The Coin Cradle pendulum: physics, fit to the real readings, and discrimination

Buyandsellbullion.com technical report

25 September 2026

Labels: [derived] follows from Maxwell's equations and mechanics; [fitted] estimated from the 31 readings; [inserted] an assumption chosen here, with a test; [memory] a handbook or mint value quoted from memory and not re-verified.

Summary

The dial reading is proportional to the pendulum's total travel, which goes as one over the braking the coin provides:

$$R = \frac{A}{\Gamma}, \qquad \Gamma = \sigma \int_{-\infty}^{\infty} dx \int_0^t d\zeta\, \kappa(x,\ z_1 + (S-1)\Delta z - t + \zeta),$$

where κ is the eddy-current drag per unit conductivity, per unit velocity and per unit depth, computed for the coin's actual outline. With conductivities held at the workbook values, this three-parameter model fits the 31 readings with known dimensions to 13.8% rms under leave-one-out cross-validation. The paper's four-parameter power law does no better (13.9%). The coin's edge and the coin's thickness bringing its top face closer to the magnet carry the fit. Air drag, pivot friction and the arc of the swing do not earn a place.

A gold-plated tungsten fake of a fine-gold coin is 6.3 noise units away from the formula and about 19 away from a same-type reference coin. The 22-carat coins are the weak point: a same-size, same-mass tantalum–tungsten stack reads within 11% of a Krugerrand, and adding rhenium gives an exact match. The ±5% gate used by the paper and the site leaves 61–63% of the genuine coins in this dataset outside it.

1 What is in the data

Source: the Coin Cradle #0001 calibration workbook. The *Tested Coins* sheet has 44 rows; the *Coin Types* sheet lists 13 alloys with conductivity and density.

- **Metals:** fine gold (Type A, 8 rows); 22 ct Au–Cu (D, 4); 22 ct Au–Ag–Cu (E, 5); fine silver (H, 12); 90% silver (K, 8); thaler (J, 1); 40% silver clad with no type (6).
- **Spacers:** 1 (8 rows), 2 (12), 3 (24). There are none at 0 or at 4–6.
- **Recorded:** name, spacer count, "Average Degrees" (a mean; the individual releases are absent), "grams (actual)", condition and type. "Error Margin" is a constant 0.05. It is the assumed tolerance, not a measured spread.
- **Not recorded:** individual releases, release angle, the gap in mm, spacer thickness, coin diameter or thickness, temperature, date and specimen identity. **There are no counterfeits.** "Grams (actual)" is nominal: every fine-silver 1 oz coin is 31.2 g, and the worn and cull halves are both 12.5 g.

- **Replicates:** two gold Kangaroos read 47 and 48 (2% apart); two Morgan dollars read 59 and 82 (33% apart).

Diameters come from the site's coin specifications and from US Mint specifications [memory]. The effective thickness is $t = m/(\rho\pi R^2)$. The five generic rounds and two bars have unsourced dimensions and are held out, as are the six clad coins, which have no conductivity. That leaves a fit set of 31. The dial is ratcheted and read cumulatively (calibration spec; a forum user: "the dial just kept going"), so the reading is taken to be proportional to total travel [inserted].

2 Governing physics

2.1 Equation of motion

Let θ be the pendulum angle, I its moment of inertia about the pivot, M its mass and ℓ_c the pivot-to-centre-of-mass distance. The magnet face sits a distance L below the pivot and moves along $x = L\sin\theta$. With the coin at the bottom of the arc,

$$I\ddot{\theta} + Mg\ell_c \sin\theta = -L\,F_{\text{eddy}}(x,\dot{x}) - \tau_{\text{air}}(\dot{\theta}) - \tau_f\,\text{sgn}(\dot{\theta}). \tag{1}$$

The right-hand side holds the three dissipative terms: eddy-current braking, air drag, and pivot (Coulomb) friction. The device specification and the 2026 paper write the same thing in Rayleigh form, $\ddot{q} + \omega_0^2 q + \partial\mathcal{D}/\partial\dot{q} = 0$; that is equivalent but leaves $\mathcal{D}$ unspecified, which is where all the physics is.

2.2 Eddy-current force in a thin coin at low magnetic Reynolds number

In the magnet's rest frame the coin moves with velocity $-\mathbf{v}$. Ohm's law in the moving conductor is $\mathbf{J} = \sigma(\mathbf{E} + \mathbf{v}\times\mathbf{B})$ with $\mathbf{E} = -\nabla\phi$ (quasi-static, induced field neglected: see §2.3). For a lamina of thickness $d\zeta$ lying in the xy-plane, only B_z drives in-plane current: $\mathbf{v}\times B_z\hat{\mathbf{z}} = -vB_z\,\hat{\mathbf{y}}$ for $\mathbf{v} = v\hat{\mathbf{x}}$. Charge conservation, $\nabla\cdot\mathbf{J} = 0$, and the edge condition $\mathbf{J}\cdot\hat{\mathbf{n}} = 0$ fix ϕ. Writing $\mathbf{J} = \sigma\,v\,\nabla\times(\psi\hat{\mathbf{z}})$ and taking the curl of Ohm's law gives

$$\nabla^2\psi = \frac{\partial B_z}{\partial x} \quad \text{inside the coin,} \qquad \psi = 0 \text{ on the edge.} \tag{2}$$

The power dissipated equals the braking power, $Fv = \int J^2/\sigma\,dV$, so the force is viscous, $F = c(x)\,v$, with

$$c(x) = \sigma\int_0^t d\zeta\;\kappa(x,\,g+\zeta), \qquad \kappa(x,h) = \int_{\text{coin}} |\nabla\psi|^2\,dA = -\int_{\text{coin}} \psi\,\frac{\partial B_z}{\partial x}\,dA, \tag{3}$$

where g is the gap between the magnet face and the coin's top face, h the depth of a lamina below the magnet face, and x the magnet's horizontal offset from the coin centre. Three consequences follow directly:

1. $F \propto \sigma$ exactly (at low R_m). The conductivity exponent is -1 by derivation, not by fit.
2. For an infinite sheet, the Fourier form of (2) gives $\kappa = \frac{1}{2}\int B_z^2\,dA$: the factor $\frac{1}{2}$ is the solenoidal fraction $\langle k_x^2/k^2\rangle$ of an axisymmetric field. So $F \approx \frac{1}{2}\sigma tv\int B_z^2\,dA$ times a geometry factor.
3. *A uniform field induces no current.* Only $\partial B_z/\partial x$ drives ψ. A coin lying wholly under the flat central pole of a pot magnet is barely braked; braking comes from the region where the field changes, mostly the pole edges and the coin's own edge.

Gap dependence. For a point dipole of moment m at height h, $B_z \sim \mu_0 m/(4\pi h^3)$ over an area $\sim h^2$, so for an infinite thin sheet $\int B_z^2 dA \propto h^{-4}$. The force therefore goes as z^{-4} for a coin much wider than the gap. For a coin much smaller than the gap the field over it is nearly uniform, and since a uniform field drives no current (point 3), only the gradient $\partial_x B_z \propto h^{-4}$ counts: $\psi \sim \partial_x B_z\,R^3$ and $\kappa \propto (\partial_x B_z)^2 R^4 \propto h^{-8}R^4$. Real coins sit between these limits, so the effective power is between

z^{-4} and z^{-8} (the brief's z^{-6} is a mid-range value, not a limit). A pot magnet, whose face carries equal and opposite magnetic charge, has a short-range near field ($B_z \sim e^{-kh}$ with $k \sim 2.4/a$ for pole scale a), which falls faster still. The observed factor per spacer, about 2.6–3.1, sits in that range; the model below does not assume any of these limits but computes B_z from the magnet's face charge.

Magnet field. The device specification names a 25 mm pot magnet (F4ME25M5-EYE, N42) with the N pole at the centre and S on the rim. Modelled as a face charge +1 on $r < a_1$ and a balancing negative charge on $a_2 < r < a_3$ (net zero), with $a_1 = 9.5$, $a_2 = 10.5$, $a_3 = 12.5$ mm (inserted; not measured):

$$B_z(\rho, h) \propto \sum_i q_i a_i \int_0^\infty J_1(ka_i) J_0(k\rho) e^{-kh}\, dk. \tag{4}$$

A plain axially magnetised cylinder ($r = 12.5$ mm, length 10 mm) is the alternative tested in the ablation.

Finite radius. Equation (2) is solved numerically (five-point finite differences, 0.25 mm grid, sparse LU) inside the actual coin outline, for each magnet offset x and lamina depth h. Checks: for a 50–80 mm disc with the magnet centred the solver reproduces $\frac{1}{2}\int B_z^2 dA$ to 0.1%; halving the grid changes κ by under 1%. The edge suppresses braking strongly for coins not much larger than the magnet: for a 16.5 mm coin with the magnet centred, the edge-constrained damping is 1–9% of the infinite-sheet value (Fig. 3).

2.3 Skin depth and magnetic Reynolds number

The thin-sheet, low-R_m approximation needs the induced field to be small compared with the applied one. For a sheet the controlling ratio is v/w with $w = 2/(\mu_0 \sigma t)$:

coin	σ (MS/m)	t_{eff} (mm)	w (m/s)
1 oz fine silver	61	2.6	10
1 oz fine gold	44.7	1.9	18
1 oz 22 ct	9.7	1.9	86

Released from 90°, the bob passes the bottom at $v_0 = \sqrt{2gL} = 1.3$–2.0 m/s for $L = 8$–20 cm, so $v/w \le 0.2$ even for silver on the fastest pass. For a thin sheet the drag is $F = F_{\text{low}}\,[1 - \frac{3}{4}(v/w)^2 + \dots]$ (Reitz 1970, dipole over a sheet), a correction of at most 3% on the first pass and less thereafter. In frequency terms, a pass over a 25 mm field region at 2 m/s is a half-cycle at $f \approx 40$ Hz, where the skin depth $\delta = (\pi f \mu_0 \sigma)^{-1/2}$ is 10 mm in silver and 12 mm in fine gold, against coin thicknesses of 1–3.3 mm ($t/\delta \le 0.3$). The approximation holds; this is **derived**, and it also means the test reads the whole thickness of the coin, not the plating.

The thin-disc approximation in the other sense, $t \ll g$, does *not* hold: the coin is 1–3 mm thick and the fitted gap is of the same order. The model therefore integrates over depth (equation 3) rather than using a single sheet.

2.4 Baseline damping

Air drag on the bob and rod is quadratic at these Reynolds numbers ($F \propto \rho_{\text{air}} C_d A v^2$) and pivot friction is a constant torque τ_f. The calibration spec requires more than 50 swings without a coin, against a few swings with one, so both are expected to be small. They enter the fit as one baseline constant b (below) and are tested for whether they earn their place.

2.5 Amplitude decay and the reading

Uniform viscous damping $(I\ddot{\theta} + c\dot{\theta} + k\theta = 0)$ gives exponential decay, $\theta_n = \theta_0 e^{-n\delta}$ with log decrement $\delta = \pi c/\sqrt{Ik}$. The number of swings to fall below a detection threshold θ_{th} is $N = \ln(\theta_0/\theta_{\text{th}})/\delta \propto 1/c$.

Coulomb friction gives linear decay, $\Delta\theta = 4\tau_f/k$ per cycle, and $N = (\theta_0 - \theta_{\text{th}})k/(4\tau_f)$, independent of σ.

Localised eddy damping (this device). The coin brakes the bob only while the magnet is over it, near the bottom of the arc where the path is nearly horizontal. Integrating $M_{\text{eff}}\, dv/dt = -c(x)v$ over one pass gives $M_{\text{eff}}\, dv = -c(x)\, dx$, so the speed drops by a *fixed amount* per pass, independent of speed:

$$\Delta v = \frac{\Gamma}{M_{\text{eff}}}, \qquad \Gamma = \int_{-\infty}^{\infty} c(x)\, dx. \tag{5}$$

Between passes energy is conserved, so the bottom speed after n passes is $v_n = v_0 - n\Delta v$ and the amplitude is $\theta_n = 2\arcsin[(v_n/v_0)\sin(\theta_0/2)]$. The number of passes before the bob stops is

$$N = \left\lceil \frac{v_0}{\Delta v} \right\rceil = \left\lceil \frac{M_{\text{eff}}\sqrt{2gL(1-\cos\theta_0)}}{\Gamma} \right\rceil \propto \frac{1}{\sigma}. \tag{6}$$

With a detection threshold θ_{th}, $N_{\text{th}} = (v_0 - v_{\text{th}})/\Delta v$. The Coin Cradle's dial is ratcheted ("turn clockwise only"; a forum user's report that "the dial just kept going" over base-metal coins), so the reading is taken to be proportional to the total angle travelled:

$$R = G\Big[\theta_0 + 2\sum_{n\geq 1}\theta_n\Big] \xrightarrow{\Delta v \ll v_0} \frac{2G\, M_{\text{eff}} v_0}{\Gamma}\int_0^1 2\arcsin\!\left(s\sin\tfrac{\theta_0}{2}\right) ds \;=\; \frac{A}{\Gamma}. \tag{7}$$

Adding the baseline, the model fitted to the readings is

$$\boxed{R = \frac{A}{\Gamma_{\text{coin}} + b}, \qquad \Gamma_{\text{coin}} = \sigma\int_{-\infty}^{\infty} dx \int_0^t d\zeta\; \kappa(x,\; z_1 + (S-1)\Delta z - t + \zeta)} \tag{8}$$

with κ from equation (3) for the coin's own outline. The gap to the coin's top face is $z_1 + (S-1)\Delta z - t$: the coin lies on a flat bed, so a thicker coin comes closer to the magnet. The exact finite sum in (7) ("saturation", which matters only when the bob stops within a pass or two) is tested as an alternative to A/Γ.

2.6 Known and fitted parameters

quantity	status	source
σ per alloy	held fixed	workbook *Coin Types* sheet
coin diameter	held fixed	site's coin specifications; US Mi
effective thickness $t = m/(\rho\pi R^2)$	held fixed	workbook mass and density
magnet outline a_1, a_2, a_3	inserted	specification (25 mm pot magne
A (magnet strength, bob mass, dial gearing, release angle)	fitted	
z_1 (gap at one spacer, zero-thickness coin)	fitted	
Δz (thickness of one spacer)	fitted	
b (air and pivot), $1/L$ (arc), saturation	tested, dropped	ablation

3 Fit to the real readings

The fit is least squares on $\ln R$ over the 31 readings, with multiplicative errors [inserted].

model	k	in-sample rms	LOO rms	BIC
M1 physics (edge, gap $= z_S - t$, depth integral)	3	0.129	**0.138**	**−116.8**
paper power law $\ln R = k + a \ln \sigma + b \ln V + cS$ (same 31)	4	0.125	0.139	−115.4
paper power law, all 38 typed rows (reproduces $R^2 = 0.985$)	4	0.139	0.151	
M1 without edge currents (infinite-sheet rule)	3	0.265	0.280†	−71.9
M1 without thickness–gap coupling	3	0.175	0.192	−97.8
M1 with a single mid-plane lamina	3	0.130	0.139	−116.3
M1 with a plain cylinder magnet	3	0.130	0.149	−116.3
M1 + air/pivot baseline b	4	0.129	0.137	−113.4
M1 + arc of swing $1/L$	4	0.129	0.138	−113.4
M1 + exact finite-pass sum	4	0.123	0.133	−116.1

† one of 31 folds failed (gap ≤ 0).

Parameters (M1, 300-sample bootstrap, 95%): $\Delta z = 3.36$ mm (3.11–3.59); $\ln A = 7.78$ (7.56–8.26); $z_1 = 18.7$ mm (17.2–19.3). The absolute gap z_1 is *not identified*: the cylinder magnet fits equally well with $z_1 = 8.2$ mm and $\Delta z = 4.2$ mm. The spacer thickness is roughly 3–4 mm under either magnet, and the site's code assumes 3 mm.

Which terms earned their place:

- Edge currents: essential; removing them doubles the error.
- Thickness–gap coupling: essential; removing it raises the out-of-sample error by 40%. This is most of what $V^{1.431}$ absorbs.
- Depth integral: no measurable effect. It is kept because it costs no parameter.
- Baseline $b \to 0$ and arc $1/L \to 0$: dropped.
- Finite-pass sum: dropped. It improves leave-one-out error by 0.005 and worsens BIC, which is within noise.

Residuals. The two largest are specimen effects. The second Morgan dollar reads 33% high and the cull Walking Liberty half 49% high; a worn coin carries less metal than its nominal mass, so it brakes less. Without them the rms is 9.6%. There is a systematic offset by metal: fine silver reads 16% below the model, and 90% silver about 10% above. This points to the workbook's silver conductivities or to the magnet outline; a test is to measure σ with a meter.

Held-out checks. The generic silver rounds read 6–18% below prediction, the same offset as the other fine silver. The bars read 20–33% above; their shape is guessed. The clad coins infer $\sigma = 38$–59 MS/m.

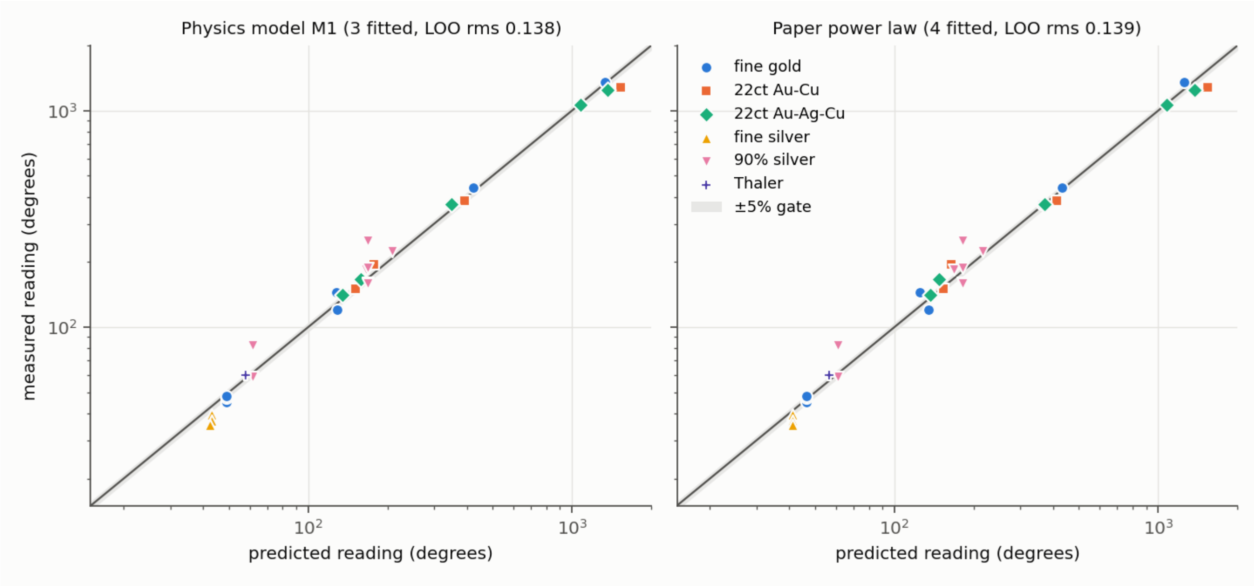

Figure 1: Measured against predicted readings. Grey band: the ±5% gate.

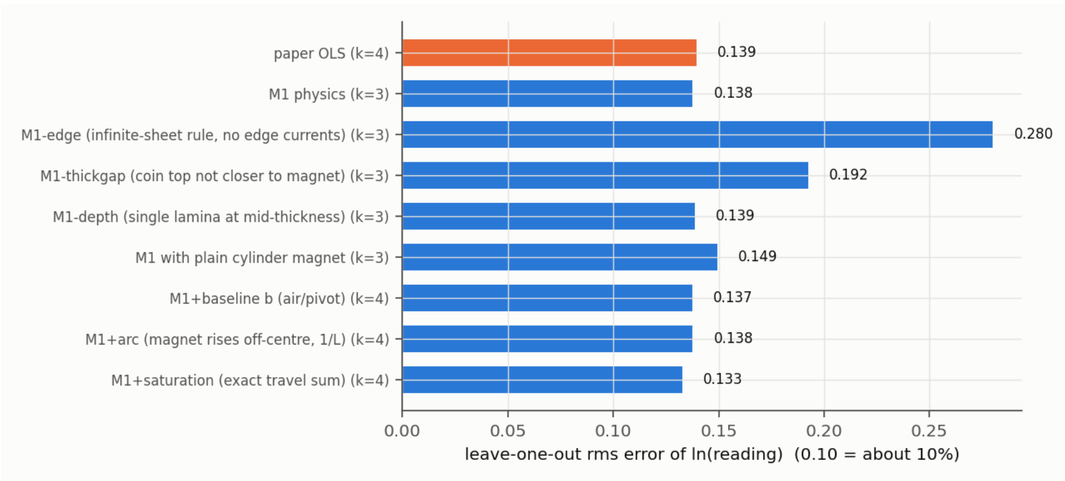

Figure 2: Ablation: leave-one-out error for each model variant.

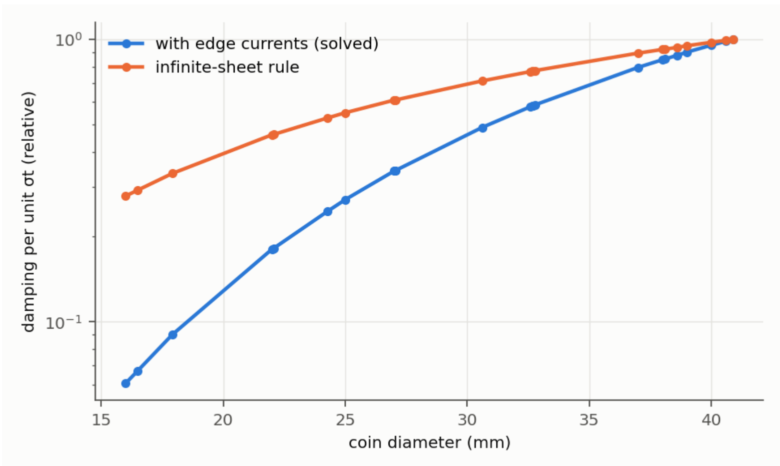

Figure 3: Braking per unit σt against coin diameter at a fixed gap. The edge suppresses small coins far more than the infinite-sheet rule assumes.

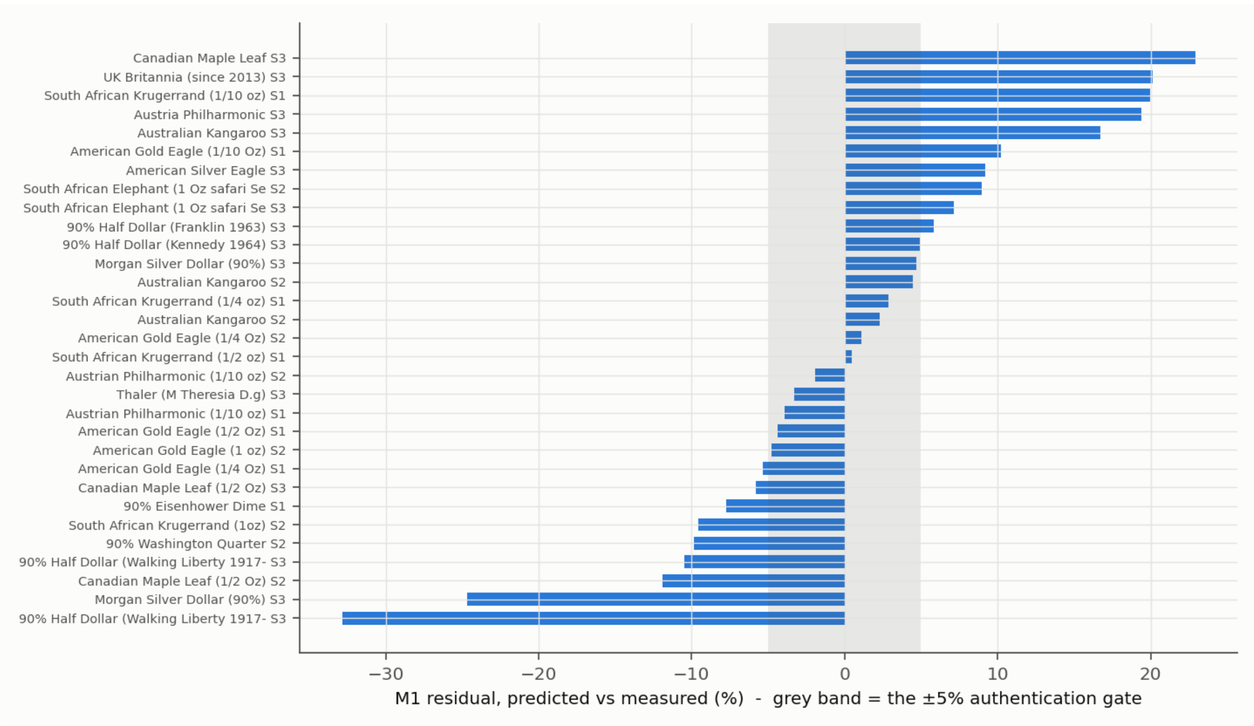

Figure 4: Residuals by coin, M1. Grey: the ±5% gate.

4 Discrimination

Physics [derived]. For a same-size layered coin the braking is $\sum_i w_i \sigma_i t_i$, with depth weights w_i from the fitted model: the top sixth of a Krugerrand carries 22% of the weight and the bottom sixth 12%. For a symmetric stack,

$$\ln \frac{R_{\text{fake}}}{R_{\text{genuine}}} = -\ln \frac{\sigma_{\text{eff}}}{\sigma_{\text{genuine}}},$$

which to first order is *independent of gap and magnet.*

A counterfeiter who matches size and mass (to within 0.5%) can reach any point in the convex hull of their materials in the (ρ, σ) plane. The closest reachable point is found by linear programming using handbook values [memory].

Noise. Against the formula the noise is 0.138, the leave-one-out error. Against a same-type reference coin it is 0.045, the scatter of like fine-silver coins; the Kangaroo pair suggests about 0.02.

target	fake (same size and mass)	$R_{\text{f}}/R_{\text{g}}$	vs formula	vs reference
fine gold 44.7	gold-plated W	2.37	6.3	19
fine gold	Ag 0.28 + Ir 0.72	1.36	2.2	6.8
22 ct Krugerrand 9.7	gold-plated W (10% heavy)	0.51	4.8	15
22 ct Krugerrand	W + Cu, mass-matched	0.38	7.1	22
22 ct Krugerrand	**Ta 0.72 + W 0.28**	**0.90**	**0.8**	**2.3**
22 ct Krugerrand	W 0.41 + Bi 0.26 + Re 0.34	1.00	0	0
22 ct Eagle 11.1	Ta 0.60 + W 0.40	0.92	0.6	1.9
fine silver 61	W 0.15 + Cu 0.85	1.14	1.0	2.9

- **Fine gold is safe** in this model. No cheap metal is both denser and more conductive than W. To pass a $\pm 2.5\sigma$ gate, a same-size, same-mass fake must be at least 48% gold by volume (formula) or 80% (reference coin).
- **22 ct coins are the weak point**, because 9.7–11.1 MS/m lies inside the hull of cheap dense metals.
- Flipping the coin separates asymmetric stacks. For the Ta/W stack, Ta on top gives ln ratio -0.08 and W on top -0.24, a difference of 3.3 reference-noise units. A symmetric sandwich defeats the flip.
- The acoustic test separates Ta, W and Re, which are two to five times stiffer. That is outside this model.

Design.

- Compare with a same-type reference coin on the same unit, not with the formula. This cuts the noise about threefold.
- Set the threshold at about $\pm 2.5\times$ the measured replicate noise per coin type, in place of $\pm 5\%$.
- Test both faces.
- Use the largest gap at which the coin still dominates the no-coin baseline and the reading stays on the dial. Separation is gap-independent to first order, and a larger gap slightly flattens the depth weighting.
- A magnet whose field changes across a 16 mm coin would brake fractional coins more strongly than the 25 mm pot magnet does.

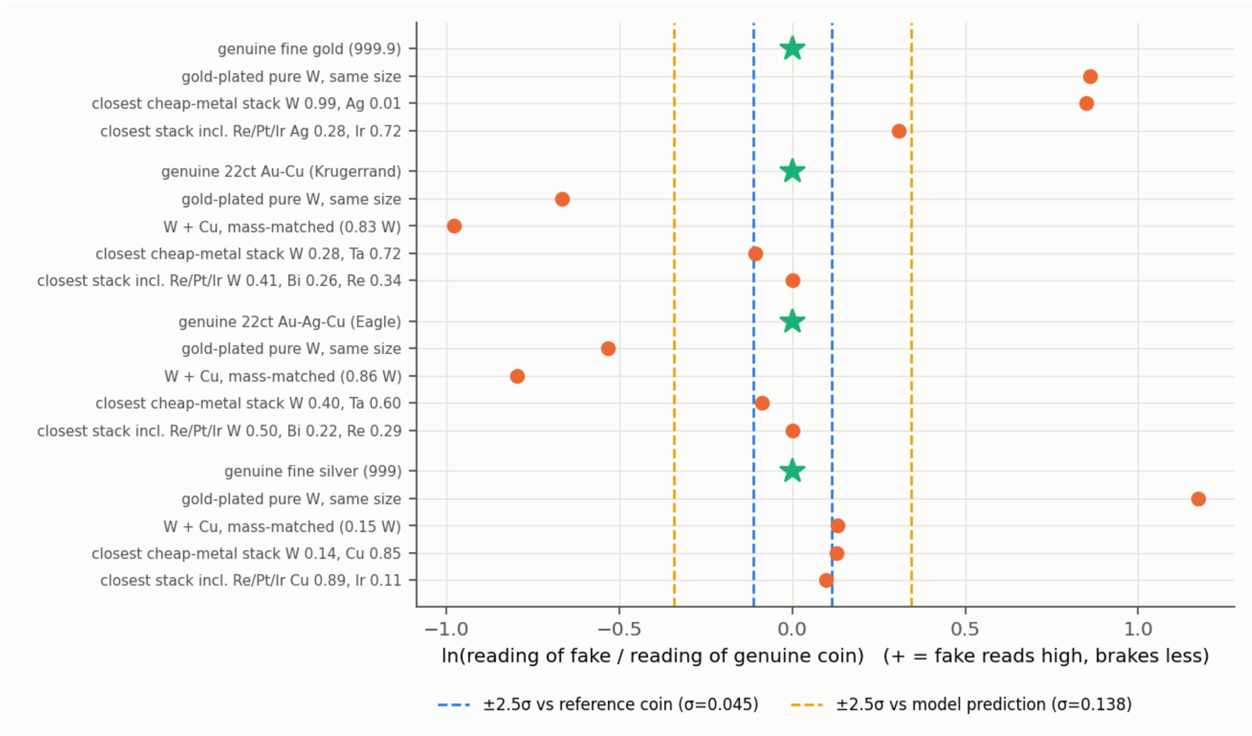

Figure 5: Where fakes land relative to the genuine coin. Dashed lines: $\pm 2.5\sigma$ gates for the two noise scales.

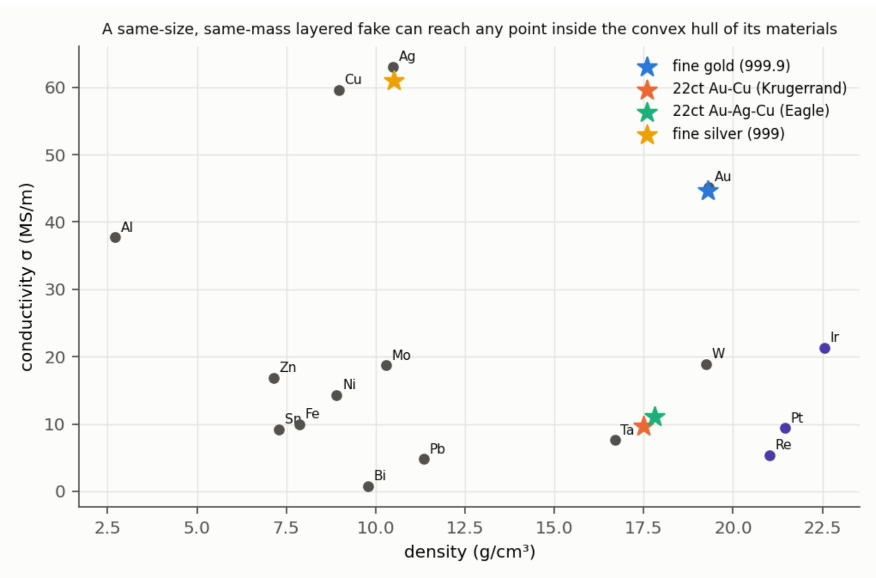

Figure 6: Density and conductivity of candidate fake materials (handbook values) and the target alloys.

5 The existing documents and site (as of 25 September 2026)

The site has since been corrected on each of the points below. **What is established, and how:**

- Proven by the data: the reading falls roughly as $1/\sigma$; it rises by a factor of 2.4–3.1 per spacer; it falls with coin size.
- Proven by physics: $F \propto \sigma$; the whole thickness is read; a fine-gold fake cannot be made from cheap metals.
- Fitted: 488, 2.90 and 1.431.
- Conjecture: all fake performance, since no fake was measured.

Overclaims:

1. **The ±5% gate** (the 2026 paper, the site's code and the dial's green band). The paper's own fit leaves 63% of its 38 genuine coins outside it, and the calibration spec records 53% within. As deployed, it would reject most genuine coins.
2. **"$R^2 = 0.985$, highly accurate".** This is R^2 on $\ln R$ across a 40-fold range. The residual that matters is 15% out of sample.
3. **"$V^{1.431}$ disproves the wire-loop theory ... exponentially increasing the drag".** It is a power law, not an exponential. V merges thickness, the coin's top face moving closer to the magnet, and diameter, and coins of equal volume but different shape read differently.
4. **"38-sample dataset of authentic coins and known counterfeits"** (site articles). The workbook contains no counterfeits.
5. **"The Mathematical Impossibility of Counterfeiting"** (the February note and the live proof page, headed "Formal Proof" with a QED mark). The algebra is right *given* the fitted law, but:
   - it assumes a homogeneous fake, when layers make σ_{eff} and ρ_{eff} continuous mixtures;
   - it uses equalities where the test has tolerances;
   - it asserts "no known alloy" without showing it;
   - the EM half is spoofable for 22 ct targets, although it holds for fine gold;
   - the acoustic half makes the same homogeneity assumption, and a composite's bending stiffness is set by its outer layers.

   It argues why two orthogonal tests help. It is not a proof.
6. **"Tungsten 17.9 vs gold 44.0, more than 2:1, impossible to fake"** (site). This is true for fine gold only. For 22 ct targets, W is *more* conductive than the coin.
7. **Spacer geometry is inconsistent.**
   - The site code has 3 mm per spacer.
   - The kit design has 1 mm spacers placed *under the coin*, which closes the gap, the opposite direction to the Cradle.
   - The forum post says the spacers go under the magnet.
   - 488 and 2.90 belong to Cradle #0001. "Same physics, same accuracy, no calibration" for the kit is untested.
8. **Simulator inputs.**
   - $\sigma_{\text{Eagle}} = 8.5$ in the simulator against 11.1 in the workbook.

- Fine gold 44.0 against 44.7.
- Its tungsten fakes are not mass-matched for 22 ct targets.
- Its spacer-0 predictions are extrapolations; there is no reading below one spacer.

6 Assumption ledger

assumption	status	test
reading $\propto$ total travel	inserted	film releases; compare frame-counted travel with the dial
low R_m, full-thickness penetration	derived ($v/w \leq 0.2$, $t/\delta \leq 0.3$)	readings independent of release angle beyond the A/Γ scaling
drag $\propto \sigma$	derived	two same-size coins of different alloy
coin top closer by t	derived (flat bed)	shim a thin coin up 1 mm and compare
pot-magnet outline 9.5/10.5/12.5 mm	inserted	Hall-probe scan of the face
spacer thickness	fitted, 3.4 mm	calipers
workbook σ table	inserted (source not stated)	eddy-current meter per alloy
round and bar dimensions	inserted	calipers
layered fake = $\sum w\sigma t$	derived	test the plated-W Eagle, both faces
handbook σ, ρ (Ta, Re, W, ...)	memory	check against CRC
noise 13.8% and 4.5%	fitted, few replicates	10 releases each on 3 coins

7 What to do next

1. Record individual releases: 10 per coin, including the plated-tungsten Eagle already in hand. That gives the first measured fake and a real noise figure.
2. Measure the spacer and, if possible, map the magnet face. That fixes z_1 and removes the inserted geometry.
3. Before selling, change the gate to a reference-coin comparison, stop calling the note a proof, and remove "known counterfeits" from the site. (These site changes have since been made.)