

Metastable nanoparticle cores amplify single-ion quantum level crossings

Xiang Zhao

sungchao@gmail.com

Australian National University <https://orcid.org/0000-0002-1168-3439>

Shuangchi Liu

Institute of Geology and Geophysics, Chinese Academy of Sciences

Tongwei Zhang

Institute of Geology and Geophysics, Chinese Academy of Sciences

Physical Sciences - Article

Keywords: Quantum level crossings, Nanomagnetism, First-order reversal curves (FORC), Single-ion anisotropy, Zero-field splitting, Mesoscopic physics, Superparamagnetism, Magnetoferritin, Spin dynamics, Quantum switching resonance

Posted Date: October 15th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10767658/v2>

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

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Metastable nanoparticle cores amplify single-ion quantum level crossings

Xiang Zhao^{1*†}, Shuangchi Liu^{2†}, Tongwei Zhang^{3†}

¹ Research School of Earth Sciences, The Australian National University, Canberra, ACT 2601, Australia

² State Key Laboratory of Lithospheric and Environmental Coevolution, Institute of Geology and Geophysics, Chinese Academy of Sciences, Beijing 100029, China

³ Key Laboratory of Planetary Science and Frontier Technology, Institute of Geology and Geophysics, Chinese Academy of Sciences, Beijing 100029, China

† These authors contributed equally.

Present address (X.Z.): Independent researcher, Canberra, ACT, Australia.

* Correspondence: sungchao@gmail.com

Abstract

Classical size dispersion and experimental artefacts have long obscured the discrete quantum energy levels of individual magnetic sites in disordered nanoparticle ensembles. We report a field-periodic, phase-coherent switching resonance—robust under instrumental and processing controls—in the first-order reversal curve (FORC) response of biogenic magnetoferritin at 2 and 5 K. The period, $\mu_0\Delta H = 11.9 \pm 1.6$ mT, is invariant and phase-coherent ($p \approx 10^{-3}$), and the resonance disappears above the blocking temperature. Cross-correlation shows the ridges tilt in proportion to the core remanence slope ($\rho \approx 0.78$): the structure tracks the internal hysteretic state of the core, not the applied field, excluding instrumental and processing origins. The spacing gives an axial zero-field splitting $D = 0.016 \pm 0.002$ K, identifying the active unit as an individual surface Fe^{3+} ion ($S = 5/2$), whose relaxation is exponentially amplified by cores thermally poised at their switching threshold ($\Gamma t \approx 1$). The resulting size–temperature window ($T_B/9.0 \leq T \leq T_B$) accounts for the thirty-year absence of resolved resonances in conventional magnetometry and predicts the geometric signature recovered in independent data at 6.9σ . This establishes a paradigm where macroscopic metastability serves as a quantum amplifier, enabling benchtop spectroscopy of single-ion quantum states in inherently disordered nanomaterials.

Resolving the discrete energy levels of individual magnetic ions within disordered nanoparticle ensembles would expose local quantum structure without requiring the monodisperse lattices or orientational order of molecular nanomagnets. The barrier is not merely small signal size.

Twice-differentiated magnetometry readily manufactures spurious periodicities from instrumental noise and numerical artefacts, and the history of macroscopic quantum effects in ferritin is a testament to this difficulty^{1,2}. Three decades ago, resonant switching was inferred from isolated zero-field anomalies in magnetic viscosity³ and thermal magnetization peaks². Yet, in each instance, classical relaxation models over broad size distributions were argued to reproduce the observations^{4,5}. The original authors disputed that reading⁶, and because an isolated anomaly near zero field carries no internal metric, the debate stalled: no subsequent measurement could definitively separate quantum mechanics from classical polydispersity.

That record sets the standard a new claim must meet. It also identifies what would constitute progress: not an isolated anomaly, but a structure carrying an internal metric: features whose spacing is itself the measurement, and whose geometry in the two-dimensional measurement space can be tested independently of their amplitude.

FORC analysis gives access to such a structure. The FORC distribution $\rho(H, H_r) = -\frac{1}{2} \partial^2 M / \partial H \partial H_r$ is a mixed second derivative^{7,8} and is therefore selective, not amplifying: it returns near zero on a smooth switching background and retains only systematic curvature. A localised, field-periodic modulation of switching probability that is buried beneath the classical distribution in a bulk $M(H)$ loop survives this operation. There is precedent for reading quantum structure this way. The quantum-FORC framework has been applied to Mn_{12} -acetate to separate thermally activated from tunnelling contributions⁹. More directly relevant, level-crossing resonances are resolved in dilute rare-earth systems such as $\text{LiYF}_4:\text{Ho}^{3+}$ by stepped, dwell-based magnetometry of the same character as ours¹⁰, where the features at the crossing fields are reproduced by cross-relaxation and phonon-bottleneck effects and not by coherent passage¹¹. That system has a non-Kramers ground doublet, whereas the ion considered here is a Kramers one; the precedent we take from it is the measurement protocol and the incoherent character of the mechanism (in the relaxation dynamics sense, as opposed to Landau–Zener coherent passage), not the level structure.

Beyond resolving the ferritin paradox, this work establishes that macroscopic metastability—often viewed as a source of decoherence—can be harnessed as a quantum amplifier, providing a non-destructive, table-top route to single-ion quantum spectroscopy in materials that are neither

crystalline nor monodisperse.

Results

The material and the observation. Magnetoferritin (M-HFn) was prepared by biomimetic synthesis of magnetite within recombinant human heavy-chain ferritin cages, whose 8 nm interior cavity and 12 nm outer shell enforce physical separation between cores¹²; the same study gives a mean physical core diameter of 5.3 ± 1.2 nm and an interaction field $H_{b,1/2} = 2.1$ mT. Langevin inversion of the room-temperature magnetisation gives a median magnetic diameter of 4.5 nm (Fig. 1a), and the blocking behaviour is conventional (Fig. 1). Every quantity we use is measured on the specimen that produced the FORC data (Methods): the saturation remanence loses half its moment by 8 K, with two thirds still blocked at 5 K and 2.0% at 50 K. No anomaly appears in any of these.

FORC measurements at 2 K show a different picture (Fig. 2). Beneath the main switching distribution the residual carries curved, sign-alternating ridges fanning across the low-field region, with inclination increasing towards small $|H_r|$. The same structure is present at 5 K. At 50 K the ridges are gone, and the residual contains a single vertical feature at the field zero crossing; the full series through 300 K is in Extended Data Fig. 6.

That the ridges are resolved follows from the variable-resolution acquisition protocol of ref. 13, which concentrates field steps where the distribution carries structure. Over the 0–90 mT interval the median step is 2.21 mT; the acquisition parameters, and the reason this grid was preferred to a regular one of the same total cost, are given in Methods.

The periodic component accounts for only a few tenths of one percent of the bulk switching amplitude, but is resolved in ρ through the selective curvature sensitivity of the mixed second derivative ($(W/\Delta H)^2 \approx 29\text{--}37$), which suppresses the broad classical switching background while preserving rapid oscillatory modulations.

Measuring the ridge geometry. A ridge in a FORC distribution is a locus in the (H, H_r) plane, not a feature at a single field, so any test for periodicity must first specify that locus. Fitting the locus and the period together is possible in principle but not on data of this size: the coherence statistic is far more sensitive to the locus parameters than a coarse parameter grid suggests, and

the effective number of independent locus trials is of order 10^5 – 10^6 , not the number of grid points (Methods). We therefore measure the locus first.

For each pair of adjacent reversal curves we cross-correlate their residual FORC profiles in H and take the displacement that maximises the correlation. This returns the local slope of whatever structure the two curves share, with no template and no free parameters; we report it as $-dH_n/dH_r$, positive in the sense the mechanism predicts (Methods). The median displacement is +0.476 at 2 K over 41 curve pairs and +0.277 at 5 K over 34 (Fig. 3a). At 50 K the median displacement over nineteen adjacent pairs is 0.000, with an interquartile range of 0.000: 63% of pairs return exactly zero and the mean is $+0.17 \pm 0.17$ (s.e.m.), consistent with no displacement. The structure there is vertical, recurring at the same applied field from one reversal curve to the next.

Periodicity along the measured locus. Projecting along the measured locus, we compute for each block of reversal curves the complex Fourier amplitude of the residual at trial period ΔH on the native grid points, normalise it to unit modulus, and form the resultant of the K unit vectors (Fig. 3b). For independent blocks the resultant follows the Rayleigh distribution, so $p = \exp(-KR^2)$ in closed form and no surrogate construction is required. Only ΔH is free. Blocks are fixed by the smoothing factor without tuning ($K = 15$ at 2 K and $K = 12$ at 5 K across the three channels; Methods):

T	$\mu_0 \Delta H$ (mT)	R	K	p
2 K	13.10	0.879	15	9.5×10^{-4}
5 K	10.80	0.921	12	3.9×10^{-3}
50 K	26.70	0.993	6	0.275
100 K	12.40	0.983	6	0.310
150 K	15.70	0.946	6	0.473
200 K	28.60	0.909	6	0.715
300 K	13.20	0.925	6	0.603

The two temperatures below blocking return $p \approx 10^{-3}$, and none of the five above returns better than 0.27. However, the control datasets carry only 20 to 22 reversal curves ($K = 6$) against 42 at 2 K. Subsampling the 2 K dataset to $K = 6$ yields $R = 0.988$ with $p = 0.29$, confirming that high coherence at elevated temperatures reflects the lower statistical power of fewer blocks rather than true phase locking. Analysed by acquisition instead of pooled, no single run reaches

significance at this block size (standard series $K = 5$ at 2 K; remanence series $K = 10$, $p = 0.068$); the three channels are combined under the null independence checked by the 300 K calibration (Methods).

We therefore require two conditions, both measured: significant phase coherence at a common period, and a tilted locus tracking the remanence branch. Above blocking the second condition fails outright: the locus is vertical (median displacement and IQR of 0.000 across nineteen curve pairs; Fig. 3a), excluding true periodicity. At 2 K and 5 K both conditions are met.

Excluding instrumental and processing origins. Extended Data Table 1 systematically evaluates twenty artifact classes; three quantitative contrasts are decisive. First, a 2.26-fold contrast in measurement dwell time (52.60 vs 23.29 s/node, ref. 14) produces only a 0.8% shift in period (13.00 vs 13.10 mT), ruling out time-locked instrument periodicities which require a 56% shift (Extended Data Fig. 3). Second, a 26% variation in node spacing shifts the period by 1.9% rather than the 139% required for aliasing (Extended Data Fig. 4). Third, across 246 ridge crests, peak positions coincide between in-field and remanence channels despite a median moment difference of 11.7 memu, excluding moment-triggered auto-ranging (Extended Data Fig. 2). Robustness against smoothing ($SF = 1-4$, Extended Data Fig. 5) and failure of data-dependent curve alignment (Extended Data Fig. 8) are detailed in Supplementary Note S6.

The tilt follows the hysteretic state. The ridge locus is not vertical, and the way it tilts is a measurement in its own right. Within a single temperature the tilt varies from row to row with the local slope of the remanence branch, and falls to zero at deep reversal fields where the loop saturates. At 2 K it falls from +0.995 near $\mu_0 H_r = -40$ mT to +0.026 near -183 mT while dM_{rem}/dH_r falls from +0.0226 to +0.0013 mT^{-1} ; row by row the two correlate with Spearman $\rho = 0.78$ over 41 pairs ($p = 1 \times 10^{-9}$), and $\rho = 0.76$ over 34 pairs at 5 K ($p = 2 \times 10^{-7}$) (Extended Data Fig. 7). The same relation holds across temperatures as well as within them: over five temperatures the slope of the remanence branch reverses sign between 5 K and 50 K, and the tilt vanishes at the same place (Fig. 3c).

The force of this does not depend on what produces the ridges: it states that the position of the structure follows the hysteretic state of the core rather than the applied field, and no artefact locked to the applied field has access to the slope of the remanence branch. We looked for the

correlation because the reading below predicts it; the exclusion it supports would survive that reading being wrong.

An independent determination confirms the geometry from a statistic in a different domain. Evaluating the complex Fourier component of the residual at each reversal-curve row, on the native non-uniform field nodes and without interpolation, returns the comb position b modulo ΔH . Its variation across rows reproduces the local tilts measured by cross-correlation, and fitting it to the remanence branch with one scale factor returns $\lambda = 29$ mT with a root-mean-square residual of 0.72 mT against 1.87 mT for a straight line (Extended Data Fig. 9). Unlike cross-correlation, which returns a displacement, the phase fixes the absolute position of the ladder, and the λ it gives agrees to within 10% with the value fitted to the cross-correlation tilt. Two statistics in different domains therefore return the same geometry.

The full result—both tilt and period—has not been reproduced in an independent sample, but the geometric signature alone has. A FORC series measured at 2 K on 7.6 nm silica-coated magnetite in another laboratory (ref. 15) gives a row-to-row displacement of $+0.68 \pm 0.10$ over 33 curve pairs, non-zero at 6.9σ (Wilcoxon $p = 1.6 \times 10^{-5}$), against $+0.17 \pm 0.17$ (s.e.m.) for our own 50 K control. The discriminant therefore returns a positive result on a different material in a different laboratory; to our knowledge no comparable structure has previously been reported in a FORC distribution. Visually, an alternating band structure is discernible in the FORC map of that sample, consistent with the ridge pattern observed here; however, the sampling density is insufficient to determine whether the bands follow a regular periodicity. Acquired on a regular grid with 3.3 samples per period against the 5.2 our variable-resolution protocol delivers, that series gives 9.9 ± 4.7 mT against 13.10 mT here, and with a single measurement channel its coherence test returns $p = 0.44$. We report the tilt as corroboration of the geometry and make no claim from it about ΔH .

A single ion, not the particle. With instrumental and processing artefacts excluded, the spacing can be interpreted directly as a physical field scale. For an ion with an axial spin Hamiltonian $\mathcal{H} = -DS_z^2 - g\mu_B\mu_0 H_{\text{eff},n} S_z$, the eigenvalues cross pairwise at $\mu_0 H_{\text{eff},n} = nD/(g\mu_B)$, so consecutive crossings are separated by $D/(g\mu_B)$ independently of n , of temperature and of particle size (Fig. 4b). Taking $g = 2$ for the ${}^6S_{5/2}$ state and $\mu_0\Delta H = 11.9 \pm 1.6$ mT, the mean and

standard deviation of the two determinations above, gives

$$\mathbf{D} = 0.016 \pm 0.002 \text{ K}$$

The interval is the spread between the two temperatures; within each, the bootstrap 68% interval on the period is [12.88, 13.30] mT at 2 K and [10.70, 11.00] mT at 5 K.

The field scale excludes a collective origin. For a 4.5-nm magnetite core the collective moment corresponds to $S \approx 1.3 \times 10^3$ and the anisotropy barrier, taken from the $K_{\text{eff}} = 1.95 \times 10^4 \text{ J m}^{-3}$ fitted to the remanence decay in Fig. 1d, is $U/k_B = 69 \text{ K}$; with $U = D_{\text{coll}} S^2$ the crossing spacing is near 0.03 mT, more than two orders of magnitude below the observation and strongly volume-dependent. The conclusion does not rest on that value: for a collective spin to give the observed spacing the blocking temperature would have to be near 1000 K, above the Curie temperature of magnetite, so the particle would cease to be ferromagnetic before it could block (Supplementary Note S2). A spacing of order 10 mT can only come from a unit with S of order unity, and the composition then fixes the species: the material contains no cation but iron, Fe^{2+} has integer spin $S = 2$ with no Kramers protection, and the only $S = 5/2$ species available is Fe^{3+} . **Neither step uses a literature value of D .**

No comparable literature value exists in principle. Every published D for Fe^{3+} is measured on an ion dilutely substituted in a diamagnetic host, the condition under which a single-ion splitting is defined and measurable; the ion here sits in an exchange-coupled lattice, where exchange averaging removes the fine structure and the same observable is not available. The range spanned by Fe^{3+} in octahedral oxygen coordination, and why agreement at this level would not identify which term of the ${}^6S_{5/2}$ spin Hamiltonian is being compared, are set out in Supplementary Note S11. Only D enters the crossing condition.

Exchange bias and core coupling. A surface ion experiences a longitudinal bias from its exchange coupling to the core^{16,17}, so the crossing appears in applied field at $\mu_0 H_n = nD/(g\mu_B) - \mu_0 H_{\text{bias}}$. H_{bias} follows the core magnetisation, which below blocking depends on H_r , which is what displaces the locus with H_r ; above blocking the core retains no memory of H_r and the locus is vertical. Only sites with small longitudinal bias place their crossings in the accessible range. This is not a claim that the mean surface exchange field is small. A fully coordinated site experiences a molecular field of order 10^2 T , consistent with the 850 K ordering temperature

of magnetite¹⁸. The requirement is only that the disorder distribution carries non-zero weight near zero. In the bulk-inverted spinel structure, the (111) facet termination can host antiferromagnetically coupled B-site pairs or geometrically frustrated coordination where the net local exchange field projects near zero onto the anisotropy axis (Supplementary Note S3). Both the small compensation ratio ($\sim 5 \times 10^{-4}$) and the narrow bias dispersion ($\sigma_{\text{bias}} \lesssim 3$ mT) are direct empirical constraints required for the phase-coherent comb to survive ensemble averaging. These rare sites carry the ridges; the broadly distributed majority contributes only to the smooth background that the second derivative suppresses. The spacing is independent of H_{bias} , which cancels between consecutive crossings. Constructing the locus on this reading reproduces the model-free result: the horizontal excursion of the integrated cross-correlation across the H_r window is 50.8 mT at 2 K and 31.4 mT at 5 K, against $\lambda \Delta M_{\text{rem}} = 53.6$ and 31.4 mT for the construction (Fig. 4c). The correlation between tilt and remanence slope is significant at both temperatures, $\rho = 0.78$ and 0.76. What is not constant is the constant of proportionality, 31.9 at 2 K against 20.6 at 5 K, so we do not treat λ as a temperature-independent quantity.

The nature of the process. The single-ion barrier implied by the measurement is $DS^2 \approx 0.11$ K, far below the measurement temperatures: $k_B T / DS^2 \approx 18\text{--}45$, all six m states are thermally accessible, and the ion is not the metastable two-well system of a molecular nanomagnet. The process therefore cannot be coherent Landau–Zener tunnelling. For half-integer S, time-reversal symmetry forbids tunnelling by the crystal-field terms alone¹⁹, and the gap opened by hyperfine and dipolar perturbations is far too small to support coherent passage before thermal decoherence. What survives is an incoherent resonant enhancement of the transition rate at the degeneracy fields (in the relaxation dynamics sense, as opposed to coherent quantum passage), where otherwise energy-non-conserving transitions become allowed, the mechanism established in dilute rare-earth systems measured with stepped protocols^{10,11}.

What makes a single ion observable. The single ion acts as a trigger; the macroscopic core provides the amplifier. An individual Fe^{3+} spin lacks the energy to reverse the $\sim 10^3$ -spin core, falling short by three orders of magnitude. Instead, it modulates a switching process that is already critically poised. When the applied field brings the core to the brink of thermal reversal, its escape rate²⁰ depends exponentially on the remaining energy barrier (Fig. 4a). A resonant

acceleration of the surface ion's spin relaxation at the level-crossing fields H_n perturbs this barrier, translating a sub-kelvin quantum crossing into a macroscopic modulation of the core switching rate. The mixed second derivative then extracts this sharp, periodic rate modulation while discarding the broad classical distribution beneath it.

That mechanism is available only to cores at the margin of blocking, which makes it a two-sided condition on size (Fig. 4d). Below the lower edge the core is superparamagnetic: it reverses freely on the measurement timescale and carries no irreversible switching for the resonance to modulate. Above the upper edge the field at which the core first becomes thermally switchable lies beyond the highest tooth of the resonance comb, and its escape rate there is suppressed by a factor that is exponential, not gradual. The window is 4.1–8.5 nm at 2 K, holding 83% of the moment of this specimen, and 5.5–11.5 nm at 5 K, holding 62%. At 50 K it is 11.9–24.8 nm. The window has not closed, but it has slid past the distribution, and 1.9% of the moment remains inside. This, and not any failure of the resonance condition, is why the ridges are absent at 50 K. The condition itself is electronic and present at every temperature.

The Kramers crossing theorem for an axial $S = 5/2$ spin restricts the fundamental crossing index to $n \leq 2S - 1 = 4$, predicting that the primary resonance comb terminates near $4\Delta H \approx 50$ mT (Supplementary Note S6.8). At the present field resolution, whether the resolved profile represents four single-sign crossings or two bipolar lobes cannot be formally distinguished (second-harmonic amplitude 0.018 versus non-harmonic median 0.037; bootstrap phase interval $[-131^\circ, +97^\circ]$). Residual extrema observed at 59.6 and 78.0 mT persist under sixfold variations of background filtering width without shifting (Extended Data Fig. 1); these higher-field features may reflect higher-order transverse crystal-field mixing or secondary surface environments, which higher-density grid acquisitions can systematically resolve.

Discussion

A testable pattern across nanomaterials. The size window turns that single corroboration into a testable pattern. Because T_B is proportional to volume the window is equivalent to $T_B/9.0 \leq T \leq T_B$, which assigns each preparation its own temperature and removes the diameter from the comparison; the product of window occupancy and per-core modulation (falling as $1/T$) peaks

between $T_B/6$ and $T_B/5$. Of the seven preparations of ref. 15, exactly one was measured near its own optimum — the 7.6 nm sample at 2 K against an optimum of 2.2 K — and it is the one in which the tilt is recovered (Fig. 5). The other six sit 7 to 21 K from their optima, where eighteen to two hundred times the integration would be needed, so a null at 2 K is a limit of detection and not an absence. Two conditions govern whether they would be seen: sampling density, which is hard and which averaging cannot repair, and sensitivity, which must be bought with sample rather than with time.

The window slides with temperature but does not widen: its width in $\ln d$ is the same at every temperature, so the fraction of moment an optimally sized specimen could place inside it is a constant, while the modulation each core contributes falls as $1/T$. Whatever the particle size, the periodic component available at 50 K is bounded at about a twenty-fifth of that available at 2 K. That ceiling, and not the absence of the effect, is what has kept the structure out of conventional low-temperature magnetometry (Supplementary Note S10).

One structural caveat bears on the prediction. The coarser preparations that place the window near 50 K are large enough for the Verwey transition to return, and charge order there could distribute D. The transition is suppressed at the sizes studied here and none appears in the present material (Extended Data Table 1, row 17); what is untested is the charge-ordered state itself. Supplementary Note S10 sets out both routes, the preparations and temperatures the account prescribes, and what each would cost.

Open questions. Two targeted questions define natural avenues for future work. First, the temperature independence of ΔH is constrained but not yet definitively established: the two determinations at 2 K and 5 K agree to within 19%, and the paired resampling test returns no significant difference (15.1%), but the bootstrap distributions overlap only in their tails. A third determination at an intermediate temperature would be required to settle the constancy—a measurement that remains for future work. Second, the fine waveform remains to be resolved. At the current ~ 5 samples per period, four single-sign crossings cannot be formally distinguished from the lobes of two; higher-density sampling grids with 15–20 samples per period represent a natural next step to resolve the harmonic structure and test the predicted termination at $n = 4$. Beyond this system, the method is general and inexpensive: row-to-row cross-correlation mea-

sures ridge geometry directly, and the geometry separates structures that follow the hysteretic state from those locked to the applied field. Our own control temperatures show why that separation is worth making: there the vertical feature alone reaches high coherence, and a periodicity would have been claimed from the statistic had the geometry not been measured. We suggest the test be applied wherever periodic structure is reported in a FORC distribution.

Methods

Sample

The magnetoferritin (M-HFn) studied here is the material prepared and characterised in ref. 12, to which the reader is referred for the full characterisation; no structural characterisation is repeated in the present work. Magnetite cores were formed by biomimetic synthesis within recombinant human heavy-chain ferritin cages. Transmission electron microscopy (JEOL JEM2100, 200 kV) gives a mean core size of 5.3 ± 1.2 nm over $N = 509$ particles, where size is defined as $(L + W)/2$ from the major and minor axes of each crystal; crystallinity was confirmed by high-resolution electron microscopy and electron diffraction. FORC analysis of the same material gives a magnetostatic interaction field $H_{b,1/2} = 2.1$ mT, six times smaller than the periodicity reported here, confirming that the protein shell suppresses inter-particle coupling. Throughout, the magnetic diameter and every quantity derived from the core volume, the collective spin S and the barrier $U = K_{\text{eff}} V$, use the bulk saturation magnetisation of magnetite, $M_s = 4.8 \times 10^5$ A m⁻¹ (480 emu cm⁻³). The Langevin inversion constrains the moment $\mu = M_s V$, so the diameter it returns scales as $M_s^{-1/3}$; the value of M_s must therefore be the same wherever a volume is used, and it is stated here for that reason.

Magnetic measurements

All measurements used a Quantum Design MPMS-XL SQUID magnetometer. FORC series were acquired with the variable-resolution protocol of ref. 13, which places field steps according to the major hysteresis properties and therefore concentrates points where the FORC distribution

carries structure. Acquisition parameters were read from the raw instrument records and not from the sequence files:

Parameter	2 K	5 K
reversal curves	77	76
median field step	25.2 Oe (2.52 mT)	27.1 Oe (2.71 mT)
field range	± 5022 Oe (± 502 mT)	± 12559 Oe (± 1256 mT)
time per field node, FORC	23.29 s	—
time per field node, remFORC	52.60 s	—
total acquisition, FORC / remFORC	29 h / 66 h	—

The standard FORC and remanence FORC sequences visit the same 4 529 field values in the same order; comparing the two raw records point by point gives a maximum difference of 0.05 mT, which is the field regulation noise. The remFORC protocol measures each node twice, once in field and once after returning the field to zero, which is the origin of the 2.258-fold difference in time per node exploited in the dwell-time test. The configuration is recorded in two different places, and the distinction matters for the auto-range test. The data file lists regression factor 1.825, algorithm 3, amplitude 4 cm, frequency 1 Hz, 5 cycles, 1 scan per measurement and error flag 0, all invariant across every point of both runs. It carries no field for the SQUID range or gain: the software writes those to a separate diagnostic file, which was not retained for these acquisitions. A diagnostic file recorded on the same specimen and instrument three weeks earlier shows the range and gain changing during a single measurement and tracking the magnitude of the moment, the behaviour an auto-range artefact would require. The sequence file issues a byte-identical measurement command at all 4 529 points, so any such change is the instrument's own and not a scripted one. The log therefore cannot exclude a range change, which is why the exclusion rests on the coordinate-mapping tests of Extended Data Fig. 2.

FORC distributions were computed by the moving-polynomial method of ref. 7 with smoothing factor $SF = 2$, robust fitting and a 5 Oe node buffer, on the native non-uniform grid; no interpolation onto a regular grid is performed at any stage.

Measurements used the variable-resolution acquisition protocol of ref. 13, which places the field steps unevenly at measurement time, with the step distribution specified from the major hysteresis loop. Resolution is thereby concentrated in the low-field region where the ridges lie:

the median step there is 2.21 mT, against 9.32 mT for a regular grid of the same total number of points. Processing-side optimisation of the smoothing²¹ is an alternative route to the same trade-off. The gain is in sampling density per unit of instrument time, not in anything a regular grid is barred from reaching: matching the 2.21 mT step on a regular grid requires 4.2 times the nodes, which at 23.29 s per node is five days per acquisition rather than 29 h. What that costs in practice is visible in the external comparison of the main text, where a series acquired on a regular grid at a density that was practical for its authors delivers 3.3 samples per period against our 5.2, and does not constrain the period; that is a statement about that dataset, not about regular grids.

The protocol parameters were fixed before the data existed and are recorded in the instrument sequence file, so no choice made after the FORC distribution was visible can have shaped the sampling. The principal risk to a claim of this kind is that the analysis manufactured the feature, and the record of unsuccessful approaches in Supplementary Note S6 shows that risk to be real. The same reasoning governs the two other design choices made here, fixing the ridge geometry by an independent measurement rather than by fitting, and using an analytic null rather than a surrogate construction.

Decomposition into remanent, induced and transient components followed ref. 14. Two of the four components are differences of the others; only three of the measurement operations are physically distinct, the standard sequence and the in-field and zero-field points of the remanence sequence, and only those three are used in the coherence analysis.

Residual isolation

The smooth switching distribution was removed by subtracting a Gaussian-smoothed copy: $\rho_{\text{resid}} = \rho - G_{\sigma} * \rho$ with $\sigma = 4.5$ grid nodes, applied identically at every temperature and to every component. All analysis is performed on ρ_{resid} .

The size of the periodic component quoted in the main text was obtained by removing a Savitzky–Golay baseline from each raw reversal curve over $H \in [0, 90]$ mT and taking the root-mean-square residual as a fraction of the switching amplitude. This figure depends on the width chosen for the baseline: varying it by a factor of three moves the result between 0.1 and

0.7%, which is why it is quoted only as sub-percent. The enhancement factor is more stable, because the same choice affects numerator and denominator: computing the ratio of residual to smooth part with matched definitions in the magnetisation and in ρ gives 29 at 2 K and 37 at 5 K, against the $(W/\Delta H)^2 = 29$ expected on dimensional grounds from the measured switching-peak width (Supplementary Note S9.4).

Direct measurement of the ridge locus

For each pair of adjacent reversal curves, the residual profiles in H were interpolated onto a common 0.5 mT grid over $H \in [0, 90]$ mT, normalised to zero mean and unit variance, and cross-correlated over displacements up to ± 14 mT in 0.5 mT steps. The displacement maximising the correlation, divided by the H_r interval between the two curves, gives the local slope of the ridge locus. An exchange bias that follows the core magnetisation moves the ridge to lower applied field as H_r becomes shallower, that is $dH_n/dH_r < 0$; the tilt is therefore reported throughout as $-dH_n/dH_r$, so that a positive value denotes the predicted sense and zero a vertical locus. Pairs whose maximum fell at the search boundary, and pairs separated by less than 0.5 mT in H_r , were discarded. The locus is the trapezoidal integral of the resulting slope sequence. This procedure involves no template, no model of the ridge shape, and no free parameters. Its null behaviour is given by the 50 K and 300 K data. At 50 K the median displacement is 0.000 with an interquartile range of 0.000 and a mean of $+0.17 \pm 0.17$ (s.e.m.); at 300 K the median is 0.000 with an interquartile range of $[-0.16, +0.13]$ and a mean of -0.05 ± 0.14 . Neither differs from zero. Individual pairs scatter, the pair-to-pair standard deviation being 0.75 at 50 K, so the statistic used throughout is the median.

Why the locus is measured rather than fitted

An earlier version of this analysis scanned the locus and the period jointly over a grid of a linear and a quadratic coefficient, and corrected the resulting significance by the number of grid points. That correction is wrong, and by a large factor. The coherence statistic is far more sensitive to the locus parameters than the grid spacing suggests: a change of 0.007 in the quadratic coefficient displaces the projection by a quarter of a period at the far end of the analysis window, whereas

the grid step used was 0.15. The grid was therefore coarser than the correlation length of the statistic in that direction, and the effective number of independent trials is of order 10^5 – 10^6 , not the 3.2×10^4 grid points. Corrected on that basis, the significance obtained from the joint scan is not meaningful.

Measuring the locus first removes the problem rather than correcting for it. With the geometry fixed by an independent measurement, the only free parameter is ΔH , and the search over it covers approximately 51 independent values across the scanned range, with a factor of two for the sign convention of the locus integration.

Coherence statistic

Within the analysis window ($H_r \in [-200, -20]$ mT, $H \in [-50, 150]$ mT), the reversal curves of each channel were divided into contiguous blocks. Both block parameters are set by the smoothing factor and neither was tuned:

- **Block width = $2 \cdot \text{SF} + 1 = 5$ curves.** The Pike moving-polynomial fit uses a five-node window at $\text{SF} = 2$, so five curves is the width over which the map carries one independent estimate. Narrower blocks were tried and fail the null calibration below.
- **Separation = $2 \cdot \text{SF} = 4$ discarded curves.** ρ at row r is computed from raw data at rows $r - 2$ to $r + 2$, so two blocks share no raw data only if the first row of one exceeds the last row of the other by more than $2 \cdot \text{SF}$. An earlier version used two discarded curves, at which adjacent blocks share two raw rows.

This gives five blocks per channel at 2 K ($K = 15$ over the three channels), four at 5 K ($K = 12$) and two at 150–300 K ($K = 6$), the last because those datasets contain only 20–22 reversal curves within the window.

For each block, the residual values at the native grid points were projected along the measured locus, mean-subtracted, and their complex Fourier amplitude evaluated at trial period ΔH directly on the irregular sample positions, with no interpolation onto a regular grid at any stage. Each block amplitude was normalised to unit modulus and the K unit vectors summed; R is the modulus of the resultant divided by K .

The reported maximum of $R(\Delta H)$ is the highest **resolved** maximum, defined before inspecting

the result as one spanning more than one grid point at 90% of its height. Single-point spikes are not resolved peaks and are skipped. The rule is stated in terms of peak width alone and makes no reference to period; it matters at 2 K, where a one-point spike at 8.20 mT ($R = 0.922$) would otherwise displace the resolved peak at 13.10 mT, which is 1.00 mT wide across eleven grid points.

For K blocks with independent phases under the null, R follows the Rayleigh distribution and $p = \exp(-KR^2)$ in closed form. This is the reason for adopting an analytic null: three defensible surrogate constructions applied to the same data return $p = 0.92$, 0.02 and 0.67 (Supplementary Note S6), so a surrogate-based significance is determined by the modelling choice and not by the data.

Reported p values are multiplied by 102 to correct for the 51 independent trial periods and the two-fold sign choice.

R is not comparable across temperatures. Under the null R is biased upwards as K falls, so the control datasets, with $K = 6$ against 15, return high R without significance. Compared at matched K , with the 2 K and 5 K data truncated to 20 curves per channel, R above blocking is as high as below it (0.91 to 0.99 against 0.93 to 0.99). Supplementary Note S8 gives the full comparison, and Note S6.8 records the block-count effect that a temperature dependence of R would reflect.

Independence of blocks and of channels

Within a channel, block independence follows by construction from the four-curve separation and the support of the smoothing kernel: no raw measurement enters two blocks.

Across channels it is an assumption, and the analysis depends on it. At this block size no single acquisition reaches significance: the standard series contributes one channel ($K = 5$ at 2 K, $K = 4$ at 5 K) and the remanence series two ($K = 10$ and $K = 8$, $p = 0.068$ and 0.073), so the reported result requires all three channels combined. We state this plainly instead of presenting an assumption-free alternative that the data do not support.

The evidence for the assumption is the null calibration. Estimating the effective number of independent blocks from the off-peak coherence, $K_{\text{eff}} = 1/\langle R^2 \rangle$, returns 6.4 at 300 K against a

nominal 6, so pooling three channels there does not inflate R . The same estimator is unstable at 2 K and 5 K, returning values between 3 and 22 depending on the evaluation band, because the structure itself carries power away from the peak and because $R(\Delta H)$ is strongly autocorrelated in ΔH ; we therefore do not present it as verification at the temperatures of interest. The three channels are not statistically independent in any strong sense, being three views of the same particles, and a referee is entitled to treat $p \approx 10^{-3}$ as an upper bound on the evidence from coherence alone.

This is one reason the claim does not rest on coherence alone. The geometric criterion is measured by cross-correlation on each channel separately, involves no blocks, no template and no free parameter, and returns a median of exactly zero above blocking across nineteen adjacent curve pairs.

Uncertainty on the period

Uncertainties on ΔH were obtained by resampling the K block phasors with replacement 4000 times and recomputing the resolved coherence maximum. The resulting distribution is not Gaussian: a minority of resamples lock onto the unresolved 8.20 mT spike, which inflates the standard deviation without reflecting the location of the peak. We therefore quote the 68% interval and the fraction of resamples falling in the peak band. At 2 K the interval is [12.88, 13.30] mT with 86% of resamples between 12 and 14 mT and none between 10 and 12; at 5 K it is [10.70, 11.00] mT with 90% between 10 and 12 mT and 3% between 12 and 14.

Resampling both temperatures gives a difference of +2.30 mT, with 15.1% of paired resamples returning a 2 K period no larger than the 5 K one. The difference is therefore not significant, but the two distributions overlap only in their tails.

The interval quoted on $\Delta H = 11.9 \pm 1.6$ mT, and hence on D , is the standard deviation of the two determinations and is the spread between temperatures, not the precision of either. Two values drawn independently at random from the searched range of 8–26 mT would agree as closely as the measured pair 19% of the time; the agreement therefore constrains temperature independence without establishing it.

What the contributing sites must satisfy

A site that carries the ridges has to meet two conditions at once. Its longitudinal exchange bias must be small enough to place the crossings inside the accessible field range, since the crossing appears at $\mu_0 H_n = nD/(g\mu_B) - \mu_0 H_{\text{bias}}$ and a fully coordinated site experiences a molecular field of order 10^2 T. And its axial crystal field must be weak enough to give the D inferred here. This condition is not an additional ad hoc assumption: it is quantified directly from the data via the locus excursion (50.8 mT across the H_r window at 2 K against a $\sim 10^5$ mT bulk molecular field, establishing a compensation ratio of $\sim 5 \times 10^{-4}$) and the survival of the periodicity under ensemble averaging (which requires $\sigma_{\text{bias}} \lesssim 3$ mT, Supplementary Note S3). In the bulk-inverted spinel structure, the (111) facet termination can host antiferromagnetically coupled B-site pairs or geometrically frustrated coordination where the net local exchange field projects near zero onto the anisotropy axis; an environment whose local symmetry is higher than that of the bulk B site (where the trigonal distortion along $\langle 111 \rangle$ is intrinsic) simultaneously reduces the net exchange vector and the net axial crystal-field component. The physical requirement is not that the local exchange field vanishes identically everywhere, but that the distribution of its longitudinal projection has finite weight near zero. We present this as a physically grounded candidate model supported by the data constraints; the two sums over magnetic neighbours and ligand charges are correlated through the local geometry but are not the same sum. What can be quantified is the residual coupling, and the tilt measures it. The condition is not that H_{bias} vanishes but that it is small, and because H_{bias} follows the core magnetisation the residual is what displaces the locus with H_r . The order of magnitude of the molecular field is an estimate, not a measurement, so the 5×10^{-4} ratio is an order-of-magnitude statement. Both numbers (the compensation ratio from the locus excursion and $\sigma_{\text{bias}} \lesssim 3$ mT from the survival of the periodicity) are direct inferences from the present dataset, and neither requires a further experiment.

The remanence decay and the size distribution

The curve shown in Fig. 1d is the thermal decay of the saturation remanence, $M_{\text{rs}}(T)/M_{\text{rs}}(2\text{ K})$, measured on the specimen used for the FORC work after saturation in 2.5 T at 2 K; the field

is zero throughout the warming. This is a direct measurement and involves no model. It replaces the $\Delta M(T)/\Delta M(4.9 \text{ K})$ of an earlier zero-field-cooled/field-cooled measurement used in previous versions, which was made on a different specimen and does not extend below 4.9 K. The two observables carry the same volume weighting, which is why the substitution does not change what is being measured: a blocked particle contributes 0.5μ to the saturation remanence and $\mu \tanh(\mu B/k_B T_B)$ to a low-field thermoremanence, and with $\mu = M_s V$ and $T_B = K_{\text{eff}} V/(k_B \ln(\tau_m/\tau_0))$ the argument $\mu B/k_B T_B = M_s B \ln(\tau_m/\tau_0)/K_{\text{eff}}$ carries no dependence on size at all (0.62 at $B = 1.0 \text{ mT}$). The hyperbolic tangent is therefore a common factor and both observables are exactly proportional to V . The comparison curve drawn in Fig. 1d is obtained from the Langevin size distribution of Fig. 1a with a single free parameter $K_{\text{eff}} = 1.95 \times 10^4 \text{ J m}^{-3}$, 1.4 times the bulk cubic anisotropy of magnetite; every use of K_{eff} in this paper refers to that fit. The saturation remanence is carried only by the particles that are blocked at the base temperature, so the model curve is normalised at 2.00 K exactly as the data are and both begin at unity.

No truncation of the size distribution is applied, and on this specimen the untruncated fit is the better one, each form taken at its own optimum with the model curve anchored at base temperature: r.m.s. residual 0.007 against 0.026 with a cut at the 8 nm ferritin cavity. Minimising the residual on the untruncated form alone would return $K_{\text{eff}} = 1.7 \times 10^4 \text{ J m}^{-3}$ rather than the 1.95×10^4 retained here, which leaves a residual of 0.015. All three figures are on the same anchored basis. We retain the larger value because it is the one the switching-field distribution supports: $\rho(H, H_r)$ is itself a distribution of switching fields, and superparamagnetic cores contribute no irreversibility to it, so the fraction of the blocked moment that switches at a field the comb reaches can be read from the data without K_{eff} , M_s , κ or a diameter. That fraction is 52% at 2 K (46–62% over the 48.8–59.6 mT range of the top tooth), which corresponds to $\mu_0 H_K \approx 85 \text{ mT}$; $K_{\text{eff}} = 1.95 \times 10^4 \text{ J m}^{-3}$ gives 81 mT and 1.7×10^4 gives 69 mT, at which the whole blocked moment would lie inside the window. On the untruncated distribution 23% of the moment lies above 8 nm, which exceeds the cavity that confines the cores, and we do not read it as cores of that size. The magnetic diameter is an effective quantity, not a physical dimension: it is obtained by assigning one value of K_{eff} to every core, so a core whose anisotropy exceeds that

value, or whose blocking is raised by interaction with a neighbour, appears at a larger apparent diameter. Ref. 15 measures K_{eff} distributed over $10\text{--}20 \times 10^3 \text{ J m}^{-3}$ across seven preparations of the silica-coated magnetite of that study, a spread of the magnitude required. Nothing in this paper depends on that assignment. Because $T_B \propto V$, the size window of Fig. 4d is exactly the interval $T_B \in [T, 9.0 \text{ T}]$, and the fraction of moment inside it is $M_{\text{rs}}(T)/M_{\text{rs}}(2 \text{ K}) - M_{\text{rs}}(9.0 \text{ T})/M_{\text{rs}}(2 \text{ K})$, read directly off the measured decay. That gives 82.7%, 63.5% and 1.8% at 2, 5 and 50 K, against 82.9%, 62.4% and 1.9% by the route through diameters, a difference of at most 1.1 percentage points, which is the numerical cost of the deconvolution and not a dependence on K_{eff} . The diameter axis of Fig. 4d is a relabelling that is applied identically to the distribution and to the window edges, and it cancels. What the shares do depend on is the ratio $9.0 = (1 - h_{\text{max}}/H_K)^{-\kappa}$: at $\mu_0 H_K = 70 \text{ mT}$ the 2 K share would be 97% and at 95 mT it would be 63%. But $\mu_0 H_K = 2K_{\text{eff}}/M_s$ is the measured barrier-to-moment ratio $2U/\mu$, with the numerator from the remanence decay and the denominator from the Langevin inversion, and it returns 81 mT for any M_s assumed between 2.3 and $4.8 \times 10^5 \text{ A m}^{-1}$.

Provenance of data reproduced from other work

Two sets of measurements shown here were not made for this study. *Figure 1b, and a correction not applied in the original work.* The zero-field-cooled and field-cooled curves at 1.0 mT are the measurement reported in ref. 12 on this material, replotted from the original data by their author, who is an author of the present work. **They were made in 2017 on a different specimen from the one used for the FORC measurements reported here, which was mounted and measured in 2020.** The two specimens differ: the blocking-temperature distribution of the earlier one has a median near 17 K, that of the present one near 4 K. We therefore take no quantity from Fig. 1b; it is reproduced for continuity with the published characterisation of the material. Panels a, c and d are measurements on the specimen used for this work, and every number quoted in this paper comes from them. Ref. 12 analysed these curves as measured and reported a zero-field-cooled maximum of 35 K. We apply a correction for a trapped residual field, which the original analysis did not, and obtain 31 K. The difference has no bearing on any conclusion of the present paper, since nothing here

depends on the position of that maximum, but the two numbers come from the same data and we set out the reason for the discrepancy instead of leaving it implicit.

The correction is determined by the data and is not assumed. At base temperature the two branches give -3.496×10^{-3} and $+3.552 \times 10^{-3}$ A m², a ratio of -0.9843 . If the low-field thermoremanence is linear in the cooling field, that ratio fixes the residual field at $H_{\text{res}} = r/(1 - r) \times H = -0.496$ mT: the nominal zero-field-cooling took place in -0.496 mT and the field-cooling in $+0.504$ mT, nearly equal and opposite. Each branch is then the sum of an induced response, common to both, and a thermoremanence proportional to its own cooling field and decaying with the blocked fraction $f_b(T)$,

$$M_{\text{ZFC}} = M_{\text{ind}} + k H_{\text{res}} f_b, \quad M_{\text{FC}} = M_{\text{ind}} + k (H + H_{\text{res}}) f_b,$$

so that $\frac{1}{2}(M_{\text{FC}} + M_{\text{ZFC}}) = M_{\text{ind}} + k (H/2 + H_{\text{res}}) f_b$ with $H/2 + H_{\text{res}} = 0.004$ mT. The half-sum therefore cancels the thermoremanence to better than 1% of the applied field and leaves the induced response, which is what a zero-field-cooled curve measures. Its maximum is at 31.0 K against 35.05 K for the uncorrected branch.

Two checks support the correction. The half-sum at base temperature is 2.8×10^{-5} A m², two orders of magnitude smaller than either branch, as required if the thermoremanence has been removed and only the small induced moment of blocked particles remains; this smallness is a consequence of the linear-thermoremanence assumption and not an input to it. And the blocking-temperature distribution, derived from the difference of the branches instead of their sum, is centred at 14.9 K with or without the correction.

Comparison FORC data of ref. 15. The distribution for the 7.6 nm octahedral sample was computed from the original measurement data underlying the corresponding sample of that published work, provided by H. Mamiya. Those measurements were made at 2.0 K on a Quantum Design MPMS-XL, the same temperature and instrument model as ours. In that work FORC analysis was used solely to establish that the silica-coated particles are non-interacting; the periodic structure discussed here was not reported or examined.

Statistics and reproducibility

No statistical methods were used to predetermine sample size, the experiments were not randomised, and no formal blinding was applied. Blinding in its clinical sense does not map onto a magnetometry measurement on a single specimen. What serves its purpose here is that the choices which could have shaped the result were fixed before the result was visible. The acquisition protocol and its step distribution were specified from the major hysteresis loop and written into the instrument sequence file before any FORC distribution existed. The smoothing factor and the background-subtraction width were fixed a priori, and the block scheme follows from the smoothing factor and is not chosen. In the analysis the ridge locus is measured by cross-correlation and not fitted, the null is analytic and not constructed from surrogates, and peak-aligned stacking was refused because it is self-confirming. One choice was not pre-specified: the 2–58 mT range over which the profile extrema of Extended Data Fig. 5 are counted, and the effect of relaxing it is given in that legend. All p values are two-sided where applicable and are corrected for the parameter search as described above. Analyses that failed, and the reasons, are recorded in Supplementary Note S6 and Extended Data Fig. 8.

Data availability

All raw instrument records, processed FORC matrices, and derived data supporting the findings of this study are provided in the Supplementary Materials and will be deposited on Zenodo with a permanent DOI prior to publication.

Code availability

All analysis code and processing scripts are provided in the Supplementary Materials and will be deposited with the data on Zenodo prior to publication.

Materials availability

Magnetoferritin materials can be synthesized reproducibly using the biomimetic synthesis route described in ref. 12.

Acknowledgements

We thank H. Mamiya for providing the original first-order reversal curve data underlying the silica-coated magnetite sample of ref. 15, used here for the independent comparison. This work was supported by the Australian Research Council through Discovery Project DP200100765 (X.Z.) and by the National Natural Science Foundation of China through grants 42374089 (S.L.) and 42388101 (T.Z.).

Author contributions

T.Z. designed and carried out the biomimetic synthesis of the magnetoferritin, characterised it, and provided the earlier zero-field-cooled and field-cooled data reproduced in Fig. 1b. S.L. contributed to the design of the study and performed all of the low-temperature measurements and data acquisition, on which every result reported here depends. X.Z. led the design of the study and the analysis, developed the measurement and analysis strategy, implemented the analysis, and wrote the manuscript. The three roles are distinct and none is substitutable: without the material there is nothing to measure, without the low-temperature data there is no observation, and without the analysis strategy the structure is not resolvable at all. The study could not have been carried out without any one of the three, and the three authors are accordingly listed as having contributed equally. All authors contributed to the draft.

Competing interests

The authors declare no competing interests.

Figure legends




Fig. 1 | Classical characterisation. **a**, Room-temperature magnetisation with a lognormal–Langevin fit; inset, the fitted size distribution (median 4.53 nm, mean 4.79 nm, mode 4.05 nm). **b**, Zero-field-cooled and field-cooled curves at 1.0 mT, from the measurement of ref. 12 on this material. Faint lines, the two branches as measured and as analysed in that work; bold line, the half-sum $\frac{1}{2}(\text{FC} + \text{ZFC})$, in which the two thermoremanences cancel and the induced response survives, which is what a zero-field-cooled curve measures. The residual-field correction is derived in Methods. **c**, In-phase AC susceptibility at 1 Hz and 1 kHz, peaking at 45 K and 57 K. **d**, Thermal decay of saturation remanence acquired in 2.5 T after field-cooled and zero-field-cooled preparation; the field is zero throughout the warming, and the two preparations give normalised curves agreeing to better than 1%. Two thirds of the remanence survives to 5 K and

2.0% to 50 K. The single free parameter $K_{\text{eff}} = 1.95 \times 10^4 \text{ J m}^{-3}$ used elsewhere in this paper is fitted here (Methods). Inset, the same measurement as the cumulative unblocked fraction $1 - M_{\text{rs}}(T)/M_{\text{rs}}(2 \text{ K})$: half the remanence is lost by 8 K. No derivative is taken. Panels **a**, **c** and **d** are measurements on the specimen used for the FORC work; panel **b** is an earlier measurement on a different specimen and no number is taken from it (Methods). No anomaly appears in these data, and no claim rests on this figure.




Fig. 2 | FORC distributions and residuals. Top, FORC distributions at 2 K, 5 K and 50 K; bottom, the residuals after removal of the smooth background, each normalised to its own maximum. At 2 K and 5 K the residual carries curved, sign-alternating ridges whose inclination increases towards small $|H_r|$. At 50 K the ridges are absent and the residual contains only the vertical superparamagnetic feature at the field zero crossing.




Fig. 3 | Measurement of the ridge geometry and of the periodicity. **a**, Row-to-row displacement $-dH_n/dH_r$ from the cross-correlation of adjacent reversal curves, with no fitting and no template; the sign convention is set in Methods, positive denoting the sense the mechanism predicts. Points are individual curve pairs and the lines a running median over seven pairs. At 2 K and 5 K the displacement is non-zero and grows as $H_r \rightarrow 0$, where the descending branch is steepest; at 50 K the median and the interquartile range are both exactly 0.000 over nineteen curve pairs, with 63% of pairs returning exactly zero; individual pairs scatter about that median (pair-to-pair s.d. 0.74, one outlier off scale). **b**, Phase coherence R against trial period, evaluated along the measured locus with ΔH as the only free parameter. Shading marks 11.9 ± 1.6 mT. The resolved maxima are at 13.10 mT (2 K) and 10.80 mT (5 K); the one-point spike at 8.20 mT is not a resolved peak (Methods). R above blocking is comparably high but rests on $K = 6$, and R is not comparable across different K . The criterion applied requires coherence **and** a tilted locus, and it is the locus that separates the regimes. **c**, Tilt and remanence slope dM_{rem}/dH_r against temperature; the tilt vanishes where the remanence slope reverses sign.




Fig. 4 | Interpretation. **a**, Switched fraction $f = 1 - \exp(-\Gamma t)$ and its sensitivity to a barrier perturbation, $\Gamma t \exp(-\Gamma t)$, peaking at $\Gamma t = 1$ with the value $1/e$, Γ being evaluated at the resonance field. A perturbation δU shifts it by $\delta U / (ek_B T)$, 0.29% per core at 2 K. The two ends are the edges of the window in **d**. Inset, the geometry. **b**, Level crossings for $S = 5/2$ with $D = 0.016$ K and no transverse term, at $nD/g\mu_B$, spacing 11.9 mT. Dots above the axis count the pairs crossing at each tooth, two at $n = 1$ and 2 and one at $n = 3$ and 4: a property of the level structure, not a prediction of amplitude. No gap is extracted. **c**, Lines, the locus $H_n = nD/g\mu_B - \lambda M_{\text{rem}}(H_r)$ built from the measured remanence branch. Points, the same locus from integrating the row-to-row cross-correlation, which uses no model. Cross-correlation returns a displacement, so this fixes the tilt, not the absolute field (Extended Data Fig. 9). Shading, the analysis window. **d**, The size window at three temperatures; edges in Supplementary Note S10, thin bars the 48.8–59.6 mT range of the top tooth. Grey, the size distribution from the remanence decay of Fig. 1d, hatched below 4.07 nm. Filled, the moment inside the 2 K window: 83%,

against 62% at 5 K and 1.9% at 50 K, where the window has slid past the distribution. The distribution is measured and the window derived.




Fig. 5 | The kinetic window over the size–temperature plane, and what a core is worth inside it. Band, the two-sided window, with edges $d_B(T)$ and $2.08 d_B(T)$ derived in Supplementary Note S10; both scale as $T^{1/3}$, so the window translates without changing width in $\ln d$. Dotted, the logarithmic centre $d = 5.9 (T/2)^{1/3}$ nm, the size best matched to each temperature. Colour, the modulation a single core contributes, $\delta U / (e k_B T)$, which depends on temperature alone and not on size; the band therefore dims as it climbs and carries no vertical variation. **The colour is per core and does not include the fraction of the distribution lying inside the window: for a fixed specimen the two fall together, and their product drops by 1083 between 2 and 50 K where the colour alone drops by 25.** Teal, the seven preparations of ref. 15 at their magnetic diameters, each drawn across the temperature interval over which it lies inside the window; O7 is the only one measured near its own optimum (2.2 K), and it is the one in which the tilt is recovered. Left panel, two size distributions on the same axis: this specimen in blue, from the measured remanence decay, and O7 in teal, from the lognormal of ref. 15. This specimen intersects the band at 2 K and has left it by 50 K. Right-hand axis, the equivalent reading $T_B \in [T, 9.0 T]$, which involves no fitted K_{eff} or M_s .

References

1. Awschalom, D. D., Smyth, J. F., Grinstein, G., DiVincenzo, D. P. & Loss, D. Macroscopic quantum tunneling in magnetic proteins. *Phys. Rev. Lett.* 68, 3092–3095 (1992). □□: *Phys. Rev. Lett.* 71, 4279 (1993)
2. Gider, S., Awschalom, D. D., Douglas, T., Mann, S. & Chaparala, M. Classical and quantum magnetic phenomena in natural and artificial ferritin proteins. *Science* 268, 77–80 (1995). doi:10.1126/science.7701343
3. Tejada, J., Zhang, X. X., del Barco, E., Hernández, J. M. & Chudnovsky, E. M. Macroscopic resonant tunneling of magnetization in ferritin. *Phys. Rev. Lett.* 79, 1754–1757 (1997).
4. Garg, A. Does macroscopic quantum coherence occur in ferritin? *Science* 272, 424–425 (1996). doi:10.1126/science.272.5260.424-b
5. Hanson, M., Johansson, C. & Mørup, S. Comment on "Macroscopic resonant tunneling of magnetization in ferritin". *Phys. Rev. Lett.* 81, 735 (1998).
6. Tejada, J., Zhang, X. X., del Barco, E., Hernández, J. M. & Chudnovsky, E. M. Tejada et al. Reply. *Phys. Rev. Lett.* 81, 736 (1998).
7. Pike, C. R., Roberts, A. P. & Verosub, K. L. Characterizing interactions in fine magnetic particle systems using first order reversal curves. *J. Appl. Phys.* 85, 6660–6667 (1999). doi:10.1063/1.370176
8. Roberts, A. P., Heslop, D., Zhao, X. & Pike, C. R. Understanding fine magnetic particle systems through use of first-order reversal curve diagrams. *Rev. Geophys.* 52, 557–602 (2014). doi:10.1002/2014RG000462
9. Béron, F., Novak, M. A., Vaz, M. G. F., Guedes, G. P., Knobel, M., Caldeira, A. & Pirota, K. R. Macroscopic quantum tunneling of magnetization explored by quantum-first-order reversal curves. *Appl. Phys. Lett.* 103, 052407 (2013). doi:10.1063/1.4816131
10. Giraud, R., Wernsdorfer, W., Tkachuk, A. M., Mailly, D. & Barbara, B. Nuclear spin driven quantum relaxation in $\text{LiY}_{0.998}\text{Ho}_{0.002}\text{F}_4$. *Phys. Rev. Lett.* 87, 057203 (2001).

11. Malkin, B. Z., Vanyunin, M. V., Barbara, B. & Bertaina, S. Relaxation rates of magnetization in $\text{LiYF}_4\text{:Ho}^{3+}$ crystals. *J. Alloys Compd.* 451, 473–476 (2008). doi:10.1016/j.jallcom.2007.04.193
12. Zhang, T. & Pan, Y. Constraining the magnetic properties of ultrafine- and fine-grained biogenic magnetite. *Earth Planets Space* 70, 206 (2018). doi:10.1186/s40623-018-0978-2
13. Zhao, X., Heslop, D. & Roberts, A. P. A protocol for variable-resolution first-order reversal curve measurements. *Geochem. Geophys. Geosyst.* 16, 1364–1377 (2015). doi:10.1002/2014GC005680
14. Zhao, X., Roberts, A. P., Heslop, D., Paterson, G. A., Li, Y. & Li, J. Magnetic domain state diagnosis using hysteresis reversal curves. *J. Geophys. Res. Solid Earth* 122, 4767–4789 (2017). doi:10.1002/2016JB013683
15. Mamiya, H., Fukumoto, H., Cuya Huaman, J. L., Suzuki, K., Miyamura, H. & Balachandran, J. Estimation of magnetic anisotropy of individual magnetite nanoparticles for magnetic hyperthermia. *ACS Nano* 14, 8421–8432 (2020). doi:10.1021/acsnano.0c02521
16. Meiklejohn, W. H. & Bean, C. P. New magnetic anisotropy. *Phys. Rev.* 102, 1413–1414 (1956). doi:10.1103/PhysRev.102.1413
17. Nogués, J. & Schuller, I. K. Exchange bias. *J. Magn. Magn. Mater.* 192, 203–232 (1999). doi:10.1016/S0304-8853(98)00266-2
18. Dunlop, D. J. & Özdemir, Ö. *Rock Magnetism: Fundamentals and Frontiers* (Cambridge Univ. Press, 1997). doi:10.1017/CBO9780511612794
19. Abragam, A. & Bleaney, B. *Electron Paramagnetic Resonance of Transition Ions* (Clarendon Press, Oxford, 1970).
20. Brown, W. F. Thermal fluctuations of a single-domain particle. *Phys. Rev.* 130, 1677–1686 (1963). doi:10.1103/PhysRev.130.1677
21. Egli, R. VARIFORC: an optimized protocol for calculating non-regular first-order reversal curve (FORC) diagrams. *Glob. Planet. Change* 110, 302–320 (2013). doi:10.1016/j.gloplacha.2013.08.003
22. Krebs, J. J., Stauss, G. H. & Milstein, J. B. EPR of Fe^{3+} and Mn^{2+} in single-crystal

- 769 ZnGa_2O_4 spinel. *Phys. Rev. B* 20, 2586–2587 (1979).
- 770 23. Bleaney, B. & Trenam, R. S. Paramagnetic resonance spectra of some ferric alums,
771 and the nuclear magnetic moment of ^{57}Fe . *Proc. R. Soc. Lond. A* 223, 1–14 (1954).
772 doi:10.1098/rspa.1954.0096
- 773 24. Bleaney, B. & Trenam, R. S. Paramagnetic resonance spectra of some ferric alums,
774 and the nuclear magnetic moment of ^{57}Fe . *Proc. R. Soc. Lond. A* 223, 1–14 (1954).
775 doi:10.1098/rspa.1954.0096

Extended Data

Extended Data Table 1 | Alternative mechanisms and controls

Each alternative carries a quantity the data can measure. The discriminant column gives that quantity; the full evidence, with cross-references, is set out in the notes below the table.

#	Alternative	Predicted signature	Discriminant (measured)	Status
A. Instrumental				
1	Time-locked periodicity	Period $\propto 1/\text{dwell}$; $\times 2.258$ requires a 56% shift	0.8% (13.00 vs 13.10 mT)	Excluded
2	Auto-range / gain switching	Feature follows the recorded moment	246 crests coincide across channels differing by 11.7 memu	Excluded
3	Field overshoot / magnet charging	Vertical in (H, H_r) , independent of history	Tilt +0.476 and +0.277 ; tracks dM_{rem}/dH_r , $\rho = 0.78$ and 0.76	Excluded
4	Flux jumps	Isolated discontinuities, not periodic	Cannot give coherence across 15 blocks in three channels	Excluded
5	SQUID drift / settling	Smooth monotonic bias, no periodicity	Removed by the residual filter; carries no periodic component	Excluded
6	Dependence on measurement protocol	Result changes with field range or node distribution	ΔH 10.80 $\rightarrow$ 10.80 mT, R 0.921 $\rightarrow$ 0.921 on truncation	Excluded
B. Processing				
7	Sampling alias of sub-resolution structure	Apparent period follows the local node spacing	Node spacing 26% , period 1.9% ; an alias requires 139%	Excluded
8	Smoothing-factor artefact	Result changes with SF	CV 0.04 (SF2) against 0.47 (SF3); coherence peaks at SF2	Excluded
9	Self-confirming stacking	Alignment derived from the data manufactures a period	Not used; on surrogates it exceeds the data (0.306 vs 0.266)	Excluded†
10	Rendering artefact	Apparent segmentation where grid spacing changes	Segments vanish under interpolation; R = 0.225 against 0.183	Excluded

#	Alternative	Predicted signature	Discriminant (measured)	Status
11	Look-elsewhere in the geometry search C. Physical alternatives	Significance inflated by fitting the locus	Locus measured, not fitted; only ΔH free, $\times 102$ applied	Excluded†
12	Size sub-populations	Continuous features, no fixed periodicity	Coherence at one period across 15 and 12 independent blocks	Excluded
13	Multi-domain states	Domain-wall signatures	Sub-10 nm magnetite is far below the single-domain threshold	Excluded
14	Inter-particle dipolar interactions	Features depend on packing density	$H_{b,1/2} = 2.1 \text{ mT}$, six times below the spacing; bias cancels	Excluded
15	Spin-glass surface freezing	Broad disordered features, no regular period	No single well-defined period, and no locus following the remanence	Excluded
16	Collective (whole-particle) level crossings	Spacing scales as U/S^2 , strongly volume-dependent	0.03 mT predicted against 11.9 mT observed	Excluded
17	Verwey transition	Features near 100–120 K	None in M-HFn; the structure reported here is at 2 and 5 K	Excluded
18	Superparamagnetic response at the field zero crossing D. Not excluded	Vertical at $H \approx 0$, no H_r dependence	Vertical at $H = 0$; the low-temperature structure is tilted and offset	Separated
19	Waveform: four same-sign crossings vs the lobes of two	Relative phase of the second harmonic	$A_2 = 0.018$ below the 0.037 non-harmonic median; $\Delta\varphi$ CI $[-131^\circ, +97^\circ]$	Undetermined
20	Exact temperature independence of ΔH	Common period at many temperatures	13.10 against 10.80 mT, agreeing to 19% ; resampling 15.1%	Not established

† Rows 9 and 11 are excluded by construction and by design respectively, the procedure in question being one we do not use and the search it would inflate one we do not perform, and not by a measurement. Row 18 is not an artefact but a real response that is identified and kept separate.

Notes to Extended Data Table 1

1. The two sequences visit the same 4 529 field values point by point (all agree to $\leq 0.05 \text{ mT}$) with mean times per node of 23.29 s and 52.60 s, a factor of **2.258**. A time-locked period would

appear at 5.76 mT in the remFORC channels; the coherence there is 0.53 against 0.86 at the measured period, and at $K = 10$ is not significant. The measured periods are 13.00 mT (FORC) and 13.10 mT (remFORC): a shift of **0.8%** against the **56%** a time-locked artefact requires (the main text, "Excluding instrumental and processing origins", ED Fig. 3)

2. The remanence sequence measures every field node twice, in field and after removing the field, so the two channels record different moments at the same node, a median difference of **11.7 memu**, with 85% of nodes differing by more than 5 memu. Yet **246 ridge crests occur at the same (H, H_r) in both channels**: a moment-locked feature cannot appear at the same field in two channels whose moment differs by that much. The data file records no SQUID range or gain field, because the software writes those to a diagnostic file that was not retained for these runs, and a diagnostic file from the same specimen shows range and gain changing with the magnitude of the moment, so the log cannot address auto-ranging in either direction, and the exclusion rests on the two measurements above (ED Fig. 2)

3. Such features are vertical in (H, H_r) . The measured row-to-row displacement is +0.476 at 2 K and +0.277 at 5 K; the structure is tilted. Moreover the tilt varies systematically **within** each dataset, tracking dM_{rem}/dH_r across 41 (2 K) and 34 (5 K) independent row pairs, Spearman $\rho = 0.78$ and 0.76, behaviour no fixed-field artefact can reproduce (the main text, "Measuring the ridge geometry", the main text, "Excluding instrumental and processing origins", ED Fig. 7)

4. Isolated events cannot produce phase coherence across 15 independent blocks drawn from three measurement channels (the main text, "Periodicity along the measured locus")

5. A monotonic bias contributes to the smooth background removed by the residual filter and carries no periodic component

6. Truncating the 5 K dataset to the 2 K field range and reprocessing leaves $\Delta H = 10.80 \rightarrow 10.80$ mT and $R = 0.921 \rightarrow 0.921$, unchanged to the resolution of the period scan (the main text, "Excluding instrumental and processing origins")

7. Restricting to the densely sampled sub-window changes the node spacing by **26%** at 5 K (2.71 $\rightarrow$ 2.01 mT) and the measured period by **1.9%** (10.80 $\rightarrow$ 11.00 mT); an alias would shift by 139%. At 2 K, 12% and **1.5%** against a predicted 71% (the main text, "Excluding instrumental and processing origins", ED Fig. 4)

8. At $H_r = -122.6$ mT the SF2 standard FORC profile gives spacings 6.6 / 6.6 / 6.2 / 6.1 / 6.0 / 6.0 mT (CV **0.04**); at SF3 two extrema merge and the sequence becomes 6.6 / 8.4 / 10.2 / 2.0 / 4.0 mT (CV **0.47**). The statistic must depend on SF, since SF is the bias–variance parameter of the estimator, and the dependence has the expected shape. Holding the block scheme fixed at $K = 15$, the coherence in the 12–14.5 mT band is 0.63 (SF1), **0.88 (SF2)**, 0.83 (SF3) and 0.35 (SF4), while point-to-point fluctuation relative to overall amplitude falls monotonically 0.78 $\rightarrow$ 0.44 $\rightarrow$ 0.27 $\rightarrow$ 0.19. At SF1 noise dominates and the resolved maximum is captured by a spike at the band edge; at SF4 the ridges are gone. **SF = 2 is the maximum of the trade-off**, and was fixed a priori. SF = 2 is used throughout (the main text, "Excluding instrumental and processing origins", ED Fig. 5)

9. Peak-aligned stacking is not used. On phase-randomised surrogates it produces oscillation amplitudes $\geq$ those of the measured data (2 K: 0.306 vs 0.266; 5 K: 0.199 vs 0.155). Only rigid projection along the independently measured locus is used (the main text, "Excluding instrumental and processing origins", Supplementary Note S6.1)

10. Segmentation of the steepest part of the main peak disappears under interpolated shading while the ridges do not; segment boundaries coincide with grid nodes; the region carries no phase coherence ($R = 0.225$ against a random expectation of 0.183 for $K = 30$) (the main text, "Excluding instrumental and processing origins")

11. The locus is measured by row-to-row cross-correlation, with no fitting and no scanning. Only ΔH is free, and the correction covers ~ 51 independent values (the main text, "Measuring the ridge geometry" and "Periodicity along the measured locus", Supplementary Note S6.4)

12. A continuous distribution cannot produce phase coherence at a single period across 15 and 12 independent blocks ($R = 0.879$ and 0.921 ; $p = 9.5 \times 10^{-4}$ and 3.9×10^{-3})

13. Impossible in sub-10 nm magnetite, far below the single-domain threshold

14. FORC analysis of this same material gives a magnetostatic interaction field of $H_{b,1/2} = 2.1$ mT [Zhang & Pan 2018], six times smaller than the measured spacing, because the protein shell enforces core separation. The spacing is in any case set by the crossing condition and is independent of the bias offset, which cancels between consecutive crossings (Supplementary Note S3)

15. A frozen disordered state produces no single well-defined period, and no coherent locus that follows the remanence (the main text, "Exchange bias is the reading we give it")

16. For a 4.5 nm core, $S \approx 1.3 \times 10^3$; with the $K_{\text{eff}} = 1.95 \times 10^4 \text{ J m}^{-3}$ fitted in Fig. 1d the barrier is $U/k_B = 69 \text{ K}$, so $D_{\text{coll}} = U/S^2 = 4.3 \times 10^{-5} \text{ K}$ and the crossing spacing is near 0.03 mT, **two orders of magnitude** below the observation. Independently of K_{eff} , matching the observation would require a barrier of $U/k_B = 2.5 \times 10^4 \text{ K}$ and a blocking temperature near 1000 K, above the Curie temperature of magnetite (the main text, "A single ion, not the particle")

17. No Verwey transition occurs in ultrafine magnetite of this size. In the same material, [Zhang & Pan 2018] observe the transition near 100 K in the coarser magnetosome magnetite of the comparison sample and none in M-HFn. The structure reported here is at 2 and 5 K

18. This is what the residual contains above blocking. It sits at the field zero and is vertical by construction, whereas the low-temperature structure is tilted and lies away from zero (the main text, "Excluding instrumental and processing origins", ED Fig. 6). It is distinct from the near-vertical striping seen at deep H_r at 2 and 5 K, which is the ridge locus itself flattening as the loop saturates (ED Fig. 7)

19. Evaluated along the measured locus at the reported period, the second-harmonic amplitude is $A_2 = 0.018$, below the median amplitude of 0.037 obtained at non-harmonic frequencies; the bootstrap confidence interval on $\Delta\varphi$ is $[-131^\circ, +97^\circ]$. At ~ 5 samples per period the waveform is unresolved (the main text, "The nature of the process", Supplementary Note S6.5)

20. The two periods differ by **2.30 mT** (13.10 against 10.80), agreeing to **19%**. Paired resampling returns a 2 K period no larger than the 5 K one **15.1%** of the time, so the difference is not significant; but the two bootstrap distributions overlap only in their tails, and the sign of the difference is the direction a kinetically set field scale would predict. The observation constrains temperature independence; it does not establish it (the main text, "What these data do not decide")


Independent corroboration of the geometry. The tilt discriminant of row 3 has been applied to an independent dataset: 7.6 nm silica-coated magnetite, measured at 2 K on a Quantum Design MPMS-XL in another laboratory, provided by the authors of ref. 15. It returns +0.68

875 ± 0.10 over 33 curve pairs, non-zero at 6.9σ (Wilcoxon $p = 1.6 \times 10^{-5}$), against $+0.17 \pm 0.17$
876 for our own 50 K control ($p = 0.87$). The geometry therefore reproduces outside this sample
877 and this laboratory. That series is on a regular field grid with 3.3 samples per period and a
878 single measurement channel, so it constrains ΔH only to 9.9 ± 4.7 mT and cannot support the
879 coherence test; no claim about the spacing is made from it (the main text, "What these data do
880 not decide").


881 **Summary.** Seventeen alternatives are excluded, each by a measurement and not by an argument.
882 Two questions remain open and are stated as such: the waveform of the individual resonance,
883 and the exact temperature independence of the spacing. Both would be settled by higher field
884 resolution and a denser temperature series (the main text, "What these data do not decide").

886 Extended Data figures

887 *Structural characterisation is not reproduced here. The magnetoferritin used in this work is*
888 *the M-HFn material characterised by Zhang & Pan, Earth Planets Space 70, 206 (2018),*
889 *doi:10.1186/s40623-018-0978-2, which reports the TEM micrographs, the core-size statistics*
890 *(mean magnetite core diameter 5.3 ± 1.2 nm) and the structural indexing. Methods should cite*
891 *that paper instead of repeating the figures.*



ED Fig. 1 | The four decomposition components at 2 K and 5 K. Residual distributions of the standard, remanent, induced and transient components over a common window. The ridges are present in all eight panels. Each panel is scaled to its own 98th percentile of $|\rho_{\text{res}}|$ because the components differ in absolute amplitude by up to a factor of 2.9; the amplitude of each panel relative to the 2 K standard FORC panel is given at lower right, from 0.85 to 2.45. Only three of the four are physically distinct measurements: the induced component is the difference of the in-field and zero-field points of the remanence sequence, and the transient component is a difference of repeated sweeps, so the coherence test uses three channels, not four (main text, "Periodicity along the measured locus"). The near-vertical striping at $\mu_0 H_r < -100$ mT is the expected behaviour of the ridge locus in the saturated part of the loop, quantified in ED Fig. 7. The near-vertical stripes crossing the main ridge are the rim left by the Gaussian high-pass on the steepest flank of the main switching peak, broken into columns by the non-uniform grid; they carry 4.5% of the variance of the map at 2 K against 24.6% for the locus-following component, they scale with the width of the background subtraction while the period does not, and removing them shifts the measured period by 0.8%. Supplementary Note S9 gives the full accounting.




ED Fig. 2 | Coordinate-mapping test for gain switching. An auto-range or gain-switching

artefact is triggered by the magnitude of the measured moment, so it is locked to M and cannot depend on field except through M . The discriminating test is made in **b** and **c**. **a**, The residual painted onto the reversal curves in (H, M) , for orientation. This projection cannot by itself discriminate: the mechanism proposed here places the ridges on a locus that follows the core magnetisation, and so does a moment-locked artefact, so the two predict row-to-row displacements too close to separate on this projection. The discrimination is made in **b** and **c**. **b**, **c**, The decisive test. The in-field and zero-field channels of the remanence sequence record different magnetisations at the same applied field, because one is measured in field and the other after removing it. 246 ridge crests are located at the same (H, H_r) in both channels; the moment recorded at those crests differs by a median of 11.7 memu (IQR 7.2–15.8), and 85% differ by more than 5 memu. A feature locked to the measured moment cannot appear at the same field in two channels whose moment differs by ~ 12 memu there. Points in **b** would have to lie on the red diagonal for a gain-locked origin; they do not.


Why the test is done this way. One might instead appeal to the instrument log, in which every recorded acquisition parameter (regression factor, algorithm, amplitude, frequency, cycles, scans per measurement, error flag) is invariant across all 4 529 points of the 29 h run. We report that, but do not rest the exclusion on it, and the reason is specific, not precautionary: the data file has no field for the SQUID range or gain at all. The software writes those to a separate diagnostic file, which was not retained for these acquisitions; a diagnostic file recorded on the same specimen and instrument three weeks earlier shows the range and gain changing within a single measurement and tracking the magnitude of the moment. The log is therefore silent on exactly the quantity at issue, and auto-ranging during these runs is likely and not excluded,

which is what makes the tests below necessary, not confirmatory. The two logged variables that do move are the eddy-current compensation flag, whose field-periodic component at the ridge period has normalised amplitude 0.007 against 0.64 for a fully field-locked binary, and the RSO waveform fit quality, whose binned medians vary by 0.15% across the ridge window with no field structure. The tests in this figure are used instead because they do not depend on the log: they exclude any moment-locked effect, recorded or not.



ED Fig. 3 | Dwell-time control. **a**, The field sequences of the standard FORC and remFORC measurements, overlaid: the same 4 529 nodes in the same order, agreeing point by point to ≤ 0.05 mT. **b**, Time per field node from the raw instrument timestamps: 23.29 s for the FORC sequence and 52.60 s for the remFORC sequence, which measures in-field and zero-field points at every node. The ratio is 2.258. **c**, Phase coherence R against trial period, computed separately for the two dwell groups along the same measured locus. The FORC channel peaks at 13.00 mT and the remFORC channels at 13.10 mT, a shift of 0.8%. A structure locked to elapsed time would place the remFORC period at $13.00/2.258 = 5.76$ mT (shaded), where the measured coherence is 0.53 against 0.86 at the peak, not significant at $K = 10$. The scan is extended down to 4 mT for this figure so that the predicted position lies inside the searched range.


954




955

ED Fig. 4 | Sampling and rendering controls. **a, b**, Coherence spectra at 5 K and 2 K computed on the full field window and on the densely sampled sub-window, along the same measured locus. At 5 K the node spacing falls from 2.71 to 2.01 mT (26%) and the period moves from 10.80 to 11.00 mT (1.9%); at 2 K the spacing changes by 12% and the period by 1.5%. If the structure were an alias of sub-resolution modulation, the apparent period would follow the node spacing: the predicted shifts are 139% and 71% respectively. **c, d**, The same region of the 2 K map rendered with flat shading (one cell per node, node positions overlaid) and with interpolated shading. The blocky segmentation of the steep flank is a property of the rendering and disappears under interpolation; the ridges are present in both.


965



ED Fig. 5 | Smoothing factor. **a**, Standard FORC residual profile at $\mu_0 H_r = -122.6$ mT at SF = 2 and SF = 3, with extrema located over $\mu_0 H = 2$ –58 mT. Within that range SF = 2 resolves seven extrema at CV = 0.04, half the period, as expected for sign-alternating ridges; at SF = 3 two merge and the sequence becomes irregular (CV = 0.47). Both spacing sequences are given in note 8 to Extended Data Table 1; the count is unchanged for peak prominences between 0.4 and 0.7. That range is narrower than the 0–90 mT window used elsewhere: over the full window the two coefficients are 0.34 and 0.83, a ratio of 2.4 and not 11.3 (Supplementary Note S9). **b**, Extrema resolved per row over $-165 \leq \mu_0 H_r \leq -60$ mT: the median falls from 7 to 5 to 4 as SF goes $2 \rightarrow 3 \rightarrow 4$. **c**, Coherence against trial period at SF = 2, 3 and 4 with the **block scheme scaled to the kernel** (width $2 \cdot \text{SF} + 1$, separation $2 \cdot \text{SF}$), the less favourable of the two ways of varying SF: the maximum moves from 13.10 mT to 9.50 mT between SF = 2 and SF = 3, only K = 9 blocks survive, and the result is no longer significant. Holding the block scheme fixed instead, so that K = 15 throughout, the shift is 11% and the coherence peaks at SF = 2, which is the basis for using it (Supplementary Note S8.7).




ED Fig. 6 | Temperature series. Residual FORC distributions at 2, 50, 100, 150, 200 and 300 K on a common field window, each scaled to its own 98th percentile. The tilted ridges are present only at 2 K (and, not shown here, at 5 K). From 50 K upward the residual is dominated by a feature that is vertical in (H, H_r) and sits at the field zero crossing, the superparamagnetic response of the unblocked cores (the main text, "Two conditions, both measured"). Its verticality is the discriminant: a structure carried by the exchange bias must follow the core magnetisation and therefore tilt.




ED Fig. 7 | The tilt tracks the remanence slope within a single dataset. The mechanism

of the main text, "Exchange bias is the reading we give it" says the ridge tilt is not a fixed geometric property but is set by how fast the core remanence changes with the reversal field: $d(\Delta H_n)/dH_r \propto dM_{rem}/dH_r$. This can be tested inside one measurement. **a, b**, Row-to-row tilt measured by cross-correlation (blue, left axis) and the slope of the remanence branch (red, right axis) against $\mu_0 H_r$, at 2 K and 5 K. Both fall towards zero together as the reversal field deepens: once the loop is saturated the core remanence no longer records H_r , the exchange bias stops moving, and the ridges become vertical. This is the origin of the near-vertical striping at low H_r in ED Fig. 1, and it is the opposite of the high-temperature vertical feature of ED Fig. 6, which sits at the field zero and arises from unblocked cores. **c, d**, The same data as a row-by-row correspondence, coloured by H_r : Pearson $r = 0.76$ (2 K, $n = 41$ row pairs) and 0.76 (5 K, $n = 34$), Spearman $\rho = 0.78$ and 0.76 ($p = 1 \times 10^{-9}$ and 2×10^{-7}). The dashed line is a fit through the origin. The relation is monotonic but not a fixed proportionality, the fitted slope differing between the two temperatures (31.9 and 20.6), so the correspondence is reported as qualitative, as in the main text, "Exchange bias is the reading we give it".



ED Fig. 8 | Failed analyses. The three approaches documented in SI Note S6, each of which was tried and abandoned. **a**, Peak-aligned stacking. Aligning each row on its strongest residual peak before stacking produces an oscillatory super-profile from the measured 2 K data, and an equally oscillatory one from phase-randomised surrogates that contain no periodicity by construction. The autocorrelation sideband is 0.266 for the data against a surrogate median of 0.306; at 5 K, 0.155 against 0.199. The surrogates score *higher*, so the method cannot support the claim. **b**, Small-window autocorrelation. The best sideband in the 10–20 mT band, scanned over the projection angle, plotted against $1/\sqrt{N_{\text{cells}}}$ for six temperatures. The points fall on a line through the origin ($A \cdot \sqrt{N} = 3.4$, $r = 0.75$): the statistic is measuring the number of degrees of freedom in the window, which is why the smallest window (50 K, 165 cells) returns the largest value. **c**, Null construction. Three surrogate schemes applied to the identical 2 K statistic. Bars show the 5–95th and 25–75th percentiles of each null, normalised to the observed value. Row-wise phase randomisation gives $p = 0.92$, a two-dimensional FFT surrogate gives $p = 0.02$, and a spectral notch gives $p = 0.67$. When the answer is set by a modelling choice and not by the data, the design must be replaced, which is why the main text uses an analytic Rayleigh null on a directly measured locus instead.



ED Fig. 9 | The comb position measured row by row, from the phase of the residual. An

independent determination of the ladder geometry, from a statistic in a different domain from the cross-correlation of Fig. 3a. For each reversal-curve row the complex Fourier component of the residual at period $\Delta H = 12.12$ mT is evaluated on the native, non-uniform field nodes with no interpolation, and its phase gives the comb position b modulo ΔH (Supplementary Note S10).

a, The positions, with local slopes fitted over seven rows, each labelled with $-db/dH_r$ from the phase and, in parentheses, the median cross-correlation displacement $-dH_n/dH_r$ over the same rows: +0.117 (+0.122), +0.321 (+0.268) and +0.376 (+0.574) as the reversal field becomes shallower. The two agree on the value at deep H_r and on the direction of the trend throughout. Grey, a single line fitted across the whole window returns +0.250; it understates the tilt because the slope is not constant. **b**, The same positions compared with the remanence branch $M_{\text{rem}}(H_r)$ measured on this specimen, with λ the only free parameter. The fit returns $\lambda = 29$ mT with a root-mean-square residual of 0.72 mT, against 1.87 mT for the straight line: the comb position follows the shape of the remanence branch, not a constant slope. Unlike the cross-correlation, this fixes the absolute position of the ladder and not only its tilt; λ here may be compared with the 31.9 mT fitted to the cross-correlation tilt at 2 K (main text).

Supplementary Information

Metastable nanoparticle cores amplify single-ion quantum level crossings

Xiang Zhao^{1*†}, Shuangchi Liu^{2†}, Tongwei Zhang^{3†}

¹ Research School of Earth Sciences, The Australian National University, Canberra, ACT 2601, Australia

² State Key Laboratory of Lithospheric and Environmental Coevolution, Institute of Geology and Geophysics, Chinese Academy of Sciences, Beijing 100029, China

³ Key Laboratory of Planetary Science and Frontier Technology, Institute of Geology and Geophysics, Chinese Academy of Sciences, Beijing 100029, China

† These authors contributed equally.

Present address (X.Z.): Independent researcher, Canberra, ACT, Australia.

* Correspondence: sungchao@gmail.com

Companion to the main text. Notes S1–S5 carry the physical framework; Note S6 records the methodological findings, including the analyses that failed and why.

Note S1 — The FORC operator and what it selects

A first-order reversal curve measurement saturates the sample, reduces the field to a reversal field H_r , then sweeps back towards positive saturation while recording $M(H, H_r)$ for $H \geq H_r$. Repeating over many H_r gives a two-dimensional magnetisation array, from which

$$\rho(H, H_r) = -\frac{1}{2} \frac{\partial^2 M}{\partial H \partial H_r}$$

The first derivative $\partial M / \partial H$ is the susceptibility, containing both reversible and irreversible contributions from the whole ensemble. The mixed second derivative retains only the part of the susceptibility that depends on magnetic history. Three properties follow.

Reversible processes vanish. Any process independent of H_r (paramagnetism, diamagnetism, reversible rotation) contributes zero.

Smooth backgrounds are suppressed, not amplified. Differentiation is *selective*: a broad classical switching distribution produces a low, smooth feature in ρ , while a sharp, localised, systematically curved modulation survives. It is important not to describe this as amplification. The operator does not increase signal; it removes what is smooth, and it increases high-frequency noise by the same token. This is why the analysis of Note S6 is required.

This selectivity can be quantified on the present data. A component of period ΔH is enhanced relative to a switching distribution of width W by a factor of order $(W/\Delta H)^2$; with the measured switching-peak width $W = 70.2$ mT (Note S9.4) and $\Delta H = 13.1$ mT this is 29. Measured directly, as the ratio of the residual to the smooth part with matched definitions in the magnetisation and in ρ , the enhancement is 29 at 2 K and 37 at 5 K. The periodic component amounts to a few tenths of one percent of the switching amplitude in M , a figure that varies between roughly 0.1 and 0.7% with the width chosen to separate the background, and is therefore quoted only as sub-percent. Two consequences follow: the structure cannot be seen in a bulk hysteresis loop, and the same structure measured with a sampling density too coarse to resolve the period would not be recoverable by this operator either (main text, "What these data do not decide").

Different H_r select different sub-populations. The two-dimensional map resolves the switching behaviour of sub-populations and not the ensemble mean, which is what makes the ridge locus a measurable object.

The decomposition of ref. 14 separates the response into remanent (irreversible), induced (reversible) and transient (stochastic) components. The three measurement channels used in the main text, the standard FORC sequence and the in-field and zero-field points of the remanence sequence, are physically distinct measurement operations, which is the property the coherence test exploits.

Note S2 — Level crossings of a single surface ion

S2.1 Spin Hamiltonian

A surface Fe^{3+} ion ($S = 5/2$) in the low-symmetry environment of the particle surface is described by

$$\mathcal{H} = -DS_z^2 + E(S_x^2 - S_y^2) - g\mu_B\mu_0 H_{\text{eff}}S_z + \mathcal{H}_{\text{int}}$$

with D the axial zero-field splitting, E the rhombic term, $\mu_0 H_{\text{eff}} = \mu_0 H + \mu_0 H_{\text{bias}}$, and $\mathcal{H}_{\text{int}}$ the local micro-perturbations (hyperfine coupling to ^{57}Fe , dipolar fields from neighbouring ions).

S2.2 Crossing condition and spacing

For $E \approx 0$ the eigenvalues are $E_m = -Dm^2 - g\mu_B\mu_0 H_{\text{eff}}m$. Setting $E_m = E_{m'}$ for $m \neq m'$,

$$D(m'^2 - m^2) = g\mu_B\mu_0 H_{\text{eff}}(m - m') \implies \mu_0 H_{\text{eff},n} = \frac{nD}{g\mu_B}, \quad n = -(m + m')$$

so consecutive crossings are separated by

$$\mu_0 \Delta H = \frac{D}{g\mu_B}$$

independent of the crossing index, of temperature, and of particle size. The bias enters as a constant offset and cancels in the difference between consecutive crossings, so ΔH measures the crystal field alone.

S2.3 The ladder is finite

Because $S = 5/2$ is half-integer and a crossing requires $m \neq m'$, the positive crossing index is restricted to $n \in \{1, 2, 3, 4\}$: there is no $n = 5$, since $m + m' = -5$ would require $m = m' = -5/2$. Hence $n_{\text{max}} = 2S - 1 = 4$.

For $S = 5/2$, the Kramers counting constraint restricts the crossing index to $n \leq 2S - 1 = 4$, confining the periodic structure to a field range near $4\Delta H \approx 50$ mT and not extending

indefinitely, which higher field resolution would test.

S2.4 Spin parameters

From the periodicity measured at the two temperatures where the core is blocked, $\mu_0 \Delta H = 11.9 \pm 1.6$ mT, and with $g = 2$ for the ${}^6S_{5/2}$ state,

$$D = g\mu_B\mu_0\Delta H = 0.016 \pm 0.002 \text{ K}$$

Per-temperature values: 0.0176 K (2 K), 0.0145 K (5 K). The nominally significant coherence returned at 50 and 100 K arises from the vertical superparamagnetic feature at the field zero crossing (main text, "Two conditions, both measured") and is not a measurement of ΔH .

There is no directly comparable literature value. A single-ion zero-field splitting is defined and measured for a paramagnetic ion dilutely substituted in a *diamagnetic* host, where it is isolated from exchange; in a magnetically ordered lattice, exchange averaging and the resulting linewidth remove the fine structure. Fe^{3+} in magnetite is neither dilute nor in a diamagnetic host, so its zero-field splitting has not been, and cannot readily be, measured by electron paramagnetic resonance. For scale only, the dilute-limit value for Fe^{3+} at the octahedral B site of the diamagnetic spinel ZnGa_2O_4 is $D = -0.2442 \text{ cm}^{-1} = 0.351 \text{ K}$ (ref. 22), twenty-two times larger than the value implied here. The two are not the same observable, and we draw no inference from the comparison in either direction. The identification of the active unit rests on the field scale, which excludes a collective origin by two orders of magnitude, and on the composition of the sample, which admits no $S = 5/2$ species other than Fe^{3+} .

Note S3 — Why the exchange bias does not destroy the ladder

The crossing appears in *applied* field at

$$\mu_0 H_{\text{applied},n} = \frac{nD}{g\mu_B} - \mu_0 H_{\text{bias}}$$

so only sites whose longitudinal bias is small place their crossings in the accessible low-field

range.

This is not a statement about the mean surface exchange field. A fully coordinated Fe^{3+} site experiences a molecular field of order 10^2 T, consistent with the 850 K ordering temperature of magnetite. The requirement is weaker: that the disorder distribution of local biases carries non-zero weight near zero, $P(H_{\text{bias}} \approx 0) \neq 0$, i.e. that a minority of geometrically compensated or frustrated surface sites exists. These rare sites carry the ridges; the broadly distributed majority contributes only to the smooth background that the second-derivative operator suppresses.

Constraint from the observation. For the periodic structure to survive ensemble averaging, the dispersion of the bias among contributing sites must be small compared with the spacing. A Gaussian offset dispersion σ suppresses the oscillation amplitude by $\exp(-2\pi^2\sigma^2/\Delta H^2)$, so $\sigma \lesssim 3$ mT is required at $\Delta H \approx 12$ mT. We present this as a constraint *derived from* the observation and not as an independent assumption, and note that symmetry-compensated sub-surface B sites are a natural candidate for a sub-population satisfying it.

Note S4 — Orientation averaging

In a randomly oriented ensemble the crossing field of an active ion depends on the angle θ between its local axis and the applied field, $\mu_0 H_n(\theta) = nD/(g \mu_B \cos \theta)$. Orientation averaging does not wash out the periodicity: each crossing produces an axial powder lineshape with a sharp turning-point edge at the low-field limit $nD/(g\mu_B)$ and a high-field tail, so spectral weight accumulates at the crossing fields and the spacing of the edges is preserved.

Because irreversible switching appears as steps in M, the sharp component of the FORC second derivative scales as the derivative of this resonance density: the steep low-field edges survive as sharp features while the smooth tails are suppressed. This is the same mechanism that renders the resonances visible in ρ while leaving them buried in the bulk M(H) loop.

Note S5 — Energy scales and the nature of the transition

The single-ion barrier implied by the measurement is $DS^2 \approx 0.11$ K, far below the measurement temperatures: $k_B T / DS^2 \approx 18$ at 2 K and ≈ 45 at 5 K. All six m states are thermally accessible and the ion is **not** a metastable two-well system.

Two consequences, stated as limitations and not glossed.

The process is not coherent tunnelling. For half-integer S, time-reversal symmetry forbids tunnelling by the crystal-field terms alone; the gap opened by hyperfine (~ 1 mT) and dipolar (~ 1 – 10 mT) perturbations is far too small to support coherent passage before thermal decoherence. Furthermore, the measurement protocol steps the field, settles, and measures, with a dwell of tens of seconds per point; there is no continuous coherent sweep, and a single-passage Landau–Zener probability is not an observable of this experiment. What survives is an **incoherent resonant enhancement of the transition rate at the degeneracy fields**, where otherwise energy-non-conserving transitions become allowed. Precedent exists in dilute rare-earth systems measured with stepped protocols of the same character, notably $\text{LiYF}_4:\text{Ho}^{3+}$, where the features at the crossing fields are reproduced by cross-relaxation and phonon-bottleneck effects. That system has a non-Kramers ground doublet, unlike the ion considered here; what carries over is the measurement protocol and the incoherent character of the mechanism, not the level structure.

Magnitude of the symmetry-breaking gap. For completeness we note the scale of the gaps that a transverse perturbation would open, computed by exact diagonalisation with $D = 0.016$ K and $E = 0.002$ K. With the rhombic term alone the minimum adjacent separation at the four crossings is 1.7, 0.2, 10.5 and 0.5 mK for $n = 1$ – 4 . Adding a transverse field of 1 mT, a plausible scale for hyperfine and dipolar perturbations, raises these to 1.7, 1.5, 11.7 and 3.7 mK, the effect being largest at the crossings that the rhombic term alone cannot mix.

These numbers are given as an order-of-magnitude orientation only. The transverse field is assumed, not measured; no gap is extracted from our data, and none is required by the interpretation, which rests on the crossing *positions*. Figure 4a is therefore computed without a transverse field, so that it shows the crossings themselves and not gaps whose size would be set by an assumed parameter.

Positions are temperature-independent; amplitudes are not. The degeneracy condition fixes the crossing fields; the ability of the FORC operator to register them depends on the existence of irreversible switching, which disappears above blocking. The persistence of a common spacing at the two temperatures where the irreversible channel is open, the monotonic decline of coherence as that channel closes, and the change of locus geometry at the blocking transition, are precisely this division.

Note S6 — Methodological findings, including analyses that failed

This note records approaches that appear reasonable, that we applied, and that we found to be invalid on data of this character. Each is reported with the diagnostic that revealed the problem. We include it because the failures are not obvious a priori and because the same procedures are in common use.

S6.1 Peak-aligned stacking is self-confirming

Aligning each reversal curve on its strongest residual extremum before stacking is a natural way to raise signal-to-noise. It is invalid when the alignment path is determined from the same data that are then tested.

We applied the procedure to phase-randomised surrogates of our own data, which by construction contain no cross-curve phase structure. At both temperatures where we report a period, the stacked super-profiles of the surrogates oscillate **as strongly as or more strongly than** those of the measured data; above blocking the ordering is not systematic, which is itself the point: the statistic tracks the alignment, not the data:

	measured	surrogate (median of 50)
2 K	0.266	0.306
5 K	0.155	0.199
50 K	0.241	0.192
300 K	0.092	0.107

1209 The mechanism is straightforward: aligning on the maximum builds a peak at zero lag by con-
 1210 struction, and the autocorrelation reads its sidelobes as periodicity. The output is independent of
 1211 whether the input contains a real period. We do not use this procedure because peak alignment
 1212 artificially builds a peak at zero lag.

1213 **S6.2 Small-window autocorrelation sidebands measure degrees of freedom, not signal**

1214 The amplitude A of the first autocorrelation sideband, computed on a projected residual profile
 1215 within a restricted window, scales with the number of cells N in that window as $A \propto N^{(-1/2)}$.
 1216 Across six temperatures we find

$$A\sqrt{N} = 3.6, 3.8, 4.3, 3.7, 3.1, 2.7 \quad (2, 5, 50, 100, 200, 300 \text{ K})$$

1218 i.e. constant. After accounting for degrees of freedom the statistic carries no temperature in-
 1219 formation. The diagnostic that settles it: decimating the data *increases* A , which cannot be the
 1220 behaviour of a statistic measuring physical signal. Comparisons of sideband amplitude between
 1221 datasets of different size are therefore uninterpretable unless normalised for N .

1222 **S6.3 Surrogate-based significance is dominated by null design**

1223 The natural null for a periodicity test is a surrogate that preserves the noise character while
 1224 destroying the structure of interest. We constructed three defensible variants and applied them
 1225 to identical data:

null construction	p (2 K)
row-wise phase randomisation, total-variance normalisation	0.92
the same, with out-of-band power high-passed before normalisation	0.02
spectral notching with two-dimensional phase randomisation	0.67

1226 The results span the full range. The cause is that row-wise phase randomisation destroys **all** row-
 1227 to-row correlation, including the correlation imposed by the $SF = 2$ smoothing kernel, which
 1228 by construction spans ± 2 reversal curves in H_r . Real FORC rows are therefore always "more
 1229 coherent" than such surrogates, regardless of whether ridges are present; whether this manifests

as suppression or enhancement of the observed statistic depends on how the normalisation treats out-of-band power.

This is why the main text uses an **analytic** Rayleigh null. Phase coherence among K independent blocks follows a known distribution, so no surrogate construction, and therefore no null-design freedom, enters the analysis.

S6.4 The effective number of geometry trials, and why the locus must be measured

An analysis that fits the ridge locus and the period simultaneously must correct for the search over locus parameters. Correcting by the number of grid points is not conservative here, because the grid is *coarser* than the correlation length of the statistic in those parameters.

For the quadratic locus coefficient, a change of 0.007 displaces the projection by one quarter of a period at the far end of the analysis window; over the scanned range this corresponds to roughly 170 independent values, against 9 grid points. The linear coefficient similarly supports about 77 independent values against 23 grid points. The effective number of trials is of order 10^5 – 10^6 , not 3×10^4 , and a significance estimate corrected by grid count is optimistic by more than an order of magnitude.

Two further symptoms of the same problem were observed. Solutions maximising coherence were repeatedly found at the boundary of the scanned range, where the quadratic term displaces the projection by more than 200 mT across the window, a shear sufficient to convert a smooth background into an apparent periodicity, and corresponding to no feature visible in the FORC map. And the locus parameters obtained by free fitting were inconsistent with those required by the measured descending branch: fitted tilt and curvature implied coupling constants differing in sign.

The main text therefore measures the locus first, by row-to-row cross-correlation, which involves no fitting, no template and no scanning, and then evaluates the periodicity along the measured locus with ΔH as the only free parameter. The correction then covers about 51 independent values.

S6.5 The waveform is unresolved

Whether the alternating bands in the FORC distribution correspond to four same-sign crossings or to the positive and negative lobes of two cannot be decided from these data. The discriminating quantity is the phase of the second harmonic relative to the fundamental: a train of same-sign peaks has its harmonics in phase ($\Delta\varphi = \varphi_2 - 2\varphi_1 \approx 0$), whereas a train of derivative-like dipoles has $\Delta\varphi \approx \pm 90^\circ$.

Evaluated along the measured locus at $\Delta H = 13.10$ mT, the second-harmonic amplitude is $A_2 = 0.018$. The comparison level is the distribution of $|c|$ at frequencies away from the fundamental and its harmonics, whose median is 0.037 and maximum 0.060: A_2 lies below even the median, so no second harmonic is detected. A bootstrap over reversal curves gives a confidence interval on $\Delta\varphi$ of $[-131^\circ, +97^\circ]$, spanning both hypotheses. At roughly five field samples per period the waveform is simply not resolved, and the counting argument of Note S2.3 cannot be tested.

S6.6 Smoothing factor

SF = 3 removes alternating extrema and must not be used for periodicity analysis. At $H_r = -122.6$ mT the 2 K remanence FORC computed with SF = 3 gives extremum spacings of 6.6 / 8.4 / 10.2 / 2.0 / 4.0 mT (CV **0.47**), six extrema against seven, one having been lost to smoothing. The same row of the standard FORC computed with SF = 2 gives the regular sequence 6.6 / 6.6 / 6.2 / 6.1 / 6.0 / 6.0 mT. All periodicity analysis in the main text uses SF = 2.

S6.7 Summary of recommended practice

For periodic structure in FORC distributions we recommend: measuring the ridge locus directly rather than fitting it; using an analytic null where one exists; reporting the row-to-row displacement, which distinguishes structures following the hysteretic state from those locked to the applied field; normalising any amplitude statistic for the number of cells in the analysis window; and avoiding any stacking procedure whose alignment is derived from the data being tested.

S6.8 Two analyses superseded during this work

Both are recorded here because each was reached by a route that looks reasonable and fails for a reason worth stating; neither survives into the main text.

A monotonic decline of R with temperature. Evaluating the coherence statistic at each temperature on all available reversal curves returns values that fall as temperature rises, which invites reading R as a discriminator between the blocked and unblocked regimes. It is not one. Under the null R is biased upwards as the number of blocks K falls, and the control datasets carry 20 to 22 reversal curves against 42 at 2 K, so they yield $K = 6$ against 15. Truncating the 2 K data to the same size raises R from 0.879 to 0.988 while p goes to 0.29, on the same data and with the same physics. Compared at matched K, R above blocking is as high as below it. Any temperature dependence of R on these data is a block-count effect, and the discrimination is carried by the geometry instead (Note S8.3).

Four resolved bands matched to the four Kramers crossings. For $S = 5/2$ the crossing index is restricted to $n \leq 2S - 1 = 4$, and the residual profile shows alternating bands whose count is of that order, so the match is tempting. It cannot be made: at roughly five field samples per period the waveform is unresolved, and four same-sign crossings cannot be distinguished from the positive and negative lobes of two (Note S6.5). The counting result is therefore carried in the main text as a prediction only. It is also in tension with the profile, which has extrema at 59.6 and 78.0 mT beyond $4\Delta H$ that survive a sixfold variation of the background-subtraction width without moving by a grid point.

Note S7 — Code and Data Availability

All raw reversal-curve data, processed FORC matrices, and analysis scripts are provided with this submission and will be deposited on Zenodo with a permanent DOI prior to publication.

Note S8 — The block scheme, and what depends on it

The coherence statistic divides the reversal curves of each channel into blocks and treats the block phases as independent draws. Two parameters control that division: block width and curve separation. This note details their specification and reports what the conclusions depend on.

S8.1 Both parameters are set by the smoothing factor

Block width = $2 \cdot \text{SF} + 1 = 5$ reversal curves. The Pike moving-polynomial fit uses a $(2 \cdot \text{SF} + 1)$ -node window, so five curves is the span over which the map carries one independent estimate. **Separation = $2 \cdot \text{SF} = 4$ discarded curves.** The blocks themselves are contiguous and non-overlapping, but ρ at row r is computed from raw data at rows $r - 2$ to $r + 2$. Two blocks therefore share no raw measurement when the first row of one exceeds the last row of the other by $2 \cdot \text{SF}$.

This gives five blocks per channel at 2 K ($K = 15$ across the three channels), four at 5 K ($K = 12$) and two at 150–300 K ($K = 6$), the last because those datasets contain only 20–22 reversal curves inside the analysis window.

S8.2 Parameter consistency across channels

The block scheme specified in S8.1 fixes all parameters directly in terms of the smoothing factor SF , ensuring consistent block counts across measurement channels.

S8.3 Null calibration

The effective number of independent blocks, estimated from the off-peak coherence as $K_{\text{eff}} = 1/\langle R^2 \rangle$, is 6.4 at 300 K against a nominal 6. Pooling the three channels therefore does not inflate R under null conditions. At 2 K and 5 K the same estimator is unstable, returning values between 3 and 22 depending on the evaluation band, because the structure itself carries power away from the peak and $R(\Delta H)$ is strongly autocorrelated; we do not present it as verification at those temperatures.

S8.4 The resolved-peak rule

The reported maximum of $R(\Delta H)$ is the highest maximum spanning more than one grid point at 90% of its height. The rule refers only to peak width, involves no period, and was fixed before the result was inspected. It is not cosmetic: at 2 K a one-point spike at 8.20 mT reaches $R = 0.922$ and would otherwise be reported in place of the resolved peak at 13.10 mT, which is 1.00 mT wide across eleven grid points.

S8.5 What the period does not depend on

Across seven block configurations the 2 K period stays within 12.70–13.20 mT, a spread of 4%:

width / separation	2 K	5 K	lowest p among the five controls
5 / 4 (adopted)	13.10 mT, $p = 9.5 \times 10^{-4}$	10.80 mT, $p = 3.9 \times 10^{-3}$	0.275
7 / 4	12.70 mT, $p = 1.6 \times 10^{-3}$	12.60 mT, $p = 2.4 \times 10^{-2}$	0.312
5 / 6	12.90 mT, $p = 2.5 \times 10^{-3}$	11.40 mT, $p = 4.3 \times 10^{-2}$	0.274
5 / 2	13.00 mT, $p = 1.5 \times 10^{-5}$	10.80 mT, $p = 1.7 \times 10^{-4}$	0.033
5 / 0 (continuous)	12.90 mT, $p = 5.8 \times 10^{-8}$	11.00 mT, $p = 2.9 \times 10^{-6}$	0.003
4 / 4	12.80 mT, $p = 6.4 \times 10^{-5}$	10.80 mT, $p = 2.2 \times 10^{-3}$	0.033
3 / 4	13.20 mT, $p = 2.1 \times 10^{-5}$	10.80 mT, $p = 2.4 \times 10^{-4}$	0.027

The configurations in the lower half return much smaller p at 2 K and 5 K (reaching $p = 5.8 \times 10^{-8}$ under the continuous 5/0 scheme), but each also drives at least one control temperature below 0.05, and a control dataset should return none. Specifically, while the continuous block scheme (5/0, no unused curves between blocks) yields an ostensibly stronger significance ($p \approx 10^{-8}$), it fails the strict independence requirement for the Rayleigh null because adjacent blocks share raw measurement curves across their boundaries through the smoothing kernel ($SF = 2$). We therefore adopt the conservative 5/4 scheme throughout the main text.

Across the configurations and the five control temperatures, combinations reaching $p < 0.05$ occur only under the rejected settings with smaller separation: their resolved maxima lie outside

the interval reported here, $\mu_0\Delta H = 11.9 \pm 1.6$ mT. **No configuration, at any control temperature, returns a significant result at the reported period.** The rejected settings inflate the false-positive rate; they do not reproduce the claimed periodicity above blocking. We do not adopt the settings that give the best p, and we report them here so that the choice is visible.

S8.6 What the significance does depend on

p depends strongly on K, and K depends on how many reversal curves a dataset contains. This is the main limitation of the coherence test as applied here.

dataset	curves per channel	K	R	p
2 K, full	42	15	0.879	9.5×10^{-4}
2 K, truncated to 20	20	6	0.988	0.29
5 K, truncated to 20	20	6	0.978	0.33
50 K	20	6	0.993	0.27
300 K	20	6	0.925	0.60

Truncated to the size of the control datasets, the 2 K data return the same p as the controls. Under the null R is biased upwards as K falls, so **R is not comparable between datasets of different size**, and the high R above blocking reflects small K and not a real periodicity. The control temperatures therefore do not establish that the coherence is absent above blocking; they lack the power to address it. What they establish is geometric: the locus is exactly vertical, with an interquartile range of 0.000 across nineteen adjacent curve pairs. That measurement uses no blocks, no template and no free parameter, and it does not weaken with dataset size.


S8.7 Dependence on the smoothing factor

SF is the bias–variance parameter of the FORC estimator, and the result must depend on it: with too little smoothing the residual is dominated by point-to-point noise, with too much the ridges are averaged away. The question is not whether there is a dependence but whether $SF = 2$ sits at the optimum. Holding the block scheme fixed at 5 / 4 so that $K = 15$ throughout and R is comparable:

SF	point-to-point fluctuation / overall amplitude	R in the 12–14.5 mT band	resolved maximum
1	0.78	0.63 @ 13.60 mT	8.20 mT (a band-edge spike; noise dominates)
2	0.44	0.88 @ 13.10 mT	13.10 mT
3	0.27	0.83 @ 14.50 mT	14.60 mT
4	0.19	0.35 @ 13.40 mT	19.10 mT (structure smoothed away)

The noise proxy falls monotonically, as smoothing requires, while the band coherence has a single maximum at SF = 2. SF = 2 was fixed a priori as standard practice for FORC distributions of this density; the scan confirms it and does not select it.

At SF = 3, with the block scheme held at 5 / 4 so that only the smoothing changes, the resolved maximum is displaced to 14.60 mT, an 11% shift. If the block scheme is also scaled to the larger kernel (width 7, separation 6), only K = 9 blocks remain, the resolved maximum moves to 9.50 mT, and the result is no longer significant (p = 0.07). Extended Data Fig. 5c shows the latter case. Visual peak-picking degrades far faster than the statistic: the extremum spacings go from CV = 0.04 at SF = 2 to CV = 0.47 at SF = 3 (Note S6.6), which is why peak-picking is not used.



Supplementary Fig. S8 | The block scheme. **a**, Phase coherence R against trial period along the measured locus, block scheme 5 / 4. Blue, 2 K (K = 15); red, 5 K (K = 12); grey band, the range spanned by the five control temperatures, each with K = 6. Triangles mark the resolved maxima at 13.10 and 10.80 mT. The control temperatures reach R of 0.91 to 0.99 without reaching significance, because R is biased upwards as K falls; R is therefore not a quantity that can be compared between datasets of different size (S8.6). **b**, The 2 K period across six block schemes

(shading, 12.70–13.20 mT), with the corrected p at 2 K and the lowest p among the five control temperatures. The three schemes returning the smallest p each drive a control temperature below 0.05 and are not used; the adopted scheme, $5/4$, is the one set by the smoothing factor (S8.1).

S8.8 The null, its range of validity, and why we do not sharpen it

The reported significance is $p = \exp(-KR^2)$, multiplied by 102 for the 51 independent trial periods and the two sign choices of the locus. Two things about this expression need stating, because the second is more important than the first and they point in opposite directions.

It is an asymptotic form, and here it is conservative. For K unit vectors with independent, uniformly distributed phases, the tail of the resultant obeys a large-deviation law, $P(R \geq r) \asymp \exp(-K I(r))$ with $I(r) = r\kappa - \ln I_0(\kappa)$ and $I_1(\kappa)/I_0(\kappa) = r$. The familiar $\exp(-KR^2)$ is the small- r expansion of that rate function: as $r \rightarrow 0$, $\kappa \rightarrow 2r$ and $I(r) \rightarrow r^2$. The two agree to 1% at $r = 0.2$ but not at the values reached here: $I(r)/r^2$ is 1.42 at $r = 0.879$ and 1.56 at $r = 0.921$. Carried into the exponent at $K = 15$ and $K = 12$ this is a factor of order 10^2 in the probability, and always in the same direction: the quoted p is an upper bound. Direct Monte Carlo, calibrated first against the $K = 2$ case where $P(R \geq r) = 2 \arccos(r)/\pi$ is exact, agrees with the rate-function estimate in order of magnitude at every point tested.

But the independence assumption fails at small K , so the sharper form is not the better one. Truncating the 2 K dataset to successively fewer reversal curves changes nothing but the block count, since the data, the measured locus and the processing are identical, and R rises monotonically as blocks are removed: 0.879 at $K = 15$, 0.948 at $K = 12$, 0.956 at $K = 9$ and 0.988 at $K = 6$. The 5 K dataset behaves the same way, from 0.921 at $K = 12$ to 0.978 at $K = 6$. A null of independent block phases does not predict this; fewer blocks should make the estimate noisier, not more coherent. Under independent uniform phases $R = 0.988$ at $K = 6$ would carry an exact tail probability near 10^{-6} , which it plainly does not, being the same measurement differently partitioned. The blocks share one field grid, one background-subtraction kernel and, above blocking, one localised feature, and at $K = 6$ that shared structure dominates.

Where the truncation starts makes little difference: holding $K = 6$ and stepping the first retained curve through the 2 K dataset returns R between 0.978 and 0.993. That interval covers the five control values (0.909 to 0.993) entirely, and two of the starting positions reproduce the 50 K control value of 0.993 exactly. The recovered period is also unstable there, moving from 13.10 to 11.80 mT at 2 K as the dataset is truncated. Coherence at $K = 6$ therefore carries no information about whether periodic structure is present.

The exact tail is therefore not more trustworthy than the asymptotic one; it is only more precisely wrong. We keep $\exp(-KR^2)$ because it is conservative, we report it as an index rather than as a calibrated probability, and we answer the question it cannot, whether the control temperatures mean anything, with the truncation test on our own data instead of with a null model. This is the quantitative content of the statement in the main text that R does not separate the two regimes: at matched K it does not, and the separation is carried entirely by the geometry of the locus.

Note S9 — The residual rim on the main ridge, and the width of the background subtraction

Extended Data Fig. 1 shows near-vertical stripes crossing the main ridge. This note identifies them, states their magnitude, and shows what the reported result does and does not depend on.

S9.1 How large they are

Decomposing the 2 K residual into a component that depends on H alone, constant from one reversal curve to the next, and a component that follows the measured ridge locus:

	fraction of the variance of the map
H-only ("vertical")	4.5%
following the measured locus	24.6%

The vertical component has an amplitude of 0.20 of the map standard deviation. At 5 K the corresponding figures are 7.3% and 18.3%.

S9.2 They are not the superparamagnetic feature, and not an instrumental common mode

Above the blocking transition the H-only component is large and has a characteristic shape: 40–67% of the variance, concentrated at the field zero crossing (48–75% of its amplitude within $|\mu_0 H| < 15$ mT, integrated over the field axis rather than counted per node, the axis being non-uniform), peaking between +3.5 and +8.8 mT, and shared between the three temperatures (mutual correlation 0.52 to 0.88). That is the superparamagnetic response discussed in the main text.

The 2 K component has none of those properties. It peaks at $\mu_0 H = +39.6$ mT, only 15% of its amplitude lies within $|\mu_0 H| < 15$ mT, and it is uncorrelated with the high-temperature profile ($r = 0.12$ against 50 K, -0.10 against 300 K). An instrumental effect fixed in applied field would give the same profile at every temperature; it does not.

S9.3 They are a residual of the background subtraction

The smooth switching distribution is removed as $\rho_{\text{resid}} = \rho - G_\sigma * \rho$. A Gaussian high-pass leaves a rim wherever the curvature of the main switching peak exceeds what the kernel follows, and that rim lies on the steepest flank of the peak, which is where the stripes are. The column-to-column granularity follows from the non-uniform grid: adjacent ρ columns are fitted from different subsets of raw points, so the rim is broken and not smooth (cf. Note S6 and Extended Data Fig. 4).

This makes a falsifiable prediction: the rim should scale with σ , a physical periodicity should not. Sweeping σ over a factor of six:

σ (grid nodes)	2.0	3.0	4.5	6.0	8.0	12.0
H-only component (% of variance)	1.3	2.1	4.5	7.7	12.5	21.4
resolved period at 2 K (mT)	13.10	13.10	13.10	13.10	13.10	13.10
ridge tilt at 2 K	0.277	0.348	0.397	0.400	0.400	0.402

The tilt in this table and in S9.4 is the locus tilt recovered under the σ sweep. It is not the same statistic as the pair-by-pair median displacement of +0.476 quoted in the main text, which is taken over the full H_r window at the adopted σ . | R at 2 K (K = 15) | 0.862 | 0.870 | **0.879** | 0.890 | 0.905 | 0.928 |

The rim grows by a factor of sixteen while the period does not move by one grid step. At 5 K the period is 10.80 mT for $\sigma \leq 6$, with a single excursion to 11.50 mT at $\sigma = 8$.

S9.4 Why $\sigma = 4.5$

The background remover is a high-pass, and its width is bounded on both sides by quantities we have measured.

Below, it must pass the ridges. A Gaussian high-pass transmits a component of wavelength λ with weight $1 - \exp(-2\pi^2\sigma^2/\lambda^2)$. At the ridge period $\Delta H = 13.1$ mT this reaches 94.6% for $\sigma = 5.0$ mT and exceeds 99.99% for $\sigma \geq 11$ mT. A narrower kernel begins to follow the ridges and subtract them as though they were background.

Above, it must follow the switching distribution. The smooth component has a measured full width at half maximum of 70.2 mT along H, so a Gaussian copy of width σ reproduces it to within a fractional broadening of $\sqrt{\sigma_W^2 + \sigma^2}/\sigma_W - 1$, where $\sigma_W = 29.8$ mT. That is 7.0% at $\sigma = 11.3$ mT and 42.4% at $\sigma = 30.2$ mT. The mismatch is what is left on the flank of the peak, the rim of S9.1.

The two bounds are separated by a factor of 6.0, and $\sigma = 4.5$ grid nodes = 11.3 mT sits in the **window**: the ridge period is transmitted without measurable attenuation while the switching distribution is reproduced to 7%.

σ (grid nodes)	2.0	3.0	4.5	6.0	8.0	12.0
σ (mT)	5.0	7.6	11.3	15.1	20.2	30.2
transmission at ΔH	94.6%	99.86%	>99.99%	>99.99%	>99.99%	>99.99%
broadening of the switching peak	1.4%	3.2%	7.0%	12.1%	20.7%	42.4%
ridge tilt	0.277	0.348	0.397	0.400	0.400	0.402

The measured tilt confirms the lower bound independently: it is 30% smaller at $\sigma = 5.0$ mT, where the transmission calculation says 5% of the ridge amplitude has been removed, and it saturates once the transmission does. The width of the main switching distribution is itself consistent with the enhancement factor quoted in the main text: $(W/\Delta H)^2 = 29$ gives $W = 70.5$ mT against the 70.2 mT measured here.

S9.5 Effect on the reported result

Removing the H-only component from the map entirely leaves the resolved coherence maximum at 12.90 mT with $R = 0.960$, against 13.00 mT and $R = 0.958$ for the unmodified map (standard FORC channel alone, $K = 5$; the three-channel values are those of the main table), a shift of 0.8%, below the scan step.

The geometric criterion is likewise unaffected, and the direction of any residual bias is conservative. Injecting synthetic H-only structure at $0.5\times$, $1\times$ and $2\times$ the observed amplitude leaves the measured tilt at +0.397; only at $4\times$ does it move, and then **downwards, to +0.266**. Structure at constant H contributes a zero-displacement component to the row-to-row cross-correlation and can only bias the measured tilt towards zero, towards the null. It cannot manufacture a tilt.

S9.6 Why the criterion is conjunctive

Any structure that depends on H alone is automatically phase-coherent across blocks, because every block sees the same profile at the same field. The 4.5% H-only component therefore contributes to R, and R alone cannot separate it from a genuine comb. Note also that R rises monotonically with σ in the table above, in step with the rim; R should not be read as a measure of signal strength. This is the limitation set out in the main text under "Two conditions, both measured", and it is why the criterion is conjunctive. The half that does the discriminating, the locus measured by row-to-row cross-correlation, is immune to it, for the reason given in S9.5. The scan is reproduced by 03_Code/figures/sigma_sweep.py.

S9.7 Alternative background estimators

A Gaussian high-pass is not the only way to remove the smooth switching distribution, and several standard alternatives follow a steep flank better. We compared four of them on the same three quantities: the rim, the amplitude of the ridges, and the period.

background estimator	rim (% of variance)	ridge amplitude	period (mT)	tilt
Gaussian, $\sigma = 4.5$	4.6	100%	13.10	0.397
median along H	2.3	71%	14.90	0.268
Savitzky–Golay along H	0.8	58%	12.70	0.261
Gaussian then median	1.8	69%	12.90	0.268
grey opening along H	33.8	89%	8.20	0.348

Filters along H use a window of 11 grid nodes = 28 mT, or 2.1 ridge periods. Ridge amplitude is the standard deviation of the residual averaged along the measured locus, normalised to the Gaussian case.

Three of the four do reduce the rim, the best of them by a factor of six. **Each does so by removing between 29% and 42% of the ridge amplitude**, and the measured tilt falls with it, from 0.397 to 0.26–0.27. This is not a defect of those filters but a consequence of the scales involved: the rim lies on a flank whose curvature scale is 20–30 mT and the ridges have a period of 13 mT, so any one-dimensional filter along H narrow enough to follow the flank is also narrow enough to absorb the ridges. A window of 28 mT spans 2.1 periods, and a quadratic or a median over that span tracks part of the oscillation.

The Gaussian at $\sigma = 4.5$ is the one estimator in the set that leaves the ridge amplitude intact, for the reason given in S9.4: at that width its transmission at the ridge period exceeds 99.99%, so it does not see the ridges at all. **We keep the rim because the alternative is to lose part of the signal, and because the rim biases the discriminating measurement towards the null and not away from it (S9.5).**

The comparison also strengthens the result. Across four background estimators built on different principles (linear smoothing, rank filtering, local polynomial fitting, and a combination) the

period is 12.70, 12.90, 13.10 and 14.90 mT. All four lie inside the interval $\mu_0\Delta H = 11.9 \pm 1.6$ mT quoted in the main text. The periodicity therefore does not depend on the width of the background subtraction (S9.3) or on the method by which the background is estimated.

Reproduced by 03_Code/figures/background_estimators.py.

Note S10 — Kinetic accessibility, and the size window it implies

The amplification that makes a single ion visible is borrowed from the core, and it works only while the core sits at the margin of thermal blocking. That is a two-sided condition on size, and it is kinetic, not energetic: the ion never supplies the energy to reverse the core, falling short by three orders of magnitude; it only modulates a rate that is already poised.

Below the window the core is superparamagnetic. It reverses freely on the measurement timescale, carries no irreversible switching, and leaves nothing in the mixed second derivative for the perturbation to modulate. With $T_B = K_{\text{eff}} V / (k_B \ln(\tau_m/\tau_0))$, $K_{\text{eff}} = 1.95 \times 10^4$ J m⁻³ and $\ln(\tau_m/\tau_0) = 25$ (Methods), this edge lies at 4.07 nm at 2 K, 5.53 nm at 5 K and 11.91 nm at 50 K.

Above the window the core is too deeply blocked where the resonances are. The field at which a core first becomes thermally switchable is set by $\Gamma t = 1$, equivalently $U(H) = 25 k_B T$: $H(\Gamma t = 1) = H_K [1 - (T/T_B)^{1/\kappa}]$, $\mu_0 H_K = 2K_{\text{eff}}/M_s = 81$ mT.

A core registers the modulation only if that field lies at or below the highest tooth of the resonance comb. Both the comb spacing and its position are measured: evaluating the complex Fourier component of the residual at each reversal-curve row, on the native non-uniform field nodes and without interpolation, returns a spacing $\Delta H = 12.12$ mT and a position b modulo ΔH that tracks the tilt of the locus (Extended Data Fig. 9). Rows are kept where the comb amplitude exceeds 40% of its median, and the positions are unwrapped across rows. Fitted as a straight line the positions scatter by 1.87 mT about it, but the scatter falls to 0.72 mT when they are compared instead with the remanence branch of the same specimen, with the single scale factor $\lambda = 29$ mT, close to the 31.9 mT obtained from the cross-correlation tilt, and an independent determination of the same ratio. With $n \leq 2S - 1 = 4$ the highest tooth therefore

lies between 48.8 and 59.6 mT over the H_r window. At 2 K a 4.5 nm core reaches $\Gamma t = 1$ at 12.0 mT, on the first crossing; a 14.6 nm core not until 69 mT, where its barrier is still $U/k_B T = 187$ against a threshold of 25, that is $\Gamma t = e^{-162}$. The suppression is exponential, not gradual.

The windows. At 2 K the window is 4.1–8.5 nm and contains 83% of the measured moment of this specimen; at 5 K it is 5.5–11.5 nm and contains 62%. At 50 K the window has not closed. It is 11.9–24.8 nm, and cores inside it escape freely, with Γt running from 10^9 at the lower edge to unity at the upper, but the window has slid past the distribution, and only 1.9% of the moment now lies inside it. Multiplying each share by the per-core modulation $\delta U / e k_B T$ gives 0.244%, 0.073% and 0.0002% of the switching amplitude at 2, 5 and 50 K: a factor of 1083 between 2 K and 50 K. This is why the ridges are absent at 50 K while the resonance condition, which is electronic, is present at every temperature.

The window translates with temperature; it does not widen. Its width in nanometres grows as $T^{1/3}$, from 4.4 nm at 2 K to 12.9 nm at 50 K, but its width in $\ln d$, the only width that matters for a log-normal size distribution, is $\ln 2.08 = 0.733$ at every temperature, because the escape rate depends on size and temperature only through $d/d_B(T)$. The fraction of moment that an optimally sized specimen could place inside the window is therefore a constant: 73% for the present width $\sigma_{\ln} = 0.335$, and 93% for a preparation as narrow as $\sigma_{\ln} = 0.20$. Since the per-core modulation falls as $1/T$, the periodic component of an ideal specimen is bounded above by that constant divided by temperature. No choice of particle size recovers at 50 K what is available at 2 K; the ceiling is 25 times lower there and 150 times lower at room temperature. This, rather than any failure of the resonance condition, is why conventional low-temperature magnetometry has not encountered the structure.

Inverted, the window assigns each preparation its own temperature. Because $T_B \propto d^3$, the size condition $d \in [d_B(T), 2.08 d_B(T)]$ is exactly $T \in [T_B/9.0, T_B]$, which removes the diameter as an intermediate quantity. Applied to the seven silica-coated magnetite preparations of ref. 15, whose moments and blocking temperatures are tabulated there:

preparation	T_B (K)	window (K)	at 2 K
O7	6–14	0.7–14	inside
O11	23–36	2.5–36	above

preparation	T_B (K)	window (K)	at 2 K
C10	32–47	3.5–47	above
O15	37–49	4.1–49	above
S17	47–61	5.2–61	above
O17	69–85	7.6–85	above
O23	79–107	8.7–107	above

The table is written in terms of the window alone. Where inside the window a preparation is best measured is a separate question, answered below, and it is not at the warm end.

The ranges span three routes to T_B , the published magnetic diameter, that diameter rescaled to a common M_s , and the measured blocking temperature, none of which is free of bias, the first because each diameter is derived with its own M_s , the last because a ZFC peak is weighted towards the coarse tail. At 2 K all three routes place all seven samples identically. O7 is the one preparation that falls inside the 2 K window, and it is the one in which the tilt of the switching locus is recovered at 6.9σ (Extended Data, "Independent corroboration of the geometry"), an independent specimen measured in another laboratory, whose diameter we take from the published table and have no freedom to adjust.

Each preparation has its own optimum, and it is not where occupancy alone is greatest.

The observable is the product of the window occupancy and the per-core modulation, and the second falls as $1/T$. Taking the blocking temperatures and lognormal widths reported for the seven preparations, and weighting by moment, the product peaks between $T_B/6$ and $T_B/5$, that is at the **low-temperature end** of each window and not in its middle:

preparation	T_{Bm} (K)	σ_m	in window at 2 K	optimum T (K)	signal, relative	integration
O7	14	0.26	76%	2.2	0.944	1 ×
O11	36	0.60	4.0%	7.1	0.218	21×
O15	42	0.45	1.0%	7.5	0.234	18×
C10	42	0.62	2.4%	8.1	0.181	31×
S17	58	0.61	0.6%	11.5	0.133	57×
O17	69	0.27	0.0%	11.0	0.188	28×
O23	107	0.66	0.0%	21.2	0.067	225×

Signal is relative to this specimen at 2 K; integration scales as its inverse square. An earlier version of this note predicted ridges in O17 and O23 remeasured between 30 and 60 K. That

was obtained from the occupancy alone and is withdrawn: including the $1/T$ factor moves every optimum below 22 K, and the two coarsest preparations are the two least favourable, not the most. **O7 is the only one of the seven measured near its own optimum, and it is the one in which the tilt is recovered.** We make no prediction that requires anyone to remeasure a sample we do not hold. Two things would defeat the pattern, and they are of different kinds.

The first is that the ridges could be present and not be seen, which is a statement about the measurement and not about the sample. Sampling density is the hard part of it: at four samples per period no amount of averaging establishes a periodicity, and that is where the published series sit. Signal-to-noise is the soft part, and the table above gives it: at its own optimum each of the six carries 0.07 to 0.23 of the signal this specimen gives at 2 K, so 18 to 225 times the integration would be needed, against 4 529 field nodes at 23.29 s for a single acquisition here. That route is therefore not available. But the periodic component is a fixed fraction of the switching amplitude, which scales with the amount of material, whereas the noise floor of the magnetometer does not: the same deficit is covered by 4 to 15 times the specimen. It is linear in sample where it is quadratic in time, and only the first is practical.

The second is structural. Every optimum in the table lies below the Verwey transition, where bulk magnetite is monoclinic; if charge order makes the B sites inequivalent and distributes D, the comb will not form. Two things bear on this. The transition is suppressed in particles of the size considered here, and no Verwey signature appears in the present material (Extended Data Table 1, row 17). And the structure reported here is itself measured at 2 and 5 K, below the transition, so the concern is not that a comb cannot form below it. What is untested is the charge-ordered state itself: because the present material does not enter it, our own data bear on the temperature and not on the structural phase. Working above the transition would remove the question altogether, and the window there opens at 16 nm, but the modulation at 120 K is sixty times below its 2 K value: the sensitivity that must be recovered grows with the structural certainty that is gained. We set both out and recommend neither.

What remains assumed. The expression above is the aligned-easy-axis switching field. For randomly oriented uniaxial particles the Stoner–Wohlfarth switching field runs from $H_K/2$ to H_K with orientation, so a core at 45° reaches $\Gamma t = 1$ at half the field and the upper edge is a

steep decline, not a hard zero. Taking $\kappa = 3/2$ instead of 2 moves the 2 K upper edge from 7.5–9.8 nm to 6.5–7.9 nm. And $n \leq 2S - 1$ follows from $S = 5/2$, which the main text establishes from the field scale and the composition and not from the ridge count; the four strongest extrema of the profile at $\mu_0 H_r = -122.6$ mT do sit at 23.2, 35.6, 47.6 and 59.6 mT, one full period apart, which is the pattern four crossings would produce, but the waveform is not resolved (Note S6.5) and we do not present that as confirmation.

The table above places each preparation by its median core, which is a fair representative for the preparations of ref. 15, whose widths are $\sigma_{\text{ln}} = 0.09$ to 0.22, and a poor one for this specimen, whose $\sigma_{\text{ln}} = 0.335$. Our own median has $T_B = 2.75$ K and by that criterion sits near the edge of the 2 K window, while the distribution as a whole places 83% of its moment inside it; the two statements are not in conflict, but only the second is the one the measurement responds to. Finally, the $1/T$ ceiling is derived over temperatures at which the distribution is measured, $T \geq 2$ K. Below 2 K the window slides under 4.07 nm, where a SIRM decay is insensitive because those cores are already superparamagnetic, and we make no claim about whether a lower temperature would serve better.

Note S11 — The zero-field splitting of octahedral Fe^{3+} , for scale only

This note collects the literature comparison referred to in the main text and in Methods. **No inference in the paper depends on it:** the identification of the active unit rests on the field scale, which excludes a collective origin by two orders of magnitude, and on the composition of the sample, which admits no $S = 5/2$ species but Fe^{3+} .

For an $S = 5/2$ ion in octahedral oxygen coordination the axial zero-field splitting D is set by the degree of local distortion and varies over two orders of magnitude with it. The diluted ferric alums, in which the ion sits in a nearly regular $[\text{Fe}(\text{H}_2\text{O})_6]^{3+}$ octahedron, span most of that range on their own: from the summary table of ref. 23, $D = +31 \times 10^{-4} \text{ cm}^{-1}$ in rubidium sulphate at 20 K, $-115 \times 10^{-4} \text{ cm}^{-1}$ in potassium selenate at 20 K, and $(-)-1880 \times 10^{-4} \text{ cm}^{-1}$ in methylamine sulphate at 90 K, with $(?-)-160 \times 10^{-4} \text{ cm}^{-1}$ quoted there for ferric ammonium sulphate at 4 K from the earlier analysis of Meijer. The strongly distorted B site of the diamagnetic spinel ZnGa_2O_4 gives $D = -0.2442 \text{ cm}^{-1}$ (ref. 22). The value implied here,

1648 0.0111 cm^{-1} , lies in the lower part of that span, closest to potassium selenate.

1649 Two things limit any such comparison. All of these are ions dilutely substituted in dia-
1650 magnetic hosts, which is the situation in which a single-ion zero-field splitting is defined and
1651 measurable. Fe^{3+} in magnetite is neither dilute nor in a diamagnetic host: exchange averaging
1652 and the resulting linewidth remove the fine structure, and the quantity inferred here, an effective
1653 splitting read from the field spacing of switching resonances in an exchange-coupled lattice, is
1654 not the same observable. And the axial term D and the cubic fourth-degree term a of the ${}^6S_{5/2}$
1655 spin Hamiltonian are numerically similar in these salts (in potassium selenate $a = -127 \times 10^{-4}$
1656 cm^{-1} against $D = -115 \times 10^{-4} \text{ cm}^{-1}$), so numerical agreement at this level does not identify
1657 which term is being matched. Only D enters the crossing condition $\mu_0 H_{\text{eff},n} = nD/(g\mu_B)$.