Supplementary Information

Fixed laboratory-axis readout of retained qubit information near a Floquet crossing

Matthias Jakob
Kolping-Schulen Neckarpark, Lenore-Volz-Strasse 1, 70372 Stuttgart, Germany
mathias.jakob@kbw-gruppe.de | 7 November 2026

S1 Model convention and reconstruction

The system, bath, encoding, and measurement protocol are those of the main text. We use hierarchical equations of motion (HEOM) and distinguish quantum Fisher information (QFI) from measurement-specific classical information. All frequencies are in units of $\omega_0 = 1$ and times in ω_0^{-1}. The drive frequency is $\omega = 1$. Throughout the primary operating-point comparison, the same input direction $\boldsymbol{v}_0$ and secondary axis are used:

$$\boldsymbol{v}_0 = \begin{pmatrix} 0.9444375376251152 \\ -0.07970011008844596 \\ -0.3188818432844661 \end{pmatrix}, \qquad \boldsymbol{n}_* = \begin{pmatrix} 0.9999643407660554 \\ 0.008427672102985097 \\ -0.0005399437312560307 \end{pmatrix}. \tag{S1}$$

Both are unit vectors to numerical precision. The laboratory-x axis is $(1, 0, 0)$. These values are prescribed inputs of this benchmark. Their preliminary time-convolutionless calibration is not repeated or optimized using any of the new HEOM results.

Provenance of the prescribed directions

The earlier time-convolutionless, second-order (TCL2) study used the perturbative framework of Ref. [1]. It selected the input at the crossing and reference bath parameters by numerically maximizing finite-horizon trace-distance backflow [2]. If $F_{I,\Omega}(t)$ propagates the contrast in the interaction picture, its objective and search domain were

$$\max_{\|\boldsymbol{d}\|=1} \int_0^{150} \left[\frac{d}{dt} \|F_{I,\Omega_c}(t)\boldsymbol{d}\| \right]_+ dt, \qquad [u]_+ = \max(u, 0). \tag{S2}$$

The resulting candidate was converted to the prescribed laboratory input and frozen. The search does not establish global optimality. It was not a maximization of the present HEOM Fisher information.

For that input, the historical secondary axis $\boldsymbol{n}_*$ was obtained by numerical maximization of the TCL2 functional $\mathcal{J}(\boldsymbol{n})$ over the unit sphere, with $T = 150$ and the reference drive and bath. The retained confirmation record uses spacing 0.00625 and 1920 Matsubara terms. It gives $\mathcal{J}(\boldsymbol{n}_*) = 0.33706685$, $\overline{F_Q} = 0.37119935$, and upper envelope 0.33715789. These historical TCL2 values are not substituted for HEOM values. Online Resource 2 includes the retained calibration summary and table with their provenance. The complete original input-direction search history is not part of this package; the full vectors above, rather than a rerun of that search, define the present reproducible benchmark. The new HEOM axis optimization below is a separate measurement comparison with the input unchanged.

Table S1 separates inherited calibration from the present dynamical calculations and diagnostic measurement optimization.

Table S1 Provenance and role of the prescribed and optimized quantities. A known input or operating point may change between comparisons; a constant measurement axis never changes with delay.

Quantity	Origin	Role in this work
Input $\boldsymbol{v}_0$	Preliminary TCL2 backflow search	Frozen input of the primary HEOM comparison
Axis $\boldsymbol{n}_*$	Preliminary TCL2 information search	Frozen secondary axis; not retuned to HEOM
Laboratory-x	Cartesian laboratory axis	Fixed readout in all studies
Controls $\Omega_\pm$	Isolated Floquet spectrum	Selected before their HEOM propagations
Axis $\boldsymbol{n}_{\mathrm{num}}$	Present information optimization	Diagnostic feasible axis for each known setting
Mean QFI	Present HEOM trajectory	Allows a locally optimal measurement at each delay

Bath normalization

The manuscript uses a Lorentz–Drude spectral density [1, 3] with the explicit normalization

$$J_B(\nu) = \frac{\alpha\Lambda\nu}{\nu^2+\Lambda^2}, \quad C_B(u) = \frac{1}{2\pi}\int_0^\infty J_B(\nu)\left[\coth\!\left(\frac{\nu}{2T_B}\right)\cos(\nu u) - i\sin(\nu u)\right]d\nu. \tag{S3}$$

The QuTiP Drude spectral density is $J_Q(\nu) = 2\lambda\gamma\nu/(\nu^2+\gamma^2)$, with the corresponding correlation prefactor $1/\pi$ [4, 5]. Therefore

$$\lambda = \alpha/4, \qquad \gamma = \Lambda, \qquad T_Q = T_B \quad\Longrightarrow\quad J_Q = J_B/2, \quad C_Q = C_B. \tag{S4}$$

The system coupling operator is σ_x without a further factor of $1/2$. This normalization is retained throughout the computations.

Hierarchy and propagations

We use the bosonic HEOM solver in QuTiP 5.3.1 [3–5]. The bath correlation uses the Padé spectrum decomposition [6, 7], represented with `DrudeLorentzPadeBath`; its implementation follows Ref. [7]. The parameter $N_k = 4$ produces five combined bath exponents, including the Drude contribution; $N_k = 5$ produces six. The baseline hierarchy depth is $L = 6$. The numbers of auxiliary density operators are 462 at $(N_k, L) = (4, 6)$, 792 at $(4, 7)$, 924 at $(5, 6)$, and 1716 at $(5, 7)$.

The bath-expansion tail terminator returned by the bath object is included in the system Liouvillian [4, 5]. It is distinct from the finite hierarchy-depth truncation. Its coefficients and the full exponent lists are stored in each run's metadata. The initial root is $\rho_\pm(0)$; every non-root auxiliary operator starts at zero. The two signs represent separately prepared system–bath copies, not two stages of one experiment.

The output states are not renormalized. Their laboratory Bloch vectors give $\boldsymbol{b} = (\boldsymbol{r}_+ + \boldsymbol{r}_-)/2$ and $\boldsymbol{w} = (\boldsymbol{r}_+ - \boldsymbol{r}_-)/2$. This reconstructs the entire one-parameter family by linearity. This two-endpoint construction alone does not reconstruct a complete qubit channel. The supplementary input study below instead uses six basis-state endpoints to obtain the full affine map.

Batches and their counts

The local-variation batch, labelled `hr`, has 11 physical points and 38 completed settings. The spectrum-control batch, labelled `fc`, has 3 physical points and 12 completed settings. Their central physical point is shared, and its baseline and depth-refinement calculations were repeated. The additional record `rv` contains one joint refinement at this same reference. The primary record therefore contains 13 distinct physical points, 51 completed solver settings, and 102 endpoint-state propagations. Repeated central calculations are retained as reproducibility checks rather than counted as new physical operating points.

S2 Operating points and time windows

The reference values are $\Omega_c = 9.001856389545972$, $\alpha = 0.1$, $\Lambda = 1.2$, and $T_B = 1$. Each row in Table S2 changes only the displayed quantity; all other parameters, the input, and the measurement axes remain fixed.

Table S2 All 13 distinct primary points at $T = 150$. The upper gap is $100(U - \mathcal{J}_x)/\overline{F_Q}$ in percentage points. For the shared reference, both batches give identical baseline values.

Change from reference	$\overline{F_Q}$	$\mathcal{J}_x$	$100\eta_x$	upper gap
$\Omega = 8.887487$ (left)	0.179641	0.151229	84.1838	0.07932
$\delta = -0.001$	0.357159	0.324312	90.8034	0.03179
$\delta = -0.0005$	0.357261	0.324413	90.8056	0.03179
Reference	0.357351	0.324502	90.8075	0.03179
$\delta = +0.0005$	0.357429	0.324578	90.8091	0.03179
$\delta = +0.001$	0.357494	0.324642	90.8105	0.03179
$\Omega = 9.117649$ (right)	0.195672	0.166950	85.3213	0.07533
$\alpha = 0.08$	0.426729	0.388063	90.9390	0.02041
$\alpha = 0.12$	0.303988	0.275684	90.6892	0.04556
$\Lambda = 0.96$	0.335745	0.304018	90.5504	0.03945
$\Lambda = 1.44$	0.379430	0.345260	90.9943	0.02597
$T_B = 0.8$	0.412885	0.375645	90.9806	0.02976
$T_B = 1.2$	0.311001	0.281917	90.6483	0.03437

Selection of the two separated controls

The closed-system propagator satisfies $\dot{U} = -iH_S(t)U$, with $U(0) = I$. It is integrated over 2π using DOP853 [8], as implemented in SciPy [9], with relative tolerance 2×10^{-12}, absolute tolerance 2×10^{-14}, and maximum step 0.025. The phases of its two eigenvalues determine the quasienergy difference, reduced modulo $\omega = 1$.

Starting at Ω_c, each direction is scanned in steps of 0.0025 until $T\Delta_F$ first reaches 2π. A bracketed root search uses amplitude tolerance 10^{-13}. The selected amplitudes are

$$\Omega_- = 8.887487160959518, \qquad \Omega_+ = 9.117649351028664. \tag{S5}$$

At all saved scan points the gap increases monotonically up to the selected root. The search is local and numerical; it is not a proof that an unresolved degeneracy cannot occur between

samples. The branch selection, full coordinates, tolerances, and acceptance criteria were saved before the HEOM runs. No dissipative observable enters the selection.

A tighter integration, with relative/absolute tolerances 3×10^{-13} and 3×10^{-15} and maximum step 0.0125, gives gaps 0.04188790204761833 and 0.04188790204811637. The change from the selection integration is below 3×10^{-13}. At the reference, the residual gap is 1.5003037×10^{-7}. The refined one-period unitarity errors are below 8.1×10^{-15}. These are numerical consistency checks of the same propagator equations.

Table S3 gives the cumulative information at all five reported horizons.

Table S3 Cumulative laboratory-x information from the primary control trajectories. The final-time values are separately reported in the main text.

T	$\mathcal{J}_x^-$	$\mathcal{J}_x^c$	$\mathcal{J}_x^+$	$\mathcal{J}_x^c/\mathcal{J}_x^-$	$\mathcal{J}_x^c/\mathcal{J}_x^+$
30	0.553975	0.709553	0.601935	1.2808	1.1788
60	0.352289	0.565775	0.387069	1.6060	1.4617
90	0.247851	0.461139	0.273229	1.8605	1.6877
120	0.188390	0.383297	0.207903	2.0346	1.8436
150	0.151229	0.324502	0.166950	2.1458	1.9437

S3 Numerical convergence and consistency

This section establishes stability under the tested hierarchy, bath-expansion and integration refinements, and checks the physical states and information formulas. These checks use one HEOM implementation; they do not establish rigorous errors for the infinite hierarchy.

All primary HEOM runs use relative tolerance 10^{-11}, absolute tolerance 10^{-13}, maximum internal step 0.025, and output spacing 0.00625 up to $T = 150$. Each trajectory is evaluated at the five cumulative horizons $30, 60, 90, 120, 150$. Baseline integrations use Adams; the integration control uses the backward differentiation formula (BDF) with the same tolerances.

In `hr`, depth $L = 7$ is tested at all 11 points. Padé $N_k = 5$ and BDF are separately tested at the two outer offsets $\delta = \pm 0.001$ and all six changed-bath points. In `fc`, all three refinements are tested independently at all three points. The depth check holds $N_k = 4$; the Padé check holds $L = 6$. These are one-variable refinements, not an extrapolation to infinite depth and expansion order.

For each variant and horizon, relative information changes are computed against the corresponding baseline as $|X_{\mathrm{var}}/X_{\mathrm{base}} - 1|$ for $X \in \{\overline{F_Q}, \mathcal{J}_x, \mathcal{J}_*\}$. Fractions are compared by $100|\eta_{\mathrm{var}} - \eta_{\mathrm{base}}|$, in percentage points. The frozen working criteria are 10^{-3} for the relative information change and 0.05 percentage points for either fraction.

Table S4 summarizes these refinements.

The working criteria pass in both batches. The largest observed relative change in the mean information quantities is 2.3928601×10^{-6}. In the control batch, the largest relative changes of the instantaneous final QFI and final laboratory-x information are 1.90×10^{-5} and 1.991×10^{-5}, respectively, across the reported horizons. These endpoint diagnostics are not the primary averaged quantities.

Table S4 Maxima over all tested points and five horizons within each batch and refinement. The information column takes the maximum over the three absolute information quantities; the fraction column takes the maximum over both fixed axes.

Batch	Refinement	Relative information change	Fraction change (pp)
`hr`	Depth $L = 7$	1.794×10^{-8}	4.162×10^{-8}
`hr`	Padé $N_k = 5$	1.917×10^{-6}	5.604×10^{-5}
`hr`	BDF integrator	9.450×10^{-7}	2.105×10^{-7}
`fc`	Depth $L = 7$	3.521×10^{-9}	1.073×10^{-7}
`fc`	Padé $N_k = 5$	2.393×10^{-6}	6.040×10^{-5}
`fc`	BDF integrator	3.402×10^{-7}	2.319×10^{-6}

Combined depth and bath-expansion refinement

An additional calculation at the reference raises both truncation parameters to $(N_k, L) = (5, 7)$, keeping the input, bath and integrator settings fixed. This calculation was specified in its own saved plan after the original batches. It gives, at $T = 150$, $\overline{F_Q} = 0.35735180336$, $\mathcal{J}_x = 0.32450225874$, and $\mathcal{J}_* = 0.32452655907$. Over all five horizons and the three information quantities, the largest relative change from $(4, 6)$ is 1.439981×10^{-6} when rounded outward. The largest fraction change is 3.648395×10^{-5} percentage points. Both pass the same criteria. The largest endpoint Bloch-component change is below 9.697×10^{-7}. Endpoint trace and Hermiticity errors remain below 6.22×10^{-14} and 3.24×10^{-15}, respectively; the smallest eigenvalue is -2.776×10^{-17} to the shown precision. Full values and all five comparisons are in `repro/rv/results/joint.json`. This supplements the separate refinements; it is not an infinite-order extrapolation.

Additional consistency diagnostics are listed in Table S5.

Table S5 Additional checks of the two original batches, using their larger deviation where applicable. The joint refinement is documented separately above.

Quantity	Maximum deviation
Baseline central trajectories between batches	0 on the stored grid
Relative information change after time-grid thinning by 2 or 4	5.529×10^{-7}
Spectral versus Bloch QFI	1.444×10^{-15}
Born-probability versus projective information	6.662×10^{-16}
Trace error of propagated endpoint states	3.064×10^{-13}
Hermiticity error of propagated endpoint states	4.350×10^{-15}
Smallest endpoint-state eigenvalue	-2.776×10^{-17}
Largest endpoint Bloch-component change under solver variants	1.537×10^{-6}

All tabulated bounds in Table S5 are rounded outward. The analysis also checks the initial Bloch vectors, time grids, stored bath and solver parameters, positivity of the Fisher denominators, $j_{\boldsymbol{n}} \leq F_Q$, and $F_Q \leq 1$ to the stated numerical tolerances. No clipped negative eigenvalues or state renormalizations are used to obtain the reported information. Positivity checks cover the propagated endpoint states and their convex mixtures; they do not certify complete positivity of an approximate map on arbitrary inputs.

The three primary records use the same QuTiP HEOM implementation. These are refinement and consistency checks, not independent second-code validation or rigorous continuum error bars. The endpoint Bloch trajectories, solver metadata, and detailed comparisons remain available for external review.

S4 Information formulas and an analytical control

Recorded-time likelihood

For $p_\theta(s,t) = p(t)[1 + s(a(t) + \theta c(t))]/2$, where $p(t) = 1/T$, $a = \boldsymbol{n}^\mathsf{T}\boldsymbol{b}$, and $c = \boldsymbol{n}^\mathsf{T}\boldsymbol{w}$, direct differentiation at zero gives

$$\sum_{s=\pm1} \int_0^T p_0(s,t)\left[\partial_\theta \log p_\theta(s,t)|_0\right]^2 dt = \frac{1}{T}\int_0^T \frac{c(t)^2}{1-a(t)^2} dt. \tag{S6}$$

Thus the classical information is averaged after conditioning on the recorded time. If time were discarded, the information would instead be $\langle c\rangle_T^2/[1 - \langle a\rangle_T^2]$, which is a different measurement record.

Table S6 makes this distinction explicit. Keeping the timestamp does not require changing the measurement axis. The stated Cramér–Rao bounds assume a locally unbiased estimator and the regular local model at $\theta = 0$ [10]; they are not a simulated estimator performance.

Table S6 Laboratory-x information at $T = 150$ with and without the timestamp, and the standard-deviation lower bound for 10^4 copies with recorded times.

Point	$\mathcal{J}_x$ (recorded)	$\mathcal{J}_x$ (discarded)	$1/\sqrt{10^4\mathcal{J}_x}$
left	0.15122886	0.08472368	0.025715
ref	0.32450179	0.28548986	0.017555
right	0.16694990	0.09544771	0.024474

The QFI used here is defined by the symmetric logarithmic derivative (SLD), not a skew-information variant [11]. For $\rho_0(t) = (I + \boldsymbol{b}\cdot\boldsymbol{\sigma})/2$ with eigenvalues p_a and derivative $D(t) = \partial_\theta\rho_\theta(t)|_{\theta=0} = \boldsymbol{w}\cdot\boldsymbol{\sigma}/2$, the independent spectral calculation is

$$F_Q(t) = \sum_{a,b} \frac{2|\langle a|D(t)|b\rangle|^2}{p_a + p_b}. \tag{S7}$$

All evaluated $\rho_0(t)$ are full rank. This expression is checked against the Bloch formula in the main text. The Born check evaluates $p_\pm = \mathrm{Tr}(\rho_0 P_\pm)$ and their derivatives directly for $P_\pm = (I \pm \boldsymbol{n}_*\cdot\boldsymbol{\sigma})/2$.

Constant projective axes

For each unit $\boldsymbol{n}$, $1 - \|\boldsymbol{b}\|^2 \le 1 - (\boldsymbol{n}^\mathsf{T}\boldsymbol{b})^2 \le 1$. Multiplication by the nonnegative numerator and integration give

$$\boldsymbol{n}^\mathsf{T} G\boldsymbol{n} \le \mathcal{J}(\boldsymbol{n}) \le \boldsymbol{n}^\mathsf{T} H\boldsymbol{n}. \tag{S8}$$

Rayleigh's bound and the quantum information bound give $\mathcal{J}(\boldsymbol{n}) \le U = \min\{\lambda_{\max}(H), \overline{F_Q}\}$. Since laboratory-$x$ is an allowed axis,

$$0 \le \max_{\|\boldsymbol{n}\|=1} \mathcal{J}(\boldsymbol{n}) - \mathcal{J}_x \le U - \mathcal{J}_x. \tag{S9}$$

For $\boldsymbol{b} = 0$, $H = G$, $\overline{F_Q} = \operatorname{tr} G$, and the best constant axis is the principal eigenvector of G. In the present nonunital trajectories, the full H envelope is used; the direction of G is an interpretation of the contrast geometry, not a replacement for that bound.

Direct numerical optimization

We evaluate the original nonlinear $\mathcal{J}(\boldsymbol{n})$ with trapezoidal weights on all 24 001 stored times at $T = 150$, separately at all 13 primary physical points. Antipodal axes define the same information, so the search can be restricted to a hemisphere. A 1024-direction Fibonacci hemisphere supplies the best 12 starts and eight dispersed starts. The three Cartesian axes and the principal eigenvectors of G and H bring the total to 25. Sequential least-squares programming (SLSQP) [12] uses the analytic gradient, the constraint $\boldsymbol{n}^{\mathsf{T}}\boldsymbol{n} = 1$, objective tolerance 10^{-13}, and at most 200 iterations. A separate angular differential-evolution calculation [13] uses seed 20261103, population multiplier 12, tolerance 10^{-10}, absolute tolerance 10^{-13} and at most 250 generations, followed by polishing. Both optimizers use SciPy [9].

The two methods agree in objective to better than 2.2×10^{-12} at every point. The largest objective spread over the 25 SLSQP results is below 3.1×10^{-14}, and the selected axes have tangent-gradient norm below 3.0×10^{-9}. Some starts stop without a success flag; those flags are counted in the archive, and the independent optimizer and objective agreement are retained as checks. These observations support numerical optimization, not a formal global certificate. Writing $\mathcal{J}_{\mathrm{num}} = \mathcal{J}(\boldsymbol{n}_{\mathrm{num}})$, the valid bracket remains

$$\mathcal{J}_{\mathrm{num}} \le \mathcal{J}_{\max}^{\mathrm{const}} \le U. \tag{S10}$$

The remaining envelope gap $100(U - \mathcal{J}_{\mathrm{num}})/\overline{F_Q}$ is at most 0.04554 percentage points. The lower bound $\mathcal{J}_x/\mathcal{J}_{\max}^{\mathrm{const}} \ge \mathcal{J}_x/U$ exceeds 99.9058% across all points; the comparison with the numerical candidates is above 99.9579%. These ratios concern the constant binary Pauli class only, not the time-adaptive QFI benchmark.

Here `dm1`, `dm5`, `dp5`, and `dp1` denote drive offsets $-0.001, -0.0005, +0.0005, +0.001$. The labels `am/ap`, `cm/cp`, and `tm/tp` denote the lower/higher coupling, cutoff, and temperature values in Table S2.

Exact loss and a near-unital bound

For full-rank ρ_0, set $K = I - \boldsymbol{b}\boldsymbol{b}^{\mathsf{T}}$ and $\boldsymbol{\ell} = K^{-1}\boldsymbol{w}$. The SLD is $L = -(\boldsymbol{b}^{\mathsf{T}}\boldsymbol{\ell})I + \boldsymbol{\ell}\cdot\boldsymbol{\sigma}$, as follows by inserting it into $\partial_\theta\rho = (\rho L + L\rho)/2$ [10, 11, 14]. Consequently

$$F_Q = \boldsymbol{\ell}^{\mathsf{T}} K \boldsymbol{\ell}, \qquad j_{\boldsymbol{n}} = \frac{(\boldsymbol{n}^{\mathsf{T}} K \boldsymbol{\ell})^2}{\boldsymbol{n}^{\mathsf{T}} K \boldsymbol{n}}. \tag{S11}$$

Completing the square with $c_{\boldsymbol{n}} = (\boldsymbol{n}^{\mathsf{T}}\boldsymbol{w})/(\boldsymbol{n}^{\mathsf{T}} K \boldsymbol{n})$ yields the exact nonnegative loss

$$F_Q - j_{\boldsymbol{n}} = (\boldsymbol{\ell} - c_{\boldsymbol{n}}\boldsymbol{n})^{\mathsf{T}} K (\boldsymbol{\ell} - c_{\boldsymbol{n}}\boldsymbol{n}). \tag{S12}$$

Table S7 Direct optimization at all 13 primary points. The last column gives the angle from laboratory-x in degrees. Full three-component axes and optimizer diagnostics are archived.

Point	$\mathcal{J}_{\mathrm{num}}$	$100\mathcal{J}_x/\mathcal{J}_{\mathrm{num}}$	$100\mathcal{J}_{\mathrm{num}}/\overline{F_Q}$	$100\eta_x$	angle
Left	0.15129248	99.95795	84.21924	84.18383	1.22751
Reference	0.32452615	99.99249	90.81429	90.80747	0.50845
Right	0.16700820	99.96509	85.35109	85.32130	1.11773
dm1	0.32433680	99.99248	90.81023	90.80341	0.50879
dm5	0.32443763	99.99249	90.81239	90.80557	0.50860
dp5	0.32460233	99.99250	90.81592	90.80911	0.50832
dp1	0.32466618	99.99250	90.81729	90.81048	0.50821
am	0.38808126	99.99538	90.94325	90.93905	0.39877
ap	0.27571467	99.98883	90.69930	90.68917	0.62056
cm	0.30404475	99.99116	90.55837	90.55037	0.55222
cp	0.34528175	99.99361	91.00008	90.99427	0.46891
tm	0.37566389	99.99493	90.98522	90.98061	0.41781
tp	0.28194682	99.98940	90.65787	90.64826	0.60456

Equivalently, let $\chi_{\boldsymbol{n}}(t)$ be the angle between $K^{1/2}\boldsymbol{n}$ and $K^{-1/2}\boldsymbol{w}$. Wherever $F_Q > 0$, $j_{\boldsymbol{n}}/F_Q = \cos^2\chi_{\boldsymbol{n}}$. Thus $\eta_{\boldsymbol{n}}$ is the QFI-weighted mean of $\cos^2\chi_{\boldsymbol{n}}$ over time. A uniform bound $|\sin\chi_{\boldsymbol{n}}| \le \sin\chi_0$ at all informative times implies $\eta_{\boldsymbol{n}} \ge \cos^2\chi_0$. This gives a sufficient directional condition, rather than an inference from the ordinary Euclidean angle alone when $\boldsymbol{b} \neq 0$.

The unital loss $\|\boldsymbol{w}\|^2 - (\boldsymbol{n}^{\mathsf{T}}\boldsymbol{w})^2$ also has a controlled correction:

$$\left|(F_Q - j_{\boldsymbol{n}}) - \left[\|\boldsymbol{w}\|^2 - (\boldsymbol{n}^{\mathsf{T}}\boldsymbol{w})^2\right]\right| \le \frac{\|\boldsymbol{b}\|^2\|\boldsymbol{w}\|^2}{1 - \|\boldsymbol{b}\|^2}. \tag{S13}$$

To see this, subtract the two unital terms from the respective QFI and classical formulas. Both corrections are nonnegative and at most the right-hand side; their difference therefore obeys the stated bound. If $\|\boldsymbol{b}\| \le \epsilon < 1$ throughout the window, the averaged correction is at most $\epsilon^2\langle\|\boldsymbol{w}\|^2\rangle_T/(1-\epsilon^2)$. At the reference, the tighter time-averaged right-hand side evaluated along the trajectory is 9.7770×10^{-5}, compared with mean loss 0.03284964 for laboratory-x. The exact identity is numerically checked at every stored time for laboratory-x and the optimized candidate at all 13 points, with maximum absolute discrepancy below 6.7×10^{-16}. These are consequences and checks of standard qubit information formulas, not a new general estimation theorem.

Known symmetric outcome errors

As an illustrative detector model, independently flip the recorded sign with known probability $q < 1/2$, identically at every time and operating point. The time is still recorded. Then

$$\mathcal{J}_x(q) = \left\langle \frac{(1-2q)^2 w_x^2}{1 - (1-2q)^2 b_x^2} \right\rangle_T. \tag{S14}$$

For nearly unital trajectories this is close to $(1-2q)^2\mathcal{J}_x(0)$, explaining why common errors mainly rescale absolute information in Table S8. No noise value was inferred from an experiment. Unknown or asymmetric errors, missed events, timing and axis uncertainty require different likelihoods and potentially nuisance-parameter estimation.

Table S8 Illustrative calibrated outcome-flip model at $T = 150$. The last two columns are the reference-to-control information ratios.

q	$\mathcal{J}_x^-$	$\mathcal{J}_x^c$	$\mathcal{J}_x^+$	$\mathcal{J}_x^c/\mathcal{J}_x^-$	$\mathcal{J}_x^c/\mathcal{J}_x^+$
0.00	0.15122886	0.32450179	0.16694990	2.14577	1.94371
0.01	0.14524003	0.31165150	0.16033847	2.14577	1.94371
0.05	0.12249470	0.26284637	0.13522854	2.14578	1.94372

A high fraction without Floquet driving

Consider a different, time-homogeneous Markovian model,

$$\dot{\rho} = \frac{\gamma}{2}(\sigma_x \rho \sigma_x - \rho), \qquad \gamma > 0, \tag{S15}$$

with the same prescribed input. It gives $\boldsymbol{b} = 0$ and $\boldsymbol{w} = (v_x, e^{-\gamma t} v_y, e^{-\gamma t} v_z)$. Therefore

$$\mathcal{J}_x = v_x^2, \qquad \overline{F_Q} = v_x^2 + (1 - v_x^2)\frac{1 - e^{-2\gamma T}}{2\gamma T}, \tag{S16}$$

$$\eta_x = \frac{v_x^2}{v_x^2 + (1 - v_x^2)(1 - e^{-2\gamma T})/(2\gamma T)} \xrightarrow{\gamma T \to \infty} 1. \tag{S17}$$

The fraction can approach unity because transverse contributions disappear. This calculation is an analytical counterexample to a broad inference from the fraction alone. It is not a matched Markovian simulation of the driven spin–boson Hamiltonian, and its performance is not used as a baseline for the claimed operating-point ratios.

S5 Dependence on known input directions

This section reconstructs the input-dependent information matrices, establishes the full-cone advantage within those numerical maps, and quantifies the known y counterexample. It then documents uniform refinement bounds, diagnostic axis searches and the resolved output-grid failure.

Affine-map reconstruction and information matrices

This study retains the three amplitudes $\Omega_-, \Omega_c, \Omega_+$, the reference bath, drive phase and recorded-time protocol. It varies the known unit encoding direction in $\rho_\theta(0) = (I + \theta \boldsymbol{v} \cdot \boldsymbol{\sigma})/2$, locally at $\theta = 0$. All six basis-state endpoints $(I \pm \sigma_i)/2$ are propagated separately from identical factorized thermal baths and zero non-root auxiliary operators. For each basis axis, the half-difference gives column i of $M(t)$. The mean of the three pair midpoints gives $\boldsymbol{b}(t)$. No state normalization or eigenvalue clipping is applied.

The reconstructed family is $\boldsymbol{r}_\theta(t) = \boldsymbol{b}(t) + \theta M(t)\boldsymbol{v}$. With $K = I - \boldsymbol{b}\boldsymbol{b}^\mathsf{T}$, the information matrices are

$$A_x = \left\langle \frac{M^\mathsf{T} \boldsymbol{e}_x \boldsymbol{e}_x^\mathsf{T} M}{1 - b_x^2} \right\rangle_T, \qquad Q = \left\langle M^\mathsf{T} K^{-1} M \right\rangle_T, \tag{S18}$$

$$\mathcal{J}_x(\boldsymbol{v}) = \boldsymbol{v}^\mathsf{T} A_x \boldsymbol{v}, \qquad \overline{F_Q}(\boldsymbol{v}) = \boldsymbol{v}^\mathsf{T} Q \boldsymbol{v}. \tag{S19}$$

The laboratory-x measurement remains unchanged. For each known direction and amplitude, a separately optimized constant axis is diagnostic. The measurement can depend on the known preparation and operating point but not on delay. This comparison is not an estimator analysis with unknown preparation as a nuisance parameter.

The largest component disagreement between pair midpoints is 2.11×10^{-8}. The original $\boldsymbol{v}_0$ endpoints reconstructed from the map agree with their independently stored primary propagations within 1.10×10^{-8}. Separately propagated endpoints for the held-out direction $\boldsymbol{u} = (1, 2, 3)/\sqrt{14}$ agree within 3.23×10^{-11}. The reconstructed normalized Choi matrices [15] have minimum eigenvalue no lower than -8.50×10^{-12} across the baseline, joint-refinement and dense-output checks. These are floating-point consistency and physicality diagnostics on saved times, not a rigorous complete-positivity certificate.

Named directions and the fixed hemisphere grid

The 13 named inputs comprise $\boldsymbol{v}_0$, the three laboratory axes, $\boldsymbol{u}$, and eight tilts of $\boldsymbol{v}_0$: 5° and 15°, each at four azimuths separated by 90° in a fixed transverse frame. Their full vectors are saved in the frozen plan. We also use 512 deterministic equal-area hemisphere samples,

$$x_k = \frac{k + 1/2}{512}, \quad \phi_k = k\pi(3 - \sqrt{5}), \quad \boldsymbol{v}_k = (x_k, \sqrt{1 - x_k^2}\cos\phi_k, \sqrt{1 - x_k^2}\sin\phi_k), \tag{S20}$$

for $k = 0, \ldots, 511$. Antipodal inputs have identical local scalar information. The total is $(13 + 512) \times 3 = 1575$ input–amplitude candidates. The directions and tilts were fixed before the new endpoint propagations.

Table S9 compares five named inputs at the reference. Table S10 resolves the reversal for $\boldsymbol{e}_y$ across the original three amplitudes.

Table S9 Known inputs at the reference, $T = 150$. The last two columns are percentages and compare different measurement resources. $\boldsymbol{u} = (1, 2, 3)/\sqrt{14}$.

Input	$\overline{F_Q}$	$\mathcal{J}_x$	$\mathcal{J}_{\mathrm{num}}$	$100\eta_x$	$100\mathcal{J}_x/\mathcal{J}_{\mathrm{num}}$
$\boldsymbol{v}_0$	0.357351	0.324502	0.324526	90.807	99.992
$\boldsymbol{e}_x$	0.339682	0.307615	0.307650	90.560	99.989
$\boldsymbol{e}_y$	0.025156	0.000128	0.013647	0.507	0.935
$\boldsymbol{e}_z$	0.047805	0.022765	0.022828	47.621	99.724
$\boldsymbol{u}$	0.027085	0.002674	0.013462	9.873	19.863

Table S10 Known y input at the original comparison amplitudes. The reference has both lower retained QFI and much poorer laboratory-x accessibility.

Point	$\overline{F_Q}$	$\mathcal{J}_x$	$\mathcal{J}_{\mathrm{num}}$
Left	0.037388	0.011925516	0.0142040
Reference	0.025156	0.000127594	0.0136467
Right	0.037590	0.012210653	0.0134484

All eight tilted inputs retain the reference advantage over both controls. At the reference, their $\mathcal{J}_x/\mathcal{J}_{\mathrm{num}}$ ratios exceed 99.9894%. On the hemisphere grid, the reference exceeds both

controls for 335 directions in $\mathcal{J}_x$ and 334 in mean QFI. Numerical constant-axis optima favor the reference for 325 directions; in 324 of these, its feasible value also exceeds both control upper bounds. These are grid counts, not statistical success probabilities. Table S11 gives a grid-independent diagnosis within the numerical maps: the smallest eigenvalues of both reference-minus-control difference matrices are negative, so neither information advantage is uniform over unit inputs.

Table S11 Smallest eigenvalues of the reference-minus-control information matrices at $T = 150$.

Control	$\lambda_{\min}(A_x^c - A_x^{\mathrm{ctrl}})$	$\lambda_{\min}(Q^c - Q^{\mathrm{ctrl}})$
Left	-0.020530	-0.021349
Right	-0.021273	-0.022006

A bound over the entire preparation cone

Let D denote either difference matrix and let the columns of B form an orthonormal basis perpendicular to $\boldsymbol{v}_0$. Define

$$a = \boldsymbol{v}_0^{\mathsf{T}} D \boldsymbol{v}_0 > 0, \qquad g = \|B^{\mathsf{T}} D \boldsymbol{v}_0\|, \qquad c = \lambda_{\min}(B^{\mathsf{T}} D B). \tag{S21}$$

Write any direction in the cone as $\boldsymbol{v} = \boldsymbol{v}_0 \cos\gamma + B\boldsymbol{z} \sin\gamma$, with $\|\boldsymbol{z}\| = 1$ and $0 \le \gamma \le \gamma_0 \le 45^\circ$. The cross term is bounded below by $-g \sin(2\gamma)$ and the transverse term by $c \sin^2\gamma$. Monotonicity on this interval gives

$$\boldsymbol{v}^{\mathsf{T}} D \boldsymbol{v} \ge a \cos^2\gamma_0 - g \sin(2\gamma_0) + \min(c, 0) \sin^2\gamma_0. \tag{S22}$$

The positive lower bounds in Table S12 establish a reference advantage in both absolute information measures throughout the 15° cone of the numerical maps. They remain positive under the joint hierarchy/bath and dense-output checks. They do not certify constant-axis near-optimality throughout the cone or give interval error bounds for the infinite hierarchy.

Table S12 Downward-rounded lower bounds in the full 15° cone at $T = 150$.

Control	Reference-minus-control $\mathcal{J}_x$	Reference-minus-control $\overline{F_Q}$
Left	0.146515	0.150079
Right	0.118676	0.121617

Uniform refinement and the resolved output-grid diagnostic

The initial batch contains 21 antipodal pairs: nine basis-direction pairs at $(N_k, L) = (4, 6)$, nine at $(5, 7)$, and three held-out pairs. All use output spacing 0.00625. Joint refinement is evaluated over all unit inputs and five horizons $T = 30, 60, 90, 120, 150$. For a positive-definite baseline information matrix A_0, the maximum relative change over all inputs is the largest absolute generalized eigenvalue of $(A_1 - A_0, A_0)$. The same calculation applies to Q. The largest values are 1.27×10^{-5} for $\mathcal{J}_x$ and 4.86×10^{-6} for mean QFI.

A uniform bound on a constant-axis information change also follows from matrix norms. Set $m_i = \|M_i\|_2$, $\beta_i = \|\boldsymbol{b}_i\|$, $d_M = \|M_1 - M_0\|_2$, and $d_b = \|\boldsymbol{b}_1 - \boldsymbol{b}_0\|$, with $i = 0, 1$ denoting baseline and refinement. For every unit input and measurement axis,

$$|\mathcal{J}_1 - \mathcal{J}_0| \le E, \tag{S23}$$

$$E = \left\langle \frac{d_M(m_0 + m_1)}{1 - \beta_1^2} + \frac{m_0^2 d_b(\beta_0 + \beta_1)}{(1 - \beta_1^2)(1 - \beta_0^2)} \right\rangle_T . \tag{S24}$$

The bound also applies after optimizing over the constant axis. If $\epsilon_Q < 1$ bounds the relative QFI change, the optimal-axis fraction changes by at most $100[E/\lambda_{\min}(Q_1) + \epsilon_Q/(1 - \epsilon_Q)]$ percentage points. Here this bound is 0.003930 percentage points. It is a bound on the change between two computed maps, not on the error of either map relative to exact dynamics.

All 1575 baseline candidates use seven-start constrained axis optimization with an analytic gradient. The 39 named input–amplitude combinations are also checked by differential evolution; the largest objective discrepancy is 3.54×10^{-11}. Across all candidates, $100(U - \mathcal{J}_{\mathrm{num}})/\overline{F_Q} \le 0.04619$ percentage points. Feasible optima remain lower bounds and the algebraic envelope remains an upper bound; matching searches do not establish a global optimum.

Fourfold thinning of the original output grid gives a maximum relative information change 1.83×10^{-4}, exceeding the preset 10^{-4} criterion. Nine further antipodal pairs therefore use output spacing 0.003125 at the baseline hierarchy. The direct original-to-dense comparison changes $\mathcal{J}_x$ and mean QFI by at most 9.52×10^{-6} over every unit input and the five horizons. Fourfold thinning of the dense grid at $T = 150$ gives 4.57×10^{-5}, and the 39 named constant-axis optima change by at most 1.94×10^{-5} relatively. The original failed diagnostic is preserved and its resolution recorded. The complete input study contains 30 pairs and 60 endpoint propagations.

S6 Broader amplitude scan and performance widths

Frozen design and additional spectral anchors

The amplitude scan retains $\boldsymbol{v}_0$ and $\boldsymbol{e}_y$, the reference bath, phase, $T = 150$ and laboratory-x measurement. All non-root auxiliary operators start at zero; no output normalization is used. The initial grid is $\Omega = 6, 6.1, \ldots, 12$, supplemented by the exact old reference and the undriven point. The old control trajectories are reused. All design choices and acceptance thresholds are stored in `ext/ds/r/in/study.json`, with its checksum fixed before new propagation.

The adaptive targets are the largest observed $\mathcal{J}_x$ and QFI for each input in $[6, 12]$, plus the $\boldsymbol{v}_0$ $\mathcal{J}_x$ maximum and the $\boldsymbol{e}_y$ $\mathcal{J}_x$ minimum in $[8.7, 9.3]$. Ties choose the smaller amplitude and duplicate targets are merged. The first refinement uses radius 0.1 and spacing 0.025; the second uses radius 0.025 and spacing 0.00625. Points are clipped to the main interval, and both inputs are always propagated. Other local extrema are not exhaustively refined.

An independent closed-system spectrum calculation uses DOP853 over one period, relative/absolute tolerances 10^{-11} and 10^{-13}, and maximum internal step 0.02. The spectral grid has spacing 0.01. It locates minima near 7.426269, 9.001856 and 10.576197; selected gaps are rechecked more tightly. The two outer spectral minima are explicitly post-plan HEOM anchors, chosen from the spectrum before their dissipative calculations. The old reference is not retuned. The largest one-period unitarity defect in this spectrum scan is 2.64×10^{-12}.

Spectral correlation alone does not isolate the gap from changes in modes, micromotion or bath coupling.

Observed maxima and the undriven comparison

Table S13 gives the largest computed values in the final 112-amplitude main interval and the separate undriven values. For $\boldsymbol{v}_0$, the best sampled $\mathcal{J}_x$ exceeds its old reference by 6.5535×10^{-5} at the baseline hierarchy; joint refinement gives 6.5557×10^{-5}. This supports the ordering of those two checked amplitudes but does not locate an exact continuous peak. For $\boldsymbol{e}_y$, the undriven $\mathcal{J}_x$ and QFI exceed every sampled point in $[6, 12]$.

Table S13 Largest sampled information in $[6, 12]$ at $T = 150$. The last columns give the unchanged reference and undriven comparison for the same input.

Input	Quantity	Ω_{best}	Best value	At Ω_c	At $\Omega = 0$
$\boldsymbol{v}_0$	$\mathcal{J}_x$	9.008106	0.324567	0.324502	0.026991
$\boldsymbol{v}_0$	$\overline{F_Q}$	9.008106	0.357417	0.357351	0.058401
$\boldsymbol{e}_y$	$\mathcal{J}_x$	6.606250	0.022848	0.000128	0.030984
$\boldsymbol{e}_y$	$\overline{F_Q}$	6.606250	0.048313	0.025156	0.063978

Width definition and crossing brackets

For each observed maximum of $\mathcal{J}_x$ or mean QFI, define the connected component around it on which $I(\Omega) \geq 0.9 I_{\text{peak}}$. The threshold is frozen at the second-refinement peak value. Subsequent endpoint propagations bisect the threshold-crossing brackets until each is at most 0.00625 wide, using at most five rounds. All crossings here are resolved in four rounds and none is truncated by a boundary. The near-reference $\boldsymbol{v}_0$ target coincides with its largest observed main-interval peak.

Table S14 separates linear interpolation estimates from the width range implied by the two crossing brackets. If the left bracket is $[a, b]$ and the right is $[c, d]$, the width lies between $c - b$ and $d - a$, conditional on the observed single crossing in each bracket. Table S15 gives those brackets, and Fig. S1 shows the normalized sampled profiles. These are amplitude-resolution statements, not statistical confidence intervals or certified HEOM error bars.

Table S14 Interpolated 90% boundaries and widths, with width ranges from the numerical crossing brackets.

Input	Quantity	Left	Right	Width	Width range
$\boldsymbol{v}_0$	$\mathcal{J}_x$	8.967444	9.043524	0.076080	[0.06875, 0.08125]
$\boldsymbol{v}_0$	$\overline{F_Q}$	8.965778	9.045250	0.079472	[0.06875, 0.08125]
$\boldsymbol{e}_y$	$\mathcal{J}_x$	6.252835	6.984412	0.731576	[0.72500, 0.73750]
$\boldsymbol{e}_y$	$\overline{F_Q}$	6.080035	7.185776	1.105741	[1.10000, 1.11250]

Table S15 Threshold-crossing brackets. Displayed endpoints are rounded outward to seven decimals; the full saved values set the 0.00625 stopping rule.

Input	Quantity	Left crossing bracket	Right crossing bracket
$\boldsymbol{v}_0$	$\mathcal{J}_x$	[8.9643563, 8.9706064]	[9.0393563, 9.0456064]
$\boldsymbol{v}_0$	$\overline{F_Q}$	[8.9643563, 8.9706064]	[9.0393563, 9.0456064]
$\boldsymbol{e}_y$	$\mathcal{J}_x$	[6.2500000, 6.2562500]	[6.9812500, 6.9875000]
$\boldsymbol{e}_y$	$\overline{F_Q}$	[6.0750000, 6.0812500]	[7.1812500, 7.1875000]

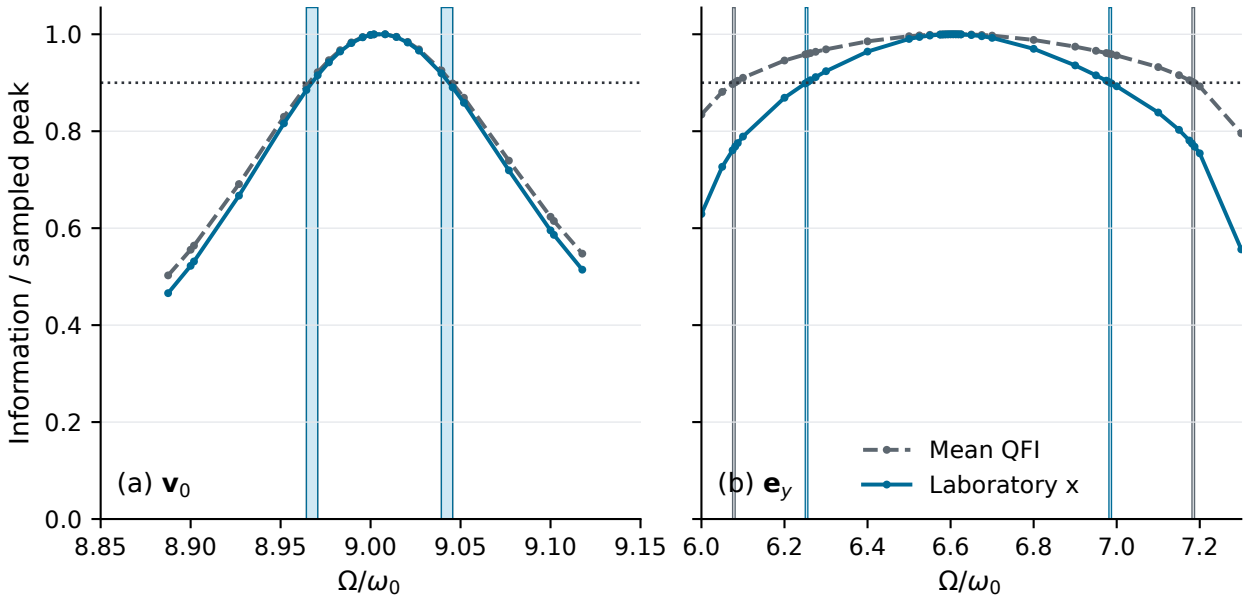


Fig. S1 Profiles normalized to their frozen sampled peaks for (a) $\boldsymbol{v}_0$ and (b) $\boldsymbol{e}_y$. Solid blue curves show laboratory-x information; dashed gray curves show mean QFI. Markers are evaluated amplitudes. The horizontal dotted level is 0.9, and vertical shaded strips show the crossing brackets. The horizontal scales differ

Refinement, optimizer exceptions and counts

The scan contains 234 new antipodal pairs, or 468 endpoint propagations. Baseline settings are $(N_k, L) = (4, 6)$ and output spacing 0.00625. Joint validation uses $(5, 7)$ and spacing 0.003125 at six amplitudes: 0, 6, 12, the old reference, 6.60625, and 9.008106389546, always for both inputs. The largest relative information change at $T = 150$, over QFI, $\mathcal{J}_x$ and the numerical constant-axis value, is 1.30×10^{-5}. Across the five horizons, the largest relative $\mathcal{J}_x$ or QFI change is 1.28×10^{-5}. The largest accessible-fraction change is 0.000463 percentage points. These checks apply to the selected amplitudes, not uniformly to the entire continuum.

Twofold output-grid thinning gives a largest relative change 8.54×10^{-5} and passes the preset 10^{-4} criterion. Maximum trace and Hermiticity errors are 4.57×10^{-13} and 4.36×10^{-15}; the smallest endpoint eigenvalue is -2.78×10^{-17}. A fresh propagation of the exact old reference with the same HEOM implementation reproduces the archived endpoints within 6.94×10^{-18} componentwise. This deterministic reproducibility check is not an independent physical validation.

Every drive-pair analysis uses seven-start constrained optimization and a separate differential-evolution search. At seven evaluations, the original hemisphere-chart differential evolution stopped below a known feasible local-search value, with largest discrepancy 5.39×10^{-6}. Each failed comparison is retained. Additional independent searches use two full-sphere coordinate charts, fixed seeds, larger populations and tighter tolerances, without inserting the local-search candidate. The agreement criterion is not relaxed. The final largest discrepancy is 3.00×10^{-11}. The remaining algebraic envelope gap $100(U - \mathcal{J}_{\mathrm{num}})/\overline{F_Q}$ is at most 0.17993 percentage points in the main interval and 3.45331 percentage points at the undriven points. Thus matching searches are not treated as a global optimality proof.

All final working criteria are satisfied, with the initial optimizer disagreements preserved separately. Combined with the primary 51 pairs and 30 input-study pairs, the numerical record contains 315 antipodal runs and 630 endpoint propagations. Repeated references are retained as actual runs; reused copies of historical controls are not counted again. The studies do not test unknown directions, a different bath family, equal drive-energy costs, a matched Markovian reference or an independent HEOM implementation.

S7 Files and reproduction

Online Resource 2 (`ESM2.zip`) contains the complete numerical record under a short top-level directory `r/`. The two original research batches reside in `repro/hr/` and `repro/fc/`. Each includes:

Directory	Contents
`in/`	Model, prescribed vectors, frozen plans, and prior central comparison trajectories
`data/`	Time arrays and full endpoint Bloch vectors, plus metadata for each setting
`logs/`	Available grid-job output; the separately launched `hr` central baseline has metadata and trajectories but no retained console log
`results/`	Information, refinements, quadrature, response matrices and summaries
`code/`	Solver, grid driver, analysis and plotting scripts; `fc` also contains the spectrum-selection script

Table S16 distinguishes reanalysis of stored trajectories from fresh propagation. Commands are run in the stated working directory with the pinned environment below. Input and output paths are relative to that directory. Use a separate copy for any command that regenerates results.

Table S16 Main reproduction routes. The HEOM column states whether new physical trajectories are computed.

Goal, working directory and command	Inputs and outputs	New HEOM
Primary tables and figures; `r/` `python code/prepare.py`	Input: primary trajectories, plans and calibration records. Output: `numbers.tex`, `extra.tex`, `fig/`, `revision/` and checks.	No
Input/scan integration; `r/` `python code/prepare_extended.py`	Input: frozen `ext/iv/` and `ext/ds/` checkpoints. Output: `paper/` assets and `evidence/extended_checks.json`.	No
Original batch analysis; `r/repro/hr/` or `r/repro/fc/` `python code/analyze.py`	Input: the batch's stored trajectories and frozen plan. Output: regenerated `results/`.	No
Fresh primary dynamics; a separate batch copy `python code/grid.py --workers 6`	Input: retained `in/` with empty `data/` and `logs/`. Output: new trajectories, metadata and logs.	Yes

Recomputing tables and figures

The tested Python environment is 3.12.14 with NumPy 2.3.5, SciPy 1.18.1, QuTiP 5.3.1 and Matplotlib 3.10.8. Install `requirements.txt`. From `r/`, run `python code/prepare.py`. The script checks the original source-copy hashes and frozen plans, recalculates the three averaged information quantities at all five horizons for the 50 original settings, and regenerates `numbers.tex` and the manuscript figures. It also calls `code/extend.py`: this verifies and evaluates the additional joint refinement, optimizes the constant axis at all 13 primary points, checks the exact loss identity, and recalculates outcome-error and timestamp comparisons. It writes `extra.tex`, `revision/analysis.json` and associated CSV tables. It performs no new HEOM propagation.

The joint-refinement dynamics, metadata, saved plan and exact rerun command are under `repro/rv/`. The retained historical calibration records are under `calibration/`; these establish provenance without replacing the full original TCL2 archive. The input and drive research checkpoints are preserved unchanged under `ext/iv/` and `ext/ds/`, each with its original directory layout and checksum manifest. Their `r/` subdirectories contain the physical data and code. Historical checkpoint prose records its stage of work; the present manuscript and SI provide the integrated interpretation.

Reanalysis and fresh dynamics

Work in a separate copy to preserve the delivered record. From either batch directory, run `python code/analyze.py` to reproduce its original analysis, including the response matrices and consistency checks. Results and summaries will be regenerated; the original checksums then refer to the delivered versions.

For fresh dynamics, retain the frozen plan and use empty `data/` and `logs/` directories in the working copy. Run `python code/grid.py --workers 6`, reducing the worker count if required by memory. Complete existing runs are otherwise skipped after parameter checks, and the individual solver refuses to overwrite an existing metadata filename. The preliminary calibration search is outside this package; the prescribed full-precision vectors are sufficient to reproduce this benchmark.

The `fc/code/select.py` script derives the controls from `in/model.json`. It also hashes `in/ref_hr.npz` for provenance, without using its trajectories in the selection. To repeat selection, retain both inputs in a separate copy with empty `data/` and `results/` folders and without `in/plan.json` or `in/plan.sha256`. A new timestamp changes the plan hash.

Integrated tables, figures and audit

Run `python code/prepare_extended.py` from `r/` after installing the same pinned dependencies. This verifies both complete checkpoint manifests, reconstructs the three baseline affine maps from their basis-state endpoints, and reevaluates QFI, laboratory-x information, the stored feasible measurement axes and algebraic upper envelopes for all 1575 input candidates. It also reevaluates all 238 rows of the drive analysis (234 new pairs and four reused control pairs), including their five-horizon $\mathcal{J}_x$ and QFI values, and checks the signs at every width bracket. No new propagation or optimization is performed. Outputs are `paper/extension.tex`, `paper/Fig4.*`, `paper/FigS1.*`, and `evidence/extended_checks.json`. Copy these assets, together with the original generated tables and Figs. 1–3, into the manuscript source directory before running `build.sh`.

To repeat the original input-axis searches, use the commands in `ext/iv/README.txt` in a separate copy. To regenerate all frozen drive trajectories in a new directory, use `ext/ds/r/code/fresh.py`; `ext/ds/README.txt` gives the command and resource settings. This regeneration retains the supplied amplitude lists and width thresholds rather than adaptively selecting a different experiment. Repeated references are intentional runs; copied historical controls do not increase the total of 315 antipodal records.

Provenance

`evidence/provenance.json` records the original copied research files. The source archives for this integration are identified in `evidence/source_archives.json`; the two unchanged checkpoint manifests cover their complete file contents. `SHA256.txt` checks the delivered data archive. Historical model metadata and source READMEs have been retained; any older release DOIs in those files do not identify the new data. Editorial changes are documented in `evidence/editorial_review.json`; the preceding delivery check is retained as `evidence/integration_delivery_qa.json`. The present record is supplied directly as Online Resource 2.

Zenodo data and code record. https://doi.org/10.5281/zenodo.24026283.

Version note. This journal version adds bibliographic references to the archived methods document. Equations, numerical results and the data package supplied as Online Resource 2 are unchanged.

References

[1] Heinz-Peter Breuer and Francesco Petruccione. *The Theory of Open Quantum Systems*. Oxford University Press, Oxford, 2007. https://doi.org/10.1093/acprof:oso/9780199213900.001.0001.

[2] Heinz-Peter Breuer, Elsi-Mari Laine, and Jyrki Piilo. Measure for the Degree of Non-Markovian Behavior of Quantum Processes in Open Systems. *Physical Review Letters*, 103:210401, 2009. https://doi.org/10.1103/PhysRevLett.103.210401.

[3] Yoshitaka Tanimura. Numerically "exact" approach to open quantum dynamics: The hierarchical equations of motion (HEOM). *The Journal of Chemical Physics*, 153:020901, 2020. https://doi.org/10.1063/5.0011599.

[4] Neill Lambert, Tarun Raheja, Simon Cross, Paul Menczel, Shahnawaz Ahmed, Alexander Pitchford, Daniel Burgarth, and Franco Nori. QuTiP-BoFiN: A bosonic and fermionic numerical hierarchical-equations-of-motion library with applications in light-harvesting, quantum control, and single-molecule electronics. *Physical Review Research*, 5:013181, 2023. https://doi.org/10.1103/PhysRevResearch.5.013181.

[5] N. Lambert, E. Giguère, P. Menczel, B. Li, P. Hopf, G. Suárez, M. Gali, J. Lishman, R. Gadhvi, R. Agarwal, A. Galicia, N. Shammah, P. Nation, J. R. Johansson, S. Ahmed, S. Cross, A. Pitchford, and F. Nori. QuTiP 5: The quantum toolbox in Python. *Physics Reports*, 1153:1–62, 2026. https://doi.org/10.1016/j.physrep.2025.10.001.

[6] Jie Hu, Rui-Xue Xu, and YiJing Yan. Communication: Padé spectrum decomposition of Fermi function and Bose function. *The Journal of Chemical Physics*, 133:101106, 2010. https://doi.org/10.1063/1.3484491.

[7] Jie Hu, Meng Luo, Feng Jiang, Rui-Xue Xu, and YiJing Yan. Padé spectrum decompositions of quantum distribution functions and optimal hierarchical equations of motion construction for quantum open systems. *The Journal of Chemical Physics*, 134: 244106, 2011. https://doi.org/10.1063/1.3602466.

[8] Ernst Hairer, Syvert P. Nørsett, and Gerhard Wanner. *Solving Ordinary Differential Equations I: Nonstiff Problems*, volume 8 of *Springer Series in Computational Mathematics*. Springer, Berlin, Heidelberg, second revised edition, 1993. https://doi.org/10.1007/978-3-540-78862-1.

[9] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3):261–272, 2020. https://doi.org/10.1038/s41592-019-0686-2.

[10] Matteo G. A. Paris. Quantum estimation for quantum technology. *International Journal of Quantum Information*, 7(supp01):125–137, 2009. https://doi.org/10.1142/S0219749909004839.

[11] Wei Zhong, Zhe Sun, Jian Ma, Xiaoguang Wang, and Franco Nori. Fisher information under decoherence in Bloch representation. *Physical Review A*, 87:022337, 2013. https://doi.org/10.1103/PhysRevA.87.022337.

[12] Dieter Kraft. A software package for sequential quadratic programming. Technical Report DFVLR-FB 88-28, DFVLR, 1988.

[13] Rainer Storn and Kenneth Price. Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11(4): 341–359, 1997. https://doi.org/10.1023/A:1008202821328.

[14] Samuel L. Braunstein and Carlton M. Caves. Statistical distance and the geometry of quantum states. *Physical Review Letters*, 72:3439–3443, 1994. https://doi.org/10.1103/PhysRevLett.72.3439.

[15] Man-Duen Choi. Completely positive linear maps on complex matrices. *Linear Algebra and its Applications*, 10(3):285–290, 1975. https://doi.org/10.1016/0024-3795(75)90075-0.