

WestQuant: An evidence-aware system for reproducible quantum-memory modelling and network integration

David Vesterlund

david@vesterlundventures.se

Vesterlund Ventures Holding AB <https://orcid.org/0009-0000-6455-1141>

Method Article

Keywords: quantum memories, reproducible simulation, provenance, open experimental data, spin coherence, quantum channels, quantum networks.

Posted Date: October 13th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11337112/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

WestQuant: an evidence-aware system for reproducible quantum-memory modelling and network integration

David Vesterlund*

WestQuant, Vesterlund Ventures, Sweden

11 October 2026

Abstract

Quantum-memory simulations often combine coherence, retrieval efficiency and conditional fidelity measured under different operating conditions. This makes a reproducible calculation insufficient to establish that its memory model applies to a particular experiment. We present WestQuant, a software system that connects versioned source data, typed observations, condition-matched references, explicit quantum channels and native network-simulator adapters. Its contribution is an executable evidence contract across these interfaces, with retained failure outcomes and reproducible study identities. Open rare-earth and diamond-spin datasets demonstrate source reanalysis, operating-point calibration and rejection of an inadequate Ramsey model. A complete effective write-store-read and detection instrument agrees with an independent numerical reference across 336 state probes. Six refinement controls characterize synthetic spin baths up to eight spins before network export. A larger SeQUeNCe/QuNetSim study completes 128 paired comparisons below the declared 10^{-10} numerical tolerance. All results are regenerated from one clean tagged software release in one locked runtime. The system supports transparent model integration and audit; matched experimental protocol tomography, material-specific bath validation and native physical-latency comparison remain explicit requirements for broader predictions.

Keywords: quantum memories; reproducible simulation; provenance; open experimental data; spin coherence; quantum channels; quantum networks.

1 Introduction

The useful lifetime of a quantum memory depends on the task being performed. Spin coherence, optical coherence, storage efficiency and conditional qubit fidelity are different observables. A material may preserve a spin superposition under a specialized pulse sequence while lacking an efficient or sufficiently quiet optical write/read interface. Consequently, a numerical value labeled T_2 cannot by itself determine whether a device supplies useful states to a quantum network.

Published experiments illustrate these distinctions. Teller et al. demonstrated individually addressable memory cells and measured path- and time-bin-qubit fidelities with weak coherent probes [1]. Ortu et al. reported protocol-dependent effective spin-coherence times together with time-bin storage [3]. A recent Pr-memory experiment distinguishes extended classical echoes from nonclassical heralded-photon storage [4]. Molecular-qubit studies supply complementary information about phase memory and chemical environments [6, 7], whereas integrated defect experiments report distinct Ramsey and Hahn observables [9]. These data support different model components and should not be pooled as interchangeable lifetime labels.

*Corresponding author: david@vesterlundventures.se. ORCID: 0009-0000-6455-1141.

Existing computational tools already solve much of the relevant physics. PySCF provides electronic-structure calculations [10]; PyCCE evaluates spin-bath dynamics [11]; QuTiP supplies open-system evolution [12]; and MaxwellBloch treats optical propagation [13]. SeQUeNCe supplies event-driven network simulation [17], while QuNetSim provides a layered quantum-network framework [16]. General workflow systems such as AiiDA address computational provenance [18]. Recent cross-validation of network simulators emphasizes that matching parameter names is insufficient when event and connection semantics differ [20, 19]. WestQuant reuses these tools. It does not replace their solvers, and the implementation evaluated here does not use AiiDA as its execution engine.

The research question is whether a common executable workflow can move from heterogeneous experimental and materials data to appropriately scoped memory simulations while retaining the evidence needed to assess each result. We address it with a typed catalog, condition-sensitive reference queries, explicit model dispatch, an acceptance-and-erasure channel contract and archived validation outputs. The same representation permits controlled candidate variations, but a generated candidate is not assumed to be a stable species or a functioning memory.

WestQuant contributes (i) a source-to-model workflow that keeps units, evidence type and operating conditions attached to observations; (ii) a shared channel contract linking conditional state evolution, acceptance, background and erasure; and (iii) a reproducible evaluation pipeline spanning independent numerical references and two native network stacks. Section 2.3 shows a complete user workflow, and Table 3 locates these contributions relative to existing tools. The historical project name is *Memory Discovery Lab*: “discovery” denotes candidate exploration, while the evaluated release establishes modelling and verification capabilities. The repository retains its original `memory-discover-lab` slug for stable links.

2 System architecture and evidence model

Figure 1 summarizes the implemented workflow. Each source artifact retains its URL, revision or record identifier, raw bytes, content hash, retrieval metadata and license information. Bounded downloads check source checksums when supplied. Offline execution checks the cached SHA-256 values and fails on missing or corrupted artifacts. A retrieval timestamp alone is not treated as an immutable source revision: live reference pages are frozen by their retrieved bytes and hashes.

2.1 Candidates, observations and references

A candidate record contains elemental composition, explicit charge, electronic spin, supplied geometry, host and any site-resolved isotope substitutions. Candidate enumeration checks electron-count/spin parity and preserves the parent configuration. These checks remove inconsistent specifications; they do not establish a bound electronic state, chemical stability, accessible optical transitions or an experimentally realizable device.

An observation contains a subject identifier, quantity, value, declared unit, conditions, source location, provenance and evidence kind. The seven allowed evidence kinds are `experimental`, `curated_experimental`, `fit_to_experiment`, `evaluated_reference`, `computed`, `hypothetical` and `unknown`. Result metadata separately records input and output evidence. For example, a computed channel constrained by a published mean remains a modeled channel rather than an experimental process-tomography result.

The catalog uses SQLite with immutable record identities. The coherence-reference extension stores a definition, a positive time in seconds, reported uncertainty convention, explicit conditions, primary-source location and unresolved conditions. Reference matching requires material identity and coherence definition, then compares the supplied reported conditions. Missing, mismatched or unverifiable conditions cannot silently yield an applicable numerical estimate. Even a reported-




Figure 1: **System architecture and evidence boundaries.** Implemented data and model boundaries. Unresolved observations remain visible in the catalog. Reviewed table values and specified physical models can produce scoped outputs; optical channel construction additionally requires write/read and loss parameters. SeQUeNCe is an optional downstream workload. A table reference is not automatically converted into a full channel.

condition match retains the source’s unresolved microscopic or operating details. No temperature, isotope, host, charge or pulse-protocol interpolation is performed by the reference-query path.

This restriction differs from the separate *pure-coherence table* model: a reviewed complex-coherence curve may be interpolated linearly within its own time interval under matching temperature and control conditions. Its initial value must be unity, its magnitude must remain within the unit disk, and it must explicitly exclude population relaxation. Extrapolation is prohibited. Neither table mechanism converts a bare formula or an isotope mass into T_2 .

2.2 Execution and failure semantics

Study execution enumerates configuration identities before launching bounded worker processes. Hashes include configuration and source information; random streams are derived from explicit master seeds and identities. Atomic JSON writes, checkpoints and resume logic preserve completed records. Invalid geometry, unsupported model families, missing units, incompatible temperature domains and numerical nonconvergence are failures or unresolved outcomes rather than substituted defaults. Acceptance gates evaluate declared numerical and evidence requirements; unresolved physical assumptions remain attached to accepted computational results.

The implemented dispatcher groups candidates into atomic ensembles, trapped atoms/ions, rare-earth solids, defect spins, molecular spins and a restricted negative-ion branch. This is a software taxonomy with six labels, not an exhaustive classification of quantum memories. The allowed model combinations are listed in Table 1. Quantum dots and optomechanical devices do not have independently validated microscopic adapters in this release.

2.3 Study definition and user workflow

A user selects a source lock and reviewed mapping, specifies the observable and operating conditions, then chooses the solver and downstream workload. The source lock pins file hashes and published record identifiers. The memory profile declares evidence, temperature, control, time convention

Table 1: Implemented model domains. Availability of an adapter and numerical execution are separate from experimental validation of that domain.

Family or component	Implemented scope
Atomic ensembles; trapped atoms/ions	Supplied Markov model or condition-matched pure-coherence table; no general first-principles device identification.
Rare-earth, defect and molecular spins	Above descriptions plus finite nuclear-bath evolution or conventional PyCCE coherence with supplied parameters.
Negative ion	Supplied Markov/table storage and a specified photodetachment environment-overlap model; no inferred bound state or coherent optical interface.
Electronic and propagation backends	Molecular PySCF; ORCA/QE input and coupling-import hooks; a separate MaxwellBloch propagation adapter. These do not automatically identify a memory channel.
Experimental operating point	A reviewed rare-earth point with exact time/protocol/temperature matching; lifetime extrapolation rejected.

and optical acceptance. Export serializes Kraus operators, exact ages and evidence metadata; the adapter consumes this object at retrieval rather than reconstructing a lifetime from a table label.

The bundled Teller study [1, 2] provides a complete example. The first command verifies licensed source files, populates the observation catalog, reconstructs the reviewed operating point and produces analytic and QuTiP channels. The second command exports that profile and executes the SeQUeNCe event workload:

```
python scripts/reproduce.py --skip-tests --out /tmp/wq-paper
python scripts/example_source_to_network.py \
  --baseline /tmp/wq-paper/lab/baselines/teller2025 \
  --out /tmp/wq-network
```

The central API calls in the executable example are:

```
profile = read_json(baseline / "measured_point_profile.json")
channels = evaluate(profile, [18e-6], backend="qutip")
export(channels, output / "device.json")
device = load_device(output / "device.json")
result = run_link(device, device, workload, 20261014)
```

Here `workload` explicitly specifies attempts, herald probability, interval, storage age and the Bell-fidelity threshold. The release replay executes this example and checks its fresh confirmation. The profile is an effective isotropic fit at a single measured age; requesting another age fails. Its result is stored as computed evidence even though source observations are experimental.

Table 2 summarizes the schema. SQLite stores canonical JSON payloads keyed by immutable identities. Required observation fields and evidence labels are checked directly; physical applicability is evaluated by model and reference contracts rather than by database storage alone.

These comparisons concern documented responsibilities, not benchmarked superiority (Table 3). Users can implement equivalent contracts in other workflow systems or scripts. QuISP [22] is another network simulator; together with the Chung studies [20, 19], it motivates keeping state/trace agreement separate from native timing and resource equivalence.

Table 2: Implemented record schema and evidence flow.

Record	Fields retained	Enforcement and interpretation
Artifact	Source URL, record/revision, raw-file hash, license, retrieval metadata	Offline replay verifies licensed bytes; a timestamp does not identify a source revision.
Observation	<code>subject_id</code> , quantity, value, unit, conditions, provenance, <code>evidence_kind</code> , <code>source_location</code>	All fields required; seven evidence kinds; content identity is immutable.
Coherence reference	Material, definition, seconds, uncertainty convention, conditions, source location, unresolved conditions	Identity and definition must match; supplied conditions are checked; reference-query interpolation is absent.
Profile and result	Family, control, temperatures, durations, rates, IO parameters; input/result evidence and source/configuration hashes	A fitted or hypothetical input remains distinguishable from the computed output.
Network export	Serialized channel samples, exact ages, acceptance, optical and evidence metadata	Missing ages fail; original input is evolved once at final retrieval.

Table 3: Complementary scope of existing systems and WestQuant.

Tool	Existing capability	Role of WestQuant in this study
AiiDA [18]	Workflow automation, scalable execution and computational provenance	Adds memory-specific evidence, observable and condition contracts; does not claim a new general provenance engine.
NetSquid [21]	Discrete-event quantum-network modelling, timing and component/protocol stacks	Supplies a possible future consumer of memory models; no NetSquid adapter or benchmark is evaluated here.
SeQUeNCe [17]	Customizable network components and discrete-event kernel	Imports source-linked conditional channels and explicitly sampled acceptance through native density/event hooks.
QuTiP scripts [12]	Open-system solvers, quantum objects and channel representations	Supplies study manifests, typed source records, applicability guards, serialized exports and full-study replay around solver calls.
WestQuant	Source-to-channel study orchestration and evidence contracts	Integrates solvers and native adapters with retained negative outcomes and release-bound reproduction.

3 Methods: physical and operational model

3.1 Stage-resolved conditions and Markov dynamics

A modeled memory specifies writing, storage and reading temperatures (T_w, T_s, T_r) and durations (t_w, t_s, t_r). Each stage is isothermal in the evaluated implementation. An unspecified thermal ramp is rejected; a temperature-dependent coefficient is used only inside its declared validity interval. Phenomenological rate functions can be supplied as

$$\Gamma(T) = \Gamma_0 + aT + b(T/T_{\text{ref}})^n + \frac{c}{\exp(\Theta/T) - 1}, \quad (1)$$

with nonnegative coefficients and an explicit reference temperature. These coefficients are inputs or calibrated model quantities. Equation (1) does not establish a material-specific relaxation mechanism merely because terms resemble direct, Raman or activated processes.

For stage j , the Markov reference uses population relaxation $\Gamma_{1,j}$, independent pure dephasing $\Gamma_{\phi,j}$, excited-state equilibrium population $p_{e,j}$ and detuning ν . With $H = h\nu|1\rangle\langle 1|$ and $\mathcal{L}_A(\rho) =$

$$A\rho A^\dagger - \{A^\dagger A, \rho\}/2,$$

$$\dot{\rho} = -\frac{i}{\hbar}[H, \rho] + \Gamma_1(1 - p_e)\mathcal{L}_{\sigma_-}(\rho) + \Gamma_1 p_e \mathcal{L}_{\sigma_+}(\rho) + \frac{\Gamma_\phi}{2}\mathcal{L}_Z(\rho), \quad (2)$$

$$\rho_{01}(t) = \rho_{01}(0) \exp[-(\Gamma_1/2 + \Gamma_\phi)t + i2\pi\nu t]. \quad (3)$$

Thermal detailed balance is optional and, when declared, uses a supplied qubit frequency to determine p_e . A separately declared activity contribution can add a phenomenological rate λa to storage dephasing; it is not an independently validated model of network-induced heating or pulse errors.

Under this stationary, independent-noise description,

$$T_1 = \Gamma_1^{-1}, \quad T_2 = (\Gamma_1/2 + \Gamma_\phi)^{-1}. \quad (4)$$

An experimentally reported *total* T_2 must not be inserted as $\Gamma_\phi = 1/T_2$ and combined with population relaxation again. The model requires an explicit rate definition to prevent this double counting. The familiar $T_2 \leq 2T_1$ statement is conditional on this model and is not used as a universal filter for protocol-dependent fitted times.

3.2 Finite nuclear-bath evolution

Hamiltonians in the spin module are frequency operators G in Hz; physical energies are $\hbar G$. The supplied bath contains nuclear spins I_i , laboratory-frame Zeeman vectors $\boldsymbol{\omega}_i$, hyperfine tensors A_i , optional symmetric quadrupole tensors Q_i and pair tensors J_{ij} :

$$G_B = \sum_i \boldsymbol{\omega}_i \cdot \mathbf{I}_i + \sum_i \sum_{ab} Q_{i,ab} I_{i,a} I_{i,b} + \sum_{i < j} \sum_{ab} J_{ij,ab} I_{i,a} I_{j,b}, \quad (5)$$

$$G_m = G_B + m \sum_i \sum_a A_{i,za} I_{i,a}. \quad (6)$$

The central-spin pure-dephasing approximation retains the central S_z coupling and the two specified logical projections m_0, m_1 . Transverse central-spin transitions and phonon relaxation are not generated by this approximation. The original controls use Gibbs initialization,

$$\rho_B(T) = \frac{\exp[-\hbar G_B/(k_B T)]}{\text{Tr} \exp[-\hbar G_B/(k_B T)]}, \quad L(t) = \text{Tr}[U_0(t) \rho_B(T) U_1^\dagger(t)]. \quad (7)$$

The extension also supports explicitly maximally mixed initialization, $\rho_B = I/d$, used for its PyCCE/exact convergence comparisons. The zero-temperature Gibbs implementation uses the normalized projector onto the numerically degenerate ground-state subspace. Free propagators contain $\exp(-i2\pi G_m \Delta t)$. Ideal CPMG pulses swap the two branch Hamiltonians at $t(j + 1/2)/N$, $j = 0, \dots, N - 1$. A finite Hilbert-space budget bounds exact evolution.

Conventional CCE is available through PyCCE [11]. Its backend units are explicitly converted: seconds to milliseconds, hyperfine Hz to kHz and field tesla to gauss. Generalized CCE density elements are not silently normalized and exported as pure coherence. The extension explicitly passes supplied pair tensors to PyCCE and checks Zeeman conventions when required. Bath-size, cluster-order, pair-cutoff, propagation-resolution and time-grid convergence are examined for synthetic controls in Section 8; none validates a new material.

The operational descriptor returned by exact evolution is the first *sampled* $1/e$ crossing of $|L|$, represented by adjacent time-grid endpoints. Revival detection and the finite observation window accompany it. This descriptor is not declared to be a unique exponential T_2 . Absence of a sampled crossing is not proof that none occurs between samples, and is not reported as an infinite measured lifetime.

3.3 Electronic structure, optical propagation and negative ions

The molecular electronic adapter evaluates HF or DFT at supplied geometry, charge and spin, optionally with scalar relativistic treatment and Tamm–Dancoff excitations. Convergence is checked and recipes are included in the cache identity. A finite-basis energy does not prove an anion bound state; a vertical excitation energy does not specify T_2 . Couplings from external calculations require explicit site mapping, tensor-frame alignment and units before entering Eq. (6). The evaluated workflow does not automatically calculate a complete spin–phonon spectral density from a molecular structure.

The MaxwellBloch adapter evolves a supplied multi-level optical recipe separately from the qubit-channel wrapper. It can provide propagation fields and states, but a finite-field solution is not automatically a single-photon acceptance instrument. Protocol calibration and an observable mapping remain necessary.

For photodetachment, neutralization leaves an electron or environmental record. For a modeled superposition

$$|\Psi\rangle = a|0\rangle|E_0\rangle + b|1\rangle|E_1\rangle, \quad \rho_{01} \mapsto \chi\rho_{01}, \quad \chi = \langle E_1|E_0\rangle, \quad (8)$$

tracing out that record reduces coherence unless the records overlap sufficiently. The restricted negative-ion branch therefore requires a supplied complex χ with $|\chi| \leq 1$. It does not infer χ from a charge change or assert that photodetachment produces a stable quantum memory.

3.4 Write-store-read channel and erasure

Let Φ_j be the conditional stage channel from Markov dynamics, and let $\mathcal{D}_z$ multiply ρ_{01} by complex z , with $|z| \leq 1$. The storage map is

$$\Phi_s(t) = \mathcal{D}_{L(t)} \circ \Phi_{s,\text{Markov}}(t), \quad \Phi(t) = \Phi_r \circ \Phi_s(t) \circ \Phi_w. \quad (9)$$

Write/read depolarization is separately supplied. Composition with bath or table coherence is allowed only when its independence from the Markov contribution is explicitly declared. The optical acceptance parameters η_w, η_r and independent erasure rates yield

$$\eta(t) = \eta_w \eta_r \exp[-\Gamma_{e,w} t_w - \Gamma_{e,s} t - \Gamma_{e,r} t_r]. \quad (10)$$

This separable expression is a model assumption, not a universal relation between optical efficiency and spin coherence.

The implemented complete map sends a qubit to a qubit-plus-erasure output:

$$\mathcal{E}_t(\rho) = \eta(t)\Phi_t(\rho) \oplus (1 - \eta(t)) \text{Tr}(\rho)|e\rangle\langle e|. \quad (11)$$

If $\{K_a\}$ is a Kraus representation of Φ_t , accepted operators are the 3×2 matrices obtained by embedding $\sqrt{\eta}K_a$ in the logical output block. Two loss operators are $\sqrt{1 - \eta}|e\rangle\langle 0|$ and $\sqrt{1 - \eta}|e\rangle\langle 1|$. Their adjoint products sum to the identity. The unnormalized, column-vectorized Choi matrix is

$$J(\mathcal{E}) = \sum_a \text{vec}(K_a) \text{vec}(K_a)^\dagger, \quad J \succeq 0, \quad \text{Tr}_{\text{out}} J = I_{\text{in}}. \quad (12)$$

Kraus compression uses Choi eigendecomposition with a declared numerical cutoff. Complete positivity, trace preservation and state validity are checked numerically. These tests establish a physical map under its assumptions; they do not establish that the map is the correct device model.

Equation (11) assumes state-independent acceptance and ideal loss flags. State-dependent acceptance requires a general trace-decreasing quantum instrument; dividing its output by a state-dependent success probability would produce a nonlinear conditional transformation. That broader instrument is outside the implemented export contract.

3.5 Readout background and a complete effective instrument

The signal-only erasure map is extended with a named background source independent of the noise already in Φ_t . The event model accepts a retrieved signal first; background is accepted only if signal retrieval failed. If $r(t) = r_0 + at$ is the supplied background rate and Δ the detector window,

$$b(t) = 1 - e^{-r(t)\Delta}, \quad q(t) = [1 - \eta(t)]b(t), \quad p_{\text{acc}}(t) = \eta(t) + q(t). \quad (13)$$

For a fixed background qubit state σ , the complete linear instrument is

$$\mathcal{D}_t(\rho) = [\eta(t)\Phi_t(\rho) + q(t)\sigma \text{Tr } \rho] \oplus [1 - p_{\text{acc}}(t)] \text{Tr } \rho |e\rangle\langle e|. \quad (14)$$

Signal efficiency η , noise-click probability q and total acceptance p_{acc} remain separate. A fixed-state background has six-Pauli-state average target overlap $1/2$, giving

$$F_{\text{det}}(t) = \frac{\eta(t)F_{\text{signal}}(t) + q(t)/2}{p_{\text{acc}}(t)}, \quad O_{\text{uncond}}(t) = p_{\text{acc}}(t)F_{\text{det}}(t). \quad (15)$$

Conditional fidelity is undefined when $p_{\text{acc}} = 0$ and is stored as missing. Duplicate noise-source identities are rejected. This is an effective state-independent first-click instrument; it excludes signal/background double clicks, multiphoton physics and state-dependent detection, and is not itself an experimentally identified detector model.

3.6 Absolute age and network semantics

The exported memory device stores channels at explicitly requested ages. Retrieval applies the complete age-dependent channel once to the original input state; incremental application is rejected. This avoids applying the same decay twice and avoids a semigroup assumption in revival-containing dynamics. It supports fixed, declared control histories, not arbitrary interventions on a non-Markovian bath. The network adapter performs no age interpolation or extrapolation: an absent sample is a missing-data error.

The SeQUeNCe adapter acts on the selected key in its complete joint density matrix, with explicit tensor ordering. It samples loss before applying conditional channels, resets consumed or erased memory states and does not add SeQUeNCe’s default memory decoherence a second time. Event durations are converted to integer picoseconds with a representability check. The evaluated workload loads an ideal Bell pair, waits for a fixed age and attempts retrieval from two nodes. It uses the actual event kernel and density-matrix manager, not a full routing, swapping or realistic photon-generation stack.

4 Observable definitions and inference gates

4.1 Fidelity and useful output

For an input $|\psi\rangle$, conditional target fidelity is $F_\psi = \langle\psi|\Phi(|\psi\rangle\langle\psi|)|\psi\rangle$. A uniform six-Pauli-state mean equals Haar-average qubit fidelity only for one linear channel and appropriate state-independent conditioning. For that model,

$$F_{\text{av}} = \frac{2F_e + 1}{3}, \quad F_e = \frac{1}{4} \sum_a |\text{Tr } K_a|^2. \quad (16)$$

The framework separately reports F_{av} , success probability and unconditional target overlap ηF_{av} . These quantities should not be merged into one unspecified “fidelity.” A classical benchmark for weak coherent probes also cannot be replaced by a universal $2/3$ threshold without the relevant photon-statistics and efficiency model.

In the two-node workload, a chosen Bell target defines F_B after retrieval. For a specified threshold $F_{\min}$, the quality-qualified delivery rate is the number of accepted pairs with $F_B \geq F_{\min}$ divided by total simulated attempt time. The threshold is a benchmark parameter, not a QKD security guarantee. The implementation does not claim that this closed workload supplies a realistic end-to-end network rate.

4.2 Protocol-specific coherence inference

The data representation distinguishes Ramsey T_2^* , Hahn coherence, DD-dependent effective spin T_2 , molecular phase-memory time T_m and effective optical AFC coherence. For example, an AFC-efficiency expression with exponent factor four and a spin-wave intensity model with exponent factor two cannot be interchanged [3]. A single fidelity point cannot determine a lifetime.

For curve diagnostics, the Gaussian likelihood uses supplied uncertainties with

$$y(t) = b + A \exp[-f(t/\tau)^\beta], \quad (17)$$

where $\beta = 1$ for an exponential and $0.25 \leq \beta \leq 4$ for the stretched model. The positive factor f must have an explicit observable mapping. In the Pr-efficiency diagnostics below, $f = 1$ defines an *apparent retrieval-decay scale*; it is not converted to physical spin T_2 .

For supplied phase-resolved counts, $k_i \sim \text{Binomial}(n_i, p_i)$ and

$$p_i = \frac{1}{2} \{1 + b + A \exp[-f(t_i/\tau)^\beta] \cos \phi_i\}. \quad (18)$$

Least-squares minimization of signed square-root binomial-deviance residuals implements the corresponding likelihood fit. The evaluated count control has known phases; phase drift is not fitted away as decay.

Fits use three starting scales and bounded log-time parameters. Every fourth unique time group is held out, including all phases at that time. Gaussian standardized residuals or binomial Pearson residuals give a held-out RMSE. Model mismatch is flagged when the approximate full-data goodness-of-fit diagnostic has $p < 0.01$ or held-out RMSE exceeds three. Missing mapping review, boundary estimates, a scale outside the sampled interval, or unstable/wide bootstrap intervals produce an inconclusive outcome. These are explicit screening diagnostics under the declared error model, not calibrated general-purpose hypothesis tests.

Parametric bootstrap intervals condition on the fitted model and supplied likelihood. They exclude shared systematics and model uncertainty. A negative diagnostic is retained; the software never selects an unreviewed numerical scale simply because a solver returned a value. In particular, an efficiency-decay fit and a correlation $g^{(2)}$ measurement do not supply qubit process tomography.

5 Methods: evaluation design and source audit

All analyses are retrospective. Splits and gates are explicit in the configuration and code, but were not prospectively preregistered. We distinguish four forms of evidence: data-consistency/reanalysis checks, within-dataset model diagnostics, mathematical/software verification, and optional simulator integration. No new hardware experiment or independent experimental replication was performed.

The first experimental case uses Teller et al.'s released processed data [1, 2]. Source identifiers and hashes are retained; the offline release replay verifies the bundled licensed numerical files. Ninety state-fidelity scalars cover six probe states in five path pairs and ten time-bin cells. These scalars are not ninety independent experiments. Source indexing is checked against released summary values; inconsistent headers and an ambiguous time axis are documented rather than silently trusted.

The spatial holdout calibrates a common conditional-fidelity mean on cells 1, 3, 5, 7 and 9 and assesses cells 2, 4, 6, 8 and 10. Inverse-variance weighting uses the supplied Monte Carlo uncertainty widths. The diagnostic is

$$z_i = \frac{|F_i - \hat{F}|}{\sqrt{\sigma_i^2 + \sigma_{\hat{F}}^2}}. \quad (19)$$

Cross-cell covariance and common systematics are unavailable; the gate $\max_i z_i \leq 3$ therefore supports only an internal approximation check. It does not validate a unique channel or the separately calibrated path point.

A separate Pr-memory retrieval dataset uses the seven released measurement files in Zenodo record 21532904 [4, 5]. Four classical curves provide 15 no-RF and 70 observations per RF sequence (XX, XY4, XY16), totaling 225 curve observations. Source metadata and MD5 checksums are retained alongside SHA-256 hashes. Efficiency, uncertainty and delay axes are mapped explicitly. Separate heralded-photon arrays remain separate observables: a source time of 100 μs versus an article value of 100.9 μs is an unresolved clock-offset issue, not corrected by guessing.

The reference branch contains ten primary-text-reviewed values. It covers five Eu/YSO effective coherences, three vanadium-complex phase-memory values and two integrated-SiC coherence values. It is a seed reference table, not an exhaustive or current-best material atlas. A pinned secondary molecular dataset [8] remains a historical acquisition because its units and conditions are unresolved; its audit counts are retained in Additional file 1 and excluded from the licensed offline replay.

Three deterministic specification cases exercise Eq. (4) and the finite bath. The count control uses 21 ages, four analysis phases per age, 4000 trials per phase, known $T_2 = 1$ ms and seed 20261010. Its 336,000 synthetic Bernoulli trials are not measurements. Sixty bootstrap replicates are requested for this control and thirty per experimental curve model. Eight primary curve-model fits are counted separately from bootstrap and holdout refits.

Two independent numerical descriptions are compared: the analytical Markov/channel construction and QuTiP evolution or conversion. The Teller operating-point network workload uses four seeds, 100,000 attempts each, an 18 μs age and a 1 ms attempt interval. A separate hypothetical three-age control uses 2000 attempts. Both declare ideal loaded Bell pairs and ideal heralded loss. All analyses use the unified release and its 109-test suite. The companion Supplement details the scope, manifests and commands.

6 Results

6.1 Reanalysis of the memory-array data

The uniform path and time-bin means are 0.9510824229 and 0.9140479475, reproducing the paper’s rounded summaries of 95% and 91%. Released per-state data reproduce the corresponding figure means with maximum discrepancies 2.22×10^{-16} and zero, respectively. A released output density matrix gives target- X fidelity 0.9576996483, agreeing with its reported-state counterpart to 2.70×10^{-12} . These results check released processed-data consistency, not raw coincidence reconstruction.

The odd-cell calibration gives $\hat{F} = 0.9224361$ with an approximate calibration uncertainty 0.0038484. The maximum held-out standardized discrepancy is 2.3417, below the diagnostic threshold. Figure 2 shows that passing this check is compatible with visible cell variation. It does not establish microscopic material uniformity.

We also fit fidelity versus *detection-window width* at fixed storage age. A finite Gaussian pulse with flat unpolarized background has

$$F(w) = \frac{1}{2} + (F_s - \frac{1}{2}) \frac{\text{erf}(w/(\sqrt{8}\sigma_p))}{\text{erf}(w/(\sqrt{8}\sigma_p)) + bw}, \quad (20)$$




Figure 2: **Released fidelity data and effective point model.** Reanalysis of processed data from Teller et al. [2]. Small points show the six state fidelities; dark markers show their uniform means. In panel (b), the dashed line is the inverse-variance calibration on odd-numbered cells, and error bars on even cells show supplied MC uncertainties. Conditioning, encodings and the spatial holdout are distinct. Error bars do not include shared systematic uncertainty.

with consistent time units. Calibration uses 80 widths at or below 600 ns; assessment uses 91 widths at or above 800 ns; 40 intermediate widths are unused. The three maximum held-out pointwise discrepancies are 1.729, 2.638 and 0.723. The corresponding constant-fidelity diagnostics have discrepancies 5.547, 5.279 and 4.119 and are flagged as mismatches. One fitted signal fidelity reaches its upper bound. Nested windows share counts, so these are pointwise adequacy diagnostics rather than independent statistical validation; Eq. (20) does not estimate T_2 .

6.2 A finite operating-point channel and its network control

For path pair 5,6 and a 270 ns detection window, the reviewed point has conditional six-state mean $F_0 = 0.9671968089$ and supplied MC uncertainty 0.0051670. The reported cryostat temperature is 3 K; the optical input-to-output time is 18 μ s, comprising 10 μ s AFC rephasing and 8 μ s spin storage. These are distinct from the detection-window width.

The chosen isotropic approximation is

$$\Phi_p(\rho) = (1 - p)\rho + pI/2, \quad p = 2(1 - F_0). \quad (21)$$

This is one effective model calibrated to the mean, not a channel inferred uniquely from process tomography. The derived signal-acceptance model is

$$\eta_0 = 0.16 \times 0.42 \times 0.50 \times 0.42 = 0.014112, \quad (22)$$

combining AFC efficiency, two-way optical transfer, demultiplexing and temporal acceptance. The propagated component uncertainty is 0.0015824 under independent component errors and exact temporal acceptance. Transfer measurements use classical light; acceptance uncertainty and covariance are unavailable. Fiber coupling, filters, detectors and analysis success are outside this loss model. The simulation therefore does not reproduce a measured complete single-photon network efficiency.

Analytical and QuTiP Choi matrices at this point agree to 1.21×10^{-17} . For two identical conditional depolarizing channels on an ideal Bell input,

$$F_B = \frac{1 + 3(1 - p)^2}{4} = 0.9048185746, \quad P_{\text{pair}} = \eta_0^2 = 1.99148544 \times 10^{-4}. \quad (23)$$




Figure 3: **Finite-point network workload.** Ideal-input network integration control for the calibrated, nonunique Teller operating-point model. Individual seeds each use 100,000 attempts; the pooled point uses 400,000. Bars are Wilson intervals from retrieval counts under fixed parameters. Agreement with the analytical loss model verifies implementation; uncertainty in the optical parameters and false-herald physics is not included.

SeQUeNCe produces 20, 32, 24 and 11 retrieved pairs for seeds 20261010–20261013, respectively: 87 in 400,000 attempts. The pooled fraction is 2.175×10^{-4} with Wilson 95% interval $[1.76350, 2.68249] \times 10^{-4}$, containing the fixed-loss expectation. Maximum joint-density discrepancy from the analytical result is 1.11×10^{-16} . Figure 3 reports Monte Carlo uncertainty only. The same assumptions give a simulated rate 0.2175 s^{-1} over 400 s of combined attempt time; this is not a measured network rate or a validation of its missing optics.

An additional constructive calculation makes the mean-only limitation explicit. A phase-flip channel with probability $q = 3p/4$ has the same average fidelity as Eq. (21), but two such channels yield a different Bell fidelity. The Supplement records the exact matrices and numerical values. This example concerns the mean constraint alone; it does not assert that both models fit every released state or density matrix.

6.3 Failure of simple decay models on the separate Pr experiment of Rodríguez-Moldes et al.

Figure 4 shows the released efficiency curves and the two tested envelopes. All six RF-curve/model combinations are flagged as model mismatches under the declared Gaussian-error diagnostic. XX held-out RMSE values are 3.582 (exponential) and 2.853 (stretched); XY4 values are 3.786 and 3.434; XY16 values are 1.983 and 1.417. The latter two pass the held-out RMSE screen but fail the full-data goodness-of-fit diagnostic. The source paper itself discusses nonexponential behavior [4]; our analysis verifies that the simple diagnostic path does not turn those curves into an unsupported homogeneous lifetime.

The no-RF exponential fit returns an apparent scale outside its observed interval; the stretched fit has insufficient stable bootstrap draws. Both are inconclusive. No fit in this case returns an identified physical spin T_2 . Reported tiny Gaussian goodness-of-fit probabilities condition on independence and the supplied error model; they are not evidence for a unique microscopic cause of the mismatch. Modulations, noise, field uncertainty and protocol details require additional modeling. The withheld heralded-photon arrays also do not measure qubit fidelity, so they cannot repair that missing identification by themselves.




Figure 4: **Retrieval-decay model diagnostics.** Processed open classical-efficiency data [5] and diagnostic fits with Eq. (17). Three RF curves have 70 observations each; the no-RF curve has 15. Curves are plotted against the source’s reported total delay. Error bars are the released uncertainties. Fits are apparent retrieval envelopes, not inferred spin- T_2 curves. The panels have different horizontal scales.

6.4 A distinct-family nuclear-spin observable case

Figure 5 presents the distinct-family case using Bradley et al.’s Fig. 3b and inset for a ^{13}C spin associated with an NV centre in isotopically engineered diamond [14, 15]. Storage is free Ramsey evolution in heralded NV^0 at 4 K, followed by recharging and NV^- postselection; no refocusing pulses are used. The reviewed field is $(0.3 \pm 0.1, 0.06 \pm 0.08, 46.801 \pm 0.001)$ G in the NV frame. The data include 13 transverse Bloch-length points and 11 signed X points with reported errors, without SPAM correction. Times are converted from milliseconds to seconds; no artificial zero-time point, unit normalization or clipping is introduced.

The physical context is $G = \nu_L I_z + m_s(A_{\parallel} I_z + A_{\perp} I_x)$. Hyperfine values measured in NV^- do not independently identify the NV^0 tensor or its noise process, so they are not used as a parameter-free prediction. Instead, the executed effective model uses $G_{\text{eff}} = \Delta\nu I_z$, $I_z = \text{diag}(1/2, -1/2)$, with $\Delta\nu = 79.8$ Hz fixed from a separate NMR measurement and the RF reference, not fitted to held-out Ramsey points. For $E(t) = \exp[-(t/\tau)^\beta]$,

$$\ell_{\perp, \text{model}}(t) = AE(t), \quad X_{\text{model}}(t) = E(t)[c \cos(2\pi\Delta\nu t) + s \sin(2\pi\Delta\nu t)]. \quad (24)$$

Measurement amplitudes A, c, s remain separate from the normalized channel dephasing. Independent matrix-exponential propagation verifies the specified rotation/channel equations.

The first nine envelope points calibrate A, τ, β and the last four are held out; the first five X points calibrate only c, s and the last six are held out. This split was chosen after source inspection and is not prospective independent validation. The fit gives $A = 0.53533$, $\tau = 55.247$ ms, $\beta = 1.36263$, $c = 0.48286$ and $s = 0.02318$. This τ is an effective subset-fit scale without a




Figure 5: **Bradley spin-observable diagnostics.** Bradley Fig. 3b reanalysis [15]. Calibration and held-out points are separated; the two measured observables use different time axes and measurement amplitudes. Source errors are unchanged. The envelope passes the descriptive residual screen, while signed X fails. The positive Bloch length can contain a noise floor not modeled here; this result does not overturn the original experiment’s interpretation.

new physical- T_2 confidence interval.

The maximum held-out absolute residual divided by its source error is 2.428 for the envelope and 3.158 for X , against the fixed threshold three. The X failure persists at 79.7 and 79.9 Hz. Calibration uncertainty, interpoint covariance and raw-shot likelihood are unavailable, so this is a descriptive consistency screen rather than a calibrated significance test. The observable comparison is completed with its negative result retained, while microscopic physical validation and optical network export remain false. This broadens the experimental data examined to a defect-spin memory; it does not complete independent microscopic-model validation.

6.5 Deterministic controls and condition-sensitive references

Figure 6 shows the deterministic controls used to exercise lifetime inference and pulse-dependent coherence. For the supplied Markov control, $\Gamma_1 = 10 \text{ s}^{-1}$ and $\Gamma_\phi = 20 \text{ s}^{-1}$ give $T_2 = 40 \text{ ms}$ at its declared hypothetical conditions. Independent parameter-box ranges [8, 12] and [18, 22] s^{-1} produce $T_2 \in [35.714, 45.455] \text{ ms}$. Four corner evaluations implement this range. Its relative deviation is $-10.71\% / +13.64\%$ about the nominal value; this is deterministic interval propagation, not prediction accuracy or a statistical confidence interval.

The one-nucleus bath control specifies $I = 1/2$, central projections $\pm 1/2$, longitudinal hyperfine coupling $A_{zz} = 4 \text{ kHz}$ and $T = 4.2 \text{ K}$. For the high-temperature unpolarized limit, $L(t) = \cos(\pi A_{zz} t)$. The finite-temperature implementation uses the Gibbs state exactly. Free evolution has revivals and a first sampled $1/e$ crossing bracket of 90–100 μs . An ideal Hahn pulse refocuses the commuting longitudinal coupling, with no sampled crossing up to 1 ms. These are software controls, not predictions for a newly selected molecule.

The count control recovers 0.979663 ms from the known 1 ms generator, with a conditional parametric-bootstrap 95% interval of [0.950215, 1.007737] ms. Recovery in one synthetic instance checks the likelihood implementation; it is not a coverage study or evidence for physical prediction accuracy. All phases at a held-out age remain outside the holdout fit.

The reviewed table query returns the 230 ms Eu/YSO XY16 effective-spin reference under its reported matching conditions [3]. An unknown-material query yields no applicable reference; changing a required temperature also prevents a match. Reviewed references and unresolved




Figure 6: **Deterministic coherence and rate controls.** Explicitly synthetic verification cases. (a) Supplied one-nucleus Hamiltonian under free evolution and ideal Hahn refocusing; a finite bath can revive and need not have an exponential lifetime. (b) Phase-zero projections from the four-phase binomial control and fitted probabilities. All phases, rather than only the displayed phase, enter the likelihood.

molecular entries retain distinct evidence status. The Supplement records the numerical work counts and parameter corners.

The separate hypothetical network control runs SeQUeNCe for 2000 attempts and retrieves 321 pairs. QuTiP/analytical full-channel Choi difference is 4.00×10^{-15} and maximum retrieved-density discrepancy is 5.55×10^{-17} . These are checks of supplied write/store/read parameters, distinct from the experiment-constrained point above. The combined release suite and its test count are reported in Section 10.

7 Results: evaluation of the system contracts

The study addresses four evaluation questions. E1 asks whether common channel representations agree with independent numerical references. E2 asks whether heterogeneous experimental evidence can be retained without promoting unsupported models. E3 asks whether selected contract violations are rejected while valid controls remain usable. E4 asks whether completed studies and network integrations reproduce under fixed source, configuration and random-input identities. The questions and gates are retrospective; the study introduces no new experimental design. All results in the main article use the unified release described in Section 10.

7.1 E1: numerical agreement across representations

Thirteen profiles, comprising twelve supplied controls and the experiment-constrained Teller point, yield 73 channel-age comparisons. An independent direct implementation evaluates the same Markov, finite-bath, coherence-table and photodetachment equations with identical units and conditions. Six Pauli eigenstates and the full Choi representation check channel action, conditional fidelity and acceptance. Maximum discrepancies are below 10^{-10} and pass the declared 10^{-8} representation gate (Table 4). This agreement provides a representation-level test across supported paths.

Table 4: Unified-release numerical representation comparisons.

Quantity	Result
Profiles / channel-age points	13 / 73
Processed fidelity values	90
Maximum full Choi discrepancy	1.305×10^{-14}
Maximum state-output discrepancy	1.174×10^{-14}
Maximum conditional fidelity discrepancy	3.886×10^{-15}
Acceptance / processed-value discrepancy	0 / 0
Numerical tolerance	1.000×10^{-8}

7.2 E2: retaining incomplete and inconsistent experimental evidence

The Teller point, the separate Pr efficiency curves and the Bradley spin observables exercise distinct input contracts. The source-linked point produces a network-consumable effective channel; the Pr decay models retain mismatch/inconclusive outcomes; and Bradley retains the envelope agreement and signed- X failure shown in Fig. 5. A separate synthetic molecular-spin control verifies a 41-age curve path with maximum scaled residual below 10^{-10} . Its hypothetical optical interface is labelled accordingly. A physical molecular descriptor without reviewed curve and Hamiltonian remains blocked. These outcomes demonstrate how the same study interface represents usable, diagnostic and unresolved evidence.

7.3 E3: rejecting selected contract violations

Seven constructed invalid cases test operating-point age and temperature extrapolation, total-coherence double counting, unsupported noise independence, hypothetical export without opt-in, incremental use of an absolute-age channel and modified source bytes. All seven are rejected and all seven valid controls pass. Isolated unsafe bypasses quantify selected consequences: double counting changes the specified 40 ms to 33.333 ms, a 16.667% error; applying an age map twice changes a density element by up to 0.097247. An additional ms-to-s omission changes a density element by up to 0.132201. The Supplement reports the individual cases and units.

7.4 E4: reproducible execution and network integration

Nine bounded studies use 12 jobs each across one, two and four workers. All 108 cold jobs and 108 checkpoint reuses complete with identical configuration identities and scientific channel payloads across worker counts. Missing parameters and real process-timeout probes retain explicit failure statuses; restarting after a completed checkpoint resumes exactly one result. The release confirmation compares the complete outputs from fresh directories. Timing and memory records are execution metadata in the Supplement. The following protocol, convergence and paired-simulator evaluations extend E1, E3 and E4 while preserving the same source and runtime identity.

8 Results: complete protocols, convergence and paired simulators

The pipeline executes instrument verification, spin-bath convergence, a paired simulator pilot and the larger simulator study in that order. Each numerical prerequisite must pass before the next stage starts. This extends E1, E3 and E4 under the scope in Table 7.




Figure 7: **Joint protocol instrument controls.** Signal efficiency, accepted noise and conditional fidelity vary together under supplied rates. Noise-dominated accepted output approaches average fidelity 1/2. The age scale is linear near zero and logarithmic at longer ages.

Table 5: Final spin-bath convergence diagnostics.

Control	Final comparison	Maximum error	Tolerance
Matrix exponential / spectral	4-spin reference	1.444×10^{-15}	1.000×10^{-10}
Strang substeps	128 substeps	2.332×10^{-5}	1.000×10^{-4}
CCE order	Order 4; 4 spins	4.132×10^{-10}	1.000×10^{-3}
Pair cutoff	40 angstrom; 4 spins	4.132×10^{-10}	1.000×10^{-3}
Bath size	6 to 8 spins	1.452×10^{-8}	1.000×10^{-3}
Mean $ L $ integration	129 to 257 points	1.950×10^{-6}	1.000×10^{-4}

Final convergence diagnostics over 0–0.4 ms, taking the maximum over 0, 1 and 4 ideal pulses. Solver, CCE and cutoff errors use the exact four-spin reference; bath-size and time-grid errors compare adjacent nested levels. The time-grid metric is mean absolute coherence, not every revival extremum or an inferred lifetime.

8.1 Time-dependent fidelity, signal retrieval and readout noise

Seven hypothetical instruments cover identity, complete loss, background-only acceptance, dephasing with loss, thermal/background noise, biased background and noise-dominated retrieval at eight ages spanning 0–100 ms. An independent Liouvillian and matrix-unit reference constructs the complete map in Eq. (14). Full Choi matrices, six Pauli-state outputs, fidelity, signal/noise/acceptance probabilities, positivity and trace preservation pass across 56 channel-age points and 336 probes. The maximum discrepancy is 1.443×10^{-15} at tolerance 10^{-10} (Fig. 7).

The heralded-photon arrays of Rodríguez-Moldes et al. [4, 5], labelled **pr2026** in the data bundle, supply nine ages with efficiency, noise and signal-idler $g^{(2)}$. Matching qubit fidelity is absent; a $0.9 \mu\text{s}$ source/article clock difference and the relation between noise per herald and vacuum false clicks remain unresolved. The audit therefore records joint experimental validation as inconclusive. The additional-data contract requests matched tomography, signal and background-only counts, trial denominators, clocks, detector windows and covariance.

8.2 Noncommuting spin-bath convergence

Nested synthetic baths contain 2, 4, 6 and 8 ¹³C spins, maximally mixed initialization, central projections $m = \pm 1/2$, transverse/longitudinal hyperfine terms and explicit pair tensors. Ideal CPMG control uses 0, 1 or 4 pulses over 0–0.4 ms. Matrix-exponential/spectral agreement, Strang substeps, CCE order, pair cutoff, bath size and time-grid integration yield six refinement diagnostics. All final levels pass (Table 5); lower-level failures remain visible in Fig. 8 and the archived tables.

PyCCE receives explicitly supplied pair tensors, including zeros, and its field/isotope convention




Figure 8: **Six spin-bath refinement controls.** Dashed lines mark numerical tolerances. Bath-size and time-grid errors compare adjacent levels; other panels use independent/exact four-spin references. The time-grid metric is integrated mean $|L|$. All studies use ideal pulses and deliberately decreasing outer couplings.

is checked against the exact solver. Additional file 1 gives the Hamiltonian recipe and unit conversions. A profile declaring `material_prediction` requires a content-matched certificate with the five physical refinement axes, a valid time interval and independent physical-model review; the sixth diagnostic checks solver implementation. No material certificate is issued for these synthetic controls.

8.3 A second native simulator and the larger paired study

SeQUeNCe runs its native event kernel and density manager. QuNetSim [16] runs native hosts, quantum connections, packet queues and acknowledgements with a protocol-specific QuTiP density backend that preserves joint states during transmission and consumes retrieved/lost qubits. A matrix-unit reference independently checks local and joint channel action.

The stacks share logical ages and three seeded streams for heralding and the two readouts. Sixteen combinations vary base background (0 or 5000 s^{-1}), storage age (0, 1, 10 or 30 ms) and herald probability (0.25 or 0.75), with attempts spaced by 40 ms and ensemble Bell-fidelity gate 0.8. Pilot and larger stages use distinct seeds and separate fresh processes. Table 6 gives their actual run counts.

The larger study completes 256 arm runs and 131 072 trials. All 128 paired comparisons agree in branch counts and logical trace hashes; maximum mean joint-density discrepancy is 2.465×10^{-15} , below 10^{-10} . No arm fails. Figure 9 shows why accepted-pair counts and fidelity must be considered together. Shared traces deliberately test implementation agreement; native physical timing and resource equivalence require the separate comparisons discussed by Chung et al. [20, 19].

Table 6: Executed native simulator study counts.

Study	Seeds	Pairs of runs	Arm runs	Trials/run	Trials, both arms
Pilot	3	48	96	32	3 072
Larger	8	128	256	512	131 072

Actual native simulator work counts. Each parameter/seed combination is executed once per simulator. Paired shared traces test implementation agreement; seeds are the independent Monte Carlo clusters. Pilot and larger stages are separate processes and use different seeds.




Figure 9: **Paired simulator fidelity and useful output.** Lines give model expectations and points show eight seed clusters at herald probability 0.75. Increasing background can raise acceptance while reducing fidelity. Quality-qualified rates use the ensemble gate. Both stacks agree under shared traces; seeds form independent Monte Carlo clusters per parameter combination.

9 Discussion

WestQuant makes the evidence behind a memory model available at the point where that model is used. Its practical contribution is the connection between a reviewed source mapping, a condition-sensitive profile, a shared channel/instrument representation and a reproducible network workload. This lets users distinguish an accepted numerical calculation from a supported experimental inference without reconstructing the study from separate scripts and notes.

The experimental cases show why that connection matters. A fidelity measured at one operating point can constrain an effective channel while leaving its lifetime unidentified. Nonexponential retrieval and a failed signed Ramsey observable remain useful diagnostic outputs. Joint protocol controls further show that higher accepted-event probability can coincide with lower conditional fidelity when background increases. Preserving these outcomes provides information that a single coherence number or accepted-pair count would discard.

The six bath refinements and the paired native-simulator study provide bounded verification of the chosen implementations. Their design deliberately separates propagation error, cluster truncation, bath truncation and network state/trace agreement. The resulting certificates bind declared scope and content identities; their physical meaning still depends on independently reviewed material parameters and operating conditions. Existing solvers and provenance systems provide the computational foundation (Table 3); WestQuant’s additional contribution is the memory-specific evidence contract connecting them.

Table 7: Validated scope and unresolved physical requirements.

Component	Established here	Requirement for broader use
Experimental inputs	Licensed processed-data replay and retrospective diagnostic splits	Raw counts, covariance and independent experimental/model review.
Teller point	One effective channel at matching age, protocol and reported temperature	Process identification and matched time-resolved optical fidelity, efficiency and background.
Pr and Bradley cases	Retained decay-model mismatch and Ramsey- X failure	Adequate dynamical model and independently identified parameters; no life-time promotion from these fits.
Complete instrument	Numerical agreement for supplied state-independent signal/noise/erasure controls	Matched experimental tomography; double clicks and state-dependent acceptance require new contracts.
Spin baths	Six numerical refinements of bounded synthetic baths and ideal pulses	Material-specific Hamiltonians, disorder, pulse errors and convergence over the target domain.
Network adapters	Two-node native state/branch/trace agreement using a shared logical clock	Native physical latency, multiple slots, routing, swapping and resource contention.
Coverage and review	One locked runtime, selected faults and author-operated replay	Other platforms, empirical fault coverage and independent scientific reproduction; NetSquid and QuISP were not run.

9.1 Limitations and requirements for broader prediction

Table 7 collects the boundaries relevant to interpreting all results. Subsequent validation should supply matched time-resolved tomography, signal and background-only trials with common clocks and denominators, then apply the refinement controls to independently characterized material baths. Larger network studies should separately test native timing, routing, swapping and contention. These are extensions of validated scope rather than additional claims made from the present study.

10 Reproducibility

All main-article analyses are regenerated from the clean v0.7.0 software tag [23]. A single isolated Python 3.12.14 runtime includes NumPy 2.3.5, SciPy 1.16.3, QuTiP 5.3.1, PyCCE 1.1.1, SeQUeNCe 1.2.0, QuNetSim 0.1.2, PySCF 2.14.0 and MaxwellBloch 0.12.0. The complete external-package snapshot is `requirements-release-v070.txt`; the source package is installed from the checked tag. The release manifest records the commit, clean source before and after execution, source/configuration hashes, environment, input/output hashes and test total.

One release command runs the combined test suite, recomputes every existing study and executes a fresh confirmation in separate directories. External DNS and sockets are denied in all workers after installation; loopback remains available for local tests. All 109 tests pass, with no optional skips. Full scientific payloads for E1–E4, the Bradley negative result and the executable example agree under confirmation. The extension comparison checks 359 JSON records and 11 exact random-trace arrays, including every native run. Numerical tolerance is 10^{-10} in absolute and relative comparisons. Runtime, timestamps and archive-container bytes are excluded from scientific equality. Additional file 1 supplies commands, refinement recipes and historical provenance; the main results depend on this single release, not the earlier snapshots.

11 Conclusions

WestQuant provides an executable source-to-channel-to-network workflow for quantum-memory studies. Typed evidence and condition contracts preserve the distinction between experimental observations, effective fits and hypothetical controls. A unified release reproduces open-data analyses, independent numerical comparisons, synthetic-bath refinements and native paired-simulator results while retaining model failures. This establishes a practical system for reproducible integration and audit within the declared scope, with a clear path to matched experimental and material-specific validation.

List of abbreviations

AFC: atomic frequency comb; CCE: cluster correlation expansion; CPMG: Carr–Purcell–Meiboom–Gill; CPTP: completely positive and trace preserving; DD: dynamical decoupling; IO: input/output; NV: nitrogen vacancy; RF: radiofrequency; SPAM: state preparation and measurement; YSO: yttrium orthosilicate.

Declarations

Availability of data and materials

The licensed numerical inputs, reviewed mappings, release result tables, native-run archives and confirmation manifest are available with the software and article materials [23]. The original Teller, Rodríguez-Moldes and Bradley datasets retain their published identifiers and licenses [2, 5, 15]. The molecular-repository acquisition is historical and excluded from the licensed offline replay [8]. Additional file 1 is the methodological and historical supplement, supplied as PDF and editable LaTeX.

Code availability

The source code is distributed under the repository’s MIT license. The immutable v0.7.0 tag and original commit are preserved in the distributed Git bundle; the public release branch mirrors the identical software tree [23]. The release contains the study configurations, replay runner, exact runtime snapshot and source-to-network example. The original tagged commit and exact source checkout are recoverable from the Git bundle identified in Additional file 1.

Competing interests

The author declares no competing interests. The author’s affiliation is WestQuant, Vesterlund Ventures, Sweden.

Funding

This study was supported internally by Vesterlund Ventures. No external funding was received.

Authors' contributions

DV conceived the study, specified its research questions and modelling scope, and is responsible for interpretation, manuscript review and the final scientific claims. Software implementation, analysis execution and manuscript preparation used the AI assistance described below. As sole author, DV retains accountability for the work.

Acknowledgements

The author thanks the cited experimental-data and open-source software contributors for making their work available. No additional individual acknowledgement is made.

Ethics approval and consent to participate

Not applicable. This study uses published physical-system data and computational models, with no human or animal participants.

Consent for publication

Not applicable.

AI assistance

OpenAI ChatGPT/Codex assisted with software development, study execution, numerical and source checks, drafting, editing and manuscript preparation. AI tools are not authors. The named author is responsible for checking the evidence, interpretations and final manuscript; automated checks and author-operated replay do not constitute independent scientific peer review.

References

- [1] M. Teller, S. Plascencia, S. Grandi and H. de Riedmatten, “Quantum Storage of Qubits in an Array of Independently Controllable Solid-State Quantum Memories,” *Physical Review X* **15**, 031053 (2025). doi:10.1103/z6lc-qw2d.
- [2] M. Teller et al., “Supporting data of Quantum storage of qubits in an array of independently controllable solid-state quantum memories,” Zenodo record 14929997 (2025). <https://doi.org/10.5281/zenodo.14929997>.
- [3] A. Ortu, A. Holzäpfel, J. Etesse and M. Afzelius, “Storage of photonic time-bin qubits for up to 20 ms in a rare-earth doped crystal,” *npj Quantum Information* **8**, 29 (2022). <https://doi.org/10.1038/s41534-022-00541-3>.
- [4] A. E. Rodríguez-Moldes et al., “Long-Lived Telecom-Heralded Single-Photon Storage in an Absorptive Spin-Rephased Quantum Memory,” *Physical Review Letters* **137**, 120803 (2026). <https://doi.org/10.1103/ftkb-pkvp>. Author manuscript: arXiv:2609.02579v1.
- [5] A. E. Rodríguez-Moldes et al., “Dataset from Long-lived telecom-heralded single-photon storage in an absorptive spin-rephased quantum memory,” Zenodo record 21532904 (2026). <https://doi.org/10.5281/zenodo.21532904>.
- [6] C.-Y. Lin, T. Ngendahimana, G. R. Eaton, S. S. Eaton and J. M. Zadrozny, “Counterion influence on dynamic spin properties in a V(IV) complex,” *Chemical Science* **10**, 548–555 (2019). <https://doi.org/10.1039/C8SC04122A>. Primary full text: PMCID PMC6335635.
- [7] K. R. Mullin, D. Johnson, D. E. Freedman and J. M. Rondinelli, “Systems-chart approach to the design of spin relaxation times in molecular qubits,” *Dalton Transactions* **53**, 16585–16591 (2024). <https://doi.org/10.1039/D4DT02311K>.
- [8] MTD-group, *Molecular Qubit Dataset*, source revision 5263b1d86b39243e7405dde5bf7c58bbd8fcd677, accessed 10 October 2026. https://github.com/MTD-group/Molecular_Qubit_Dataset.
- [9] H. Hu et al., “Room-temperature waveguide integrated quantum register in a semiconductor photonic platform,” *Nature Communications* **15**, 10256 (2024). <https://doi.org/10.1038/s41467-024-54606-2>.
- [10] Q. Sun et al., “Recent developments in the PySCF program package,” *Journal of Chemical Physics* **153**, 024109 (2020). <https://doi.org/10.1063/5.0006074>.
- [11] M. Onizhuk and G. Galli, “PyCCE: A Python Package for Cluster Correlation Expansion Simulations of Spin Qubit Dynamics,” *Advanced Theory and Simulations* **4**, 2100254 (2021). arXiv:2107.05843; doi:10.1002/adts.202100254.
- [12] N. Lambert et al., “QuTiP 5: The Quantum Toolbox in Python,” *Physics Reports* **1153**, 1–62 (2025). arXiv:2412.04705.
- [13] T. P. Ogden, *MaxwellBloch*, software and official documentation, version used: 0.12.0. <https://github.com/tpogden/maxwellbloch>; <https://maxwellbloch.readthedocs.io/>.
- [14] C. E. Bradley et al., “Robust quantum-network memory based on spin qubits in isotopically engineered diamond,” *npj Quantum Information* **8**, 122 (2022). <https://doi.org/10.1038/s41534-022-00637-w>.
- [15] C. E. Bradley, “Data/software underlying the publication: Robust quantum-network memory based on spin qubits in isotopically engineered diamond,” 4TU.ResearchData, version 3 (2022), CC BY 4.0. <https://doi.org/10.4121/16887658.v3>.
- [16] S. DiAdamo, J. Nötzel, B. Zanger and M. M. Beşe, “QuNetSim: A Software Framework for Quantum Networks,” *IEEE Transactions on Quantum Engineering* **2**, 1–12 (2021). <https://doi.org/10.1109/TQE.2021.3092395>.
- [17] X. Wu, A. Kolar, J. Chung, D. Jin, T. Zhong, R. Kettimuthu and M. Suchara, “SeQuENCe: a customizable discrete-event simulator of quantum networks,” *Quantum Science and Technology* **6**, 045027 (2021). <https://doi.org/10.1088/2058-9565/ac22f6>.
- [18] S. P. Huber et al., “AiiDA 1.0, a scalable computational infrastructure for automated reproducible workflows and data provenance,” *Scientific Data* **7**, 300 (2020). <https://doi.org/10.1038/s41597-020-00638-4>.
- [19] J. Chung et al., “Cross-Validation of Open-Source Quantum Network Simulators,” arXiv:2610.09322v1 (2026). <https://arxiv.org/abs/2610.09322>.
- [20] J. Chung et al., “Cross-Validating Quantum Network Simulators,” IEEE INFOCOM Workshops (2025). <https://doi.org/10.1109/INFOCOMWKSHP565812.2025.11152951>. Preprint: arXiv:2504.01290.
- [21] T. Coopmans et al., “NetSquid, a NETwork Simulator for QUantum Information using Discrete events,” *Communications Physics* **4**, 164 (2021). <https://doi.org/10.1038/s42005-021-00647-8>.
- [22] R. Satoh et al., “QuISP: a Quantum Internet Simulation Package,” IEEE International Conference on Quantum Computing and Engineering, pp. 353–364 (2022). <https://doi.org/10.1109/QCE53715.2022.00056>. Preprint: arXiv:2112.07093.
- [23] D. Vesterlund, *WestQuant Memory Discovery Lab*, software version 0.7.0 (2026). Source code and release materials: <https://github.com/VesterlundCoder/memory-discover-lab/tree/release/v0.7.0>. Accessed 11 October 2026.