Records Say What Can Be Known

Empirical access, response functionals, and severity in finite information physics

David Elliman
Neuro-Symbolic Ltd
`dave@neusym.ai`
https://neusym.ai

Draft: 6 July 2026

Abstract

Finite information models naturally invite an epistemological confusion. They specify a substrate ledger: a finite set of admissible records, repair operations, forbidden states, and internal counts. But experiments do not read an arbitrary internal ledger entry. They read calibrated response channels: currents, residues, poles, susceptibilities, cross-sections, spectra, fluxes, and likelihoods. This paper states the resulting epistemic principle: *records say what can be known; responses say what can be measured.* The claim is not offered as a new interpretation of quantum mechanics competing on the same axis as QBism, relational quantum mechanics, or Everettian quantum mechanics. It is a constraint on empirical access in any finite, record-bearing substrate theory. A record is a stable, copyable, dynamically protected fact; a response is the closed-time-path or equivalent operational map by which an instrument couples to such facts. Internal ledger structure may be real in a model while remaining unknowable in principle unless it is written into stable records and exposed by a response. The paper formalises this distinction with a record algebra, a response functional, and an empirical-access equivalence relation. It then compares the view with QBism, relational quantum mechanics, Everett, hidden-variable realism, quantum Darwinism, and engineering calibration practice. Finally it proposes a Mayo-style severity criterion for finite information physics: a numerical claim is credible only when the record object, response map, calibration constants, inherited readouts, and possible falsifiers are stated before the test. Recent applications in the finite-QEC programme—the QED $\alpha(0)$ boundary and Thomson readout, black-hole flux, CMB/halo likelihoods, and electroweak pole matching—are used as case studies. The intended contribution is methodological: the record-response split does not prove a substrate true, but it sharply limits how such a substrate may make empirical claims.

Plain summary

The phrase “finite information model” can sound as if physics is only a matter of counting bits. That is too crude. A finite model can count many internal things, but an experiment reads only what its apparatus is coupled to. The record-response split is the rule that prevents an internal count being silently promoted into a physical observable.

Records define possible knowledge. Responses define possible measurement.
Predictions require the bridge between them.

This is close to an engineering distinction. A system has internal state, but an instrument sees a transfer function. A counter inside a circuit is not the same thing as the voltage on the output pin. In a physical theory, the same discipline says that a state count, a symmetry count, or a code count is not yet a measured charge, cross-section, flux, or mass. The response map has to be supplied.

1 The issue: finite ledgers are not automatically observables

Record-based theories are attractive because many exact physical statements are discrete. A gauge anomaly either cancels or it does not. A code has a particular distance. A finite alphabet has a particular size. A closed gauge surface either admits a colour singlet or it does not. These are natural record-level facts.

They are also dangerous. A finite theory may generate many integers and ratios. Some will lie near known constants. Without a response discipline, one can mistake an internal counter for an experimental reading.

The measurement discipline paper [1] gave this rule as an audit protocol for discrete record-based physics. The present paper isolates the epistemology behind that protocol. The distinction can be put as a simple engineering analogy:

engineering language	finite substrate language	physical-observable language
internal state	record ledger	what may be stably known
sensor placement	apparatus coupling	which record the experiment probes
transfer function	response functional	how the reading is produced
calibration	pole/residue/likelihood normalisation	what number the meter reports
acceptance test	falsifier / severity test	what would kill the claim

This paper proposes that the same split should be treated as foundational for finite information physics. The substrate ledger says what can become a fact; the response functional says what a real experiment measures.

2 Records

Definition 1 (Record)**.** *A* record *is a dynamically stable, copyable, mutually distinguishable state of a monitored system. Algebraically, a record is represented by a projector P_i or by an element of the commutative algebra generated by a set of mutually orthogonal pointer projectors:*

$$P_iP_j = \delta_{ij}P_i, \qquad \sum_i P_i = I, \qquad \mathcal{R} = \mathrm{span}\{P_i\}.$$

The operational content of the definition is stability: once a record is written, later allowed operations may condition on it or copy it, but they do not change the fact that it was written.

This is intentionally close to the language of decoherence and quantum Darwinism. Pointer states are the states robust under environmental monitoring; redundant environmental records explain how classical facts become objective for many observers [2–4]. The present paper adds a ledger emphasis: once a record exists, it can be billed, copied, and referred to, but it is not thereby the same thing as every possible observable involving the system.

Proposition 1 (Record accessibility)**.** *Only record-stable distinctions can become public empirical facts for observers inside the system.*

Argument. A public empirical fact must be copyable into multiple later systems: detector state, memory, notebook, data file, and independent checks. Quantum no-cloning forbids arbitrary unknown states from being copied [5]; decoherence selects a restricted pointer algebra whose states can be redundantly recorded. Therefore public empirical facts are not arbitrary Hilbert-space distinctions but record-stable distinctions. □

The proposition does not say that unrecorded amplitudes are unreal or irrelevant. It says they are not public facts. They may still affect future responses through interference, virtual processes, or dressing. That is why the response layer is needed.

3 Responses

Definition 2 (Response functional)**.** *A* response functional *is an operational map from a record-bearing system and an apparatus coupling to a measured output distribution or correlator. In closed-time-path notation one may write the schematic form*

$$Z[J_+, J_-] = \sum_{h_+, h_- \in \mathcal{H}} \exp\{\mathrm{i}A[h_+] - \mathrm{i}A[h_-] - A_{\text{bill}}[h_+, h_-] + J_+ \cdot R[h_+] - J_- \cdot R[h_-]\} \,,$$

where $h_\pm$ are forward/backward monitored histories, A_{bill} is the record-writing or branch-mismatch cost, and $J_\pm$ encode the apparatus coupling. Measured currents, residues, susceptibilities, cross-sections and likelihood variables are derivatives or transforms of Z, not arbitrary entries in the internal ledger.

This notation is not meant to replace the standard formalisms. In different domains the same role is played by response theory [6, 7], LSZ pole residues [8], scattering cross-sections, Boltzmann likelihoods, greybody transfer functions, or renormalisation-group maps. The common point is that the observable is defined by a coupling and a readout procedure.

Principle 1 (Record-response promotion rule)**.** *A finite record count may be promoted to a physical observable only after a response map proves that the relevant apparatus reads that record object.*

Corollary 1 (Anti-numerology)**.** *An internal count near a measured constant is at most a candidate until the response map is supplied. Conversely, once the response map is fixed, a disagreement with the measured response kills the branch; a different count cannot be introduced ad hoc to repair it.*

4 Partial unknowability and empirical equivalence

The record-response split implies a modest but important form of epistemic humility. A substrate theory may contain internal bookkeeping that is not directly knowable by observers inside the universe.

Definition 3 (Empirical access)**.** *Let S be an internal substrate description, $\mathcal{R}(S)$ its record algebra, and $\mathcal{S}(S, \mathcal{I})$ the family of response functionals exposed by physically available instruments $\mathcal{I}$. Two internal descriptions S_1, S_2 are* empirically equivalent at access class $\mathcal{I}$ *when*

$$\mathcal{R}(S_1) \simeq \mathcal{R}(S_2) \quad \textit{and} \quad \mathcal{S}(S_1, I) = \mathcal{S}(S_2, I) \quad \forall I \in \mathcal{I}.$$

Internal distinctions outside this equivalence class may be meaningful in a model, but they are not empirically available at that access class.

This is not a retreat from realism. It is a statement about access. Many scientific theories already work this way. Thermodynamics does not require knowing the exact microstate of every molecule. Electronics does not expose every internal transistor state at an output pin. Quantum field theory uses virtual processes that affect amplitudes without appearing as directly recorded particles. In each case the empirically accountable object is the response, not the complete internal bookkeeping.

5 Relation to existing quantum interpretations

The record-response view is adjacent to several interpretations of quantum mechanics, but it is not identical to any of them.

view	central emphasis	agreement	difference
QBism [9]	quantum states as agent-centred expectations	measurement and knowledge are not optional afterthoughts	records are physical stabilised facts, not merely personal degrees of belief
Relational QM [10]	facts are relative to interactions	facts are interaction-bound	the mechanism is a record write plus response readout, not relation alone
Everett [11, 12]	universal unitary state and branch structure	more may exist in the accounting than one observer records	emphasis shifts from ontology of all branches to which histories become stable records and which responses expose them
Hidden-variable realism [13]	a more complete underlying state	the substrate may contain unobserved structure	unrecorded structure is not automatically empirically accessible; response equivalence matters
Quantum Darwinism [4]	redundant environmental records explain objectivity	stable records are central	the response layer is added as a separate measurement/calibration axis

The closest conceptual neighbour is quantum Darwinism, because both views make stable records central. The novelty claimed here is not the existence of pointer records. It is the promotion rule: stable records define possible knowledge, but measured quantities require a response functional. This is the step that prevents finite record models from treating every attractive ledger entry as an observable.

6 Engineering calibration as epistemology

The engineering analogy is more than a metaphor. Engineers routinely distinguish an internal state from a sensor readout. A system state becomes a measurement only through an interface:

$$\text{state} \longrightarrow \text{sensor} \longrightarrow \text{transfer function} \longrightarrow \text{calibrated output.}$$

The analogous chain for a finite information model is

$$\text{record algebra} \longrightarrow \text{monitored histories} \longrightarrow \text{response functional} \longrightarrow \text{observable.}$$

This engineering view is often implicit in experimental practice but under-explicit in speculative theory. The record-response principle makes it an audit rule: every claimed physical number must state its sensor and transfer function.

7 Severity: how the framework should make itself vulnerable

The philosopher Deborah Mayo's account of severe testing emphasises that a claim earns evidential weight by passing tests that would probably have found the error if the claim were false [14]. The record-response principle gives a version of that idea for finite information physics.

Principle 2 (Response severity criterion)**.** *A finite-substrate numerical claim is severe only if it states, before confrontation with the relevant data:*

1. *the record object being billed;*
2. *the response functional by which an instrument reads it;*
3. *which constants are calibration inputs rather than outputs;*
4. *which downstream readouts are inherited rather than independent;*
5. *which observation would kill the claim.*

This criterion is deliberately unforgiving. A framework should not receive credit for three predictions when two are inherited readouts of the first. It should not keep a failed response by introducing a new count. It should not claim a likelihood-level observable without running the likelihood-level calculation. It should publish killed predictions and branch retirements where the original claims were made.

The criterion is therefore compatible with a modest "apt metaphor" stance. A framework need not be proven true to be useful. It may earn credibility as an apt representation if it repeatedly:

- separates record facts from response observables;
- blocks tempting but illegal promotions;
- generates pre-registered numbers;
- survives severe tests that could have killed it;
- improves conceptual compression without adding escape dials.

If it fails those tests, it remains metaphorical or pedagogical rather than physical.

8 Case studies from the finite-QEC programme

This section summarises how the rule has already changed the status of several claims in the finite-QEC programme. These examples are not offered as proof that the programme is true. They are offered as demonstrations of the discipline.

8.1 QED: from count to Thomson readout

The bare service alphabet gives a record-level count

$$\alpha_0^{-1} = \mathrm{Sym}^2(16) + 1 = 137.$$

That is not the physical fine-structure constant. The response calculation requires a Wilson/Gauss endpoint contact, a second monitored endpoint vertex, and ordinary retarded running. The resulting boundary is

$$\alpha(0)^{-1} = 137.035999106904$$

in the current convention [15, 16].

The first dynamical readout is then low-energy Compton scattering. The response chain is:

endpoint record → Maxwell contact → zero-frequency residue → Thomson cross-section.

Ward conservation and the Low–Gell-Mann–Goldberger low-energy theorem [17, 18] force

$$\frac{d\sigma}{d\Omega} = \frac{r_e^2}{2}(1 + \cos^2\theta), \qquad \sigma_T = \frac{8\pi}{3} r_e^2, \qquad r_e = \frac{\alpha(0)\hbar c}{m_e c^2}.$$

Thus the Thomson readout is a theorem once the endpoint action is accepted. It is not an independent second prediction: its evidential burden is inherited from $\alpha(0)$. That distinction is precisely what the response rule is for.

8.2 Black holes: horizon counts are not fluxes

A horizon ledger may supply local service counts, but a measured Hawking flux is a response observable. It requires a local KMS state, a near-horizon Bogoliubov response, species and polarisation factors, and greybody transfer through the exterior geometry. The response layer therefore prevents a horizon alphabet count from being mistaken for an absolute flux coefficient. The count can seed the source; the flux is the response.

8.3 CMB and halos: likelihoods are responses

A pressureless zero-mode reservoir may have the right record-level shape for a dark component. But CMB acoustic peaks, lensing, and galaxy-galaxy lensing do not measure the mere existence of a ledger entry. They measure evolved spectra and likelihoods. A viable branch must pass through Boltzmann evolution, nuisance parameters, halo occupation, and lensing response. This is why the CMB/halo sector remains an empirical gate rather than a settled count.

8.4 Electroweak and QCD: matching and continuum limits

Finite record selection may identify candidate electroweak exposure factors or strong-sector Wilson surfaces. But pole masses and physical string tensions are response objects. They require fixed-scheme RG matching, Coleman–Weinberg normalisation, threshold corrections, static-potential extraction, and continuum scale setting. The record layer can supply structure; the response layer decides whether that structure computes real numbers.

9 What the principle does not claim

The record-response principle is not a proof of the finite-QEC substrate. It is not a derivation of quantum theory from pure philosophy. It is not a new hidden-variable theory. It does not say that unrecorded amplitudes are unreal. It does not say that every response has already been computed.

Its more limited claims are:

1. stable records are the only candidates for public empirical facts;
2. measured observables are response-functional outputs, not arbitrary internal ledger entries;
3. finite information models need this split to avoid numerology;
4. the split produces a concrete severity audit for predictions.

10 Conclusion

The slogan “records say what can be known” is not an invitation to retreat from experiment. It is the first half of an empirical rule. Records define which distinctions can become stable facts. Responses define which of those facts an instrument can measure, and how.

For finite information physics, this is the difference between a seductive count and a physical prediction. It also clarifies the epistemology. The substrate ledger may be richer than anything available to observers inside the universe, but science is not accountable to the inaccessible ledger as such. Science is accountable to recordable facts and calibrated responses.

The resulting view is realist but access-limited, informational but not subjectivist, relational but mechanism-bearing, and severe in Mayo's sense: claims must expose where they can fail. That is the contribution of the record-response principle. It does not make a finite substrate true. It states the conditions under which such a substrate can become empirical at all.

References

[1] David Elliman. From counts to observables: The response layer of a finite record substrate, a bridge for physicists, information scientists, and engineers. `https://doi.org/10.5281/zenodo.21210527`, 2026. Concept DOI 10.5281/zenodo.21210526.

[2] Wojciech H. Zurek. Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75:715–775, 2003. doi: 10.1103/RevModPhys.75.715.

[3] Harold Ollivier, David Poulin, and Wojciech H. Zurek. Objective properties from subjective quantum states: Environment as a witness. *Physical Review Letters*, 93:220401, 2004. doi: 10.1103/PhysRevLett.93.220401.

[4] Wojciech H. Zurek. Quantum darwinism. *Nature Physics*, 5:181–188, 2009. doi: 10.1038/nphys1202.

[5] William K. Wootters and Wojciech H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982. doi: 10.1038/299802a0.

[6] Herbert B. Callen and Theodore A. Welton. Irreversibility and generalized noise. *Physical Review*, 83:34–40, 1951. doi: 10.1103/PhysRev.83.34.

[7] Ryogo Kubo. The fluctuation-dissipation theorem. *Reports on Progress in Physics*, 29: 255–284, 1966. doi: 10.1088/0034-4885/29/1/306.

[8] Harry Lehmann, Kurt Symanzik, and Wolfhart Zimmermann. Zur formulierung quantisierter feldtheorien. *Il Nuovo Cimento*, 1:205–225, 1955. doi: 10.1007/BF02731765.

[9] Christopher A. Fuchs, N. David Mermin, and Rüdiger Schack. An introduction to QBism with an application to the locality of quantum mechanics. *American Journal of Physics*, 82: 749–754, 2014. doi: 10.1119/1.4874855.

[10] Carlo Rovelli. Relational quantum mechanics. *International Journal of Theoretical Physics*, 35:1637–1678, 1996. doi: 10.1007/BF02302261.

[11] Hugh Everett. "relative state" formulation of quantum mechanics. *Reviews of Modern Physics*, 29:454–462, 1957. doi: 10.1103/RevModPhys.29.454.

[12] David Wallace. *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford University Press, Oxford, 2012. doi: 10.1093/acprof:oso/9780199546961.001.0001.

[13] John S. Bell. *Speakable and Unspeakable in Quantum Mechanics*. Cambridge University Press, Cambridge, 1987.

[14] Deborah G. Mayo. *Statistical Inference as Severe Testing: How to Get Beyond the Statistics Wars*. Cambridge University Press, Cambridge, 2018. doi: 10.1017/9781107286184.

[15] David Elliman. Records and responses: Dressing-blind theorems for monitored quantum instruments. Zenodo preprint, 2026. https://doi.org/10.5281/zenodo.21189012.

[16] David Elliman. Records and responses in the world: A derived fine-structure boundary, certified confinement gaps, and registered discriminators for a finite record substrate. Zenodo preprint, 2026. Companion physics application paper, https://doi.org/10.5281/zenodo.21202741.

[17] Francis E. Low. Scattering of light of very low frequency by systems of spin 1/2. *Physical Review*, 96:1428–1432, 1954. doi: 10.1103/PhysRev.96.1428.

[18] Murray Gell-Mann and Marvin L. Goldberger. Scattering of low-energy photons by particles of spin 1/2. *Physical Review*, 96:1433–1438, 1954. doi: 10.1103/PhysRev.96.1433.