# Supplementary Information

Beyond sensitivity: mechanism-resolved error budgets for designing quantum sensors

Nima Leclerc, Marco Capelli, Kevin James Rietwyk, Mark Dong, Dmitry Lyakh, Geoffrey Iwata, Brandon Rodenburg, Sean Oliver, Benedikt Kloss, Jin-Sung Kim, Stefan Bogdanovic, Yunheng Chen, Meysam Sharifzadeh Mirshekarloo, Cedric Weber, Marcus Doherty, Ethan Pratt, and Joseph Hagmann

## Contents

**Overview.** This Supplementary Information provides the derivations, numerical methods, parameter tables, and validation behind the main text. Supplementary Note 1 develops the open-quantum-system model and the tangent master equation, Note 2 validates them against closed forms, Note 3 gives the first-principles rate and contrast inputs, Note 4 defines the three metrics and the Shapley attribution, Note 5 provides the time-dependent-noise model, Note 6 describes the multi-GPU implementation and its spatial output, Note 7 expands upon the cross-platform optically pumped magnetometer, and Note 8 describes the experimental extraction of the charge-state fraction from photoluminescence spectra. Time-dependent noise, per-pixel spatial structure at GPU scale, and a second sensor platform are developed here and exercised only in these Supplementary results, while the main text reports the deterministic single-point and real-device results.

# Supplementary Note 1: Extended Theoretical Derivations

## 1.1 NV Center Hamiltonian

In this paper, we consider the Hamiltonian describing the electronic structure of the negatively charged nitrogen vacancy center ($NV^-$) in diamond, accounting for the ground and excited triplet states and singlet states utilizing a 10-level system. The results in the main text are described by an abstracted Hamiltonian that still accounts for both spatial (positions corresponding to points in the field of view, $\mathbf{r}$) and temporal ($t$) dependencies,

$$\mathcal{H}_{\text{total}}(\mathbf{r},t) = \mathcal{H}_0(\mathbf{r};\{\theta_{\text{sensor}}\}) + \mathcal{H}_{\text{control}}(\mathbf{r},t;\theta_{\text{control}}) + \mathcal{H}_{\text{sense}}(\mathbf{r},t) + \mathcal{H}_{\text{noise}}(t). \tag{S1}$$

(S1) comprises a static term $\mathcal{H}_0(\mathbf{r};\{\theta_{sensor}\})$ describing the inherent electronic structure of the quantum sensor, parameterized by the set of intrinsic material constants $\{\theta_{\text{sensor}}\}$ (e.g. fine structure constants, spin-orbit coupling, etc), a control term $\mathcal{H}_{\text{control}}(\mathbf{r},t;\theta_{\text{control}})$ that captures the spatiotemporally dependent microwave or laser drive used to manipulate the quantum sensor, $\mathcal{H}_{\text{sense}}(\mathbf{r},t)$ captures the unknown field being sensed, and $\mathcal{H}_{\text{noise}}(t)$ accounts for time-dependent noise processes (e.g. microwave phase/amplitude noise, laser noise, etc.).

We ensure that our model properly accounts for phonon-induced temperature dependent rates, laser-leakage (captured in the master equation approach), and local strain fields by using a 10-level Hamiltonian in the $E_Z$ basis described by [1]. This basis spans the states $(|1_g\rangle, |0_g\rangle, |-1_g\rangle, |1_e^x\rangle, |0_e^x\rangle, |-1_e^x\rangle, |1_e^y\rangle, |0_e^y$ The states $(|1_g\rangle, |0_g\rangle, |-1_g\rangle)$ are the spin-triplet states for the ground-state manifold ($^3A_2$) of the $NV^-$ center mapping to $m_s = +1, 0, -1$ magnetic sub-levels, $\left(|1_e^{x/y}\rangle, |0_e^{x/y}\rangle, |-1_e^{x/y}\rangle\right)$ map to the $m_s = +1, 0, -1$ spin triplet states for $x$ (or $y$)-orbital branch of the excited manifold ($^3E$), $|ss\rangle$ is a shelving-state (SS) corresponding to the collapsed $^1E$ and $^1A$ singlet levels. The inclusion of these 10 states becomes relevant in high-magnetic field and high-strain environments typically encountered in quantum sensing applications. Drawing from the $NV^-$ electronic structure formulation first proposed by [2], we can expand the Hamiltonians for the ground and excited state manifolds in (S2) and (S3) (the parameter values used and their sources are found in Table S1).

$$\begin{aligned}\frac{\mathcal{H}_g(\mathbf{r},t)}{h} = D_g\left(\hat{S}_{g,z}^2 - \frac{2}{3}\mathbb{I}_3\right) + \\ \mu_b g_g\left(B_0 + B_{\text{sense}}(t,\mathbf{r}) + \delta b_{\text{noise}}(t)\right) S_{g,z} + \\ \mu_b g_g\left(1 + a_{\text{noise}}(t)\right) f_{\text{control}}(\mathbf{r},t)\left(e^{-i(\omega_{mw}t+\phi(t)+\delta\phi_{\text{noise}}(t))} S_{g,01}^+ + e^{i(\omega_{mw}t+\phi(t)+\delta\phi_{\text{noise}}(t))} S_{g,01}^-\right)\end{aligned} \tag{S2}$$

$$\begin{aligned}\frac{\mathcal{H}_e(\mathbf{r},t)}{h} &= D_e^{||}\mathbb{I}_2\otimes\left(\hat{S}_{e,z}^2-\tfrac{2}{3}\mathbb{I}_3\right)-\lambda_e^{||}\hat{\sigma}_y\otimes\hat{S}_{e,z}\\ &+D_e^{\perp}\left(\hat{\sigma}_z\otimes\left(\hat{S}_{e,y}^2-\hat{S}_{e,x}^2\right)-\hat{\sigma}_x\otimes\left(\hat{S}_{e,y}\hat{S}_{e,x}+\hat{S}_{e,x}\hat{S}_{e,y}\right)\right)\\ &+\lambda_e^{\perp}\left(\hat{\sigma}_z\otimes\left(\hat{S}_{e,x}\hat{S}_{e,z}+\hat{S}_{e,z}\hat{S}_{e,x}\right)-\hat{\sigma}_x\otimes\left(\hat{S}_{e,y}\hat{S}_{e,z}+\hat{S}_{e,z}\hat{S}_{e,y}\right)\right)\\ &+\mu_b g_e\mathbb{I}_2\otimes\left(B_0+B_{\text{sense}}(\mathbf{r},t)+\delta b_{\text{noise}}(t)\right)S_{e,z}\\ &+\mu_B g_l\left(B_0+B_{\text{sense}}(\mathbf{r},t)+\delta b_{\text{noise}}(t)\right)\hat{\sigma}_y\otimes\mathbb{I}_3+d_{es}^{\perp}E^{\perp}(\sigma_z-\sigma_x)\otimes\mathbb{I}_3+d_e^{||}E_z\mathbb{I}_2\otimes\mathbb{I}_3.\end{aligned}\tag{S3}$$

Here, $D_g$ is the ground-state fine structure constant, $S_{g,i}, i\in x,y,z$ are the spin-1 operators $(3\times 3)$, $\mathbb{I}_2$ and $\mathbb{I}_3$ are the $2\times 2$ and $3\times 3$ identity matrices, $\mu_b=\frac{e\hbar}{2m_e}$ is the Bohr magneton, $B_0$ is the static bias magnetic field, $\delta b_{\text{noise}}(t)$ is the time-dependent background magnetic field noise, $B_{\text{sense}}(t,\mathbf{r})$ is the spatiotemporally varying magnetic field being sensed, $g_g$ and $g_e$ are the ground- and excited-state spin gyromagnetic ratios, $g_l$ is the excited-state orbital $g$-factor, $f_{\text{control}}(\mathbf{r},t)$ is the spatiotemporally varying magnetic field for the microwave drive, corresponding to the quantum control pulse sequence being implemented, $a_{\text{noise}}(t)$ is the multiplicative amplitude noise, $\omega_{mw}$ is the microwave frequency driving the transition between the $m_s=0$ and $m_s=-1$ sub-levels, $\phi(t)$ is the programmable time-dependent phase from the control pulse, $\delta\phi_{\text{noise}}(t)$ is the time-dependent microwave phase noise, and $S_{g,01}^{+}=|0_g\rangle\langle -1_g|$ and $S_{g,01}^{-}=|-1_g\rangle\langle 0_g|$ are triplet state operators describing the population inversion between the $m_s=0$ and $m_s=-1$ states under the presence of a microwave drive. The excited state Hamiltonian $\mathcal{H}_e(\mathbf{r},t)$ works in a composite spin-triplet-orbital-doublet Hilbert space $\mathcal{H}_{\text{orbit}}\otimes\mathcal{H}_{\text{spin}}$, where the $\hat{\sigma}_i, i\in\{x,y,z\}$ are the Pauli operators that operate on the orbital-doublet basis, $D_e^{||}$ and $D_e^{\perp}$ are the excited state fine structure constants along the perpendicular and parallel directions, $S_{e,i}, i\in\{x,y,z\}$ are the spin-1 operators for the excited states, and $\lambda_e^{\perp}$ and $\lambda_e^{||}$ are the spin-orbit coupling strengths along the perpendicular and parallel directions.

Throughout, the Hamiltonians are written in frequency units (divided by $h$): the Zeeman couplings $\mu_b g_{g,e}B$ denote the corresponding Zeeman frequencies $\mu_b g_{g,e}B/h\equiv\gamma_{g,e}B$, with $\gamma_{g,e}=\mu_b g_{g,e}/h\approx 28\,\text{GHz/T}$, and in the effective detuning $\Delta_{\text{eff}}$ of Eq. (S9) the carrier and phase-rate terms are likewise expressed as ordinary frequencies (angular quantities divided by $2\pi$), so that every term on the right-hand side carries units of frequency. Equation (S8) is instead written as an energy Hamiltonian, with $h\Delta_{\text{eff}}/2$ on the diagonal.

In our implementation, we use a $10\times 10$ Hamiltonian constructed from the ground ($\mathcal{H}_g(\mathbf{r},t)$), excited ($\mathcal{H}_e(\mathbf{r},t)$), and the shelving state (SS) ($E_{SS}$) states given in (S4). $E_{SS}$ here is a scalar value of the singlet state energy level.

$$\mathcal{H}_{\text{NV}}(\mathbf{r},t)=\begin{pmatrix}\mathcal{H}_g(\mathbf{r},t) & 0 & 0\\ 0 & \mathcal{H}_e(\mathbf{r},t) & 0\\ 0 & 0 & E_{SS}\end{pmatrix}\tag{S4}$$

We work in the interaction picture and apply the rotating wave approximation, first separating out the time-dependent ($t$-subscript) and time-independent (0-subscript) contributions for the ground and excited states,

$$\mathcal{H}_{\text{NV}}(\mathbf{r},t)=\begin{pmatrix}\mathcal{H}_{0,g}(\mathbf{r}) & 0 & 0\\ 0 & \mathcal{H}_{0,e}(\mathbf{r}) & 0\\ 0 & 0 & E_{SS}\end{pmatrix}+\begin{pmatrix}\mathcal{H}_{t,g}(\mathbf{r},t) & 0 & 0\\ 0 & \mathcal{H}_{t,e}(\mathbf{r},t) & 0\\ 0 & 0 & 0\end{pmatrix}.\tag{S5}$$

These contributions are expanded out in (S6).

$$
\begin{aligned}
\frac{\mathcal{H}_{0,g}(\mathbf{r})}{h} &= D_g\left(\hat{S}_{g,z}^2 - \tfrac{2}{3}\mathbb{I}_3\right) + \mu_b g_g B_0 S_{g,z}, \\
\frac{\mathcal{H}_{0,e}(\mathbf{r})}{h} &= D_e^{\parallel}\mathbb{I}_2 \otimes \left(\hat{S}_{e,z}^2 - \tfrac{2}{3}\mathbb{I}_3\right) - \lambda_e^{\parallel}\hat{\sigma}_y \otimes \hat{S}_{e,z} \\
&\quad + D_e^{\perp}\left(\hat{\sigma}_z \otimes \left(\hat{S}_{e,y}^2 - \hat{S}_{e,x}^2\right) - \hat{\sigma}_x \otimes \left(\hat{S}_{e,y}\hat{S}_{e,x} + \hat{S}_{e,x}\hat{S}_{e,y}\right)\right) \\
&\quad + \lambda_e^{\perp}\left(\hat{\sigma}_z \otimes \left(\hat{S}_{e,x}\hat{S}_{e,z} + \hat{S}_{e,z}\hat{S}_{e,x}\right) - \hat{\sigma}_x \otimes \left(\hat{S}_{e,y}\hat{S}_{e,z} + \hat{S}_{e,z}\hat{S}_{e,y}\right)\right) \\
&\quad + \mu_B B_0\left(g_l\,\hat{\sigma}_y \otimes \mathbb{I}_3 + g_e\,\mathbb{I}_2 \otimes S_{e,z}\right) + d_{es}^{\perp}E^{\perp}(\sigma_z - \sigma_x) \otimes \mathbb{I}_3 + d_e^{\parallel}E_z\mathbb{I}_2 \otimes \mathbb{I}_3, \\
\frac{\mathcal{H}_{t,g}(\mathbf{r},t)}{h} &= \mu_b g_g\Bigg( \left(B_{\text{sense}}(\mathbf{r},t) + \delta b_{\text{noise}}(t)\right) S_{g,z} \\
&\quad + \left(1 + a_{\text{noise}}(t)\right) f_{\text{control}}(\mathbf{r},t)\Big(e^{-i(\omega_{mw}t + \phi(t) + \delta\phi_{\text{noise}}(t))}S_{g,01}^{+} \\
&\qquad + e^{i(\omega_{mw}t + \phi(t) + \delta\phi_{\text{noise}}(t))}S_{g,01}^{-}\Big)\Bigg), \\
\frac{\mathcal{H}_{t,e}(\mathbf{r},t)}{h} &= \mu_B\left(B_{\text{sense}}(\mathbf{r},t) + \delta b_{\text{noise}}(t)\right)\left(g_l\,\hat{\sigma}_y \otimes \mathbb{I}_3 + g_e\,\mathbb{I}_2 \otimes S_{e,z}\right).
\end{aligned}
\tag{S6}
$$

Note that time-dependent electric field noise and temperature dependent fluctuations due to laser-induced heating are neglected here, though these effects can be included in the Hamiltonian above. We now proceed to transform our lab-frame Hamiltonian to the interaction frame, first defining the unitaries $\mathcal{U}_{0,g}(t) = \exp\left(-\frac{it}{\hbar}\mathcal{H}_{0,g}(\mathbf{r})\right)$ and $\mathcal{U}_{0,e}(t) = \exp\left(-\frac{it}{\hbar}\mathcal{H}_{0,e}(\mathbf{r})\right)$. The interaction frame Hamiltonian $\mathcal{H}_I(t)$ is then given in (S7).

$$
\mathcal{H}_I(\mathbf{r},t) = \begin{pmatrix} \mathcal{U}_{0,g}^{\dagger}(t)\,\mathcal{H}_{t,g}(\mathbf{r},t)\,\mathcal{U}_{0,g}(t) & 0 & 0 \\ 0 & \mathcal{U}_{0,e}^{\dagger}(t)\,\mathcal{H}_{t,e}(\mathbf{r},t)\,\mathcal{U}_{0,e}(t) & 0 \\ 0 & 0 & 0 \end{pmatrix}
\tag{S7}
$$

We then apply the rotating-wave approximation (RWA) in the frame rotating at $\omega_{mw}$, tuned to the $|0_g\rangle \to |-1_g\rangle$ transition, giving the ground-state RWA Hamiltonian (S8). This reduced expression retains the control variables (the microwave waveform $f_{\text{control}}(\mathbf{r},t)$, carrier $\omega_{mw}$, programmable phase $\phi(t)$, and bias $B_0$) and the time-dependent noise terms (temperature $T(t)$, phase $\phi_{\text{noise}}$, amplitude $a_{\text{noise}}$, and field $\delta b_{\text{noise}}$), collected into the effective detuning $\Delta_{\text{eff}}(t)$ of (S9), in which the temperature enters through $D_g(T)$. Counter-rotating terms give small Bloch–Siegert shifts ($\ll$ kHz here) and are neglected. The $|+1_g\rangle$ sub-level remains a spectator, shifted only by the slow drifts $\delta b$ and $\delta D$.

$$
\mathcal{H}_g^{(I,\text{RWA})}(\mathbf{r},t) = \begin{pmatrix} \frac{h}{2}\Delta_{\text{eff}}(\mathbf{r},t) & \frac{\mu_b g_g}{2\sqrt{2}}f_{\text{control}}(\mathbf{r},t)\left(1 + a_{\text{noise}}(t)\right)e^{-i\phi(t)} & 0 \\ \frac{\mu_b g_g}{2\sqrt{2}}f_{\text{control}}(\mathbf{r},t)\left(1 + a_{\text{noise}}(t)\right)e^{i\phi(t)} & -\frac{h}{2}\Delta_{\text{eff}}(\mathbf{r},t) & 0 \\ 0 & 0 & 0 \end{pmatrix}
\tag{S8}
$$

$$
\Delta_{\text{eff}}(\mathbf{r},t) = \left(D_g(T(t)) - \gamma_g\left(B_0 + b_{\text{sense}}(\mathbf{r},t) + \delta b_{\text{noise}}(t)\right)\right) - \frac{\omega_{mw}}{2\pi} - \frac{1}{2\pi}\frac{\partial\phi_{\text{noise}}(t)}{\partial t}
\tag{S9}
$$

The interaction picture is defined with respect to the ground-state Hamiltonian, i.e. at microwave frequencies, whereas terms in the excited manifold evolve at optical frequencies (hundreds

of THz). Upon transforming to the interaction frame, the excited-state Hamiltonian acquires rapidly oscillating phase factors. Under the rotating-wave approximation these oscillating terms average to zero, leaving only the static diagonal splittings (orbital strain, spin–orbit coupling, Zeeman shifts) and phonon-mediated mixing terms. As a result, the effective excited-state Hamiltonian takes the same form in the interaction picture as in the Schrödinger picture. Throughout this work we therefore use the Schrödinger-picture expressions for the excited manifold, while explicitly applying the interaction-picture/RWA treatment only to the ground-state spin Hamiltonian. The final $10 \times 10$ spatiotemporal Hamiltonian used in the quantum dynamics simulations under RWA is given in (S10). We do not include coherent optical drives. The green laser excitation enters only via dissipators, so $\mathcal{H}_e$ has no explicitl time dependence.

$$\mathcal{H}^{(I,\mathrm{RWA})}(\mathbf{r},t) = \begin{pmatrix} \mathcal{H}_g^{(I,\mathrm{RWA})}(\mathbf{r},t) & 0 & 0 \\ 0 & \mathcal{H}_e(\mathbf{r},t) & 0 \\ 0 & 0 & 0 \end{pmatrix} \tag{S10}$$

This provides us with a time-dependent Hamiltonian in the interaction picture that we can use to evolve our master equation, as will be described in 1.2.

Table S1: Hamiltonian parameters used in the NV-center model, with the primary source for each parameter listed. $D_g$ is the zero-temperature zero-field splitting, and its temperature dependence enters separately in the thermal channel.

| **Parameter** | **Symbol** | **Value** | **Reference/Note** |
|---|---|---|---|
| Ground-state zero-field splitting | $D_g$ | 2.878 GHz | [3] |
| Ground/excited-state $g$-factor | $g_g$ | 2.003 | [4] |
| Excited-state $\perp$ splitting | $D_e^{\perp}$ | 0.7705 GHz | 1.541 GHz/2, nvratemodel convention [5] |
| Excited-state $\|\|$ splitting | $D_e^{\|\|}$ | 1.44 GHz | [5] |
| $\|\|$ spin-orbit coupling | $\lambda_e^{\|\|}$ | 5.33 GHz | [5] |
| $\perp$ spin-orbit coupling | $\lambda_e^{\perp}$ | 0.154 GHz | [5] |
| Excited-state orbital $g$-factor | $g_l$ | 0.1 | [6] |

## 1.2 Lindblad Master Equation

The workhorse of the digital twin is its ability to combine time-dependent noise processes, control protocols, and external fields with material-dependent energy level structures to forecast quantities of interest relevant to quantum sensing (e.g. sensitivity, accuracy) using quantum dynamics simulations. We use the Lindblad master equation ([7]) to predict the density matrix trajectories $\rho(\mathbf{r},t)$ under various time-dependent processes, working in the interaction picture $\rho_I(\mathbf{r},t)$. (S11) defines the transformations of the density matrix $\rho(\mathbf{r},t)$ and collapse operators $C_k(\mathbf{r},t)$ into the interaction frame (given for the $k$th collapse process). The collapse operators here are written a

$C_k(\mathbf{r},t) = \sqrt{\gamma_k(\mathbf{r},t)}L_k$ where $\gamma_k(\mathbf{r},t)$ is a time and spatially dependent rate for $k$th dissipative processes and $L_k$ is its corresponding non-unitary "jump" operator.

$$\begin{aligned}\mathcal{U}_0(\mathbf{r},t) = \exp\left(-\frac{i}{\hbar}\mathcal{H}_0(\mathbf{r})\right), \rho_I(\mathbf{r},t) = \mathcal{U}_0^\dagger(\mathbf{r},t)\,\rho(\mathbf{r},t)\,\mathcal{U}_0(\mathbf{r},t)\\ C_k^I(t) = \mathcal{U}_0^\dagger(\mathbf{r},t)\,C_k(\mathbf{r},t)\,\mathcal{U}_0(\mathbf{r},t)\end{aligned} \tag{S11}$$

Working in the transformed interaction frame, we solve for the density matrix $\rho_I$ solve the quantum state dynamics in (S12). In our implementation, the master equation is solved in the interaction picture and with the RWA where trajectories of $\rho_I(\mathbf{r},t)$ in time are solved up to the maximum sensing duration $T_{\text{sense}}$, then they are transformed back to the Schrödinger picture using the transformation in (S11) to recover $\rho(\mathbf{r},t)$.

$$\frac{\partial\rho_I(\mathbf{r},t)}{\partial t} = -\frac{i}{\hbar}\left[\mathcal{H}^{(I,\text{RWA})}(\mathbf{r},t), \rho_I(\mathbf{r},t)\right] + \sum_k C_k^I(\mathbf{r},t)\rho_I(\mathbf{r},t)\,C_k^{I,\dagger}(\mathbf{r},t) - \frac{1}{2}\{C_k^{I,\dagger}(\mathbf{r},t)C_k^I(\mathbf{r},t), \rho_I(\mathbf{r},t)\} \tag{S12}$$

**Remark on collapse operators in the interaction picture.** The Lindblad dissipator is invariant under multiplication of a collapse operator by a global phase, i.e. $\mathcal{D}[e^{i\theta(t)}C]\rho = \mathcal{D}[C]\rho$. When transforming to the interaction picture with respect to $H_0$, a generic collapse operator $C^{(S)}$ becomes $C^{(I)}(t) = U_0^\dagger(t)C^{(S)}U_0(t)$. The ground-state channels ($T_1$ relaxation and dephasing) are projectors or single ladders between $H_0$ eigenstates, and their interaction-picture form differs from the Schrödinger-picture form only by a phase factor $e^{i\omega t}$, which cancels inside the dissipator, so they may be written in the same algebraic form in both pictures, with only their scalar rates $\gamma_k(t)$ (e.g. laser-gating or temperature-dependent phonon exchange) carrying time dependence.

The excited-manifold channels (radiative decay, intersystem crossing, singlet relaxation, and phonon-mediated orbital exchange) are written in the bare $|E_x\rangle, |E_y\rangle$ orbital basis, which is not diagonal in $H_{0,e}$ at finite strain, spin–orbit coupling, or magnetic field. We invoke the standard secular (fast-orbital-averaging) approximation, valid above $\sim$80 K where phonon-driven mixing dominates (Methods), under which these cross terms average to zero, leaving a sum of single-frequency channels of the same algebraic form. Consistent with Sec. 1.1, we therefore use the Schrödinger-picture expressions for the excited manifold.

**Proof sketch.** Let $C^{(I)}(t) = e^{i\theta(t)}C^{(S)}$. Then

$$\begin{aligned}\mathcal{D}[C^{(I)}]\rho &= C^{(I)}\rho C^{(I)\dagger} - \tfrac{1}{2}\{C^{(I)\dagger}C^{(I)}, \rho\}\\ &= e^{i\theta}C^{(S)}\rho e^{-i\theta}C^{(S)\dagger} - \tfrac{1}{2}\{C^{(S)\dagger}C^{(S)}, \rho\}\\ &= C^{(S)}\rho C^{(S)\dagger} - \tfrac{1}{2}\{C^{(S)\dagger}C^{(S)}, \rho\}\\ &= \mathcal{D}[C^{(S)}]\rho,\end{aligned}$$

showing that global phases from the interaction-picture transformation cancel exactly. For any collapse operator that is a projector or a single ladder between $H_0$ eigenstates, the dissipator is therefore unchanged and the operator takes the same form in either picture, which gives (S13).

$$\frac{\partial\rho_I(\mathbf{r},t)}{\partial t} = -\frac{i}{\hbar}\left[\mathcal{H}^{(I,\text{RWA})}(\mathbf{r},t), \rho_I(\mathbf{r},t)\right] + \sum_k C_k(\mathbf{r},t)\rho_I(\mathbf{r},t)\,C_k^\dagger(\mathbf{r},t) - \frac{1}{2}\{C_k^\dagger(\mathbf{r},t)C_k(\mathbf{r},t), \rho_I(\mathbf{r},t)\} \tag{S13}$$

Table S2: Dissipation channels in the NV-center model.

| **Process** | **Collapse operator(s) $C_k$** | **Rate(s) / Notes** |
|---|---|---|
| Ground-state relaxation ($T_1$) | $\sqrt{\Gamma_{1,\downarrow}}\,\lvert 0_g\rangle\langle -1_g\rvert$, $\sqrt{\Gamma_{1,\uparrow}}\,\lvert -1_g\rangle\langle 0_g\rvert$ | Rates: $\Gamma_{1,\downarrow}$, $\Gamma_{1,\uparrow}$. Spin-lattice relaxation; the $0 \leftrightarrow +1$ operators are included analogously. |
| Ground-state dephasing ($T_2^*$) | $\sqrt{\gamma_\phi}\,\hat{S}_{g,z} = \sqrt{\gamma_\phi}\,(\lvert 1_g\rangle\langle 1_g\rvert - \lvert -1_g\rangle\langle -1_g\rvert)$ | Rate: $\gamma_\phi = 2/T_2^\star$. Pure dephasing channel. |
| Laser excitation (green pump) | $\sqrt{\gamma_{\mathrm{ex},m_s}(t)}\,\lvert e, m_s\rangle\langle g, m_s\rvert$ | Rate: $\gamma_{\mathrm{ex},m_s}(t)$. Time-dependent optical pumping. |
| Radiative decay | $\sqrt{\Gamma_{\mathrm{rad},m_s\to m_s'}}\,\lvert g, m_s'\rangle\langle e, m_s\rvert$ | Rates: $\Gamma_{\mathrm{rad},m_s\to m_s'}$. Optical emission with possible branching. |
| Intersystem crossing (triplet $\to$ singlet) | $\sqrt{\Gamma_{\mathrm{ISC},m_s}}\,\lvert S\rangle\langle e, m_s\rvert$ | Rates: $\Gamma_{\mathrm{ISC},m_s}$. Non-radiative, stronger for $m_s = \pm 1$. |
| Singlet repump | $\sqrt{\Gamma_{S\to 0}}\,\lvert 0_g\rangle\langle S\rvert$ | Rate: $\Gamma_{S\to 0}$. Polarization pathway back to $\lvert 0_g\rangle$. |
| Phonon orbital exchange | $\sqrt{\Gamma_{x\to y}(T)}\,\lvert E_y, m_s\rangle\langle E_x, m_s\rvert$, $\sqrt{\Gamma_{y\to x}(T)}\,\lvert E_x, m_s\rangle\langle E_y, m_s\rvert$ | Rates: $\Gamma_{x\to y}(T)$, $\Gamma_{y\to x}(T)$. One-/two-phonon processes, $T$-dependent. |

Table S3: Numerical values or scaling laws for the dissipation rates used in simulation. The radiative, intersystem-crossing, and singlet rates are the $NV^-$ rate-model values of Ref. [1] (their Table I). The ground-state $T_1$ and $T_2^\star$ are representative ensemble inputs for the sensing context, not part of that rate model; the attribution operating point instead uses $T_1 = 5\,\mathrm{ms}$ and $T_2^\star = 1.5\,\mu\mathrm{s}$ (Table S7).

| **Process** | **Symbol** | **Value / scaling** | **Source / note** |
|---|---|---|---|
| Radiative (optical emission) | $k_r$ | $55.7\ \mu\mathrm{s}^{-1}$ (17.9 ns) | |
| ISC, $m_s = \pm 1$ | $\bar{k}_{E_{12}}$ | $98.7\ \mu\mathrm{s}^{-1}$ (10.1 ns) | |
| ISC, $m_s = 0$ | $k_{E_{xy}}$ | $8.2\ \mu\mathrm{s}^{-1}$ (122 ns) | |
| Singlet decay ($T = 0$) | $\tau_{S,0}$ | 320 ns | |
| Singlet branching | $r_S = k_{S0}/k_{S1}$ | 2.26 | |
| Singlet phonon energy | $\Delta E$ | 16.6 meV | |
| Phonon orbital exchange | $\Gamma_{x\to y}(T)$ | full Debye integral (Sec. 3.1) | $\eta = 176\ \mu\mathrm{s}^{-1}\mathrm{meV}^{-3}$, $\hbar\Omega = 168\,\mathrm{meV}$ |
| Ground-state relaxation ($T_1$) | $\Gamma_1$ | $(1\ \mathrm{ms})^{-1}$ | Ensemble input |
| Ground-state dephasing ($T_2^\star$) | $\gamma_\phi$ | $(0.5\ \mu\mathrm{s})^{-1}$ | Ensemble input ($\gamma_\phi = 2/T_2^\star$, $T_2^\star = 1\,\mu\mathrm{s}$) |

### 1.3 Tangent Master Equation

Many of the quantities of interest in this work, such as Fisher information and sensitivity bounds, require not only the density matrix $\rho(t)$ but also its derivatives with respect to system parameters. For example, derivatives with respect to microwave detuning $\Delta_{\text{eff}}$, control amplitude $f_{\text{control}}(t)$, or phonon transition rates $\Gamma_{x\to y}(T)$ are needed to quantify how parameter variations affect measurement outcomes. While finite-difference methods are possible, they are computationally inefficient and prone to numerical error. Instead, we employ the *tangent master equation* (TME), the standard tangent-linear construction of open-system estimation [8], which propagates parameter derivatives directly alongside the state.

Starting from the Lindblad master equation (Sec. 1.2),

$$\dot{\rho}(t) = -\frac{i}{\hbar}[H(t), \rho(t)] + \sum_k \Big( C_k \rho(t) C_k^\dagger - \tfrac{1}{2}\{C_k^\dagger C_k, \rho(t)\} \Big),$$

we differentiate with respect to a parameter $\theta$. This yields the tangent master equation

$$\begin{aligned} \frac{\partial}{\partial t}\left(\frac{\partial \rho}{\partial \theta}\right) = &-\frac{i}{\hbar}\Big( [H(t), \tfrac{\partial \rho}{\partial \theta}] + [\tfrac{\partial H}{\partial \theta}, \rho] \Big) \\ &+ \sum_k \Big( C_k \tfrac{\partial \rho}{\partial \theta} C_k^\dagger - \tfrac{1}{2}\{C_k^\dagger C_k, \tfrac{\partial \rho}{\partial \theta}\} \Big) \\ &+ \sum_k \Big( (\tfrac{\partial C_k}{\partial \theta}) \rho C_k^\dagger + C_k \rho (\tfrac{\partial C_k^\dagger}{\partial \theta}) - \tfrac{1}{2}\{ (\tfrac{\partial C_k^\dagger}{\partial \theta}) C_k + C_k^\dagger (\tfrac{\partial C_k}{\partial \theta}), \rho \} \Big). \end{aligned} \tag{S14}$$

This governs the joint dynamics of $\rho$ and its tangent $\partial_\theta \rho$.

The TME describes the coupled evolution of the state $\rho(t)$ and its tangent vectors $\partial_\theta \rho(t)$. It preserves the Hermiticity and tracelessness of the tangent, $G_\theta = G_\theta^\dagger$ and $\text{Tr}\,[G_\theta] = 0$, with the parent state $\rho$ remaining trace-preserving and positive as in the Lindblad equation. The tangent $G_\theta$ is not itself a density matrix, being generically indefinite as the derivative of a state. This construction avoids the need for repeated simulations with slightly perturbed parameters. In our NV-center model, relevant parameters include $\Delta_{\text{eff}}$, $f_{\text{control}}(t)$, ground-state decoherence rates $(\Gamma_{1,\downarrow}, \Gamma_{1,\uparrow}, \gamma_\phi)$, and temperature-dependent orbital exchange rates $(\Gamma_{x\to y}(T), \Gamma_{y\to x}(T))$.

The tangent master equation is central to our digital-twin framework: it provides analytic parameter derivatives that are later used in Sec. 1.4 to compute Fisher information and sensitivity bounds.

#### Tangent master equation for magnetic-field sensing ($\theta = b_{\text{sense}}$)

**Setting.** In the interaction/RWA frame (Sec. 1.2) the only explicit $b$–dependence (where $b$ is the magnetic field being sensed) of the Hamiltonian is in the ground-manifold detuning,

$$\Delta_{\text{eff}}(t) \;=\; D_g(T) - \mu_B g_g \,[\, B_0 + b_{\text{sense}}\, s(t)\,] - \mu_B g_g\, \delta b(t) \;-\; \omega_{\text{mw}} - \dot{\varphi}_{\text{noise}}(t),$$

where $s(t)$ is the known sensing waveform (for DC sensing take $s(t) \equiv 1$). In the $\{|0_g\rangle, |-1_g\rangle\}$ block we use

$$H_g^{(I,\text{RWA})}(t) = \frac{\hbar}{2}\begin{pmatrix} 0 & \Omega(t) e^{-i\varphi(t)} \\ \Omega(t) e^{+i\varphi(t)} & 2\,\Delta_{\text{eff}}(t) \end{pmatrix}.$$

We take $\Omega\,(t) \;=\; \frac{\mu_B g_g}{\hbar}\,(1 + a_{\text{noise}}\,(t))\, f_{\text{control}}\,(\mathbf{r}, t)\,|\langle 0_g|S_{g,x}|-1_g\rangle$ . The excited/singlet blocks have no explicit $b_{\text{sense}}$ dependence in our model (optical processes treated via dissipators).

**Derivative operators.** Let $G_b(t) \equiv \partial_{b_{\text{sense}}}\rho_I(t)$. Since collapse *operators* $L_k$ are fixed ladders/projectors (Sec. 1.2), $\partial_{b_{\text{sense}}}L_k = 0$ and only $\partial_{b_{\text{sense}}}H$ contributes:

$$\frac{\partial H^{(I,\text{RWA})}(t)}{\partial b_{\text{sense}}} = -\frac{1}{2}\,\mu_B g_g\, s(t)\;\left(|0_g\rangle\langle 0_g| - |{-1_g}\rangle\langle {-1_g}|\right) \quad \text{(DC case: } s(t)\equiv 1\text{)}.$$

Dissipation *rates* are taken $b$-independent (no Zeeman modification of optical/phonon rates is included), so $\partial_{b_{\text{sense}}}\gamma_k(t) = 0$.

**NV-specific TME (single parameter).** Differentiating the Lindblad equation (Sec. 1.2) yields

$$\boxed{\dot{G}_b(t) = -\frac{i}{\hbar}\Big([H^{(I,\text{RWA})}(t),\, G_b(t)] + [\,\partial_{b_{\text{sense}}}H^{(I,\text{RWA})}(t),\, \rho_I(t)]\Big) \;+\; \sum_k \gamma_k(t)\,\mathcal{D}[L_k]\,G_b(t),}$$

with initial condition $G_b(0) = 0$ if $\rho_I(0)$ is $b$-independent (standard for NV initialization).

**Measurement derivative and Fisher link.** For the photoluminescence (PL) readout modeled by $M_0 = |0_g\rangle\langle 0_g|$,

$$\partial_{b_{\text{sense}}}p_0(t) \;=\; \text{Tr}[M_0\, G_b(t)]\,,$$

which we use in Sec. 1.4 to assemble the Fisher information for the chosen sensing sequence (Ramsey/echo), with $s(t)$ encoding the known modulation.

## 1.4 Poisson Statistics and Photon Shot Noise

Assume that we have $N$ independent emitters and we describe the probability of being in a correct state for that emitter to emit a photon as $s(T, b)$. We include the collection efficiency $\eta_{\text{opt}}$ and therefore the probability of detecting a photon from the single emitter is

$$p_s = \eta_{\text{opt}} s(T, b). \tag{S15}$$

Now if we have a large collection of $N$ independent emitters, the probability $p(x|b)$ of detecting $x$ photons (where $x$ is an integer in the range $[0, N]$) is given by a binomial distribution:

$$p(x|b) = \binom{N}{x} p_s^x (1-p_s)^{N-x} = \frac{N!}{x!(N-x)!} p_s^x (1-p_s)^{N-x} \tag{S16}$$

Now we can compute the Fisher Information between the measurement of photon counting $X$ and the magnetic field characterized by $b$ as

$$\begin{aligned}
I(b) &= \sum_{x=0}^{N} p(x|b)\,(\partial_b \log(p(x|b)))^2 = \sum_{x=0}^{N} p\left(\frac{1}{p}\partial_b p\right)^2 = \sum_{x=0}^{N}\frac{1}{p}\,(\partial_b p)^2 \\
&= \sum_{x=0}^{N}\frac{1}{p}\left(\frac{x}{p_s}p\partial_b p_s - \frac{N-x}{1-p_s}p\partial_b p_s\right)^2 = \sum_{x=0}^{N} p\left(\frac{x-Np_s}{p_s(1-p_s)}\partial_b p_s\right)^2 \\
&= \left(\frac{\partial_b p_s}{p_s(1-p_s)}\right)^2 \sum_{x=0}^{N} p\left[x^2 - 2Np_s x + (Np_s)^2\right] \\
&= \left(\frac{\partial_b p_s}{p_s(1-p_s)}\right)^2 \left[\left(Np_s(1-p_s) + (Np_s)^2\right) - 2(Np_s)^2 + (Np_s)^2\right] \\
&= \frac{N(\partial_b p_s)^2}{p_s(1-p_s)} = \frac{N\eta_{\text{opt}}}{s(T,b)\left[1-\eta_{\text{opt}}s(T,b)\right]}\left(\frac{\partial s(T,b)}{\partial b}\right)^2,
\end{aligned} \tag{S17}$$

where we have used $\mathbb{E}(x) = Np_s$ and $\mathbb{E}(x^2) = Np_s(1-p_s) + (Np_s)^2$. $s\,(T, b)$ and $\partial_b s$ are obtained from solving the Lindblad and tangent master equations in Sec. 1.2 and 1.3. Together, this provides a framework to numerically evaluate the impacts of photon shot noise and other time-dependent noise sources on the classical Fisher information and sensitivity of a quantum sensing protocol.

The same shot-noise treatment carries over to the optically pumped magnetometer of Sec. 7, with two substitutions. The detected observable is the transmitted-probe optical rotation rather than NV photoluminescence, and the estimated quantity is the Larmor precession frequency of the free-induction-decay signal, which maps to the field through the gyromagnetic ratio $\gamma$. The per-shot photon statistics enter the free-induction-decay Cramér–Rao bound of Sec. 7.4 in the same binomial-to-Poisson form, so the sensitivity floors of both platforms are evaluated within one photon-counting Fisher framework.

# Supplementary Note 2: Validation and benchmarks

## 2.1 Numerical checks of the assembled model

Before validating the tangent solve against closed forms, we confirm that the assembled open-system model reproduces known device physics and the full estimation chain end-to-end.

As a first check that the collapse-operator set and its rates are assembled correctly, we propagate the master equation under continuous 532 nm optical pumping from several initial spin states (Fig. S1). All initial states relax to a common steady-state polarization of $\sim$0.69 into $m_s$=0 within $\sim 2\,\mu$s, driven by the spin-selective intersystem crossing through the shelving singlet, reproducing the known optical spin-initialization mechanism of the NV$^-$ center that sets the readout contrast.

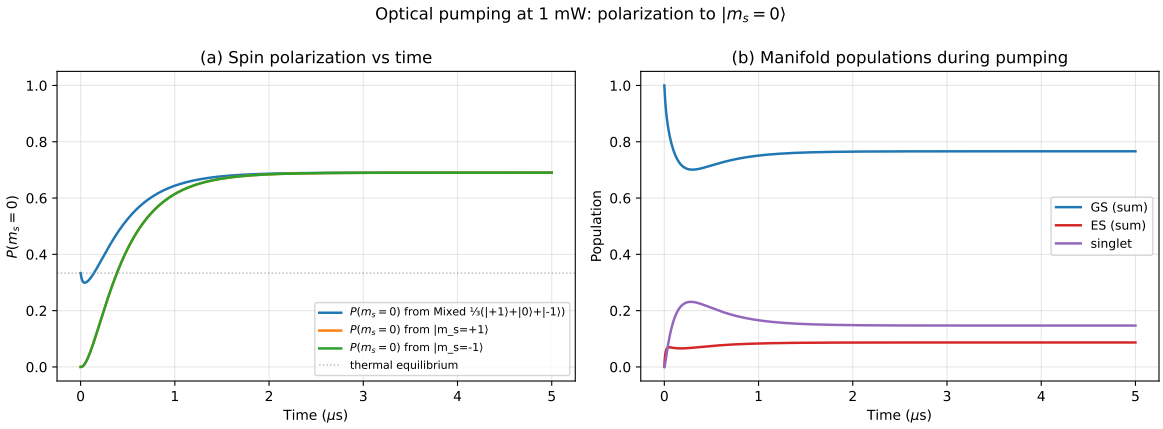


Figure S1: Optical spin polarization in the ten-level model. **(a)** Probability $P(m_s{=}0)$ under 1 mW 532 nm pumping, starting from the fully mixed state and from $|m_s{=}\pm1\rangle$. All initial states converge to a steady-state polarization of $\sim$0.69 within $\sim 2\,\mu$s, above the thermal-equilibrium value 1/3 (dotted). **(b)** Ground-state, excited-state, and singlet manifold populations during pumping. The transient excited-state population and the shelving singlet drive the spin-selective repolarization. The model reproduces the spin-initialization physics of the NV$^-$ center that sets the readout contrast of Sec. 3.2.

The full estimation chain is exercised end to end in Fig. S2, which recovers a static 2 $\mu$T field from a Ramsey sequence at $T_2^\star{=}1\,\mu$s and contrast $C{=}0.46$ under $1/f$ and photon shot noise. The

forward density-matrix solve produces the calibration fringe $p(B)$, a Poisson draw of 100 detected-photon shots at the true field yields the noisy readout, and inverting the calibration recovers the field. Repeating the draw multiple times sets the empirical field uncertainty, which sits a factor $\sim 4.0$ above the ideal projection Cramér–Rao bound because of the finite contrast and the $1/f$ background. This example exposes, in one figure, exactly which non-idealities (contrast dilution, $1/f$ dephasing, shot noise) separate the performance of a realistic device from the quantum limit.

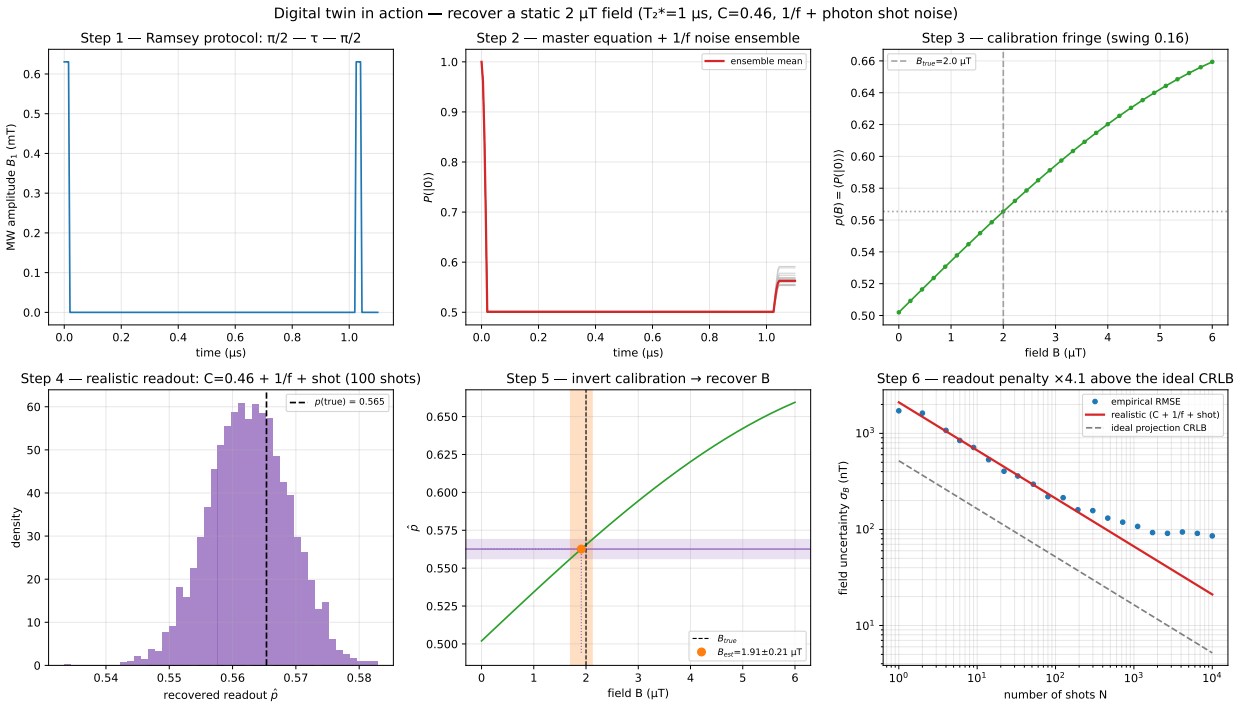


Figure S2: The digital twin recovering a static field, end to end. A 2 $\mu$T DC field is recovered from a Ramsey sequence at $T_2^\star$=1 $\mu$s and contrast $C$=0.46, under $1/f$ and photon shot noise. **(Step 1)** the $\pi/2$-$\tau$-$\pi/2$ pulse sequence. **(Step 2)** the ground-state readout $P(|0\rangle)$ propagated by the master equation (Sec. 1.2) over a $1/f$ noise ensemble. **(Step 3)** the calibration fringe $p(B)$ of swing 0.16. **(Step 4)** the Poisson-sampled readout $\hat{p}$ over 100 shots at the true field. **(Step 5)** inverting the calibration recovers $\hat{B}$=1.91 $\pm$ 0.21 $\mu$T against a true 2.0 $\mu$T. **(Step 6)** the field uncertainty $\sigma_B$ versus shot number, where the realistic readout (contrast, $1/f$, and shot noise) sits a factor $\sim$4.0 above the ideal projection Cramér–Rao bound. Together these steps exercise the full sensitivity pipeline of Secs. 1.3 and 1.4.

## 2.2 Tangent master equation versus finite difference

The tangent solve is the computational linchpin of the framework, so we validate it against independent references along two axes. First, the propagated derivative $\partial_b\rho$ agrees with a central finite-difference of two full solves to better than $10^{-4}$ across the operating range, and the resulting sensitivity is independent of the integration time-grid (a spurious $\sqrt{N_t}$ drift appears only if the Fisher information is accumulated incorrectly). The quantum Fisher information (QFI) used as the reference here is defined through the symmetric logarithmic derivative (SLD) $L_b$, the Hermitian operator solving $\partial_b\rho = \frac{1}{2}(L_b\rho + \rho L_b)$, as $F_Q(b) = \mathrm{Tr}[\rho(b)\, L_b^2]$. It is the estimator-independent upper bound on the classical Fisher information of any measurement, $F_C(b) \leq F_Q(b)$, and we assemble

it from the same tangent $G_b = \partial_b \rho$ that the framework already propagates (Sec. 1.3), so it is a check on the classical photon-counting Fisher information rather than a separate solve. Second, the assembled Fisher information reproduces closed-form limits (Fig. S3): with dissipation switched off it saturates the closed-system quantum Fisher information $F_Q = (2\pi\gamma_e\tau)^2$ to $\sim 5 \times 10^{-5}$. Under Markovian dephasing it tracks $F = (2\pi\gamma_e\tau)^2 e^{-2\tau/T_2^\star}$ to $\sim 3 \times 10^{-3}$. The symmetric-logarithmic-derivative quantum Fisher information upper-bounds the projective classical Fisher information ($F_C \leq F_Q$), as required, and the readout contrast decays as $e^{-\tau/T_2^\star}$.

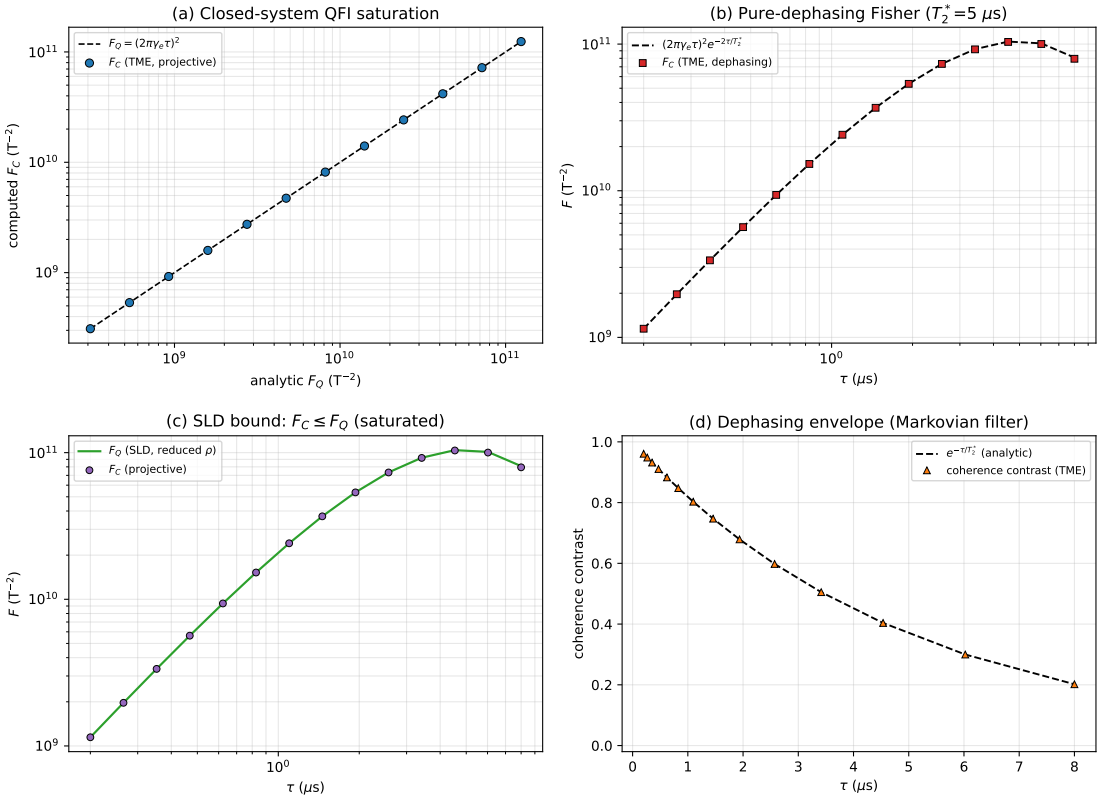


Figure S3: Validation of the tangent-master-equation Fisher pipeline against closed forms. **(a)** With dissipation switched off, the propagated classical Fisher information saturates the closed-system quantum Fisher information $F_Q = (2\pi\gamma_e\tau)^2$ (agreement $\sim 5 \times 10^{-5}$). **(b)** Under Markovian dephasing it tracks $F = (2\pi\gamma_e\tau)^2 e^{-2\tau/T_2^\star}$ (agreement $\sim 3 \times 10^{-3}$). **(c)** The symmetric-logarithmic-derivative quantum Fisher information (reduced $\rho$) upper-bounds the projective classical Fisher information, $F_C \leq F_Q$. **(d)** The readout coherence contrast follows the Markovian envelope $e^{-\tau/T_2^\star}$.

**Numerical stability of the derivative.** The $< 10^{-4}$ agreement above is attained at a well-chosen step. A finite-difference derivative is not unconditionally accurate, because it must balance truncation error (which grows as $(\Delta b)^2$ for a central difference) against floating-point cancellation of two nearly equal solves (which grows as $\varepsilon/\Delta b$ as $\Delta b \to 0$). Sweeping the step $\Delta b$ on the validated three-level NV spin model with a $T_1$ channel (the setup of the derivative test above), the relative error of the central-difference $\partial_b \langle \sigma_x \rangle$ is $U$-shaped (Fig. S4): it reaches a minimum of only $\sim 1 \times 10^{-12}$ in a narrow window near $(\Delta b)^\star \approx 5 \times 10^{-11}$ T (a fractional step $\sim 5 \times 10^{-6}$), and degrades to $\sim 5 \times 10^{-3}$ at the largest step (truncation-limited) and floors near $\sim 10^{-7}$ at the smallest step

(cancellation-limited). The best attainable finite-difference accuracy therefore sits three to four orders of magnitude above machine precision, and the optimal step is problem-dependent, so a finite-difference budget must retune $\Delta b$ at every operating point. The tangent solve carries no such step-size error: it integrates the exact variational equation alongside $\rho$, returning $\partial_b \rho$ to the integrator roundoff floor ($\sim \varepsilon = 2.2 \times 10^{-16}$) at every operating point without tuning. This is why the framework propagates the tangent state rather than perturbing the field and re-solving.

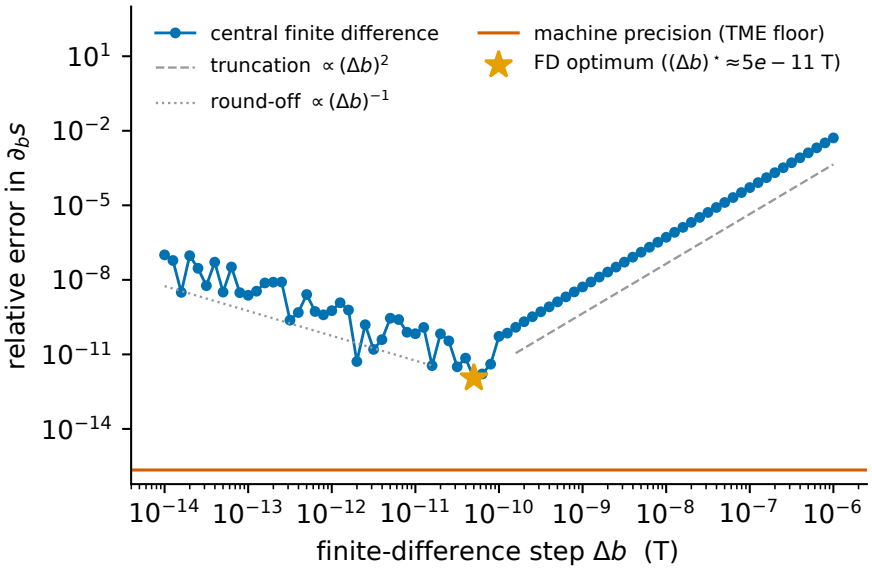


Figure S4: Numerical instability of the finite-difference field-derivative. For the validated three-level NV spin model with a $T_1$ channel, the relative error of a central finite-difference $\partial_b \langle \sigma_x \rangle$ versus step size $\Delta b$ is $U$-shaped: truncation error $\propto (\Delta b)^2$ at large $\Delta b$ (dashed guide) and floating-point cancellation $\propto (\Delta b)^{-1}$ at small $\Delta b$ (dotted guide). The optimum ($(\Delta b)^\star \approx 5 \times 10^{-11}$ T, star) reaches only $\sim 10^{-12}$, three to four orders of magnitude above machine precision. The tangent master equation returns the same derivative exactly, to the integrator roundoff floor (orange line), with no step-size dependence.

### 2.3 Benchmark table

Table S4 collects the quantitative validation benchmarks. Each row compares a digital-twin output to an independent analytic or numerical reference and corresponds to a reproducible example in the code distribution, so every headline quantity used in the main text, from the thermal apparent-field bias and the readout contrast to the Cramér–Rao sensitivity and the GPU/CPU parity, is traceable to a documented check at the stated tolerance.

Table S4: Validation benchmarks. Each check compares a digital-twin output to an independent reference.

| **Check** | **Reference** | **Agreement** |
|---|---|---|
| Tangent derivative $\partial_b \rho$ | central finite difference | $< 10^{-4}$ |
| Finite-difference $\partial_b s$ optimum | central-difference step sweep | $\sim 10^{-12}$ at $h^\star \sim 5 \times 10^{-11}$ T |
| Fisher information (no dissipation) | closed-system QFI $(2\pi\gamma_e\tau)^2$ | $5 \times 10^{-5}$ |
| Fisher information (dephasing) | $(2\pi\gamma_e\tau)^2 e^{-2\tau/T_2^\star}$ | $3 \times 10^{-3}$ |
| Sensitivity versus time grid | grid-independence | $1.000\times$ |
| Thermal apparent field | $-\Delta D_{\mathrm{gs}}/\gamma_e$ | $< 0.02\%$ |
| Synthesized $1/f$ spectral slope | target $-1$ | $-1.00$ |
| Cs Breit–Rabi spectrum | analytic Breit–Rabi | $\sim 2 \times 10^{-6}$ |
| Contrast and sensitivity | Barry *et al.* ODMR relation | within reported scatter |
| GPU versus CPU (TME, $\eta_B$) | bitwise-parity target | $\leq 10^{-6}$ |

## Supplementary Note 3: Numerical Implementations and Benchmarks

Having established in Secs. 1.3–1.4 how to propagate tangent states $G_b(t)$ and compute Fisher information, we now convert these results into sensitivity benchmarks. For a single sensing-and-readout cycle of duration $\tau_{\mathrm{shot}}$ (including initialization, control, and readout) and additional dead time $\tau_{\mathrm{dead}}$, the total number of repetitions in an experiment of duration $T_{\mathrm{tot}}$ is

$$\nu = \frac{T_{\mathrm{tot}}}{\tau_{\mathrm{shot}} + \tau_{\mathrm{dead}}}.$$

The total Fisher information is $F_{\mathrm{tot}}(b) = \nu F_1(b)$, and the Cramér–Rao bound implies where $\hat{b}$ is the estimate of $b$

$$\mathrm{Var}(\hat{b}) \;\geq\; \frac{1}{F_{\mathrm{tot}}(b)}.$$

We define the field sensitivity referred to $\mathrm{Hz}^{-1/2}$ as

$$\eta_B \equiv \sqrt{\mathrm{Var}(\hat{b})\, T_{\mathrm{tot}}} = \sqrt{\frac{\tau_{\mathrm{shot}} + \tau_{\mathrm{dead}}}{F_1(b)}}. \tag{S18}$$

Equation (S18) is the metric reported in the main text. It is determined entirely by the choice of sensing sequence $s(t)$ (Ramsey, echo, etc.), the $\mathrm{NV}^-$ Hamiltonian and dissipation parameters listed in Tables S1–S3, and the statistical detection model outlined in Sec. 1.4. In practice, we simulate $\rho_I(t)$ and $G_b(t)$, compute $p_0(b)$ and $\partial_b p_0(b)$, assemble $F_1(b)$ via Eq. (S17), and finally report $\eta_B$.

### 3.1 Temperature-Dependent Phonon Processes

Phonon-mediated orbital hopping between the excited-state branches $E_x$ and $E_y$ sets the temperature dependence of the optical contrast and the excited-state relaxation. We follow the derivation of Ernst *et al.* In this framework, we retain all the terms across the strain and temperature range of interest rather than taking a high- or low-temperature limit. The electron–phonon interaction

couples the orbital doublet to $E$-symmetric acoustic phonons through the dynamic Jahn–Teller effect, and the hopping rate between the two branches, split by $\hbar\Delta_\perp$, follows from Fermi's golden rule taken to second order in the coupling.

**Phonon spectral density.** Assuming a linear acoustic dispersion and phonon wavelengths large compared with the lattice spacing, the Debye model gives a mode density $\rho(\epsilon) \propto \epsilon^2$ and a per-mode coupling $\lambda \propto \sqrt{\epsilon}$, so the polarization-independent phonon spectral density is (in units of $\mu\mathrm{s}^{-1}$)

$$J(\epsilon) = \eta\, \epsilon^3, \tag{S19}$$

with $\eta$ the effective $E$-phonon coupling strength and $\hbar\Omega$ the Debye cutoff energy. The Bose–Einstein occupation is $n(\epsilon, T) = [\exp(\epsilon/k_B T) - 1]^{-1}$.

**One-phonon process.** At first order the upward hop $E_y \to E_x$ absorbs a single phonon of energy $\hbar\Delta_\perp$, so evaluating $J$ at the splitting gives

$$k_{\uparrow,1}(T) = 4\,\eta\,(\hbar\Delta_\perp)^3\, n(\hbar\Delta_\perp, T), \tag{S20}$$

and the downward (emission) rate follows from detailed balance,

$$k_{\downarrow,1}(T) = 4\,\eta\,(\hbar\Delta_\perp)^3 \left[\, n(\hbar\Delta_\perp, T) + 1 \,\right], \qquad \frac{k_\uparrow}{k_\downarrow} = e^{-\hbar\Delta_\perp/k_B T}. \tag{S21}$$

Writing the splitting as $\hbar\Delta_\perp \approx 2\delta_\perp h$ with $\delta_\perp$ the transverse-strain half-splitting, the low-temperature rate scales as $32\,\eta\, h^3 \delta_\perp^3\, n(2\delta_\perp h, T)$. In the high-temperature limit $k_B T \gg \hbar\Delta_\perp$ it reduces to a term linear in temperature, $k_1 \approx 16\,\eta\, h^2 \delta_\perp^2\, k_B T$.

**Two-phonon Raman process.** At second order the dominant contribution is a Raman process in which one phonon is absorbed and one emitted. Using $J(\epsilon) = \eta\epsilon^3$ with the reduced variables $x = \epsilon/k_B T$ and $x_\perp = \hbar\Delta_\perp/k_B T$, the downward Raman rate is

$$k_{\downarrow,2}(T) = \frac{64\,\hbar}{\pi}\,\eta^2\, k_B^5\, T^5\, I(T, \delta_\perp), \qquad I(T, \delta_\perp) = \int_{x_\perp}^{\hbar\Omega/k_B T} \frac{e^x\, x\,(x - x_\perp)\left[\, x^2 + (x - x_\perp)^2 \,\right]}{2\,(e^x - 1)\,(e^{x - x_\perp} - 1)}\, dx, \tag{S22}$$

with the upward rate again fixed by detailed balance. The two-phonon *emission* and *absorption* channels, in which both phonons move in the same direction, contribute below 1% for the strains ($\delta_\perp < 100$ GHz) and temperatures ($T > 10$ K) of interest and are neglected. The retained Raman rate scales as $T^5$ at low reduced temperature.

**Implementation in the digital twin.** The total downward hopping rate is $k_\downarrow = k_{\downarrow,1} + k_{\downarrow,2}$, with $k_\uparrow = k_\downarrow\, e^{-\hbar\Delta_\perp/k_B T}$. The twin evaluates the *full* rates: the one-phonon term of Eq. (S20) in closed form and the two-phonon Raman term of Eq. (S22) up to the cutoff $\hbar\Omega = 0.168$ eV, at the operating temperature and orbital splitting. It does not use a fitted polynomial. The rates are inserted into the Lindblad equation (Sec. 1.2) as collapse operators coupling $E_x \leftrightarrow E_y$, with a rate cap imposed for numerical stability at extreme strain, using the coupling $\eta = 176\ \mu\mathrm{s}^{-1}\mathrm{meV}^{-3}$ of the derivation cited above.

### 3.2 Optical readout: contrast and photon budget

The photon-counting Fisher information of Sec. 4.1 in the main text enters through two readout quantities: the effective contrast $C_{\text{eff}}$, which scales the spin-dependent fringe amplitude, and the detected-photon budget $N_{\text{det}}$. Both are computed from the open-system model rather than assumed. The optical-cycle contrast $C_{\text{optical}}$ is the normalized difference in time-integrated fluorescence between the $m_s$=0 and $m_s$=±1 spin projections, obtained by propagating the ten-level readout (Sec. 1.2) through a full initialization-and-readout cycle. It falls with temperature $T$ (phonon-activated excited-state mixing erodes spin contrast) and with transverse strain $E_\perp$ (excited-state orbital mixing). Because a per-point optical-cycle solve is expensive, we tabulate the spin-projection-averaged contrast $\bar{C}(T, E_\perp)$ on a temperature and strain grid (Supplementary Fig. S5), and interpolate it within the smooth operating region, which replaces a per-point optical-cycle solve with a table lookup at a large speed-up. The contrast that dilutes the measured signal is $C_{\text{eff}} = C_{\text{optical}}\, f_{\text{NV}^-}$, with $f_{\text{NV}^-}$ the measured NV$^-$ charge-state fraction (only NV$^-$ carries spin contrast, while NV$^0$ adds unmodulated background, Fig. 4a), extracted from photoluminescence spectra as described in Sec. 8. The detected-photon budget is $N_{\text{det}} = E_{\text{conv}}\,[N]\,V\,n_{\text{avg}}$, from the conversion efficiency $E_{\text{conv}}$, the nitrogen concentration $[N]$, and the sensing volume $V$ (so that $E_{\text{conv}}\,[N]\,V$ is the number of NV$^-$ emitters), and $n_{\text{avg}}$ the mean detected photons per NV$^-$ per readout window, which folds in the optical collection efficiency. Its fixed prefactors are representative assumed values (Sec. 6.1).

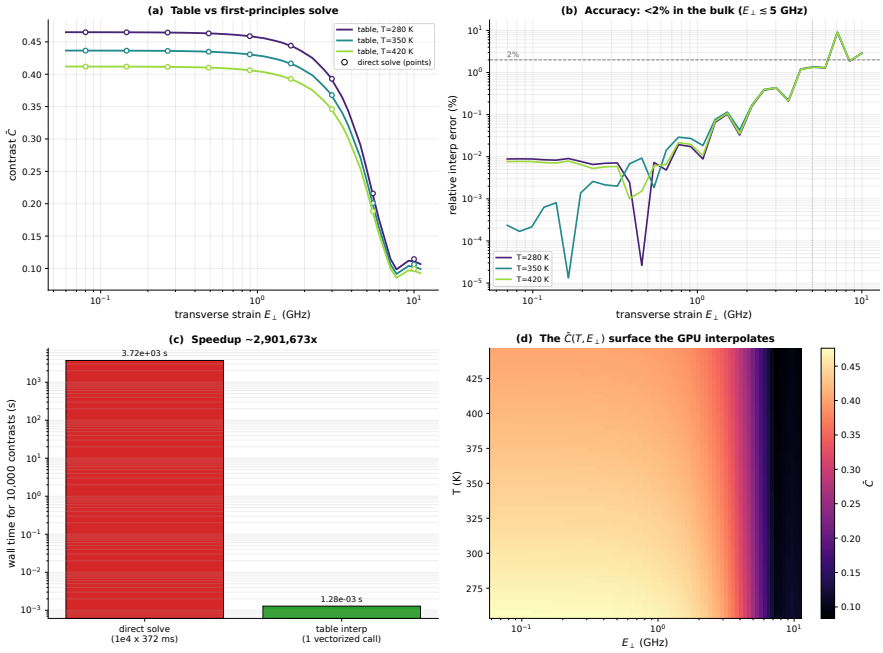


Figure S5: First-principles readout-contrast table. **(a)** The spin-projection-averaged optical-cycle contrast $\bar{C}(T, E_\perp)$ from the ten-level readout (lines) against direct per-point solves (open circles) versus transverse strain $E_\perp$ at three temperatures. **(b)** The relative interpolation error stays below 2% throughout the operating region ($E_\perp \lesssim 5\,\text{GHz}$). **(c)** Replacing the per-point optical-cycle solve with a table lookup is a $\sim 10^6\times$ speed-up over $10^4$ direct solves. **(d)** The full $\bar{C}(T, E_\perp)$ surface the GPU interpolates. The table supplies the effective contrast $C_{\text{eff}}$ used in the sensitivity and accuracy budgets.

# Supplementary Note 4: Multi-metric definitions and per-mechanism attribution

The digital twin evaluates three metrics from a *single* open-system solve: the sensitivity $\eta_B$, the accuracy (systematic bias and root-mean-square error), and the robustness $R_{\mathrm{lin}}$. The solve is parameterized by the operating-parameter vector $\theta = (\theta_{\mathrm{sensor}}, \theta_{\mathrm{control}}, \theta_{\mathrm{env}})$, collecting the material, control, and environment inputs of Eq. (S1). All three metrics are functionals of the same field-dependent density matrix $\rho(b, t; \theta)$ and its tangent $G_b(t; \theta) = \partial_b \rho(t; \theta)$ (Sec. 1.3): the sensitivity and accuracy are read at the fixed nominal $\theta_0$, and the robustness measures their stability as $\theta$ drifts about $\theta_0$. They respond to distinct features of the same solve. This section (i) defines the three metrics, (ii) proves that sensitivity and accuracy are governed by *different* quantities of the same solve through the unbiased and biased Cramér–Rao bounds, and (iii) derives the Shapley attribution that assigns each error mechanism a unique, fair share of every metric.

Throughout, $b$ denotes the sensed field, $p(b) = \mathrm{Tr}[M_0\, \rho(b)]$ the readout probability of the measurement projector $M_0 = |0_g\rangle\langle 0_g|$, and $p'(b) \equiv \partial_b p(b) = \mathrm{Tr}[M_0\, G_b]$ its tangent (Sec. 1.3). We write $C \in (0, 1]$ for the optical readout contrast and $N_{\mathrm{det}}$ for the number of detected photons per shot. For the ground-state Ramsey readout used throughout, this probability takes the fringe form

$$p(b) \;=\; \tfrac{1}{2}\Big[1 + V(t)\, \cos\varphi(b)\Big], \qquad V(t) = e^{-t/T_2^\star}, \quad \varphi(b) = \gamma_e\, b\, \tau + \varphi_0, \tag{S23}$$

with visibility $V$ set by coherence and phase $\varphi$ set by the accumulated Zeeman rotation ($\gamma_e = \mu_B g_g / h$ and $\varphi_0$ the calibration phase). Pure dephasing scales $V$ without shifting $\varphi$, the thermal shift displaces $\varphi$, and optical leakage lowers $V$ and adds a readout-baseline offset. The attribution below exploits that these mechanisms act on different factors of the same fringe.

## 4.1 The three metrics as functionals of one solve

**Sensitivity.** As derived in Sec. 1.4, the single-shot classical Fisher information for photon counting is

$$F_1(b) = \frac{N_{\mathrm{det}}\, \big(p'(b)\big)^2}{p(b)\, \big[1 - p(b)\big]} \qquad \text{(binomial readout)}, \tag{S24}$$

which reduces to $F_1 = N_{\mathrm{det}}(p')^2/p$ in the Poisson limit. The field sensitivity referred to unit bandwidth is

$$\eta_B(b) = \sqrt{\frac{\tau_{\mathrm{shot}} + \tau_{\mathrm{dead}}}{F_1(b)}}, \tag{S25}$$

the Cramér–Rao-limited field resolution per $\sqrt{\mathrm{Hz}}$. The *ideal-device* sensitivity $\eta_B^{\mathrm{sn}}$ is Eq. (S25) evaluated with all imperfections switched off. It is the shot-noise floor and serves as the baseline against which sensitivity is attributed below.

**Accuracy.** Accuracy quantifies how faithfully a field *estimate* $\hat{b}$ recovers the true $b$. The estimate is formed by inverting a calibration generated at a reference operating point. Let $\big(p_{\mathrm{ref}}, p'_{\mathrm{ref}}\big)$ be the readout probability and slope of the ideal device (all mechanisms off), and let the optical fringe seen by the detector be the contrast-diluted $p_{\mathrm{opt}} = \frac{1}{2} + C\,(p - \frac{1}{2})$. A single shot draws $k \sim \mathrm{Poisson}(N_{\mathrm{det}}\, p_{\mathrm{opt}})$. Inverting the contrast gives the probability estimate $\hat{p} = (k/N_{\mathrm{det}} - \frac{1}{2})/C + \frac{1}{2}$, and the linear calibration inverse yields the field estimate

$$\hat{b} \;=\; b_{\mathrm{ref}} + \frac{\hat{p} - p_{\mathrm{ref}}}{p'_{\mathrm{ref}}}. \tag{S26}$$

Taking the expectation over the shot noise ($\mathbb{E}[\hat{p}] = p$) gives the *systematic bias*

$$\beta(b) \;\equiv\; \mathbb{E}[\hat{b}] - b \;=\; \frac{p(b) - p_{\mathrm{ref}}}{p'_{\mathrm{ref}}} - \big(b - b_{\mathrm{ref}}\big), \tag{S27}$$

the displacement of the true fringe from the reference, referred to field units through the reference slope. Crucially, $\beta$ is a *deterministic* functional of $\rho(b)$: it involves no sampling. The statistical part is the shot-noise precision floor

$$\sigma_b^2 \;=\; \mathrm{Var}(\hat{b}) \;=\; \frac{\mathrm{Var}(\hat{p})}{(p'_{\mathrm{ref}})^2} \;=\; \frac{p_{\mathrm{opt}}\,(1 - p_{\mathrm{opt}})}{N_{\mathrm{det}}\,C^2\,(p'_{\mathrm{ref}})^2}, \tag{S28}$$

and the reported accuracy is the root-mean-square error

$$\mathrm{RMSE}(b) \;=\; \sqrt{\beta(b)^2 + \sigma_b^2}. \tag{S29}$$

The decomposition of Eq. (S29) into a deterministic bias and a shot-noise floor is exact and is the reason accuracy and sensitivity can disagree: $\sigma_b$ is (like $\eta_B$) set by the Fisher information, whereas $\beta$ is not.

**Robustness.** Robustness measures the stability of $\eta_B$ under slow drift of the operating parameters $\theta = (\theta_1, \ldots, \theta_n)$ about their nominal values $\theta_0$. We define the primitive robustness as the ratio of the root-mean-square sensitivity under drift to the nominal sensitivity,

$$R \;\equiv\; \frac{\sqrt{\langle \eta_B^2(\theta) \rangle}}{\eta_B(\theta_0)}, \qquad \theta = \theta_0 + \delta\theta, \tag{S30}$$

where the drifts are independent and zero-mean, $\langle \delta\theta_i \rangle = 0$ and $\langle \delta\theta_i \delta\theta_j \rangle = \sigma_i^2 \delta_{ij}$, with $\sigma_i$ the root-mean-square drift budget of parameter $i$ and $\langle \cdot \rangle$ the average over the drift distribution. $R = 1$ is a perfectly stable operating point. $R > 1$ is the fractional inflation of the effective sensitivity by drift.

To leading order in the drift we linearize the sensitivity,

$$\eta_B(\theta) = \eta_0 + \sum_i (\partial_{\theta_i} \eta_B)\, \delta\theta_i + O(\delta\theta^2), \qquad \eta_0 \equiv \eta_B(\theta_0), \tag{S31}$$

and insert this into Eq. (S30). Squaring, $\eta_B^2 = \eta_0^2 + 2\eta_0 \sum_i (\partial_{\theta_i} \eta_B)\, \delta\theta_i + \sum_{i,j} (\partial_{\theta_i} \eta_B)(\partial_{\theta_j} \eta_B)\, \delta\theta_i \delta\theta_j + O(\delta\theta^3)$; the linear term averages to zero and the quadratic term is diagonal ($\langle \delta\theta_i \delta\theta_j \rangle = \sigma_i^2 \delta_{ij}$), leaving

$$\big\langle \eta_B^2 \big\rangle = \eta_0^2 + \sum_i (\partial_{\theta_i} \eta_B)^2\, \sigma_i^2 + O(\sigma^3). \tag{S32}$$

Dividing by $\eta_0^2$ and using $\partial_{\theta_i} \eta_B / \eta_0 = \partial_{\theta_i} \log \eta_B$ gives the linearized robustness index,

$$\boxed{R_{\mathrm{lin}}^2 \;=\; \frac{\big\langle \eta_B^2 \big\rangle}{\eta_B^2(\theta_0)}\bigg|_{\mathrm{lin}} \;=\; 1 + \sum_i \sigma_i^2 \left(\partial_{\theta_i} \log \eta_B\right)^2.} \tag{S33}$$

Equivalently $R_{\mathrm{lin}}^2 - 1 = \langle (\delta\eta_B / \eta_0)^2 \rangle$ is the relative variance of the sensitivity in the linear-response regime, so $R_{\mathrm{lin}} = 1$ denotes a drift-immune operating point and $R_{\mathrm{lin}} = 1.1$ an environment that

inflates the effective sensitivity by 10%. The log-gradients are evaluated by centred finite differences, $\partial_{\theta_i} \log \eta_B \approx [\log \eta_B(\theta_i{+}h_i) - \log \eta_B(\theta_i{-}h_i)]/(2h_i)$, and each parameter's normalized share is

$$c_i \;=\; \frac{\sigma_i^2\,(\partial_{\theta_i} \log \eta_B)^2}{R_{\mathrm{lin}}^2 - 1}, \qquad \textstyle\sum_i c_i = 1, \tag{S34}$$

additive by construction. This form assumes the parameter drifts are uncorrelated, $\langle \delta\theta_i \delta\theta_j \rangle = \sigma_i^2 \delta_{ij}$. Correlated drifts (for example temperature and optical leakage through radio-frequency heating of the modulator) would add cross terms $\sigma_{ij}\, \partial_{\theta_i} \log \eta_B\, \partial_{\theta_j} \log \eta_B$ and make the robustness budget itself non-additive. Retaining the second-order term in Eq. (S31) adds a curvature correction $\sum_i (\partial_{\theta_i}^2 \eta_B/\eta_0)\, \sigma_i^2$ to $\langle \eta_B^2 \rangle/\eta_0^2$ that $R_{\mathrm{lin}}$ omits. When it is not negligible (large drift, or an $\eta_B$ strongly nonlinear in $\theta$) we evaluate the primitive definition Eq. (S30) directly by Monte Carlo, sampling $\delta\theta \sim \mathcal{N}(0, \mathrm{diag}(\sigma_i^2))$, recomputing $\eta_B$ per draw, and forming $R_{\mathrm{env}} = \sqrt{\langle \eta_B^2 \rangle}/\eta_B(\theta_0)$. One has $R_{\mathrm{env}} \to R_{\mathrm{lin}}$ in the small-drift limit $\sigma_i\, \partial_{\theta_i} \log \eta_B \ll 1$. The drift budget used throughout is $\{\, \sigma_b{=}1\,\mathrm{nT},\ \sigma_T{=}0.1\,\mathrm{K},\ \sigma_\varepsilon{=}10^{-5},\ \sigma_{T_2^\star}{=}2\%,\ \sigma_{T_1}{=}10\% \,\}$ (background field, temperature, optical leakage, coherence, and relaxation).

### 4.2 Unbiased and biased Cramér–Rao bounds: precision versus accuracy

The mean-squared error of any field estimator decomposes exactly as

$$\mathrm{MSE}(\hat{b}) \;=\; \mathbb{E}\big[(\hat{b} - b)^2\big] \;=\; \underbrace{\big(\mathbb{E}[\hat{b}] - b\big)^2}_{\beta(b)^2} \;+\; \underbrace{\mathrm{Var}(\hat{b})}_{\text{precision}}. \tag{S35}$$

For an *unbiased* estimator ($\beta \equiv 0$), the Cramér–Rao bound gives $\mathrm{Var}(\hat{b}) \geq 1/F_{\mathrm{tot}}(b)$ with $F_{\mathrm{tot}} = \nu F_1$ and $\nu = T_{\mathrm{tot}}/(\tau_{\mathrm{shot}} + \tau_{\mathrm{dead}})$. The sensitivity Eq. (S25) is precisely the bandwidth-referred form of this bound. For a *biased* estimator with bias function $\beta(b)$, the bound generalizes to

$$\mathrm{Var}(\hat{b}) \;\geq\; \frac{\big[1 + \beta'(b)\big]^2}{F_{\mathrm{tot}}(b)}, \qquad \mathrm{MSE}(\hat{b}) \;\geq\; \beta(b)^2 + \frac{\big[1 + \beta'(b)\big]^2}{F_{\mathrm{tot}}(b)}, \tag{S36}$$

where $\beta' = \partial_b \beta$. Equation (S36) makes the metric split explicit: the achievable *precision* is set by the Fisher information $F_1$ (hence by $\eta_B$), whereas the *accuracy* floor is set by $\beta$, an independent functional of the same solve. A device optimized to minimize $\eta_B$ (maximize $F_1$) therefore carries no guarantee on $\beta$. The two are bounded by different quantities and, as the next subsection shows, are limited by different physical mechanisms.

### 4.3 Per-mechanism attribution by Shapley values

Let $\mathcal{M} = \{1, \ldots, M\}$ index the error mechanisms. Each admits an idealized "off" operation: optical leakage $\varepsilon \to 0$, relaxation $T_1 \to \infty$, dephasing $\gamma_\phi \to 0$, and the thermal channel is switched off by recalibrating the microwave carrier to the operating $D_g(T)$. That recalibration removes the apparent-field shift a temperature offset $\delta T$ would otherwise impose: with the carrier fixed at $D_g(T_{\mathrm{cal}})$, the offset $\partial_T D_g\, \delta T$ is indistinguishable from a Zeeman shift, i.e. an apparent field

$$\beta_{\mathrm{th}} \;=\; \frac{|\partial_T D_g|\, \delta T}{\gamma_e} \;=\; \frac{|\Delta D_{\mathrm{gs}}|}{\gamma_e} \;=\; 282\ \mathrm{nT} \quad (\delta T = 0.1\ \mathrm{K}), \tag{S37}$$

using $\partial_T D_g = -79\,\mathrm{kHz/K}$ and $\gamma_e = 28.0\,\mathrm{GHz/T}$. The thermal "off" is thus a recalibration rather than the removal of a physical channel, so the thermal contribution to the accuracy budget is

equivalently the bias incurred by failing to recalibrate the carrier as $D_g(T)$ drifts. Because both the accumulated phase and the field inferred from it scale with $\tau$, $\beta_{\mathrm{th}}$ carries no interrogation-time dependence of its own. For a coalition $S \subseteq \mathcal{M}$, define the *value function*

$$v : 2^{\mathcal{M}} \to \mathbb{R}, \qquad v(S) = \big[\text{metric evaluated with exactly the mechanisms in } S \text{ active}\big], \tag{S38}$$

so $v(\varnothing)$ is the ideal device and $v(\mathcal{M})$ the full one. For sensitivity we take $v(S) = \eta_B(S)$, decomposing the reported operational metric itself; because $\eta_B \propto \mathcal{F}^{-1/2}$ is a nonlinear transform of the Fisher information, decomposing $\mathcal{F}$ instead would give different shares, and we attribute $\eta_B$ as the quantity with direct sensing meaning. For accuracy $v(S) = \beta(S)$ of Eq. (S27) with $v(\varnothing) = 0$; this holds exactly rather than by assumption because the linear calibration is referenced to the all-off device at the operating field ($b_{\mathrm{ref}} = b_{\mathrm{true}}$), so the fringe inversion is exact there and the residual cosine curvature is $O((\gamma_e b\tau)^2) \approx 10^{-10}$ at $b = 1\,\mathrm{nT}$, negligible against the nanotesla-scale attributed biases. The attribution problem is to split the total degradation $v(\mathcal{M}) - v(\varnothing)$ among the mechanisms.

**Well-posedness of the value function.** The Shapley value requires $v$ to be a function of the *set* $S$, not of the order in which mechanisms are toggled. This is manifest from the definition, Eq. (S38): each mechanism $i$ carries a nominal value $\theta_i$ and an off value $\theta_i^{\mathrm{off}}$ (for $\varepsilon$ and $T_1$ the removed dissipative rate, for $\gamma_\phi$ the removed dephasing, and for the thermal channel the microwave carrier recalibrated to $D_g(T)$), and $v(S)$ is the metric evaluated at the single parameter vector in which mechanisms in $S$ take their nominal values and those outside take their off values. That parameter vector is a function of the set $S$ alone, with no dependence on ordering, so $v$ is well defined on $2^{\mathcal{M}}$ and Eq. (S39) is unambiguous. Non-additivity of this well-defined $v$ is precisely the mechanism interaction that the Shapley partition distributes and that the pairwise index quantifies below (Supplementary Table S6). It does not enter through the readout contrast. Because the thermal "off" is a carrier recalibration at fixed temperature, $T$ and hence $C(T)$ are identical across all 16 coalitions, so $C(T)$ is a constant of the game and not a source of interaction. For accuracy the interaction arises instead from the nonlinear fringe inversion that defines $\beta$. Dephasing lowers the fringe visibility, and inverting a lower-visibility fringe recovers the thermal apparent-field offset nonlinearly, as detailed in the leave-one-out analysis below.

**Axioms and uniqueness.** A fair attribution $\phi : v \mapsto (\phi_1, \ldots, \phi_M)$ should satisfy: *(i) efficiency*, $\sum_j \phi_j = v(\mathcal{M}) - v(\varnothing)$; *(ii) symmetry*, if $v(S\cup\{i\}) = v(S\cup\{j\})$ for every $S \not\ni i, j$ then $\phi_i = \phi_j$; *(iii) null player*, if $v(S\cup\{j\}) = v(S)$ for every $S$ then $\phi_j = 0$; and *(iv) linearity*, $\phi_j(v+w) = \phi_j(v) + \phi_j(w)$. Shapley's theorem states that a *unique* map obeys all four,

$$\boxed{\phi_j = \sum_{S \subseteq \mathcal{M}\setminus\{j\}} \frac{|S|!\,(M - |S| - 1)!}{M!} \big[v(S \cup \{j\}) - v(S)\big].} \tag{S39}$$

That a *unique* map obeys all four axioms, namely Eq. (S39), is Shapley's theorem. Uniqueness follows by expanding any coalition game in the basis of unanimity games $u_T(S) = \mathbb{1}[T \subseteq S]$, fixing $\phi$ on each through the null-player, symmetry, and efficiency axioms, and combining them by linearity to reach an arbitrary game. Existence is the direct check that Eq. (S39) satisfies all four. The equivalent order-averaging form below is the one we evaluate in practice.

Equivalently, writing $\Pi_M$ for the $M!$ orderings of $\mathcal{M}$ and $P_j^\pi$ for the set of mechanisms preceding $j$ in ordering $\pi$,

$$\phi_j = \frac{1}{M!} \sum_{\pi \in \Pi_M} \big[v(P_j^\pi \cup \{j\}) - v(P_j^\pi)\big], \tag{S40}$$

i.e. $\phi_j$ is the marginal effect of switching mechanism $j$ on, averaged over every order in which the mechanisms could be activated. The weight in Eq. (S39) is the fraction of orderings whose predecessors of $j$ are exactly $S$. For each ordering $\pi$ the sum over players telescopes, $\sum_j \big[v(P_j^\pi \cup \{j\}) - v(P_j^\pi)\big] = v(\mathcal{M}) - v(\varnothing)$, so averaging over $\pi$ yields efficiency, $\sum_j \phi_j = v(\mathcal{M}) - v(\varnothing)$. With the $2^M$ coalitions enumerated exactly, efficiency holds identically for any $v$, so recovering $\big|\sum_j \phi_j - (v(\mathcal{M}) - v(\varnothing))\big| / |v(\mathcal{M}) - v(\varnothing)| \le 2 \times 10^{-16}$ checks the arithmetic of our implementation rather than the soundness of the decomposition. The substantive validation is the interaction-index analysis below. Evaluating Eq. (S39) requires the $2^M$ coalition solves. For the dissipative set plus the thermal channel ($M = 4$) this is 16 open-system solves per operating point, shared between the sensitivity and accuracy value functions.

**Leave-one-out and non-additivity.** The single-difference or "leave-one-out" (LOO) attribution, $\phi_j^{\mathrm{LOO}} = v(\mathcal{M}) - v(\mathcal{M} \setminus \{j\})$, is the special case of Eq. (S39) that keeps only the full-coalition term. It equals the Shapley value iff the marginal contribution $v(S \cup \{j\}) - v(S)$ is independent of $S$, i.e. iff $v$ is additive (no interactions). Deviations are measured by the pairwise Shapley interaction index

$$I_{ij} = \sum_{S \subseteq \mathcal{M} \setminus \{i,j\}} \frac{|S|!\,(M - |S| - 2)!}{(M-1)!} \big[v(S \cup \{i,j\}) - v(S \cup \{i\}) - v(S \cup \{j\}) + v(S)\big]. \tag{S41}$$

For *sensitivity*, the mechanisms act on $\eta_B$ through nearly independent factors (a mechanism that lowers $V$ and one that adds a photon-baseline offset combine multiplicatively near the operating point), so the $I_{ij}$ are small and LOO and Shapley agree to within a few percent. For *accuracy*, the estimate Eq. (S26) inverts the *nonlinear* fringe Eq. (S23): two mechanisms that each reshape the fringe produce a cross term $\partial^2 \hat{b} / \partial\theta_i \partial\theta_j \neq 0$. The dominant coupling is between dephasing and the thermal shift. Dephasing lowers the visibility $V$, and the maximum-likelihood inversion of a lower-contrast fringe recovers the thermal apparent-field offset nonlinearly, so the bias ascribed to the thermal channel depends strongly on whether dephasing is active. Concretely, thermal alone produces the full-visibility apparent field of Eq. (S37), $v(\{\text{thermal}\}) = 282\,\text{nT}$, well above the full-device total $v(\mathcal{M}) - v(\varnothing) = 110\,\text{nT}$. The gap is the interaction. With dephasing active the fringe visibility falls to $V \approx 0.67$, so the same thermal offset inverts to $V \cdot 282 \approx 189\,\text{nT}$ ($v(\{\text{dephasing}, \text{thermal}\}) = 189\,\text{nT}$), and adding leakage brings it to the 110 nT full-device value. The reduction from 282 to 189 nT on adding dephasing is the $-91$ nT dephasing–thermal cross term of Supplementary Table S6 made explicit, and leakage supplies the remaining reduction to 110 nT. The full set of coalition values is listed in Supplementary Table S5.

We evaluate the interaction index Eq. (S41) directly from the 16 coalition solves (Supplementary Table S6). For sensitivity every $|I_{ij}|$ is at most 4% of the total and the leave-one-out residual is $-4\%$: the budget is additive and LOO reproduces Shapley. For accuracy the dephasing–thermal interaction alone is $-91$ nT (83% of the total degradation), the leakage–thermal interaction adds a further 10%, and a single leave-one-out ordering leaves 89% of the bias in a non-additive residual. Only the order-averaged Shapley partition of Eq. (S39) returns a unique, order-independent share, which is why the accuracy budget of the main text *requires* the Shapley construction whereas the sensitivity budget does not. The pairwise indices do not capture the interaction in full: the leave-one-out residual plus the sum of the $I_{ij}$ leaves a higher-order (three- and four-body) remainder of $-2.1$ nT for accuracy, about 2% of the total, and essentially zero for sensitivity, so the non-additivity is dominated by, but not confined to, pairwise terms.

Table S5: The 16 coalition values $v(S)$ at the budget operating point of Fig. 4, from which every quantity in this subsection follows, namely the Shapley shares Eq. (S39), the leave-one-out residuals, the pairwise interaction indices of Supplementary Table S6, and the higher-order remainder. In each coalition $S \subseteq \{\varepsilon, T_1, T_2^*, \text{temp}\}$ the listed mechanisms are active and the rest are toggled off ($\varepsilon$ optical leakage, $T_1$ relaxation, $T_2^*$ dephasing, temp the thermal channel), and both metrics are read from that single solve, sensitivity $v(S) = \eta_B(S)$ and accuracy $v(S) = \beta(S)$. $v(\varnothing)$ is the ideal device and $v(\mathcal{M})$ the full one. Totals and the derived quantities of Supplementary Table S6 use the unrounded solves and can differ from a subtraction of the two-decimal entries above by up to 0.01.

| **coalition** $S$ | $\eta_B(S)$ **(pT/$\sqrt{\text{Hz}}$)** | $\beta(S)$ **(nT)** |
|---|---|---|
| $\varnothing$ | 16.59 | 0.00 |
| $\{\varepsilon\}$ | 17.33 | −70.15 |
| $\{T_1\}$ | 16.59 | −0.57 |
| $\{T_2^*\}$ | 24.74 | −0.33 |
| $\{\text{temp}\}$ | 16.71 | 282.18 |
| $\{\varepsilon, T_1\}$ | 17.33 | −70.71 |
| $\{\varepsilon, T_2^*\}$ | 25.85 | −70.46 |
| $\{\varepsilon, \text{temp}\}$ | 17.45 | 199.48 |
| $\{T_1, T_2^*\}$ | 24.75 | −0.90 |
| $\{T_1, \text{temp}\}$ | 16.71 | 281.57 |
| $\{T_2^*, \text{temp}\}$ | 24.87 | 188.82 |
| $\{\varepsilon, T_1, T_2^*\}$ | 25.86 | −71.03 |
| $\{\varepsilon, T_1, \text{temp}\}$ | 17.45 | 198.87 |
| $\{\varepsilon, T_2^*, \text{temp}\}$ | 25.98 | 110.27 |
| $\{T_1, T_2^*, \text{temp}\}$ | 24.87 | 188.22 |
| $\{\varepsilon, T_1, T_2^*, \text{temp}\}$ | 25.98 | 109.68 |

### 4.4 Uncertainty quantification of the attribution shares

The attribution shares carry two distinct uncertainties, quantified separately. The accuracy shares are the output of a photon-counting estimator, so their statistical uncertainty follows from resampling: holding the $2^M$ coalition solves fixed, we redraw the Poisson optical readout of each coalition (one $N_{\text{det}}$-photon measurement per coalition) over $R$ realizations, propagate each through the Shapley formula Eq. (S39), and report the 2.5 to 97.5 percentile interval of every share; averaging over $K$ repeated measurements narrows these by $\sqrt{K}$. Because the coalitions are redrawn independently rather than with common random numbers, the readout noise does not cancel in the marginal differences $v(S \cup \{j\}) - v(S)$, so the reported intervals are conservative; common random numbers across coalitions would tighten them without shifting the shares. The sensitivity shares are a deterministic Cramér–Rao quantity and the robustness shares a deterministic finite-difference functional, so neither carries shot-noise seed variance. Their uncertainty is parametric, and we report it by sweeping the dominant input (the measured $T_2^\star$ range for sensitivity, the assumed drift budget for robustness). Across the swept ranges the dominant mechanism of each metric is unchanged, so the metric-dependent inversion is insensitive to the precise values of the assumed inputs. For robustness this dominance is quantitative: leakage drift remains the leading contributor

Table S6: Pairwise Shapley interaction index $I_{ij}$ (Eq. (S41)) and the leave-one-out (LOO) residual, evaluated from the $2^4$ coalition solves at the budget operating point. Sensitivity is near-additive ($|I_{ij}|$ and residual $\lesssim 4\%$ of the total). Accuracy is strongly non-additive, dominated by the dephasing–thermal cross term. Percentages are of the total degradation $v(\mathcal{M}) - v(\varnothing)$.

| **Mechanism pair** $(i, j)$ | $I_{ij}$**, sensitivity (pT/**$\sqrt{\text{Hz}}$**)** | $I_{ij}$**, accuracy (nT)** |
|---|---|---|
| dephasing × thermal | 0.00 (0%) | −90.95 (83%) |
| leakage × thermal | 0.00 (0%) | −10.49 (10%) |
| leakage × dephasing | 0.37 (4%) | 2.08 (2%) |
| $T_1$× thermal | 0.00 (0%) | −0.04 (0%) |
| $T_1$× dephasing | 0.00 (0%) | 0.01 (0%) |
| leakage × $T_1$ | 0.00 (0%) | 0.01 (0%) |
| total $v(\mathcal{M}) - v(\varnothing)$ | 9.40 | 109.68 |
| sum of LOO contributions | 9.77 | 12.38 |
| LOO residual | −0.37 (−4%) | 97.30 (89%) |

provided the assumed leakage-drift budget satisfies $\sigma_\varepsilon \gtrsim 1.8 \times 10^{-6}$ (about an 18% extinction-ratio drift, roughly 5.5× tighter than the value adopted here), below which coherence ($T_2^\star$) drift takes over. The adopted $\sigma_\varepsilon = 10^{-5}$ is a representative extinction-ratio-stability assumption, and the ∼96% leakage share of $R_{\text{lin}}$ should be read against it. Because the signed accuracy contributions partially cancel, the net bias can be small or pass through zero, so shares expressed as a percentage of it would diverge. We therefore quote the accuracy contributions as absolute apparent-field offsets rather than as percentages of the net bias.

Whether the accuracy is limited by this bias or by the shot floor depends on the photon budget. At low photon number the falling readout contrast enters the accuracy only through the shot-noise precision $\sigma_b \propto 1/(C_{\text{eff}}\sqrt{N})$. Once enough photons are collected that $\sigma_b$ averages below the thermal $D_{\text{gs}}$ bias, the accuracy saturates at the bias and is flat across the operating range (Supplementary Fig. S6).

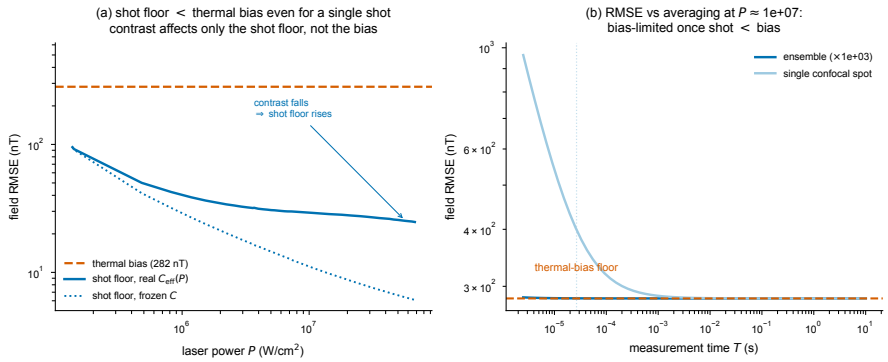


Figure S6: Shot-limited accuracy regime (companion to Fig. 4c). **(a)** Versus laser power, the shot-noise precision floor $\sigma_b \propto 1/(C_{\rm eff}\sqrt{N})$ (solid, with the measured $C_{\rm eff}(P)$; dotted, at frozen contrast) stays below the thermal $D_{\rm gs}$ bias (dashed, 282 nT), so the falling readout contrast enters the accuracy only through the shot floor, not through the systematic bias $\beta$. **(b)** Versus averaging time at fixed power, the field RMSE relaxes onto that thermal-bias floor, so once the floor averages below the bias the accuracy becomes bias-limited and nearly flat across the measured power range.

**Generality of the attribution across the design space.** The metric-dependent attribution is not specific to the single operating point of the main text. Supplementary Fig. S7 maps each metric's magnitude, together with its dominant-limiter boundary, across five two-parameter design planes. Each pixel of each panel is a full attribution. The three metrics and the $2^M$ mechanism coalitions are recomputed at that operating point, the per-mechanism shares are formed as in Sec. 4.3 (Shapley values for sensitivity and accuracy, the additive decomposition of Eq. (S34) for robustness), and the dominant limiter is the mechanism holding the largest share, whose region boundary is drawn in white. Reading down the three metric columns, the outcome is the same across all five planes: sensitivity is limited by dephasing ($T_2^\star$), robustness by optical leakage ($\varepsilon$), and accuracy by the thermal $D_{\rm gs}$ shift (temperature) over most of the space, with optical leakage taking over only in the high-leakage, low-thermal-drift corners. This inversion of the dominant mechanism between metrics persists throughout the design space, so the attribution is a property of the sensing stack and not of a particular design choice.

# Supplementary Note 5: Noise Models

This section specifies the stochastic models used for control and environment noise and how we synthesize one-sided power spectral densities (PSDs) into *time-domain* traces that are inserted into the Hamiltonian and dissipative rates in Sec. 1.1–1.2. In our model, microwave *phase noise* $\delta\phi_{\rm noise}(t)$ perturbs the control phase and contributes via its time derivative to the effective detuning $\Delta_{\rm eff}(t)$ in (S9), microwave *amplitude noise* $a_{\rm noise}(t)$ multiplies the control envelope, and *background magnetic-field noise* $\delta b_{\rm noise}(t)$ adds to the Zeeman term in (S6)–(S10). Laser/AOM noise enters as fluctuations of optical *rates* in the collapse operators (Sec. 1.2).

## 5.1 Noise sources

The model admits four stochastic channels, each specified by a one-sided power spectral density (PSD) and its insertion point, together with photon shot noise (a Poisson readout process, Sec. 1.4).

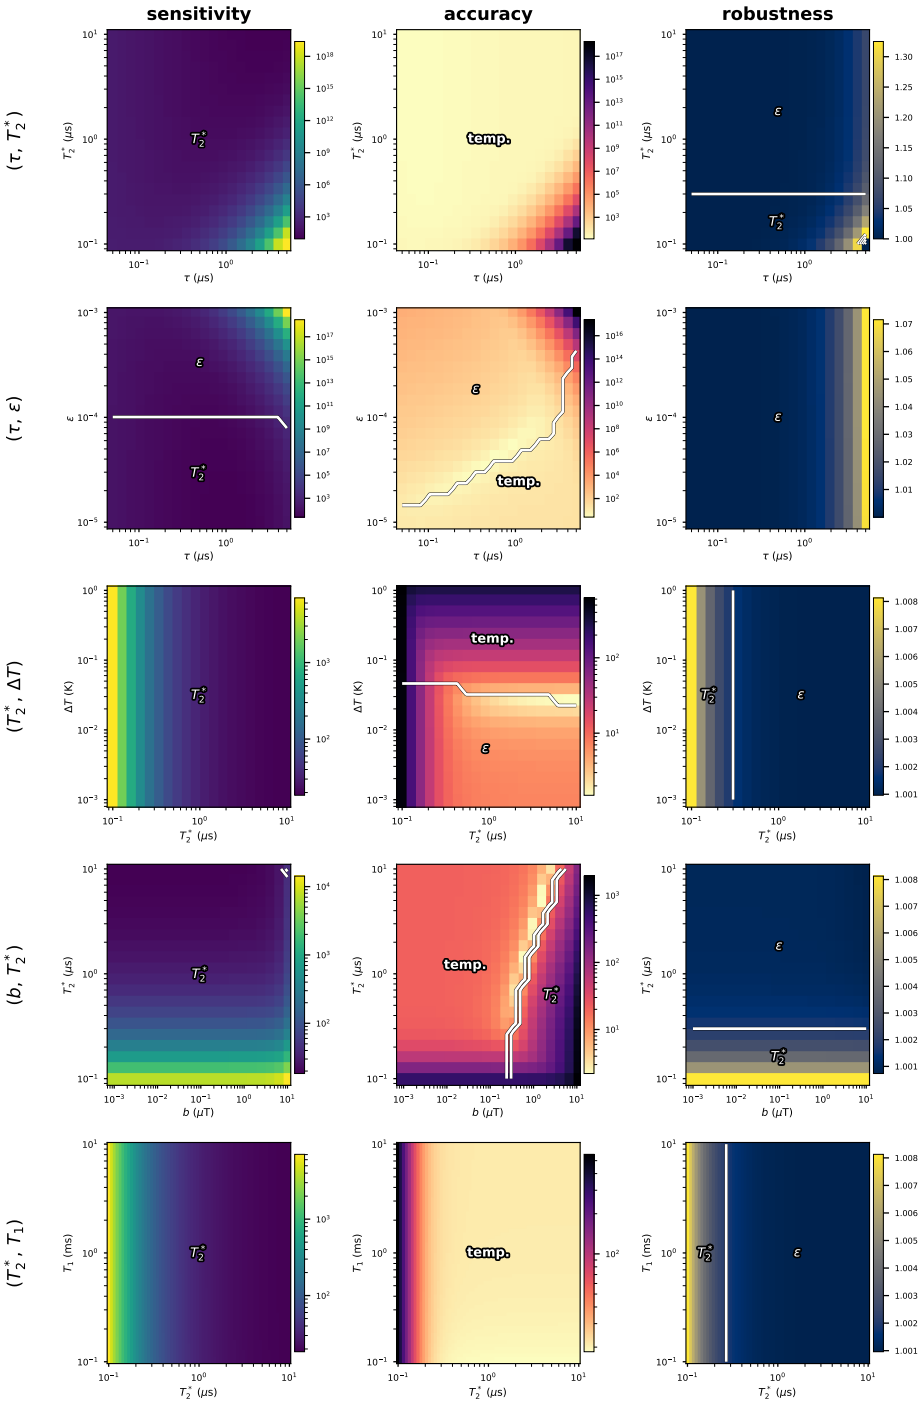


Figure S7: Generality of the metric-dependent attribution across the design space. Rows are the two-parameter design planes $(\tau, T_2^\star)$, $(\tau, \varepsilon)$, $(T_2^\star, \Delta T)$, $(b, T_2^\star)$, and $(T_2^\star, T_1)$. Columns are the three metrics (sensitivity $\eta_B$, accuracy NRMSE, robustness $R_{\mathrm{lin}}$), with color the metric magnitude. White curves are the dominant-limiter boundaries, labeled by the limiting mechanism ($T_2^\star$ = dephasing, $\varepsilon$ = optical leakage, temp. = thermal $D_{\mathrm{gs}}$ shift). Across every plane, sensitivity is limited by dephasing, robustness by optical leakage, and accuracy by the thermal shift, so the metric-dependent inversion holds throughout the design space.

Of these, the field, phase, and amplitude channels are realized as synthesized time-domain traces by the procedure of Sec. 5.2, and the optical-rate channel enters through the collapse-operator rates rather than as a synthesized trace.

**Microwave phase noise.** A stochastic phase on the control drive, $\phi(t) \to \phi(t) + \delta\phi_{\text{noise}}(t)$, with

$$S_\phi(f) = \frac{h_{-3}}{f^3} + \frac{h_{-2}}{f^2} + \frac{h_{-1}}{f} + h_0 + h_{+1} f \quad [\text{rad}^2/\text{Hz}], \tag{S42}$$

enters the effective detuning of (S9) through its derivative, $\Delta_{\text{eff}}(t) \leftarrow \Delta_{\text{eff}}(t) - \dot{\phi}_{\text{noise}}(t)$. A vendor frequency-noise PSD $S_\nu(f)$ [$\text{Hz}^2/\text{Hz}$] converts as $S_\phi(f) = S_\nu(f)/f^2$ ($f > 0$), since $\dot{\phi} = 2\pi\nu$ makes the $2\pi$ factors cancel.

**Microwave amplitude noise.** A multiplicative fluctuation on the Rabi envelope, $\Omega(t) \to (1 + a_{\text{noise}}(t))\Omega(t)$ with $S_a(f) = \alpha_{-1}/f + \alpha_0$ [$(\text{fractional})^2/\text{Hz}$] and $|a_{\text{noise}}| \ll 1$, modifies the off-diagonal control elements of (S6), (S8). These spectra are measured with a power-spectrum analyzer.

**Laser intensity noise and AOM leakage.** Optical-rate fluctuations $\Gamma_{\text{exc}} \to \Gamma_{\text{exc}} + \delta\Gamma_{\text{exc}}(t)$ with PSD $S_\Gamma(f)$ [$\text{s}^{-2}/\text{Hz}$], plus a constant leakage rate $\Gamma_{\text{leak}}$, enter the collapse-operator rates of Sec. 1.2 rather than the Hamiltonian.

**Background magnetic-field noise.** An additive Zeeman fluctuation $\delta b_{\text{noise}}(t)$ with $S_b(f) = \beta_{-1}/f + \beta_0$ [$\text{T}^2/\text{Hz}$] shifts $\Delta_{\text{eff}}(t)$ through the coupling $\mu_B g$ of (S9).

**Photon shot noise.** Photon counting is a Poisson process with mean $\lambda = \eta_{\text{opt}} R_{\text{fl}}\, T_{\text{int}}$ ($R_{\text{fl}}$ the spin-dependent fluorescence rate, $T_{\text{int}}$ the integration time) plus a background offset $\lambda_{\text{bg}}$, with the Fisher-information formulas in Sec. 1.4.

## 5.2 From PSDs to Time-Domain Noise Traces and Hamiltonian Insertion

We synthesize *real* stationary Gaussian time series $x(t)$ from assumed one-sided PSDs $S_x(f)$ and insert them into the Hamiltonian (Sec. 1.1) or dissipative rates (Sec. 1.2). The construction is spectrally faithful: naive decimation aliases the spectral slope toward $-0.64$ near the Nyquist frequency and inflates the variance, whereas the band-limited synthesis holds the target $1/f$ slope at $-1.00$ across the sensing band with a variance independent of the oversampling factor (Fig. S8).

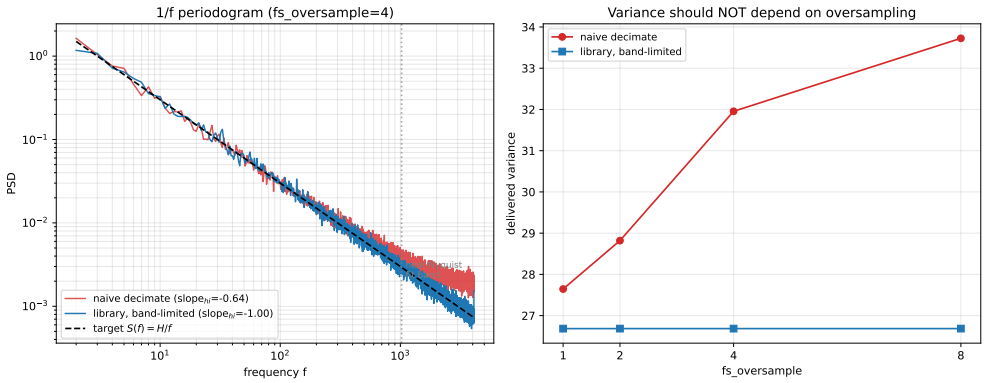


F

Figure S8: Spectral fidelity of the synthesized $1/f$ noise. **(left)** The periodogram: naive decimation aliases the spectral slope from $-1$ toward $-0.64$ near the Nyquist frequency, whereas the band-limited library synthesis holds the target slope at $-1.00$ across the sensing band. **(right)** The delivered variance: naive decimation inflates it as the oversampling factor grows, while the band-limited synthesis keeps it independent of oversampling.

**Discrete grid.** Choose duration $T$, step $\Delta t$, and $N = T/\Delta t$ (even). Define sampling rate $f_s = 1/\Delta t$ and frequency resolution $\Delta f = 1/T$. Positive frequency bins are $f_k = k\,\Delta f$ for $k = 0, \ldots, N/2$.

**Random spectrum with Hermitian symmetry.** Construct complex Fourier coefficients $\{X_k\}$:

$$X_0 = \sqrt{S_x(0)\,\Delta f}\,\mathcal{N}(0,1), \qquad X_{N/2} = \sqrt{S_x(f_{N/2})\,\Delta f}\,\mathcal{N}(0,1), \tag{S43}$$

$$X_k = \sqrt{\tfrac{1}{2}S_x(f_k)\,\Delta f}\,\left(Z_{1,k} + iZ_{2,k}\right), \quad k = 1, \ldots, N/2-1, \tag{S44}$$

$$X_{N-k} = X_k^*, \quad k = 1, \ldots, N/2-1, \tag{S45}$$

with $Z_{1,k}, Z_{2,k} \sim \mathcal{N}(0,1)$ i.i.d. The factor $1/2$ splits the one-sided PSD into conjugate bins.

**Time series and bandlimits.** Apply an upper cutoff $f_c \le f_s/2$ (hardware bandwidth). Set $S_x(f > f_c) = 0$. For $1/f^\alpha$ components, regularize low-$f$ by $S_x(f) \propto 1/\max(f, f_{\min})$ with $f_{\min} = 1/T$. Obtain the real sequence via the unnormalized inverse discrete Fourier transform

$$x[n] \;=\; \mathrm{IFFT}\{X_k\}[n] \;=\; \sum_{k=0}^{N-1} X_k\, e^{\,i2\pi kn/N}, \qquad n = 0, \ldots, N-1, \tag{S46}$$

under which the per-bin power $\mathbb{E}|X_k|^2 = S_x(f_k)\Delta f$ delivers the target time-domain variance $\sum_k S_x(f_k)\Delta f$. For control pipelines, generate at 2–4× the integrator rate, apply waveform-domain filtering if needed, then decimate.

**Insertion of the synthesized traces.** Each trace is inserted at the point defined in Sec. 5.1, namely the phase and amplitude channels into the control terms, the field channel additively in the Zeeman term, and the laser/AOM channel into the collapse-operator rates. The phase channel additionally uses the discrete derivative $\dot{\phi}_{\mathrm{noise}}[n] \approx (\phi_{\mathrm{noise}}[n{+}1] - \phi_{\mathrm{noise}}[n{-}1])/(2\Delta t)$ for the detuning shift in (S9). The model admits all four of these stochastic channels. The main-text figures use none of them, being computed from deterministic dynamics with a parametric drift budget, and the channels are exercised only in the time-dependent-noise study of Sec. 5.3. Reporting this subset is a scope choice for the paper, not a limitation of the model.

**Monte Carlo and reproducibility.** For each condition we draw $R$ independent paths $\{\delta\phi^{(r)}, a^{(r)}, \delta b^{(r)}\}_{r=1}^{R}$ with distinct RNG seeds, propagate the Lindblad/TME dynamics, and report mean $\pm 1\sigma$. We log $\{N, \Delta t, f_c, f_{\min}\}$ and fitted PSD coefficients $\{h., \alpha., \beta.\}$ with the figure artifacts.

**Implementation sketch.**

1. Fix $T, \Delta t, N$ covering the full pulse schedule (incl. idle windows).
2. For each PSD $S_x(f)$, build the one-sided grid, sample $\{X_k\}$ with variance $S_x(f_k)\Delta f/2$, enforce Hermitian symmetry, and IFFT to $x[n]$.
3. Apply masks for $f > f_c$ and low-$f$ regularization; add deterministic leakage (e.g. AOM).
4. Insert $x[n]$ into $H$ or rate channels per above; for phase, also compute $\dot{\phi}_{\text{noise}}[n]$ for (S9).

### 5.3 Sensitivity under time-dependent noise

The main-text sensitivity is evaluated at fixed dissipative rates, where dephasing enters as a Lindblad $T_2^\star$ channel. Here we document the complementary, fully time-resolved calculation, in which the dephasing is generated by an explicit stochastic magnetic-field trace rather than a rate. For each noise flavor we synthesize $\delta b(t)$ from its power spectral density (Supplementary Fig. S9a), inject the time-dependent detuning $\Delta(t) = \gamma_e[b + \delta b(t)]$ into the rotating-frame Hamiltonian, and co-propagate the density matrix and its field derivative $\partial_b \rho$ through the master equation (the tangent solve of Sec. 1.3, with the synthesized trace as the sole dephasing source). Averaging the optical readout over 4096 independent realizations (512 per seed, 8 seeds), and forming the Poisson-photon Fisher information at the measured contrast and photon budget, yields the field sensitivity $\eta_B(\tau)$ (Supplementary Fig. S9b). Because the synthesized trace is the only source of decoherence, $\eta_B(\tau)$ reflects purely the spectral content of the environment.

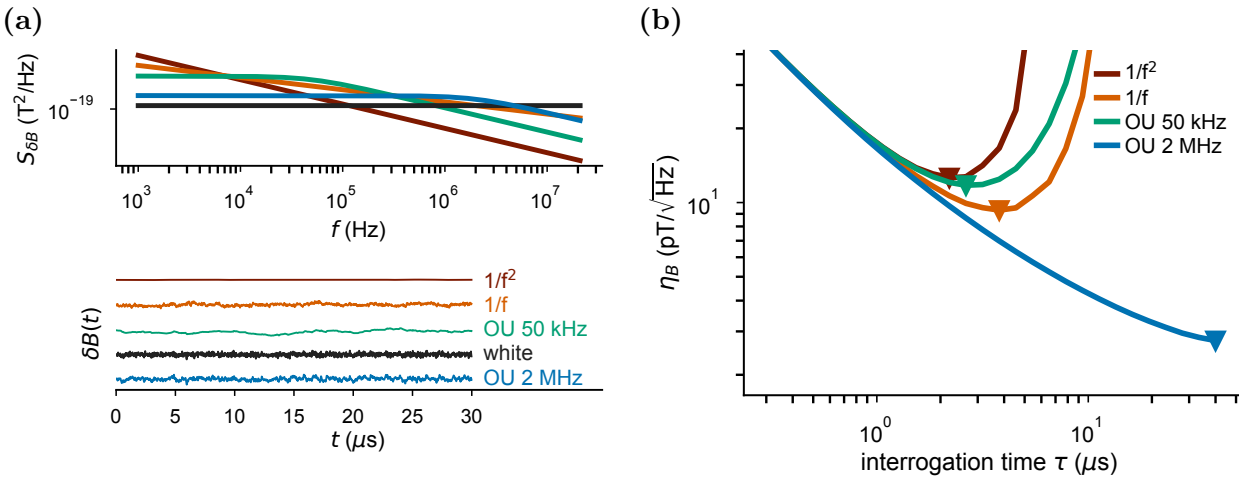


Figure S9: Time-dependent field noise and its effect on sensitivity. **(a)** One-sided power spectral density $S_{\delta B}(f)$ (top) and a representative synthesized trace $\delta b(t)$ (bottom) for the field-noise flavors used in the time-dependent study: $1/f^2$, $1/f$, an Ornstein–Uhlenbeck (OU) spin bath (corner $\sim$50 kHz), white (Johnson) noise, and a high-frequency Ornstein–Uhlenbeck process ($\sim$2 MHz). Traces are synthesized by the Hermitian-symmetric inverse-FFT construction of Sec. 5.2. **(b)** Ramsey field sensitivity $\eta_B(\tau)$ for these flavors (the near-flat white-noise curve is omitted for clarity), obtained by co-propagating $\rho$ and $\partial_b\rho$ under the injected stochastic detuning over 4096 realizations ($512 \times 8$ seeds), averaging the readout, and forming the Poisson-photon Fisher information at the measured contrast and detected-photon budget. Quasi-static ($1/f^2$, $1/f$) noise, whose weight lies below the sensing bandwidth $1/\tau$, produces the deepest, most sharply rising curves, whereas the high-frequency Ornstein–Uhlenbeck process ($\sim$2 MHz), largely averaged out within each interrogation, gives the shallowest.

## Supplementary Note 6: Simulation Details

This section details the numerical protocols used to propagate the NV dynamics and tangent-space equations under noise realizations synthesized from Sec. 5.

### 6.1 Multi-GPU CUDA-Q Backend

We implement our digital-twin dynamics on top of NVIDIA's `cuQuantum` Python bindings, specifically the `cuDensityMat` backend for mixed-state (density-matrix) simulation. Each MPI rank is bound to a single GPU and initializes a `WorkStream` object, which defines the CUDA stream and attaches to a multi-GPU multi-node (MGMN) communicator with the MPI provider. This allows domain decomposition across spatial points $r$, noise realizations $\xi$, and pulse-sequence indices $s$.

**Batched state layout.** For a given rank, we instantiate a `DenseMixedState` object with batch dimension $B = |r| \times |\xi| \times |s|$. The storage is allocated once on device and initialized to $\rho_0 = |0\rangle\langle 0|$ for each batch element. The cuDensityMat API provides local-slice information so that each GPU only stores and evolves its shard of the full batch.

**Operator Construction** The Liouvillian superoperator $\mathcal{L}(t)$ is built once using `Operator` and `OperatorTerm` objects. Its structure includes commutator contributions from the NV-center Hamiltonian and dissipators for radiative decay, intersystem crossing (ISC), phonon exchange, etc. Time dependence enters through parameters (drive envelopes, noise traces, detuning) provided at each call to `compute_action`. An `OperatorAction` handle is prepared so that the entire batch of density matrices can be updated efficiently on device.

**Time integration.** The batched state is advanced by a fixed-step RK4 integrator implemented in CuPy that applies the cuDensityMat operator action at each stage. The scalar-OPM solve of Sec. 7 instead uses Strang splitting for its stiff hyperfine term, and `torchdiffeq.odeint` with adaptive `dopri5` is available as an integration cross-check. At each step: 1. Pulses from the experimental sequence library are interpolated to $t$. 2. Noise traces $\{\delta\phi(t), a_{\mathrm{amp}}(t), \delta b(t)\}$ are indexed for each $(r, \xi)$. 3. Effective Hamiltonian coefficients are assembled and passed to `compute_action`, yielding $\dot{\rho} = \mathcal{L}(t)\rho$. 4. In the tangent master equation mode, the pair $(\rho, \partial_\theta \rho)$ is propagated by re-using the same operator action plus the additional source term $\partial_\theta \mathcal{L}\, \rho$.

**Multi-GPU distribution.** Each GPU integrates its shard of trajectories in parallel, with the batched `DenseMixedState` distributed across ranks through the cuDensityMat MGMN communicator. Observable quantities (readout probabilities $\mathrm{Tr}[M\rho]$ and their field tangents, Fisher increments) are reduced over each rank's local noise and pulse-sequence axes, then gathered to assemble the global spatial maps and noise-averaged statistics.

**Pseudocode for the batched solve.**

```
# one MPI rank per GPU; WorkStream carries the CUDA stream + MGMN communicator
WS = WorkStream(device); WS.set_communicator(comm, provider="MPI")
L_action = OperatorAction(WS, (Operator(hilbert_dims),))     # Liouvillian, built once
rho  = DenseMixedState(WS, hilbert_dims, batch_size=B_local) # B = |r|*|xi|*|s|, sharded
drho = rho.clone(zeros_like_storage=True)                    # tangent state d(rho)/d(theta)

def RHS(t, rho, drho):                          # evaluated at each RK4 stage
    H_t = assemble_H_I_RWA(t, pulses, noise)                 # Eq. (S9)
    rho_dot  = L_action.compute_action(time=t, params=H_t, state_in=rho)
    drho_dot = L_action.compute_action(time=t, params=H_t, state_in=drho) \
               - i*2*pi*[assemble_dH_db(t), rho]             # + TME source
    return rho_dot, drho_dot

rho_t, drho_t = rk4_integrate(RHS, (rho, drho), tgrid)       # fixed-step RK4 in CuPy
p, dp = Tr[M @ rho_t], Tr[M @ drho_t];  F = Fisher_from(p, dp)   # Secs. 1.4-1.5
# reduce over local (xi, s), then gather across ranks via the MGMN communicator
```

**GPU-scale spatial output.** Figure S10 shows the backend producing the full per-pixel error budget over a dense grid, the GPU-scale, spatially-resolved counterpart of the single-operating-point measured-device attribution of the main text. Each pixel carries its own coherence, relaxation, control, and stress, so a spatially varying device is a batch of independent open-system solves rather than a single one. Here a $64 \times 64$ grid with 64 noise realizations per pixel and $N_t$=1500 steps is $2.6 \times 10^5$ tangent trajectories, propagated as one batch across the eight GPUs and reduced to the six maps of Fig. S10 in a single run (the sensitivity and shot floor from the photon Fisher information of Sec. 1.4, the strain and control biases from the closed forms of Sec. 4). The metric-dependent attribution is a *local* statement. Which mechanism limits a pixel depends on that pixel's material, so a real sensor is not limited uniformly. Resolving that requires exactly this per-pixel treatment at scale, which the batched GPU formulation makes a single propagation rather than tens of thousands of serial solves. The same batched propagation scales to a realistic fabrication run: the per-pixel solves are independent, so a wafer-scale map of order $10^5$ to $10^6$ pixels is the same computation at a larger batch, with its throughput set by the aggregate GPU memory (the $B \times N_t$ trajectory-memory budget) and the number of GPUs rather than by any serial step.

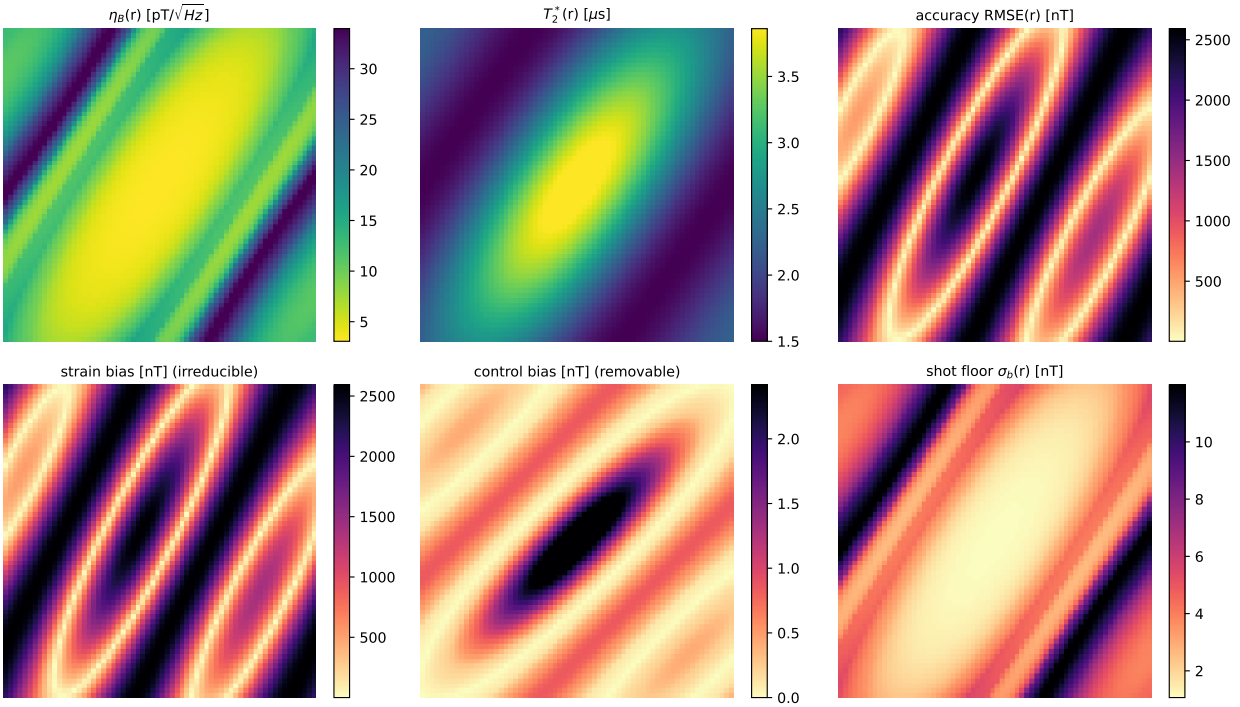


Figure S10: Per-pixel metric-dependent attribution on the multi-GPU backend, over a synthetic $64 \times 64$ grid with 64 noise realizations per pixel and $N_t$=1500 steps to $T$=4 $\mu$s. This is a pipeline demonstration on smooth synthetic material maps, the spatially-resolved GPU-scale counterpart of the single-operating-point measured-device attribution of the main text (Fig. 4). Each pixel runs the stochastic-Hamiltonian tangent solve with its own coherence $T_2^\star(r)$, relaxation $T_1(r)$, control amplitude $\Omega(r)$, and stress $\sigma_{\mathrm{diag}}(r)$, at a detected-photon budget $N_{\mathrm{det}}$=$4 \times 10^8$ and per-pixel readout contrast $C(r)$≈0.09 to 0.25. **(top row)** the sensitivity $\eta_B(r)$ (median $\sim$10 pT/$\sqrt{\mathrm{Hz}}$), the coherence $T_2^\star(r)$ it tracks (both peaking at the grid center), and the accuracy, shown as the total root-mean-square field error RMSE($r$) (median $\sim$1.3 $\mu$T, whose per-mechanism decomposition into the strain, control, and shot channels appears in the bottom row) follows a different spatial pattern. **(bottom row)** the accuracy bias decomposed into its three channels, each set by a different input map: the irreducible strain channel from the stress $\sigma_{\mathrm{diag}}(r)$ ($-\Delta D_{\mathrm{gs}}/\gamma_e$ via the measured 14.58 MHz/GPa hydrostatic coupling of Doherty *et al.*, Phys. Rev. Lett. 112, 047601, reaching $\sim$2.6 $\mu$T), the removable finite-pulse control bias from the drive amplitude $\Omega(r)$ ($\lesssim$3 nT), and the photon shot floor $\sigma_b(r)$ set by the shot-noise-limited readout at the detected-photon budget $N_{\mathrm{det}}$ and per-pixel contrast $C(r)$ ($\sim$1 to 13 nT). Thus each map is tied to a specific input: the sensitivity $\eta_B(r)$ to the coherence $T_2^\star(r)$, the strain bias to the stress $\sigma_{\mathrm{diag}}(r)$, the control bias to the drive amplitude $\Omega(r)$, and the shot floor to the photon budget. The accuracy is limited almost entirely by the irreducible strain, while the sensitivity is set by coherence. Because each metric is governed by a different input map, their spatial structures need not coincide: the control bias, for example, vanishes along the contour where $\Omega$ matches its calibrated value and so is offset from the coherence-set sensitivity. This reproduces the metric-dependent attribution of the main text, resolved per pixel at GPU scale.

**Parameter sets.** The intrinsic Hamiltonian and dissipation parameters are listed in Tables S1–S3. Table S7 collects the operating point used for the per-mechanism attribution of Sec. 4, to-

gether with the readout and drift budgets. The contrast and photon-budget prefactors are the representative values of Sec. 3.2, and the coherence times follow the measured inverse-linear law $T_2^\star\,[N] \approx 9.6\ \mu\text{s}\cdot\text{ppm}$ over $[N] = 0.08$ to 60 ppm.

Table S7: Operating point and drift budgets for the per-mechanism attribution. Intrinsic constants are listed in Tables S1–S3.

| **Quantity** | **Symbol** | **Value** |
|---|---|---|
| Sensed field | $b$ | 1 nT |
| Interrogation time | $\tau$ | 0.6 $\mu$s |
| Dephasing time | $T_2^\star$ | 1.5 $\mu$s |
| Relaxation time | $T_1$ | 5 ms |
| Optical leakage | $\varepsilon$ | $10^{-5}$ |
| Calibration temperature | $T_{\text{cal}}$ | 300 K |
| Temperature offset | $\Delta T$ | 0.1 K |
| Effective contrast | $C_{\text{eff}}$ | 0.14 |
| Detected photons per shot | $N_{\text{det}}$ | $5 \times 10^6$ |
| Zero-field splitting (0 K anchor) | $D_{\text{gs}}$ | 2.8777 GHz |
| Gyromagnetic ratio | $\gamma_e$ | 28.03 GHz/T |
| Thermal shift | $\partial_T D_{\text{gs}}$ | $-79$ kHz/K |
| Thermal apparent field | $\partial_T D_{\text{gs}}/\gamma_e$ | 2.82 $\mu$T/K |
| Drift, background field | $\sigma_b$ | 1 nT |
| Drift, temperature | $\sigma_T$ | 0.1 K |
| Drift, optical leakage | $\sigma_\varepsilon$ | $10^{-5}$ |
| Drift, coherence | $\sigma_{T_2^\star}$ | 2% |
| Drift, relaxation | $\sigma_{T_1}$ | 10% |

# Supplementary Note 7: Optically pumped magnetometer: cesium model and cross-platform test

The framework transfers to a second sensing platform: a scalar optically pumped magnetometer (OPM) run as a total-field sensor. We instantiate the same construction (an open-system solve, a Fisher-based sensitivity, and a system-level design margin) for a cesium-133 vapor and apply it to a real magnetocardiography recording. We use an established OPM model to test whether the metric-dependent, device-grounded provisioning logic carries across platforms.

## 7.1 Cesium-133 ground-state model and Breit–Rabi spectrum

The sensing states are the $^{133}$Cs ground manifold $6^2S_{1/2}$, with nuclear spin $I = 7/2$ and hyperfine constant $A_{\text{hf}} = E_{\text{hf}}/(I + \frac{1}{2}) = 2.298$ GHz (from the ground-state splitting $E_{\text{hf}} = 9.192631770$ GHz). In a bias field $\mathbf{B}$ the Hamiltonian is

$$H = A_{\text{hf}}\,\mathbf{I}\cdot\mathbf{S} + \big(g_J\mu_B\mathbf{S} + g_I\mu_B\mathbf{I}\big)\cdot\mathbf{B}, \tag{S47}$$

whose eigenvalues are given exactly by the Breit–Rabi formula for the two hyperfine manifolds $F = I \pm \frac{1}{2} = 3, 4$. In the linear (low-field) regime the Zeeman splitting of the stretched $F = 4$ manifold is $h\gamma B$ with the gyromagnetic ratio

$$\gamma = \frac{g_F \mu_B}{h}, \qquad g_F = \frac{g_J}{2I+1} \approx \frac{1}{4} \;\Rightarrow\; \gamma \approx 3.50 \text{ Hz/nT}, \tag{S48}$$

half the $^{87}$Rb value ($g_F = \frac{1}{2}$). As an internal consistency check on the Hamiltonian construction and diagonalization, our numerically diagonalized spectrum reproduces the analytic Breit–Rabi eigenvalues of the same Hamiltonian to a relative accuracy of a few parts in $10^6$ across the geomagnetic range. The leading departure from linearity, the quadratic (nonlinear) Zeeman term $\propto B^2/E_{\text{hf}}$, sets the scalar heading dependence and the systematic-error scale of the total-field readout.

Table S8: Cesium-133 ground-state parameters used in the scalar-OPM model. Atomic constants follow the standard $^{133}$Cs D-line data (Steck, *Cesium D Line Data*). The nuclear $g_I$ is in the Steck convention (relative to $\mu_B$). The lower block lists the nominal operating point.

| **Parameter** | **Symbol** | **Value** |
|---|---|---|
| Nuclear spin | $I$ | 7/2 |
| Hyperfine splitting | $E_{\text{hf}}$ | 9.192631770 GHz |
| Magnetic-dipole constant | $A_{\text{hf}} = E_{\text{hf}}/(I + \frac{1}{2})$ | 2.298 GHz |
| Electronic $g$-factor | $g_J$ | 2.00254 |
| Nuclear $g$-factor | $g_I$ | $-3.989 \times 10^{-4}$ |
| Stretched-manifold gyromagnetic ratio | $\gamma$ ($F$=4) | 3.50 Hz/nT |
| Spin-exchange cross-section | $\sigma_{\text{se}}$ | $2.0 \times 10^{-18}$ m$^2$ |
| D1 wavelength, linewidth | $\lambda_{D1}$, $\Gamma_{D1}$ | 894.593 nm, 4.56 MHz |
| Vapor pressure ($\log_{10} P_{\text{Torr}} = A - B/T$) | $A$, $B$ | 7.046, 3830 K |
| Bias field | $B_0$ | 50 $\mu$T |
| Cell temperature | $T_{\text{cell}}$ | 330 K |
| Spin-destruction rate | $R_{\text{sd}}$ | 300 Hz |
| Transverse coherence time | $T_2$ | 6.6 ms |
| Search window ($T$ optimized within) | $T_{\max}$ | 7 ms (optimum $T \approx$ 4.9 ms, Fig. S12b) |
| Readout noise on $\langle F_y \rangle$ | $\sigma_n$ | $10^{-3}$ |

### 7.2 Mean-field spin-exchange master equation and readout

At the atomic densities of a practical cell, spin-exchange collisions dominate the relaxation. Rather than a Lindblad-jump treatment, we propagate the single-atom density matrix under a *mean-field* spin-exchange term, in which the collision rate couples each atom to the ensemble-averaged spin $\langle \mathbf{F} \rangle$. This is the standard density-matrix description of a spin-exchange-broadened (as opposed to spin-exchange-relaxation-free) alkali magnetometer. The coherent Larmor precession and the spin-exchange term are advanced by Strang splitting, alternating the two half-steps to preserve second-

order accuracy. The magnetometer is read out by free-induction decay: after optical preparation the transverse spin precesses at the Larmor frequency, and is detected on *two quadratures*, so its amplitude and phase are recovered independently and the readout is insensitive to the absolute optical phase.

Explicitly, the single-atom density matrix evolves as

$$\dot{\rho} = -\,i\,2\pi\,[H,\rho] + \mathcal{V}_{\mathrm{SD}}(\rho) + \mathcal{V}_{\mathrm{SE}}(\rho) + \mathcal{V}_{\mathrm{OP}}(\rho), \tag{S49}$$

with the three mean-field dissipators, following Mouloudakis *et al.* (arXiv:2402.10746),

$$\mathcal{V}_{\mathrm{SD}}(\rho) = R_{\mathrm{sd}}\big(\mathrm{Tr}_S\rho \otimes \tfrac{1}{2}\mathbb{I}_2 - \rho\big), \tag{S50}$$

$$\mathcal{V}_{\mathrm{SE}}(\rho) = R_{\mathrm{se}}\big(\mathrm{Tr}_S\rho \otimes \varphi(\langle\mathbf{S}\rangle) - \rho\big), \qquad \varphi(\langle\mathbf{S}\rangle) = \tfrac{1}{2}\mathbb{I}_2 + 2\,\langle\mathbf{S}\rangle\cdot\mathbf{S}, \tag{S51}$$

$$\mathcal{V}_{\mathrm{OP}}(\rho) = R_{\mathrm{op}}\big(\mathrm{Tr}_S\rho \otimes |{\uparrow}\rangle\langle{\uparrow}| - \rho\big), \tag{S52}$$

where $\mathrm{Tr}_S$ is the partial trace over the electron spin, $\langle\mathbf{S}\rangle = \mathrm{Tr}(\mathbf{S}\rho)$, and $R_{\mathrm{sd}}, R_{\mathrm{se}}, R_{\mathrm{op}}$ are the spin-destruction, spin-exchange, and optical-pumping rates. Spin destruction and optical pumping are linear in $\rho$. The spin-exchange term is *nonlinear*, because the electron spin-temperature state $\varphi(\langle\mathbf{S}\rangle)$ depends on $\rho$ through $\langle\mathbf{S}\rangle$.

### 7.3 Tangent solve for the OPM

The field sensitivity requires the tangent $G_B \equiv \partial_B\rho$, propagated by the same construction as the NV platform (Sec. 1.3) but adapted to the nonlinear mean-field generator. Differentiating Eq. (S49) with respect to the field along the bias axis gives the coupled tangent equation

$$\dot{G}_B = -\,i\,2\pi\,[H,G_B] + \mathcal{D}_{\mathrm{lin}}(\rho,G_B) - i\,2\pi\,[\partial_B H,\rho], \qquad \partial_B H = \frac{g_J\mu_B S_z + g_I\mu_B I_z}{h}, \tag{S53}$$

whose last term is the source that seeds the tangent. The linearized dissipator $\mathcal{D}_{\mathrm{lin}} = \mathcal{V}_{\mathrm{SD}} + \mathcal{V}_{\mathrm{OP}} + \mathcal{D}_{\mathrm{SE,lin}}$ keeps $\mathcal{V}_{\mathrm{SD}}$ and $\mathcal{V}_{\mathrm{OP}}$ unchanged for the two linear channels, while the nonlinear spin-exchange term contributes its Fréchet derivative,

$$\mathcal{D}_{\mathrm{SE,lin}}(\rho,G_B) = R_{\mathrm{se}}\big[\mathrm{Tr}_S G_B \otimes \varphi(\langle\mathbf{S}\rangle) + \mathrm{Tr}_S\rho \otimes \delta\varphi(G_B) - G_B\big], \qquad \delta\varphi(G_B) = 2\sum_k \mathrm{Tr}(S_k G_B)\,S_k. \tag{S54}$$

The extra term $\mathrm{Tr}_S\rho \otimes \delta\varphi(G_B)$ has no NV analog. It is the derivative of the ensemble spin-temperature feedback, and it is what distinguishes the tangent of a nonlinear (spin-exchange-coupled) generator from the linear NV Lindbladian, for which $\mathcal{D}_{\mathrm{lin}} = \mathcal{D}$. Equations (S49)–(S54) are integrated together by Strang splitting, which exponentiates the fast $\sim$9.2 GHz hyperfine term exactly while resolving the $\sim$175 kHz transverse Larmor precession, so $\rho$ and $G_B$ advance on the same grid and the readout tangent $\partial_B\langle F_y(t)\rangle = \mathrm{Tr}(F_y G_B)$ is available at every step with no finite-difference reference.

### 7.4 Cramér–Rao sensitivity

$F_y = I_y + S_y$ is the $y$-component of the ground-state total angular momentum $\mathbf{F} = \mathbf{I} + \mathbf{S}$, written in the coupled hyperfine basis $\{|F, m_F\rangle\}$ (dimension 16 for $^{133}$Cs, with $F = 3$ and $F = 4$). The readout is the transverse magnetization $\langle F_y\rangle(t) = \mathrm{Tr}[\rho(t)\,F_y]$, a damped sinusoid at the Larmor frequency from which $|B|$ is estimated. The sensitivity is built from the Hamiltonian in four steps that mirror the NV sensitivity chain (Secs. 1.3–1.4). (i) The atoms are optically

pumped into the stretched state $|F{=}4, m{=}{+}4\rangle$ and tipped transverse to the bias field, $\rho_0 = e^{-i(\pi/2)F_y}\,|F{=}4,+4\rangle\langle F{=}4,+4|\,e^{+i(\pi/2)F_y}$. (ii) The state is propagated under Eq. (S49) with the pump off ($R_{\rm op} = 0$), so the transverse spin precesses at the Larmor frequency and decays. (iii) The field tangent $\partial_B\langle F_y(t)\rangle = \mathrm{Tr}(F_y G_B)$ is read from the co-propagated $G_B$ of Eq. (S53). (iv) The tangent is accumulated into the field Fisher information below. The scalar readout is a frequency estimation: the field is inferred from the Larmor frequency $f_L = \gamma B$ of the free-induction decay. The field Fisher information accumulates the squared tangent of the transverse spin over the decay,

$$\mathcal{I}_B(T) = \sum_{t\le T} \frac{\big(\partial_B\langle F_y(t)\rangle\big)^2}{\sigma_n^2}, \qquad \eta_B = \min_T \frac{\sqrt{T}}{\sqrt{\mathcal{I}_B(T)}}, \tag{S55}$$

with $\partial_B\langle F_y\rangle$ co-propagated by the same tangent solve (Sec. 1.3) and $\sigma_n$ the per-sample readout noise on $\langle F_y\rangle$. Evaluated on the recording, and optimizing the sensing time over a free-induction-decay window longer than $T_2$ (a 1 ms window truncates the optimum and inflates $\eta_B$ by $\sim 2.4\times$), Eq. (S55) gives $\eta_B \approx 4.0\,\mathrm{fT}/\sqrt{\mathrm{Hz}}$ across the cardiac band, more than four orders of magnitude below the unshielded $\sim 59\,\mathrm{pT}/\sqrt{\mathrm{Hz}}$ ambient, so the quantum layer never limits the measurement. This value is anchored to an assumed readout noise $\sigma_n$, so the sensor is modeled as readout-(technical-) noise-limited rather than projection-limited. The first-principles spin-projection floor, $\eta_B^{\rm SQL} = 1/(\gamma_{\rm rad}\sqrt{N_{\rm at}\,T_2}) \approx 1.1\,\mathrm{fT}/\sqrt{\mathrm{Hz}}$, lies a further factor of $\sim 4$ below, so the conclusion that the quantum layer is far from the system bottleneck holds under either noise model.

### 7.5 Heading-error accuracy bias

The accuracy channel of the scalar OPM is the *heading error*: the apparent field depends on the angle $\theta$ between the spin polarization (set by the pump/light axis) and the bias field. This is the cross-platform analog of the NV thermal $D_{\rm gs}$ shift, a systematic bias that is independent of the interrogation time. The dominant contribution is the nonlinear Zeeman effect. The quadratic Breit–Rabi term $\propto B^2/E_{\rm hf}$ splits the $2F$ adjacent $m{\to}m{-}1$ transitions of the $F{=}4$ manifold, so a single-sinusoid (centroid) readout returns the population-weighted mean of these lines, which shifts as the tilt reweights the $m$-coherences toward high $m$. We evaluate it by propagating the tilted state $\rho_0(\theta) = e^{-i\theta F_y}|F{=}4,+4\rangle\langle F{=}4,+4|e^{+i\theta F_y}$ and taking the power-weighted spectral centroid $f(\theta)$, with the bias referred to field units as $\beta(\theta) = [f(\theta) - f(\theta_{\rm ref})]/\gamma_{\rm cal}$, calibrated at the symmetric heading $\theta_{\rm ref} = \pi/2$. Two subdominant corrections enter the same centroid: the nuclear Zeeman term $g_I\mu_B\mathbf{I}\cdot\mathbf{B}$ shifts the two hyperfine manifolds oppositely at the $g_I/g_J \sim 2\times10^{-4}$ level, and a light shift along the pump axis adds an effective-field offset $\beta_{\rm ls} = B_{\rm ls}\cos\theta$, whose heading derivative $-B_{\rm ls}\sin\theta$ is largest at $\theta = 90°$, so this term is not stationary at the operating heading. These biases are apparent fields set by the sensor geometry rather than by the coherence, so, like the NV thermal bias, they fix the accuracy floor of the total-field readout without entering the sensitivity. Because each is a slowly-varying (near-DC) offset over the recording, it is further suppressed by the R-peak-aligned beat averaging used to form the mean cardiac field, which is the operative reason the array-level heading bias is small.

### 7.6 Common-mode rejection by principal-component analysis

The dominant environmental noise in an unshielded array is the spatially correlated ambient field, common to all sensors. We reject it by principal-component analysis. Forming the centered multi-sensor data matrix $X$ (channels $\times$ time), we take its singular-value decomposition and project out the leading $k$ spatial modes,

$$X_{\rm res} = X_c - X_c V_k V_k^{\top}, \tag{S56}$$

where $V_k$ holds the leading $k$ right-singular vectors. The residual per-channel power spectral density in the cardiac band gives the rejected floor. Projecting out a single common mode ($k$=1) suppresses the ambient by $\sim$13×, from $\sim$59 to $\sim$4.4 pT/$\sqrt{\text{Hz}}$. Retaining more components keeps improving the rejection but begins to remove genuine spatial structure, so $k$=1 is the conservative choice.

### 7.7 Equivalent-current-dipole localization

The application-level metric is the accuracy with which the cardiac source is localized. We model the source as an equivalent current dipole of moment $\mathbf{m}$ at position $\mathbf{r}_0$, with the exact point-dipole lead field

$$\mathbf{B}(\mathbf{R}) = \frac{\mu_0}{4\pi}\,\frac{3(\mathbf{m}\cdot\hat{\mathbf{R}})\hat{\mathbf{R}} - \mathbf{m}}{R^3}, \qquad \mathbf{R} = \mathbf{r}_{\text{sensor}} - \mathbf{r}_0. \tag{S57}$$

Given a trial position the moment enters linearly and is solved in closed form by least squares, so only the three source coordinates are optimized (Levenberg–Marquardt). The localization error is the root-mean-square displacement of the fitted position under Monte-Carlo per-sensor field noise and gain mismatch. It rises from a gain- and geometry-limited floor of $\sim$0.37 mm (at the spin-projection field-noise floor and a 1% gain match) to $\sim$60 mm at the raw single-channel ambient, steeply but sub-linearly in the per-sensor field noise, floor-limited at low noise and saturating at high noise. The error is also proportional to the cross-sensor gain mismatch (unit slope, from 0.037 mm at 0.1% to 3.7 mm at 10%), but at a realistic 1% match it is 0.37 mm, about 14× under the 5 mm clinical specification, so gain calibration is not the binding constraint (Fig. S11). Common-mode rejection (Sec. 7.6) carries the device onto the few-millimeter contour set by the clinical requirement. The limiting factor is therefore the array-level field-noise floor, not the atomic sensitivity, the same provisioning lesson as the NV platform.

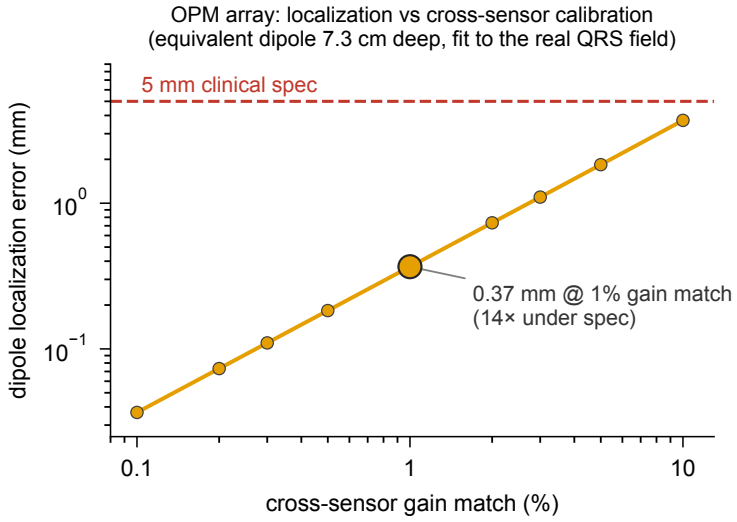


Figure S11: Localization design margin. Root-mean-square error of the equivalent-current-dipole fit to the QRS field (source $\sim$7 cm deep) versus the cross-sensor gain match, propagated by Monte Carlo. At a 1% gain match the error is 0.37 mm, about 14× under the 5 mm clinical localization specification (dashed), so cross-sensor calibration is not the limiting factor. The array-level field-noise floor is.

### 7.8 From the model to the Fig. 5 panels

The four panels of Fig. 5 are produced from the cesium model above and the single unshielded recording, by the same steps used on the bench. Supplementary Fig. S12 shows the intermediate quantity behind each panel, computed by the same core routines (the `opm_scalar` model) that populate the figure, so every reported number is reproducible from the recording.

**Panel (a), the recording.** The R-peak-aligned average of the multi-channel record gives the mean cardiac beat, with its P, QRS, and T waves (Fig. S12f), and the per-sensor field at the QRS peak, interpolated over the array geometry, gives the dipole-field map.

**Panel (b), the engine grounded in the recording.** The free-induction decay is propagated with its field tangent (Sec. 7.4). The transverse spin $\langle F_y(t)\rangle$ and its derivative $\partial_B\langle F_y(t)\rangle$ (Fig. S12a) accumulate the field Fisher information, and minimizing $\sqrt{T}/\sqrt{\mathcal{I}_B(T)}$ over the window returns $\eta_B \approx 4.0\,\mathrm{fT}/\sqrt{\mathrm{Hz}}$ at an optimal window $T \approx 4.9\,\mathrm{ms}$ (Fig. S12b). The frequency-resolved $\eta_B(f)$ and the Larmor-versus-Breit–Rabi check are read from the same solve.

**Panel (c), the sensitivity and robustness budgets.** The sensitivity is split into its spin-projection floor and the residual readout shot noise, both from the Fisher information. The robustness is the low-band (1 to 10 Hz) ambient field-noise floor before and after projecting out the leading spatial principal component (Eq. (S56)). The field amplitude spectral density before and after rejection (Fig. S12d), integrated over the band, gives the $\sim 13\times$ common-mode rejection ($58.7 \to 4.4\,\mathrm{pT}/\sqrt{\mathrm{Hz}}$), with the leading mode carrying 99.9% of the array variance. This mode is a near-uniform ambient field, nearly orthogonal to the cardiac dipole pattern, so projecting it out removes $\lesssim 0.01\%$ of the QRS amplitude while suppressing the common-mode ambient, and the cardiac signal is preserved.

**Panel (d), the localization design space.** An equivalent current dipole is fit to the QRS field across the 26 sensors (Eq. (S57)). The fitted dipole reproduces the measured field at a source depth of $\sim 7\,\mathrm{cm}$ (Fig. S12e). Sweeping the per-sensor field noise and the cross-sensor gain mismatch by Monte Carlo (Sec. 7.7) then maps the localization error over the design space against the 5 mm clinical requirement.

The accuracy channel of the model, the heading-error bias $\beta(\theta)$ (Sec. 7.5), is nearly common-mode for this array because every sensor sits within 1–2° of the same heading, $\theta \approx 90°$. The bias $\beta(\theta)$ varies approximately linearly there (slope $|d\beta/d\theta| \approx 0.12\,\mathrm{nT/deg}$, Fig. S12c), so the shared offset is absorbed into the single field calibration and only a residual spread of $\sim$265 pT survives. It therefore does not appear as a separate Fig. 5 panel.

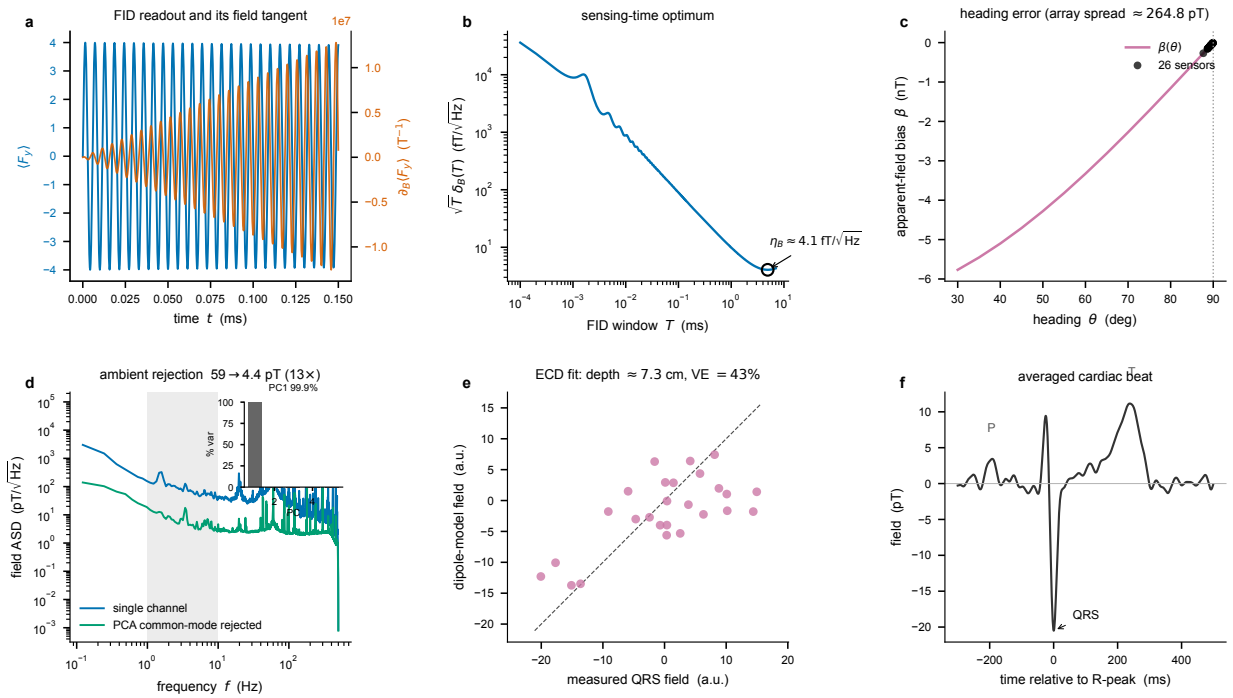


Figure S12: Intermediate quantities behind the Fig. 5 panels, each computed from the cesium model and the recording by the same routines that populate the figure. **(a)** The free-induction-decay readout $\langle F_y(t)\rangle$ (left axis) and its co-propagated field tangent $\partial_B\langle F_y(t)\rangle$ (right axis) over the first 0.15 ms. **(b)** The sensing-time curve $\sqrt{T}\,\delta_B(T)$, whose minimum is the reported sensitivity $\eta_B \approx 4.0\,\mathrm{fT}/\sqrt{\mathrm{Hz}}$ at $T \approx 4.9$ ms. **(c)** The heading-error calibration $\beta(\theta)$. The 26 real sensors cluster within 1–2° of $\theta = 90°$, where $\beta$ varies nearly linearly, so their shared offset calibrates out and the array-level accuracy bias is nearly common-mode. **(d)** The field amplitude spectral density of a single channel and of the PCA common-mode-rejected residual. Integrated over 1 to 10 Hz the ambient falls from 58.7 to $4.4\,\mathrm{pT}/\sqrt{\mathrm{Hz}}$ ($\sim$13×), the leading principal component carrying 99.9% of the array variance (inset). **(e)** The equivalent-current-dipole fit: the model field against the measured QRS field across the 26 sensors, for a source at $\sim$7 cm depth. **(f)** The R-peak-aligned averaged cardiac beat with its P, QRS, and T waves.

# Supplementary Note 8: Charge-state characterization from photoluminescence spectra

Only the negatively charged state $\mathrm{NV}^-$ carries spin contrast, so the sensitivity and photon budgets depend on the fraction of centers in that state. The effective readout contrast is $C_{\mathrm{eff}} = C_{\mathrm{optical}}\, f_{\mathrm{NV}^-}$ (Sec. 3.2), and the measured $\mathrm{NV}^-/\mathrm{NV}^{\mathrm{T}}$ charge fraction $f_{\mathrm{NV}^-}$ as a function of optical power is the starting input of the Fig. 4 analysis (Fig. 4a). This section describes how that fraction was extracted from photoluminescence (PL) spectra, following the method of Thalassinos *et al.* [9].

## 8.1 Experimental setup and samples

The PL measurements used a WITec confocal microscope with a Zeiss 100× objective (numerical aperture NA = 0.9). The excitation was a 532 nm laser with the power incident on the objective

varied from 0.1 to 50 mW. From the objective NA and the excitation wavelength, the excitation spot has an estimated lateral diameter of 306 nm and a confocal depth of 930 nm. A second laser at 355 nm, held fixed at 3 mW, recorded a reference $NV^0$ spectrum. Under ultraviolet excitation the ensemble is driven predominantly into the neutral charge state, so this spectrum gives the pure $NV^0$ emission line shape [9]. The samples were two commercial nitrogen-vacancy-doped diamond substrates (Element Six, supplied by Thorlabs, product codes DNVB1 and DNVB14), the two growths used for the charge-state and photon-budget inputs of Fig. 4.

### 8.2 Spectral features

Table S9 lists the spectral features in the 560 to 710 nm detection window. The zero-phonon lines (ZPLs) of the two charge states lie at 575 nm ($NV^0$) and 637 nm ($NV^-$) [4], each with a broad phonon sideband extending roughly 100 nm to longer wavelengths. Under 532 nm excitation the first-order diamond Raman line (1332 $cm^{-1}$) [10] falls at 572.6 nm, on the rising edge of the $NV^0$ ZPL, and is not resolved as a separate peak at the spectral resolution used here.

Table S9: Spectral features in the 560 to 710 nm detection window.

| **Feature** | **Position** | **Origin** | **Ref.** |
|---|---|---|---|
| $NV^0$ ZPL | 575 nm (2.156 eV) | $NV^0$ zero-phonon transition | [4] |
| $NV^-$ ZPL | 637 nm (1.945 eV) | $NV^-$ zero-phonon transition | [4] |
| Diamond Raman | 572.6 nm (532 nm exc.) | first-order phonon, 1332 $cm^{-1}$ | [10] |
| $NV^0$ sideband | $\sim$580 to 650 nm | phonon-assisted $NV^0$ emission | [4] |
| $NV^-$ sideband | $\sim$650 to 750 nm | phonon-assisted $NV^-$ emission | [4] |

### 8.3 Spectral decomposition

The two charge-state contributions were separated as follows. The 355 nm reference spectrum was clipped to non-negative values and normalized to unit integrated area, giving the pure $NV^0$ line shape $S_{NV^0}(\lambda)$. The pure $NV^-$ line shape $S_{NV^-}(\lambda)$ was obtained from a 532 nm spectrum by subtracting $S_{NV^0}$, scaled by least squares to match the measured spectrum over 565 to 585 nm, where $NV^-$ emission is negligible (its ZPL lies at 637 nm and its sideband extends to longer wavelengths). The non-negative residual, normalized to unit area, was taken as $S_{NV^-}$. We note that the 532 nm-excited first-order diamond Raman line at 572.6 nm falls within this 565–585 nm scaling window, whereas the 355 nm reference has no Raman feature there (its Raman line lies near 373 nm), so the least-squares scaling can absorb a small amount of Raman intensity into the $NV^0$ amplitude and bias $f_{NV^-}$ slightly low. The line is unresolved and the fit residual is below 1% of the peak, which bounds the effect. A Raman-free scaling window or an explicit Raman basis component would remove it.

Each measured spectrum $y(\lambda)$ was then decomposed as

$$y(\lambda) = c_0\, S_{NV^0}(\lambda) + c_1\, S_{NV^-}(\lambda), \qquad c_0,\, c_1 \geq 0, \tag{S58}$$

with the coefficients found by non-negative least squares [11] as implemented in SciPy [12]. Because the reference line shapes are area-normalized, $c_0$ and $c_1$ equal the integrated PL intensities of the $NV^0$ and $NV^-$ components, and these are the quantities reported as a function of excitation power. A representative decomposition is shown in Fig. S13, and across the data set the root-mean-square fit residual is below 1% of the peak intensity. The ratio of PL *intensities* maps onto a ratio of

charge-state *populations* only after correcting for the different emission and collection efficiencies of the two charge states within the detection window [9]. That corrected ratio is the $NV^-/NV^T$ fraction $f_{NV^-}$ used in the effective contrast of Sec. 3.2 and plotted against optical power in Fig. 4a. The raw spectra for both samples are available as described in Sec. **??**.

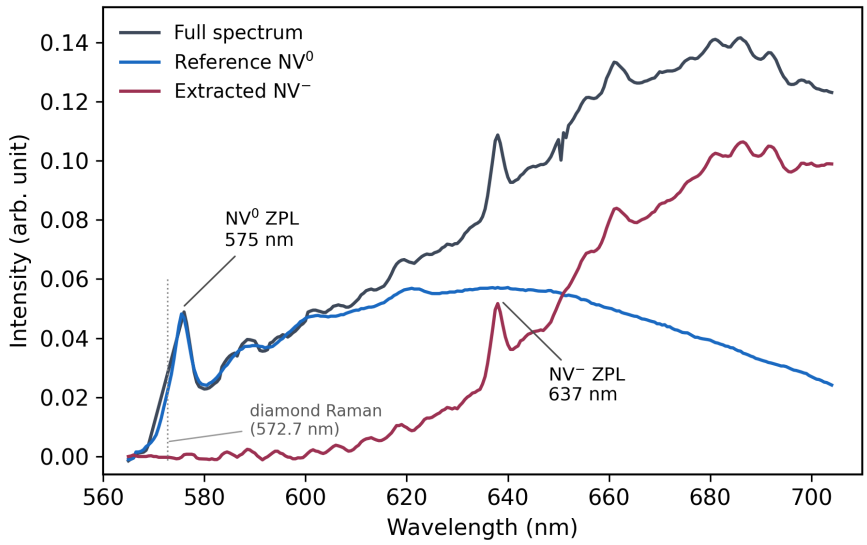


Figure S13: Representative decomposition of a measured photoluminescence spectrum (black) into its $NV^0$ (blue) and $NV^-$ (red) components. The zero-phonon lines of the two charge states are labeled, and the dotted line marks the first-order diamond Raman line under 532 nm excitation, which overlaps the rising edge of the $NV^0$ ZPL. The area under each component is the integrated photoluminescence intensity of that charge state (Eq. (S58)).

## Supplementary Note 9: Outlook: from mechanism attribution to adaptive, meta-learned control

Current practice manages device heterogeneity and drift by recalibrating the control stack on a fixed schedule, a cost that grows with processor size. Pairing the mechanism-resolved twin developed here with recent scaling laws for meta-learned adaptation [13] suggests a different route, in which calibration is budgeted per device and tracked online rather than repeated wholesale. The object being tuned is a generalized control vector $\theta \in \mathbb{R}^P$ that is agnostic to physical parameterization. Beyond gate-pulse coefficients, $\theta$ can collect any calibratable knob across the stack, from quasi-static hardware flux biases and environmental setpoints such as thermoelectric-cooler currents to compilation choices such as virtual-$Z$ offsets or dynamical-decoupling delays.

**Setup.** Let a device instance be a task $\xi \in \Xi$ drawn from a distribution $\mathcal{P}(\xi)$, where $\xi$ collects the mechanism parameters (coherence times, noise rates, strain, couplings). From a single tangent solve the twin returns the performance objective $L(\theta;\xi) = 1-\mathcal{F}$ and its gradient $\nabla_\theta L$, alongside the metrics and Shapley shares of the main text. A knob enters this construction whenever its coupling to the physical model parameters $p$ is known, since the same tangent supplies its sensitivity by the

chain rule,

$$\partial_{\theta_i}\rho \;=\; \frac{\partial\rho}{\partial p}\,\frac{\partial p}{\partial\theta_i}, \tag{S59}$$

so the framework is agnostic to what $\theta_i$ physically is, provided $\partial p/\partial\theta_i$ is characterized. The coupling through $p$ also groups the coordinates $\{1,\dots,P\}$ into mechanism-aligned blocks $S_1,\dots,S_m$, where $S_j$ collects the knobs that act on mechanism $j$. The Shapley value $\phi_j(\xi)$ of the main text quantifies how much mechanism $j$ currently degrades the objective, and $\Pi_S$ denotes the orthogonal projection that restricts an update to a block $S$.

**The twin as a data engine.** As a forward model grounded in and updated from measured device data, the twin is a generator of physically accurate synthetic training data on which a control policy $\pi_\theta$ can be meta-trained to a device family before deployment. Because the twin can activate or suppress each mechanism in the model, that data carries ground-truth per-mechanism labels no hardware measurement can supply, which is what makes it a supervised-learning source rather than a plain simulator. Producing that data at the scale meta-training demands is where the twin's batched, multi-GPU backend (Sec. 6.1) becomes essential. A single campaign co-propagates the state and its tangent across a large grid of device configurations at once, sweeping the experimental parameters $\xi$ and batching over spatial position and noise realization, so one run defines the task distribution $\mathcal{P}(\xi)$ the policy is trained on rather than solving each configuration in series (Algorithm 1).

**When and how much to adapt.** At deployment the twin supplies the inputs that set the adaptation budget. It measures the device-to-device and temporal variance $\sigma_\tau^2$ of the task parameters $\xi$, and it estimates the local landscape geometry from the tangent, so a few probe steps yield a finite budget rather than an open-ended search. The companion scaling law states that the expected gain from per-device adaptation grows with this variance and saturates with the number of adaptation steps $K$,

$$G(K) \;\approx\; A_\infty\big(1-e^{-\beta K}\big), \qquad A_\infty \propto \sigma_\tau^2, \tag{S60}$$

with an achievable ceiling $A_\infty$ set by the measured heterogeneity and an adaptation-rate constant $\beta$ that fixes a budget beyond which returns diminish. We reproduce only the form. The derivation, the constant $\beta$, and the conditions under which the bound holds are given by Leclerc *et al.* [13].

**Mechanism-targeted adaptation.** Because the twin resolves error by mechanism, the same attribution directs the budget. Rather than optimizing blindly over the full multi-layer parameter space, a runtime diagnosis restricts the update to the block that couples to the mechanism worth correcting, and the adaptation step is a masked gradient descent on the objective,

$$\theta \;\leftarrow\; \theta - \eta\,\Pi_S\,\nabla_\theta L(\theta;\xi), \tag{S61}$$

with step size $\eta$ and $\nabla_\theta L$ read from the same tangent, so the inactive parameters are held fixed and the adaptation stays within a well-behaved region of the landscape. The block is chosen by the recoverable gain, which is limited both by how much a mechanism currently costs and by how much adaptation on its block can deliver,

$$S \;=\; S_{j^\star}, \qquad j^\star \;=\; \arg\max_j \;\min\big(\phi_j(\xi),\, A_\infty^{(S_j)}\big), \tag{S62}$$

where $A_\infty^{(S_j)}$ is the ceiling of Eq. (S60) restricted to block $S_j$ and inherits its scaling with that block's task variance (Algorithm 2).

**What the readout must resolve.** Targeting a single mechanism presumes the readout can tell mechanisms apart. A scalar readout, a single contrast value or one field estimate, can respond almost identically to a thermal shift, a dephasing, and an optical leakage, leaving their Shapley shares collinear and a masked update on one channel indistinguishable from a misdirected update on another. Per-mechanism targeting is well-posed only when the Fisher information over the mechanism parameters is well conditioned, and a richer readout, a time-resolved Ramsey or free-induction trace, several pulse sequences, or a spectral or multi-axis measurement, raises its rank and separates the shares. The two platforms studied here show the effect. The NV sequence family separates coherence loss from a temperature shift because each acts on a different part of the time-resolved signal, and the cesium free-induction spectrum separates a heading bias from the true field because the two enter the line shape differently. Choosing the readout to be informative about the mechanisms one intends to control is thus the measurement-side counterpart of choosing which knobs to actuate, and the twin can rank a candidate readout by the conditioning of the mechanism Fisher matrix it would produce.

**Design time versus runtime.** The closed loop suppresses only those mechanisms whose controlling parameters are adjustable during operation. Pulse coefficients, flux biases, thermoelectric setpoints, and compilation choices can be retuned in flight. Many device properties cannot. The fabrication of the control lines, the layout of the photonic circuit, and the material itself, including its defect density and strain, are fixed at design and growth time, and no online update can change them. The attribution separates these two regimes. A large Shapley share on an online-adjustable channel is what the runtime loop is for, whereas a large share on a parameter frozen at fabrication points instead to a design-time change, a different control-line geometry, a lower-loss photonic switch, or a cleaner material. The same per-mechanism budget therefore allocates effort across the development cycle, indicating whether a given metric is best recovered by better fabrication or by better operation.

**Algorithm 1** Offline: task-distribution generation and meta-training (twin, multi-GPU).

**Require:** parameter grid $\Xi = \{\xi^{(1)}, \ldots, \xi^{(G)}\}$, policy $\pi_\theta$, meta-iterations $M$

1: solve the twin in one batched multi-GPU campaign over $\Xi$, batching position, noise realization, and parameter sweep, collecting for each $\xi$ the objective $L$, gradient $\nabla_\theta L$, and mechanism shares $\{\phi_j\}$ ▷ populates $\mathcal{P}(\xi)$

2: meta-train $\pi_\theta$ on this ensemble for $M$ iterations

3: **return** meta-initialization $\theta_0$

**Algorithm 2** Online: closed-loop, mechanism-targeted calibration (edge device).

**Require:** device, twin, meta-initialization $\theta_0$, gain threshold $\epsilon$, step $\eta$

$\theta \leftarrow \theta_0$
**loop**
  **Probe:** run $N$ characterization steps $\rightarrow$ estimate task parameters $\hat{\xi}$
  **Diagnose:** one tangent solve at $(\theta, \hat{\xi}) \rightarrow L,\ \nabla_\theta L,\ \{\phi_j\}$
  select block $S = S_{j^\star}$ by Eq. (S62); estimate $\sigma_\tau^2$ and geometry $\rightarrow$ ceiling $A_\infty$, budget $K$
  **if** $A_\infty < \epsilon$ **then**
    keep $\theta$ (non-adaptive) and **continue** ▷ near-nominal device
  **end if**
  **Adapt: for** $k = 1$ **to** $K$ **do** $\theta \leftarrow \theta - \eta\, \Pi_S\, \nabla_\theta L(\theta; \hat{\xi})$
  **Deploy** $\theta$; re-probe as drift accrues
**end loop**

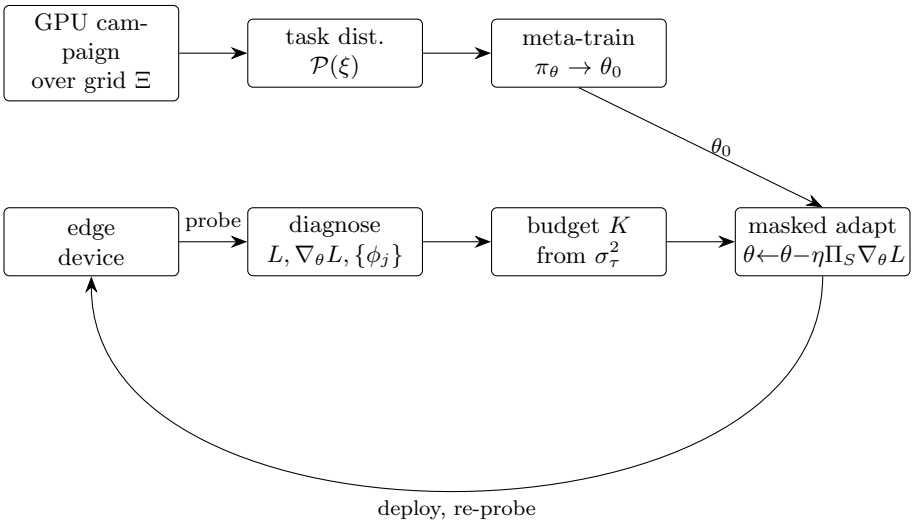


Figure S14: The proposed closed loop. *Offline*, one batched multi-GPU campaign over a parameter grid $\Xi$ defines the task distribution $\mathcal{P}(\xi)$ and meta-trains the policy to an initialization $\theta_0$ (Algorithm 1). *Online*, the twin probes and diagnoses the edge device from a single tangent solve, sets the adaptation budget $K$ from the measured heterogeneity $\sigma_\tau^2$, and applies masked updates to the diagnosed block before redeploying and re-probing as drift accrues (Algorithm 2).

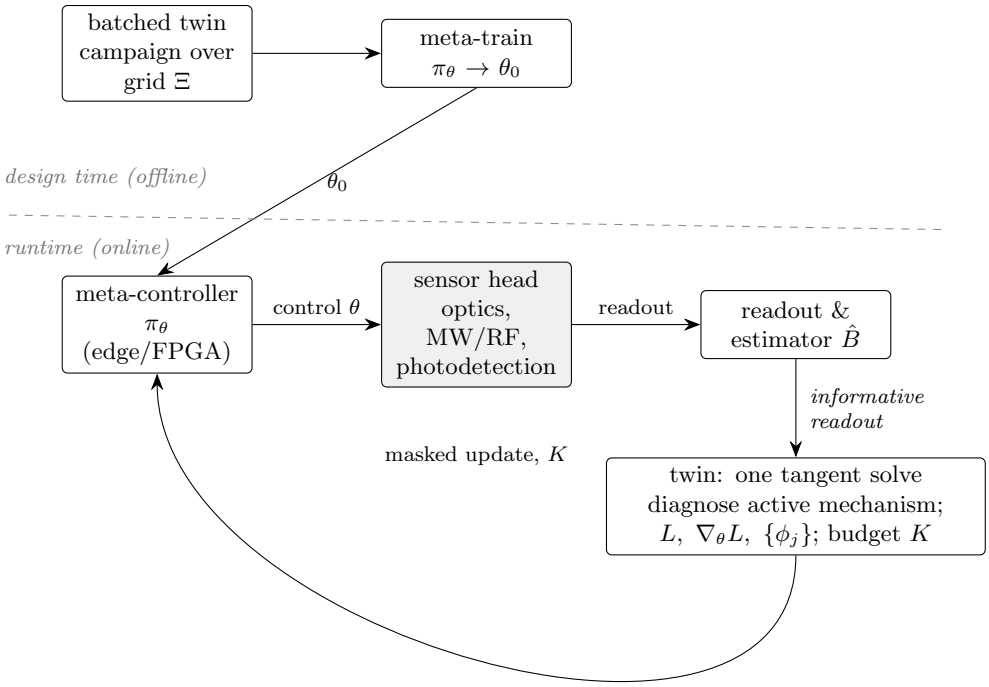


Figure S15: Proposed integration of the twin and the meta-learned controller in a sensing apparatus (high level; not implemented here). At design time, a batched multi-GPU campaign over the parameter grid $\Xi$ defines the task distribution and meta-trains the policy to an initialization $\theta_0$. At runtime, an edge controller holding $\pi_\theta$ actuates the sensor's control knobs $\theta$ (microwave phase and amplitude, bias field, optical power); the photon-count or free-induction readout is estimated and passed to the twin, which runs a single tangent solve to identify the active mechanism, read its Shapley share $\{\phi_j\}$ and sensitivities, and return a masked gradient with an adaptation budget $K$. The controller applies the targeted update, and the loop repeats as drift accrues. The readout must be informative enough to separate the mechanisms being diagnosed; a scalar readout can leave their shares unidentifiable, whereas a time-resolved or spectral readout makes them separable. Any platform exposing these control knobs and a differentiable model instantiates the sensor block; the NV ensemble and the cesium magnetometer studied here are two such instances.

**Operational sensing.** This loop matters most for quantum sensors that operate outside the laboratory, where the environment drifts and the platform itself moves. In these settings accuracy, the systematic bias emphasized throughout this work, is the binding metric, because a navigation or detection decision integrates the sensor output over time and a slowly-varying bias accumulates into a growing error rather than averaging away. A magnetometer flown for anomaly detection on a moving vehicle sees its dominant bias source change as the platform moves, from a heading-dependent shift as the vehicle turns to a thermal shift of the transition frequency as it heats or a stray field from onboard electronics. A self-adapting sensor uses the twin to diagnose which channel is active at each moment and spends a bounded adaptation budget re-nulling only that channel, so a genuine anomaly is separated from a platform artifact without returning to calibration. The

same construction supports positioning, navigation, and timing in GPS-denied conditions, where the worth of a magnetometer, gravimeter, or clock is set by how well its bias is held over an operational interval.

**Status.** Assembled, these pieces describe a closed loop: probe the edge device, diagnose the active mechanisms and read the sensitivities from one tangent solve, set the budget from the measured heterogeneity and geometry, apply targeted updates, and re-probe as drift accumulates (Fig. S14), mapped onto an apparatus in Fig. S15. We present this as a direction rather than a demonstration. The attribution framework is established here and the adaptation scaling laws are established separately, but closing the loop end to end on hardware, within the controllability and model-fidelity assumptions each component requires, remains open. Its appeal is a change in what the hardware is asked to do, from a passive subject of environmental drift toward an active participant in its own stabilization.

## Acknowledgements

This work was supported by the MITRE Independent Research and Development Program. Portions of this technical data were produced for the U. S. Government under Contract No. FA870225CB001 and W56KGU-18-D-0004, and is subject to the Rights in Technical Data-Noncommercial Items Clause DFARS 252.227-7013 (FEB 2014).



## References

[1] Ernst, S., Scheidegger, P. J., Diesch, S. & Degen, C. L. Modeling temperature-dependent population dynamics in the excited state of the nitrogen-vacancy center in diamond. *Phys. Rev. B* **108,** 085203 (2023).

[2] Lenef, A. & Rand, S. C. Electronic structure of the N-V center in diamond: Theory. *Phys. Rev. B* **53,** 13441 (1996).

[3] Chen, X.-D. et al. Temperature dependent energy level shifts of nitrogen-vacancy centers in diamond. *Appl. Phys. Lett.* **99,** 161903 (2011).

[4] Doherty, M. W. et al. The nitrogen-vacancy colour centre in diamond. *Phys. Rep.* **528,** 1 (2013).

[5] Bassett, L. C. et al. Ultrafast optical control of orbital and spin dynamics in a solid-state defect. *Science* **345,** 1333 (2014).

[6] Rogers, L. J., McMurtrie, R. L., Sellars, M. J. & Manson, N. B. Time-averaging within the excited state of the nitrogen-vacancy centre in diamond. *New J. Phys.* **11,** 063007 (2009).

[7] Breuer, H.-P. & Petruccione, F. *The Theory of Open Quantum Systems* (Oxford Univ. Press, 2002).

[8] Zhang, M. et al. QuanEstimation: An open-source toolkit for quantum parameter estimation. *Phys. Rev. Res.* **4,** 043057 (2022).

[9] Thalassinos, G. et al. Robust quantification of the diamond nitrogen-vacancy center charge state via photoluminescence spectroscopy. *APL Photonics* **10,** 101102 (2025).

[10] Solin, S. A. & Ramdas, A. K. Raman spectrum of diamond. *Phys. Rev. B* **1,** 1687–1698 (1970).

[11] Lawson, C. L. & Hanson, R. J. *Solving Least Squares Problems* (SIAM, 1995).

[12] Virtanen, P. et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* **17,** 261–272 (2020).

[13] Leclerc, N., Miller, C. & Brawand, N. When does adaptation win? Scaling laws for meta-learning in quantum control. In *Proc. 43rd International Conference on Machine Learning (ICML)* (2026).