event stream be maximally informative for every other scientific question. This Perspective proposes a distinct and falsifiable question: are the conventional variables and conditional outcome statistics used for the Bell inference information-sufficient with respect to all reproducible organization in the generating experimental process, or can an alternative representation recover nonredundant, independently verifiable information from the same observations?

Public applications of Recursive Indexing (RIX) to four NIST Bell-test event streams report stable recursive structure, reproducible run families, early diagnostic convergence, and setting-dependent path-entropy organization. These findings do not overturn Bell's theorem, invalidate reported Bell violations, establish local realism, or constitute a new theory of quantum mechanics. They establish a motivating observation: ordered event records contain structure that can be represented and measured in a way not directly displayed by conventional aggregate summaries.

We formulate the problem explicitly as two transformations of the same event stream, X : a conventional Bell representation $C(X)$ and a recursive representation $R(X)=Z$. The scientific question is not whether Z contains structure, but whether Z contains information about an independently defined experimental property Y beyond that recoverable from $C(X)$ and known experimental metadata. We propose a falsification-first program centered on probability-preserving surrogate event streams, temporal destruction controls, synthetic nulls, apparatus and preprocessing controls, independent preregistered replication, and prospective prediction. This framework converts representational completeness from a philosophical assertion into an empirical comparison of information sufficiency.

1. Two Questions, Not One

A Bell experiment is designed to answer a specific physical question. Experimental events are represented through measurement settings, outcomes, pairing and timing rules, and the conditional probabilities or related statistics required for a Bell inequality. A valid violation constrains specified classes of local hidden-variable models. Nothing in the present research program disputes that logic.

A second question can nevertheless be asked of the same experimental record: is the representation sufficient for every scientifically meaningful property of the generating process that is present in the recorded events? These questions should not be conflated. A representation can be sufficient for testing a Bell inequality while being insufficient for reconstructing temporal organization, diagnosing apparatus state, identifying regime transitions, or predicting another independently measured property.

The research program therefore does not begin from the proposition that the conventional Bell representation is incomplete. It begins from a null hypothesis: once the conventional Bell variables, relevant conditional probabilities, and known experimental metadata are preserved, an alternative representation should contain no additional independently meaningful information. The objective is to design experiments capable of rejecting or failing to reject that hypothesis.

2. Making the Computational Question Explicit: $X \rightarrow Z \rightarrow Y$

A useful criticism of the existing Bell studies is that the computational roles of input, representation, and target have not always been stated with sufficient precision. RIX should not be described simply as a supervised predictor $Y=f(X)$. Its primary operation is a representation transformation.

X - observed event stream. X is the ordered set of experimental event records retained in the public Bell-test data, including the recorded variables used by the analysis.

R - recursive representation map. R maps X into a structured state and transport representation using recursively defined basins, admissible transitions, paths, and derived structural statistics.

Z = R(X) - structural representation. Z is the resulting collection of recursive and transport features, including quantities such as basin organization, transition concentration, path structure, and path entropy.

Y - independently defined target, when prediction is tested. Y is not required to construct R(X). In downstream experiments it may be a run family, apparatus condition, regime, data-quality state, or another independently defined experimental property.

Some existing RIX studies are therefore predictive: they ask whether Z discriminates a predefined Y and how early sufficient evidence emerges. Others are descriptive or convergence studies: they ask whether properties of Z stabilize reproducibly as increasing prefixes of X are observed. These are different experiments and should be labeled accordingly.

3. The Actual Comparison: C(X) Versus R(X)

The central scientific comparison is not RIX versus Bell's theorem. It is between two representations derived from the same underlying observations.

C(X): the conventional representation used for the Bell inference, including the relevant settings, outcomes, pairings, conditional probabilities, and other prespecified conventional quantities.

R(X)=Z: a recursive representation retaining order, transitions, paths, and derived structural organization.

The decisive question is incremental information: does R(X) contain reproducible information about Y beyond C(X) and known metadata M? In predictive form, the relevant comparison is not whether R(X) predicts Y, but whether adding R(X) improves held-out prediction of Y after C(X) and M are already available. In information terms, the research target is evidence that the recursive representation contains conditional information about Y that is not exhausted by the conventional representation and known experimental variables.

This formulation prevents a common interpretive error. Demonstrating that R(X) has structure is insufficient. Demonstrating that R(X) classifies runs is also insufficient if the same classification follows

from known apparatus variables or conventional statistics. The additional representation becomes scientifically consequential only when its information is demonstrably nonredundant and independently verifiable.

4. Operational Representational Completeness

Operational definition. For a specified experimental question Y , a representation $C(X)$ is operationally incomplete if another representation of the same underlying observations contains reproducible information about Y that remains after conditioning on $C(X)$ and known experimental metadata, survives appropriate null and destruction controls, replicates independently, and cannot be attributed to the construction of the alternative representation itself.

This is deliberately narrower than philosophical or ontological notions of completeness. Completeness is always indexed to a question. $C(X)$ may be entirely sufficient for estimating a Bell statistic even if $R(X)$ proves useful for another task. Conversely, if every apparent advantage of $R(X)$ is reproduced by $C(X)$, known metadata, synthetic controls, or probability-preserving surrogates, then the stronger representational claim fails.

5. Motivating Evidence from the Existing Bell Studies

The existing studies motivate, but do not resolve, this question. Analyses of four public NIST HDF5 event streams report a stable 16-basin apparatus-level scaffold and reproducible run families. Scaffold measures stabilized by approximately 1% of event rows, while run-family classification stabilized at approximately 1-2%. The reported average reduction in event rows needed for family classification was 98.75%. An evidence-gated procedure classified 331 of 368 local samples deemed classifiable with 100% accuracy while abstaining when evidence was weak.

A second analysis reports concentration of transition structure from roughly 0.25% to roughly 40% in the top ten pathways, global path-entropy stabilization by approximately 1% of event rows, and setting-specific differences in path-entropy organization. A conservative path-evidence gate classified 21 of 28 run-prefix samples with 100% accuracy and abstained on the remainder.

These observations show that the ordered streams admit a stable recursive description. They do not establish why that description is stable. In particular, they do not yet demonstrate that the observed structure contains information beyond the conventional Bell variables, apparatus state, acquisition process, preprocessing choices, or generic sequential statistics. That distinction defines the next phase of the research.

6. The Decisive Experiment: Probability-Preserving Surrogate Streams

The strongest immediate test follows directly from the claim that the scientific information relevant to a Bell experiment is captured by the conventional setting-outcome probabilities. Rather than debate that proposition conceptually, it can be tested computationally.

Construct surrogate event streams that preserve, as closely and explicitly as possible, the conventional quantities used for the Bell inference - including the relevant setting frequencies and conditional outcome distributions - while selectively destroying event-order, recursive, and transport organization not required to preserve those quantities. The surrogate generator itself must be prespecified and validated to demonstrate that the Bell-relevant statistics have been retained.

RIX is then applied unchanged to the empirical and surrogate streams. The primary comparison is whether the empirical streams occupy a structural regime distinguishable from probability-preserving surrogates under prespecified RIX measures.

If empirical and surrogate streams produce the same RIX structure, the stronger interpretation weakens substantially. The recursive organization is then plausibly determined by the conventional probabilities, protocol, or generic properties of the data.

If the RIX structure is selectively lost or materially altered in the surrogates, the empirical stream contains organization not determined by the preserved conventional statistics alone. That result would establish additional information, but not yet its physical significance.

This experiment should become the first major gate in the research program because it directly tests information sufficiency while preserving the quantities that motivated the original Bell experiment.

7. A Falsification-First Validation Ladder

7.1 Computational specification and independent verification

The transformation from the observed event stream X to the recursive representation $R(X)$ should be specified sufficiently to establish what information enters the analysis, what transformations are performed conceptually, what quantities are produced, and how reported conclusions are derived. Where implementation details constitute protected intellectual property, independent verification can be performed through a version-locked executable, containerized implementation, controlled computational service, or blinded third-party validation rather than unrestricted release of source code. The scientific requirement is that reported claims be independently testable and falsifiable; it does not require disclosure of proprietary implementation details unrelated to evaluating those claims.

7.2 Temporal destruction controls

Apply global shuffling, within-setting shuffling, block permutation, time reversal, local temporal scrambling, and other targeted transformations. Each control should preserve specified marginal or conditional quantities while destroying a defined class of sequential relationships. The result should identify which aspects of event order are necessary for each reported RIX phenomenon.

7.3 Probability-preserving surrogates

Generate multiple null ensembles that preserve increasingly rich conventional statistics. Begin with setting and outcome frequencies, then preserve the full conditional distributions used for Bell analysis, and where feasible add known timing and apparatus distributions. This creates a hierarchy of nulls rather than a single synthetic comparison.

7.4 Apparatus and metadata controls

Test whether RIX features are explained by known settings, detector states, acquisition blocks, clocks, calibration conditions, file identity, run identity, or other available metadata. Conditional and multivariable analyses should quantify any incremental contribution of $R(X)$ after these variables are included.

7.5 Construction sensitivity

Repeat the analysis under defensible variations in pairing rules, coincidence windows, segmentation, filtering, and preprocessing. A robust phenomenon should have an interpretable relationship to these choices rather than depend on one arbitrary pipeline configuration.

7.6 Conventional-representation baselines

Build explicit baseline predictors using $C(X)$ and known metadata. When a target Y exists, compare $C(X)+M$ with $C(X)+M+R(X)$ on held-out data. This makes the nonredundancy claim quantitative rather than rhetorical.

7.7 Preregistered independent replication

Freeze the representation map, metrics, null models, thresholds, and success/failure criteria before analyzing independent Bell datasets or previously untouched runs. Validation data should not be used to tune the recursive construction.

7.8 Prospective prediction

If the preceding gates are passed, test whether $R(X)$ prospectively predicts an independently verifiable experimental property before that property is revealed or before the complete event stream is available. This is the strongest route from descriptive organization to scientific utility.

8. Decision Rules: What Would Each Outcome Mean?

Outcome	Interpretation
Independent researchers cannot reproduce $X \rightarrow R(X)$	The current evidence cannot support stronger scientific interpretation; specification and implementation must be corrected first.
Temporal controls eliminate the signal	RIX is detecting order-dependent information. This is informative about the representation but not by itself evidence of additional physics.
Probability-preserving surrogates reproduce the empirical RIX structure	The structure is plausibly entailed by conventional Bell statistics, protocol, or generic sequential properties; the information-sufficiency hypothesis is not rejected.
Known apparatus variables or metadata fully explain the signal	The structure is instrumental or procedural. RIX may still be useful diagnostically, but no stronger representational claim follows.
The signal survives surrogates and controls but fails independent replication	The phenomenon is dataset- or apparatus-specific.
The signal survives controls and preregistered independent replication	The event streams contain reproducible organization not exhausted by the tested conventional representation and known metadata.
$R(X)$ adds held-out predictive information about an independent Y beyond $C(X)+\text{metadata}$	There is evidence of operational representational incompleteness for that specified question.
The added information predicts a separately measured physical quantity and survives independent experimental replication	A physical interpretation becomes scientifically warranted, although this still would not by itself refute Bell's theorem or establish local realism.

9. What This Program Does Not Claim

The research program does not presently claim a novel theory of quantum mechanics. It does not claim that Bell's theorem is mathematically incorrect, that reported Bell violations are invalid, that local realism has been restored, or that RIX has identified a hidden physical variable. It also does not assume that every reproducible event-level pattern deserves physical interpretation.

Its claim is methodological: information sufficiency can be tested. If two representations are derived from the same observations, one can ask whether the second contains reproducible and independently useful information after conditioning on the first. Bell event streams provide an unusually demanding environment in which to perform that test.

10. Relationship to the Broader Representation Program

Recursive transport analysis has also been applied to biological fitness landscapes, learned representations, ESM-2 embeddings, and world-model latent spaces. Those studies provide methodological context by showing that representation-dependent organization can be quantified, compared across systems, and allowed to fail. Recent protein-transfer analyses further suggest that alignment of site-level sensitivity structure can be associated with cross-system transfer behavior.

None of those results constitutes evidence about Bell physics. Their relevance is narrower: they motivate a general research strategy in which representations are treated as experimentally comparable objects. Claims about Bell-test data must be established using Bell-test controls, Bell-test nulls, and independent Bell-test replication.

11. Perspective: Bell Tests as a Laboratory for Representation

The appropriate destination is not an immediate declaration that Bell was wrong. It is a better experiment about what Bell tests represent.

The conventional Bell representation was constructed for a specific and extraordinarily important inference. The proposed program asks a different question without diminishing that achievement: when an experimental record is reduced to the variables sufficient for one inference, what other reproducible information may be lost, and how can we distinguish meaningful loss from artifact?

Probability-preserving surrogate streams make this question unusually concrete. If recursive organization is reproduced whenever the Bell-relevant probabilities are preserved, the stronger hypothesis loses support. If empirical streams retain organization that carefully matched surrogates do not, the next task is to determine whether that residual information is instrumental, procedural, computational, or physical. Only prospective, nonredundant prediction of independently verifiable quantities should justify movement toward stronger interpretation.

In this form, Bell-test data become a laboratory for a broader scientific question: not whether a successful representation is wrong, but how we determine the boundaries of what it is sufficient to tell us.

Acknowledgment

The author thanks Mark White, MD, for developing the originating framework of Representation Theory, including the structural analysis of the genetic code that motivated this research program, and for contributing to discussions concerning its broader interpretation.

Public Sources

- AIIA Technologies, “The Science of Representation.”

- Allbright, J. S. “Early and Efficient Regime Detection in Bell-Test Event Streams with Recursive Indexing” (2026).
- Allbright, J. S. “Path-Entropy Convergence in Bell-Test Event Streams” (2026).
- Allbright, J. S. “Recursive Transport Topology: Predictive and Causal Organization in Protein Fitness Landscapes” (2026).
- Allbright, J. S. “Recursive Transport Topology: Diagnostics, Tunable Abstraction, and Latent Enhancement” (2026).
- Allbright, J. S. “Recursive Transport Topology Detects Structural Regime Transitions Associated with Behavioral Failure in World Models” (2026).
- Allbright, J. S. “Recovery Dynamics and Persistent Structural Excursion in Disrupted World Models” (2026).