Ping test review: what the February method did, what it got wrong, and what a phone can do now

Buyandsellbullion.com technical report, 26 September 2026. A review of the February 2026 ping work and the site's /authenticate/ping page, followed by new physics, a phone-signal study, a Bayesian inversion, and a fusion with the pendulum.

Labels: **[measured]** on our own coins; **[sourced]** from a paper, with the page; **[scraped]** detected by the February pipeline in YouTube audio; **[inferred]** my reading of evidence; **[modelled]** computed here; **[assumed]** a value I chose, with a test for it; **[memory]** a handbook figure not re-verified this session.

Short answer

- **What the old method was.** One number per coin: the (0,2) frequency of a free circular plate from Kirchhoff theory, checked against a phone FFT with a ±3% gate. No acoustic measurement of our own was made. Of the 67 database values, 14 come from two papers [sourced], 46 are marked "YouTube verified" [scraped], and 7 are theory [modelled].
- **The "YouTube verified" values are theory echoed back.** A detection counted if it lay within 5% of *any* coin's theoretical value, and a coin was "verified" if the mean lay within 3%. The ±5% windows around the 67 values cover 77% of the 2–12 kHz band. By the pipeline's own transcript labels, only 2 of the 46 coins have two or more detections, named as that coin, within 3%.
- **Accuracy against the only independent data** (Vinteler 2025, 13 bullion coins, measured with a phone): the model that set the YouTube search windows was 38% high on average. The best February model was 4.7% rms, max 9.9%, and its elastic constants were fitted to those same coins. The report's "0.90% mean error" is circular.
- **The counterfeit "signatures" are not signatures.** All four sit at 1.5–2.1 kHz, in four videos of different fakes. In the tungsten-Maple video the narrator says the fake "makes absolutely zero sound"; the detections fall during talking and the outro. Physics says a tungsten-cored Maple rings about 2.4 times *higher* (11.5 kHz), not 66% lower.
- **What is newly possible.** A phone at 48 kHz sees three to five plate modes of a bullion coin. It measures each one to about 0.001% [modelled], far below any real spread between coins. Frequency *ratios* cancel size, thickness and stiffness scale, so one ping gives Poisson's ratio to about ±0.007. Per-mode Q is measurable to 1–2% on single modes. Mode-level physics, not machine learning, carries the discrimination.
- **Fusion. One family of fake beats all three tests; nothing built from cheap metals comes close.** With at most 30% gold by volume and materials limited to W, Ta, Mo, Pb, Bi, Sn, Zn, Fe, Ni, Cu, Ag, tungsten heavy alloy and epoxy, the closest fake to a Krugerrand is 15.5 noise widths away on scale + phone-read pendulum + five-mode ping (bonded), and 11.0 as an unbonded foil deck. With platinum allowed, a bonded stack of 22 ct skin, Pb, Zn and a Pt core matches every channel to within 0.9 widths, or 2.8 once damping is added. Fine gold is safe: nothing gets within 19 widths of a Britannia [modelled].
- **Next step.** Record ten genuine Krugerrands and ten genuine Britannias, five taps each, on one phone, plus the plated-tungsten Eagle already held. That settles the one number everything above leans on: the coin-to-coin spread of the mode ratios.

1. What was done

- Read every script, data file and run of the February analysis, both papers (Manas 2015; Vinteler 2025), the LLNL smartphone report, the site's ping page and the 20 February screenshot, and the pendulum and phone-cradle notes.
- Wrote an exact Kirchhoff solver (Bessel roots) and a Rayleigh–Ritz Mindlin solver for circular plates that are radially piecewise and symmetrically layered, including an energy-consistent transverse shear stiffness for layered sections (`plates.py`). Numpy and scipy only.

- Validated both against Vinteler's own printed classical and Mindlin predictions for 13 bullion coins (`validate_plates.py`).
- Rebuilt the February provenance and accuracy numbers from the raw run files (`audit_old.py`).
- Modelled rim and relief (`relief_rim.py`), phone recording and mode extraction (`phone_sim.py`), and a Bayesian inversion (`inversion.py`).
- Searched 3,600 bonded layer stacks and 1,440 unbonded foil decks for the fake that best matches a genuine Krugerrand, and a genuine Britannia, across mass, pendulum, five modes and damping (`fusion.py`, `fusion_figs.py`).

2. What was found

2.1 The physics the old method used

A coin rings as a free thin disc. Its bending modes are labelled (s, n): s nodal circles, n nodal diameters. The lowest is (0,2), then (1,0), (0,3), (1,1) and (0,4). Kirchhoff plate theory gives

$$f_{sn} = \lambda^2_{sn}(\nu) / (2\pi\, a^2) \cdot \sqrt{D / \rho h}, \quad D = E h^3 / 12(1 - \nu^2)$$

so $f \propto (h/a^2)\cdot\sqrt{E/\rho(1 - \nu^2)}\cdot\lambda^2(\nu)$. Frequency scales linearly with thickness, inversely with diameter squared, and with the square root of stiffness over density. The edge is free; λ^2 depends on ν (5.36 at ν = 0.30, 4.95 at ν = 0.42 for (0,2)) [modelled, exact roots]. Coins are "moderately thick" (a/h = 6.6–9.2 for 1 oz), so Mindlin's correction for shear and rotary inertia lowers the higher modes by 1–6% [sourced: Vinteler §4; reproduced here].

- `coin_physics.py` fixes $\lambda^2 = 5.25$ (right only for $\nu \approx 0.33$), uses the **rim** thickness and handbook E, and drove the YouTube search windows. `kirchhoff_calculator.py` uses the mass-based effective thickness $h' = m/(\rho\pi a^2)$, which is Vinteler's choice and the right one, with per-alloy E values averaged from Vinteler's fits.
- The pipeline's "harmonic" logic looks for 2f, 3f, 4f. Plate modes are not harmonics: their ratios are about 1.75, 2.27, 3.7 and 3.9. That logic is mis-specified.
- My solvers reproduce Vinteler's printed predictions to 0.1–0.4% for 12 of 13 bullion coins [modelled]. The 13th, the gold Maple, is 8% off on every mode: his Table 3 E for that coin does not reproduce his own Table 5 frequencies. That is an inconsistency in the paper.

2.2 Where the frequencies came from

Source	Coins	What it is
Vinteler 2025	12	Measured by bouncing or pinging, recorded on a smartphone at 48 kHz, FFT in the Spectroid app [sourced, pp. 89–90]. Faithfully copied, except that the gold Maple (4,830) and Britannia (3,585) are the means of split pairs 4,590/5,060 and 3,470/3,700. Peace dollar copied from Morgan.
Manas 2015	2	"Vreneli" is the 1890 Helvetia 20 Fr (5,375/5,379 Hz) and "Napoleon" the 1839 Louis-Philippe (5,531/5,563). Same specification, different coins [sourced, Table 5].
"YouTube verified"	46	Peaks from 72 YouTube videos, accepted by closeness to theory [scraped].
"Kirchhoff calibrated"	7	Theory × a per-alloy factor [modelled].

- **Nothing was measured in-house.** No audio recording exists.
- **Misattributions.** `coin_physics.py` lists "verified measurements" (Ducat 3,460, Libertad 4,130, Panda 4,249, Silver Eagle 5,622, Washington quarter 10,328) and `kirchhoff_calculator.py` attributes them to Manas and Vinteler. Manas measured only 20 Fr coins. Vinteler gives 3,800 for the Silver Eagle and 6,370 for the quarter. These were the site's earlier values; their origin is not traceable.
- **The 20 February screenshot** shows reference values labelled verified that differ from Vinteler's measurements by up to 38% (Silver Eagle 5,234 against 3,800; Britannia gold 4,867 against 3,585).

How the YouTube verification worked. `apply_verifications.py` and `apply_cross_batch.py` kept a detection (SNR ≥ 20) if it lay within 5% of the *nearest* database value, over all 67 coins, whatever the video was about. A coin became "verified" with two such detections whose mean lay within 3% of its value. At the time that value was the theory.

- The ±5% windows round the 67 values cover **77%** of 2–12 kHz, and **61%** of all 857 detections pass [modelled from the run files, Fig. A2]. The filter is not a test.
- Reconstructing the rule on the saved runs, 444 detections support the 46 coins. **Only 20 come from a video about that coin**, and 40 of the 46 coins have none. The Buffalo is supported by Britannia, Kangaroo and Panda videos; the 1/4 oz Krugerrand by Buffalo videos; the 1 oz Philharmonic by 13 detections from the counterfeit videos [inferred from the saved run files].
- By the pipeline's own transcript labels, only **Panda 30 g and Sovereign** have two or more detections named as that coin within 3%.
- All 41 recorded deviations are ≤ 2.97%: the acceptance rule guarantees it (Fig. A2, right).
- The report's "0 narrative mismatches across 727 detections" checked coin *names* only. A genuine Eagle at 10,382 Hz, 54% from its database value, is marked "CORROBORATED". That peak is in fact Vinteler's (0,3) mode of the Eagle, 10.41 kHz: a ping's loudest peak is often not (0,2).

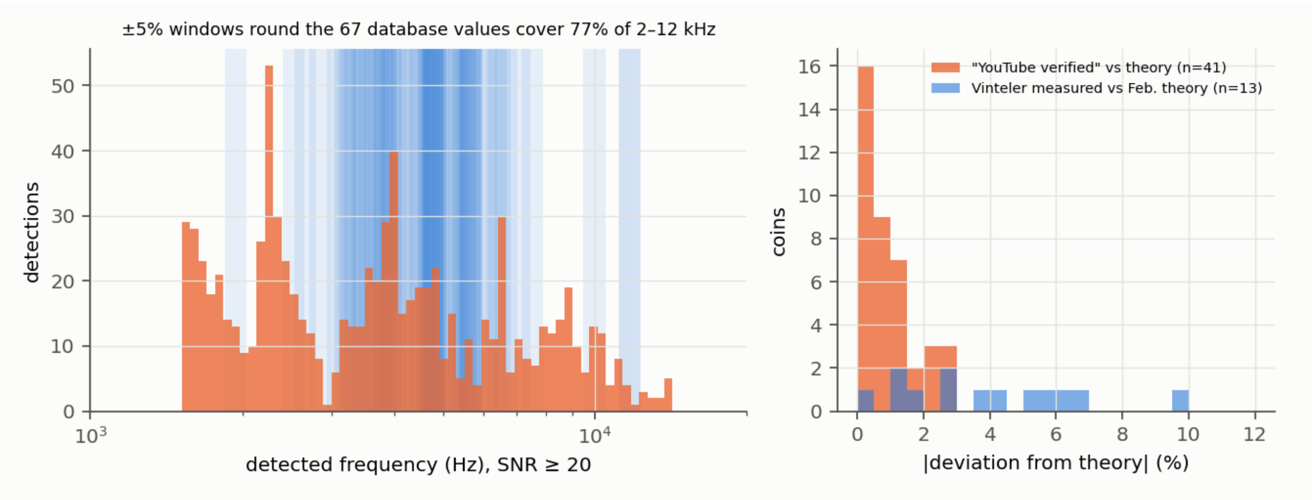

2.3 The calibration and the fudge factors

- `kirchhoff_calculator.py`: 8 per-alloy E values (4 averaged from Vinteler's fits to the very coins they are tested on, 1 from Manas, 3 "interpolated"), with ν and λ set by hand per alloy.
- `build_calibration.py`: 8 multiplicative per-alloy factors, fitted to the "verified" set, which was selected for being within 3% of the theory. The factors come out at 0.99–1.03 by construction.
- That is **16 fitted continuous numbers plus 8 hand-set eigenvalues**, validated on data selected by its agreement with them.

2.4 Prediction error against independent data

Vinteler's 13 bullion coins are the only independent measurements. Against their (0,2):

Model	Mean	rms	Max
`coin_physics.py` (rim thickness, fixed λ^2, handbook E): set the search windows	+38%	40%	59%
`kirchhoff_calculator.py` (E fitted to these coins; in-sample)	−3.8%	4.7%	9.9%
Handbook E, ν, ρ; effective thickness; exact $\lambda^2(\nu)$; Mindlin	+2.9%	4.0%	8.4%

- The limit is not the theory. Coins of one alloy have different effective moduli: Vinteler's fits give 61 and 76 GPa for two fine-gold 1 oz coins, and 92 and 104 GPa for two 22 ct Au–Cu coins [sourced, Table 3]. Minting texture, rim and relief all go into that number.

- Rim and relief alone move f02 by 3–15% at fixed mass [modelled, Fig. R1]. **A theory-only database cannot support a ±3% gate.** The reference must be measured genuine coins of the same type.
- The true out-of-sample test is the higher modes. With E and ν fitted to (0,2) and (1,0), Mindlin predicts (0,3), (1,1) and (0,4) for the 13 coins to 3.4% rms (2.6% without one outlier). Classical theory reads 2.6% high on average [modelled from Vinteler's Tables 3 and 5].
- For the Krugerrand, a rim band 0.4 mm wide cuts the higher-mode error from 1.4% to 1.0%, and the fitted E moves from 92 to 96 GPa [modelled]. The fitted "E" absorbs geometry.

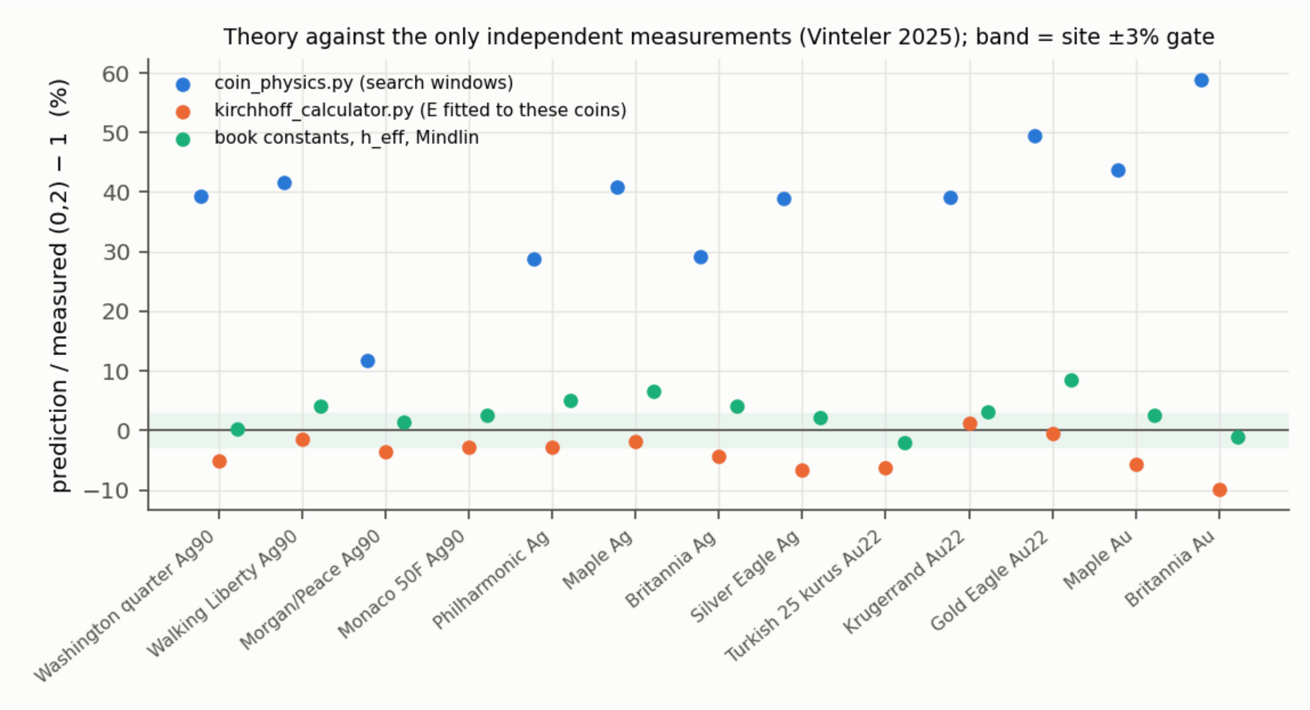

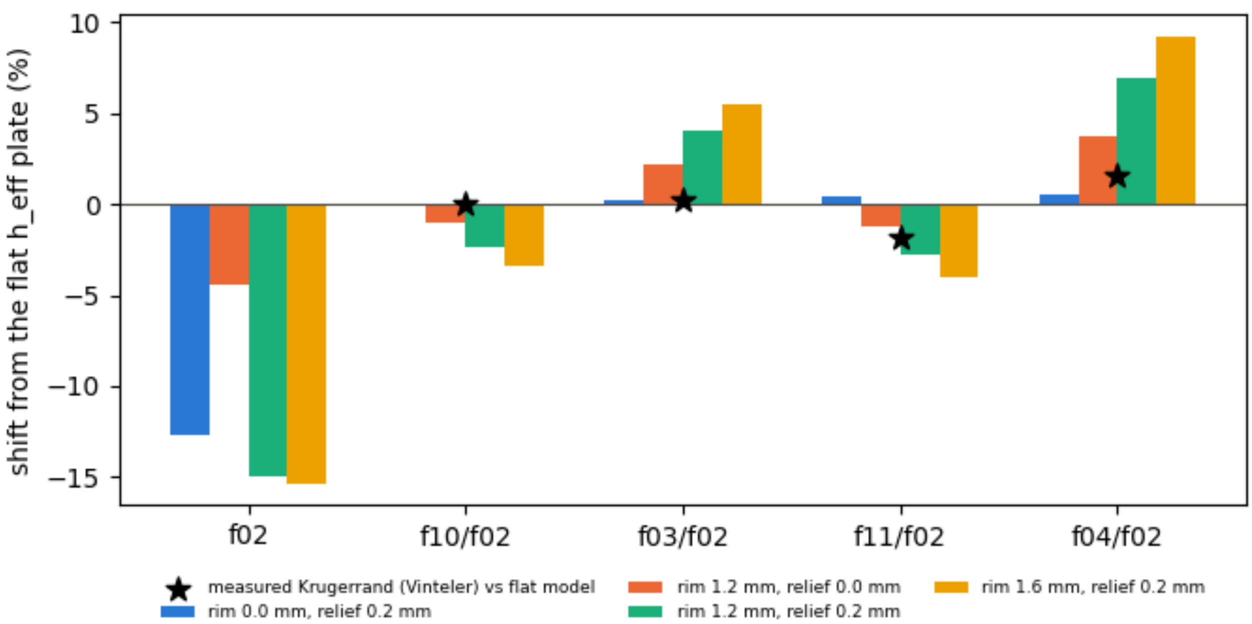

2.5 The verification report's conclusions

Claim	Holds?
"Vinteler et al. — laser vibrometry and precision microphones"	No. One author, a smartphone at 48 kHz and the Spectroid app.
"Manas & Manas"	One author.
"Strong agreement across all methods", "0.90% mean error"	Circular: agreement was the acceptance rule.
"0 narrative mismatches", "zero false identifications"	The audit matched names, not frequencies.
"Kirchhoff remarkably accurate, ~1.4%"; "±5% tolerance well justified"	No. 4–5% rms against independent data, up to 10%.
"81% verified"	14 of 67 have an independent measurement, all taken from two papers.
counterfeit_signatures.json: "counterfeits ring 40–70% lower regardless of type"; "tungsten core damping lowers frequency"	No. The 1.5–2.1 kHz detections are narration and noise. Damping does not lower frequency materially. A tungsten core raises it (Fig. A3).

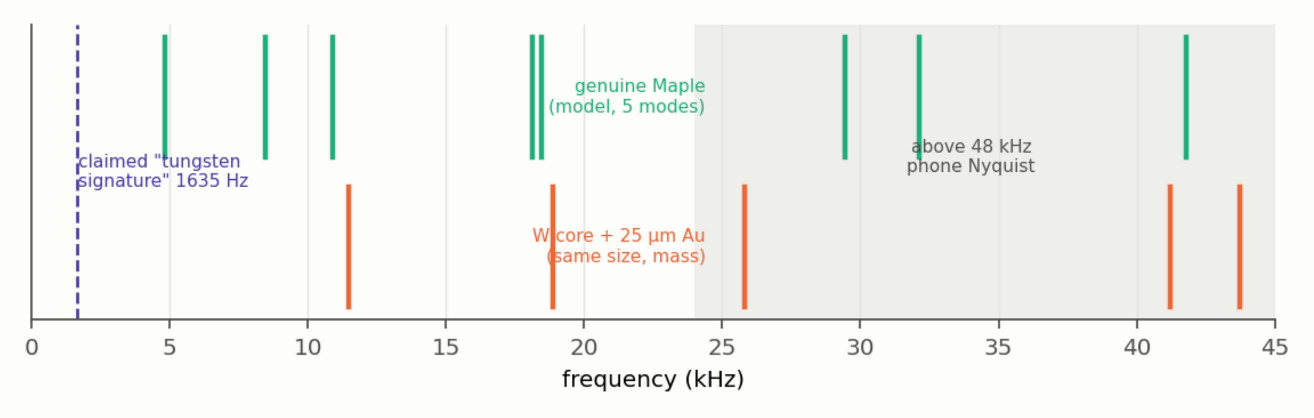

2.6 What the site page does with the microphone

- It uses a WebAudio `AnalyserNode` (FFT 16,384, smoothing 0.75), which averages spectra across animation frames rather than analysing the transient. It asks for echo cancellation, noise suppression and AGC off, which is right.
- It takes the top three peaks between 1.5 and 14 kHz and keeps whichever is closest to the expected value. It passes if that peak is within ±3%, and shows "match %" as 100 minus the deviation, so a coin 10% off shows "90% match". Decay time is counted in animation frames.
- **A genuine split Maple (4,590/5,060 Hz) fails the ±3% gate on both peaks** (−5.0%, +4.8%).
- The instructions say to balance the coin on a fingertip. Manas found that finger contact adds a second set of frequencies.

2.7 What a phone can measure now [modelled]

Modes in band. At 48 kHz (usable to about 20 kHz), 1 oz coins show all five modes, and 1/10 oz coins and the Sovereign show three:

Coin	a/h	(0,2)	(1,0)	(0,3)	(1,1)	(0,4) kHz
Krugerrand 1 oz	7.2	4.82	8.43	10.93	18.22	18.63
Sovereign	9.3	5.59	9.80	12.79	21.5	22.0
Krugerrand 1/10 oz	9.2	7.48	13.11	17.11	28.8	29.5
Britannia 1 oz Au	8.5	3.58	6.18	8.16	13.55	14.01
Philharmonic 1 oz Au	12.4	2.20	3.81	5.06	8.47	8.79

Precision. A synthetic Krugerrand ping included a split (0,2), random mode amplitudes, a click, noise, MEMS roll-off and 16-bit samples.

- An FFT with interpolation recovers each mode to 0.0001–0.002% at 20–40 dB SNR, once the search window is narrower than the (1,1)–(0,4) spacing.
- Q comes from the band-limited envelope to 0.3–2% on single modes. It is poor (15–37%) on the split (0,2), whose beating spoils the fit. Use (1,0) and (0,3) for damping.
- The phone's own clock (±50 ppm) and the temperature coefficient of E (a few × 10^{-4} per K [memory]) are negligible.
- **So measurement noise is not the limit.** Coin-to-coin variation and how the coin is held are. Both are unmeasured. I carry 1% on f02 and 0.5% on each ratio [assumed].

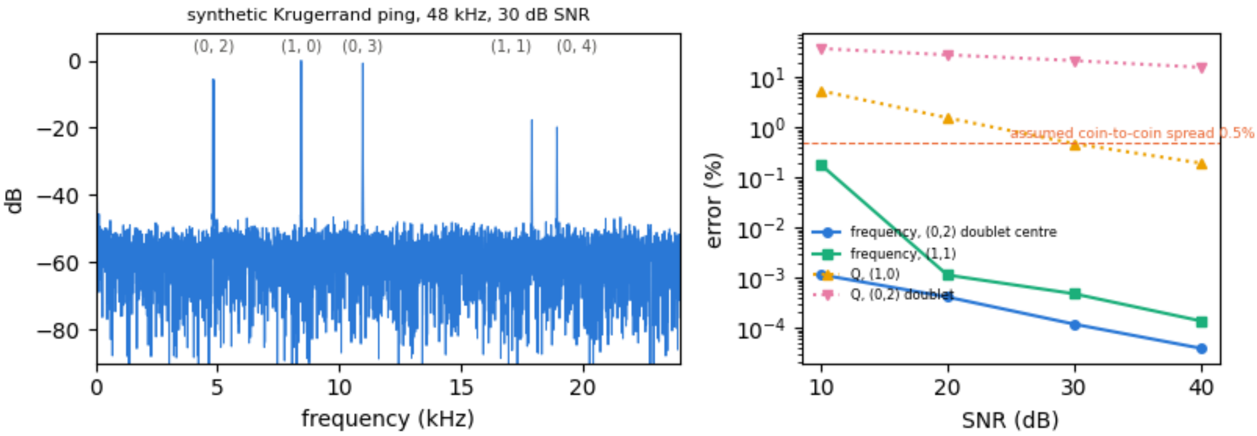

Protocol ("tap and hold"). Hold the coin lightly at its centre between two fingertips. The centre is a node of every mode with nodal diameters, as in Vinteler's pincer method. Tap the rim with a second coin or a pen barrel, keep still for one second, and repeat five times at rim points 45° apart, so that both halves of the (0,2) doublet and the (1,0) mode get excited. The (1,0) mode has its maximum at the centre, so the hold damps it; add one or two LLNL-style bounce drops as a free-edge check. Capture raw PCM at 48 kHz with processing off (WebAudio constraints; the iOS measurement mode; Android's unprocessed source [memory]), and analyse the recording offline, not in an AnalyserNode.

2.8 Bayesian inversion: what one ping can and cannot tell you [modelled]

- **Homogeneous plate, five measured Krugerrand modes plus mass and diameter.** ν = 0.331–0.361 (95%), well pinned. E spans 85–196 GPa and the implied density 17.1–22.9 g/cm^3, because thickness and stiffness trade against each other and only the small Mindlin terms separate them. **A ping does not measure density.** The scale and callipers do.
- **Any symmetric section, layered or not,** is fixed for bending by four numbers: stiffness per mass $c = \ln(D_{11}/I_0)$, $\nu_{eff} = D_{12}/D_{11}$, a shear flexibility S and a rotary inertia R. With the noise above, the ping gives c to 2.6%, ν_{eff} to ±0.007, S only to ±35% and R not at all (prior-limited; S and R correlate at −0.94).
- **So the ping sees two numbers well and a third weakly.** A counterfeiter who matches stiffness per mass and ν_{eff}, and roughly the shear flexibility, passes the ping. The fusion asks whether that is possible *while* also matching mass and $\sigma \cdot t$.

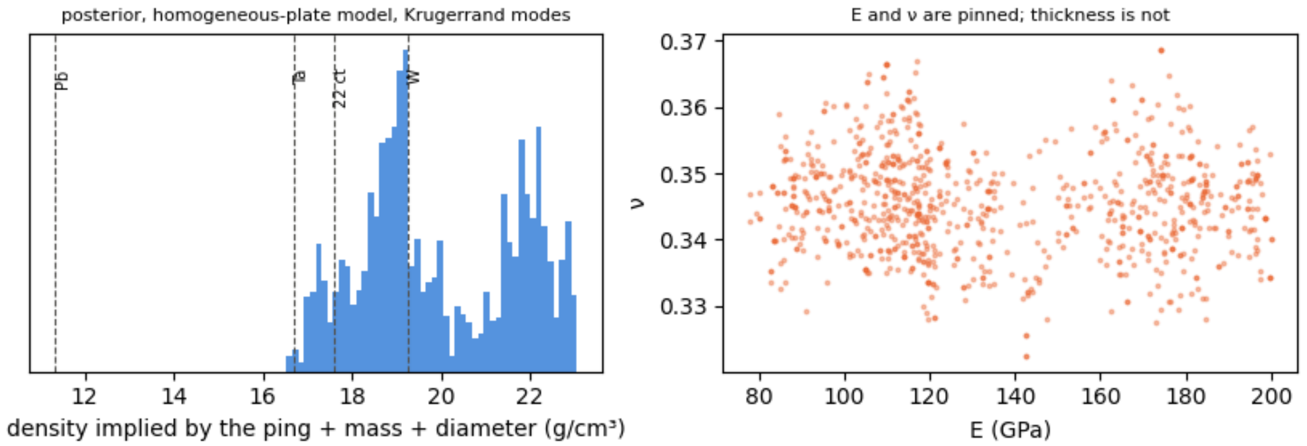

2.9 Where machine learning helps, and where it does not

- **Classical signal processing wins** on the measurement itself: five damped sinusoids with known approximate frequencies are a textbook parametric problem. FFT or Prony estimates are already about 100 times finer than coin-to-coin spread. The physics plus a Bayesian comparison with a reference library uses every bit of it, with honest uncertainty. A learned model would add error and hide it.
- **ML genuinely helps** with data hygiene in uncontrolled phone audio:
    - finding the real strikes and rejecting double taps, finger-damped taps, clipping, speech and music (a small classifier on spectrogram patches);
    - guiding the user in real time ("tap nearer the rim", "too quiet");
    - after hundreds of genuine recordings, a one-class anomaly score on spectral shape as a secondary flag.
- **General pretrained audio embeddings are a poor fit.** Most are trained on 16 kHz audio, which discards everything above 8 kHz, where most coin modes live. They also learn the phone and the room.
- **The lesson from February:** Whisper and vision models used to *label* YouTube audio produced circular ground truth. AI is useful around the measurement, not as its source of truth.

2.10 Fusion with the pendulum [modelled]

Set-up. The fake copies the genuine coin's outline and rim exactly (a die-struck shell), so diameter and thickness carry no information. It must match the rest, each compared with a same-type genuine reference and divided by a 1σ noise width:

Channel	Noise width	Status
mass	0.3%	[assumed]
pendulum σ·t, Coin Cradle	4.5%	[v1 estimate]
pendulum σ·t, phone-read cradle	1.5%	[assumed, phone-cradle note]
f02	1.0%	[assumed; Vinteler cites < 0.11 kHz sd over 300 coins]
f10/f02, f03/f02, f11/f02, f04/f02	0.5% each	[assumed]
damping of (0,2) and (1,0)	30%	[assumed]

- **Constructions.** A symmetric stack under a genuine-alloy skin (≥ 20 µm): layer A, layer B and a core, from 16 materials (3,600 ordered triples). It may sit inside a gold ring as a plug (radius 0.85 a), or be a fully plated blank. Unbonded foil decks were searched too: skin, one layer and k = 1–8 loose foils, with friction loss 0.01 added per mode [assumed]. Gold is capped at 30% of the volume, with 10% and 50% as sensitivity.
- **Model.** Mindlin plate with the laminate's own D_{11}, D_{12}, energy-consistent shear stiffness and rotary inertia. The pendulum uses the v1 depth weighting (22% in the top sixth), and a plug is area-

weighted (crude). Loss factors are handbook orders of magnitude [assumed]. For each material combination, Nelder–Mead from three starts minimises the summed squared noise widths. The reported distance D is the root of that sum.

Results (Krugerrand 1 oz unless stated):

Tests used	Cheap metals, bonded	All metals incl. Pt, Ir, Re, bonded	Unbonded foil deck (all metals)
A scale + f02 (February method, with a scale)	12.7	0.0	0.0
B scale + Coin Cradle	not run	0.0	not run
C scale + Cradle + f02	13.4	0.0	not run
D scale + five modes	13.1	0.3	6.6
E scale + Cradle + five modes	13.7	0.8	7.8
F scale + phone cradle + five modes	15.5	0.9	7.8 (11.0 with cheap metals)
G F + damping	15.5	2.8	18.3
Britannia 999.9, B / D / F	not run	4.6 / 5.2 / 23.8	F: 19.0

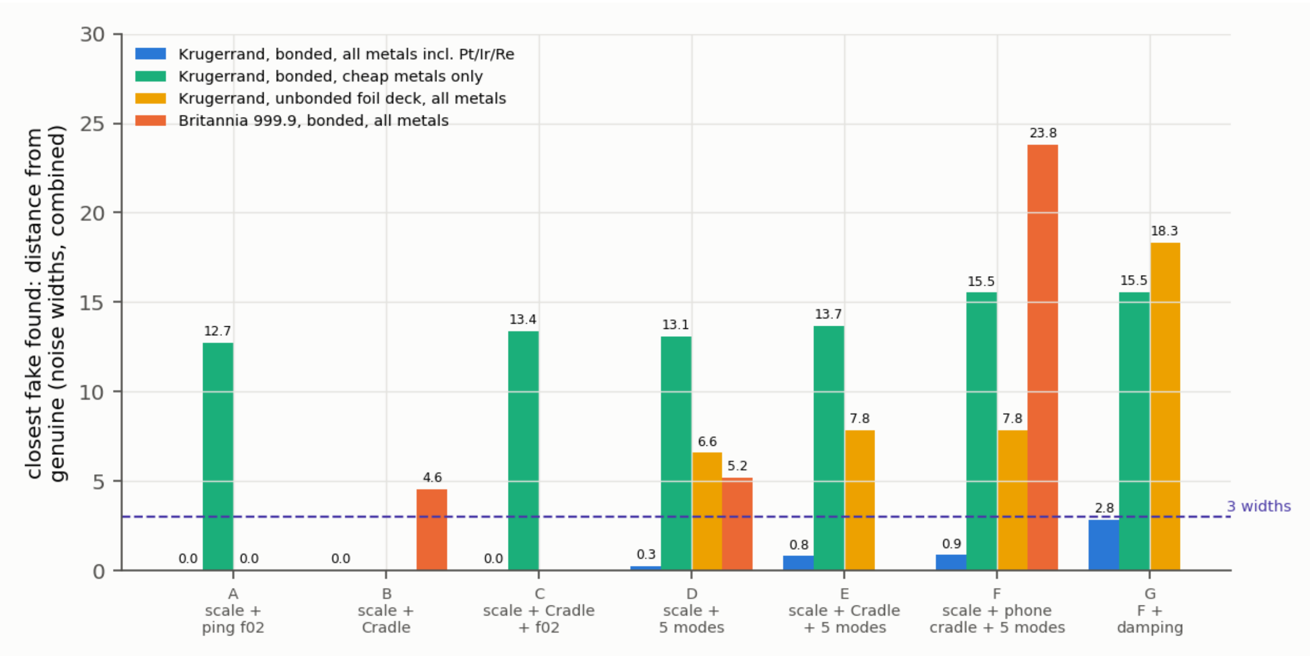

- **Why cheap fakes fail the ping.** Every cheap metal denser than about 11 g/cm³ is two to five times stiffer than gold alloy (Fig. F0). A mass-matched stack of cheap metals therefore rings at least 10% high, or comes in light. The best cheap construction (22 ct skin / Pb / Ta / W) is 10.8 widths high on f02 and 9.5 light on mass. **The tantalum–tungsten stacks that defeat the pendulum for 22 ct coins ring about 10% high**, so the ping catches them.
- **Why platinum works.** Pt is the one metal denser than gold alloy with a modulus only about twice gold's (21.45 g/cm³, 168 GPa [memory]). A soft, light Pb/Zn layer just under the skin brings the stiffness down; the dense Pt core brings the mass back. The best design:
  - 343 µm 22 ct skin / 134 µm Pb / 120 µm Zn / 1.09 mm Pt core / mirror;
  - by volume 30% gold alloy, 12% Pb, 10% Zn, 48% Pt; **by mass 30% gold alloy and 58% platinum**;
  - mass, σ·t and all five modes within 0.6 widths of genuine.
  - Platinum trades well below gold [memory], so this fake holds roughly half the genuine coin's metal value and still leaves a margin. It needs five-layer bonding at 0.1 mm tolerances. Damping moves it only to 2.8 widths, because the loss factors of bonded metals are all small.
- **Foil decks.** Letting layers slip cuts bending stiffness, so f02 can be matched with W or Ta. But a deck of thin plates loses the thick-plate (Mindlin) softening of the higher modes: f04/f02 comes out

6 widths high and f11/f02 4 high. Friction adds about 12 widths on each damped mode, if it is as large as assumed.

- **Gold budget.** With all metals, allowing only 10% gold still gives 2.6 widths; allowing 50% gives 0.4.
- **Fine gold** (44.7 MS/m) sits outside every cheap or exotic mix on conductivity. The pendulum alone gives 4.6 widths on the Cradle, and the full phone set 23.8.

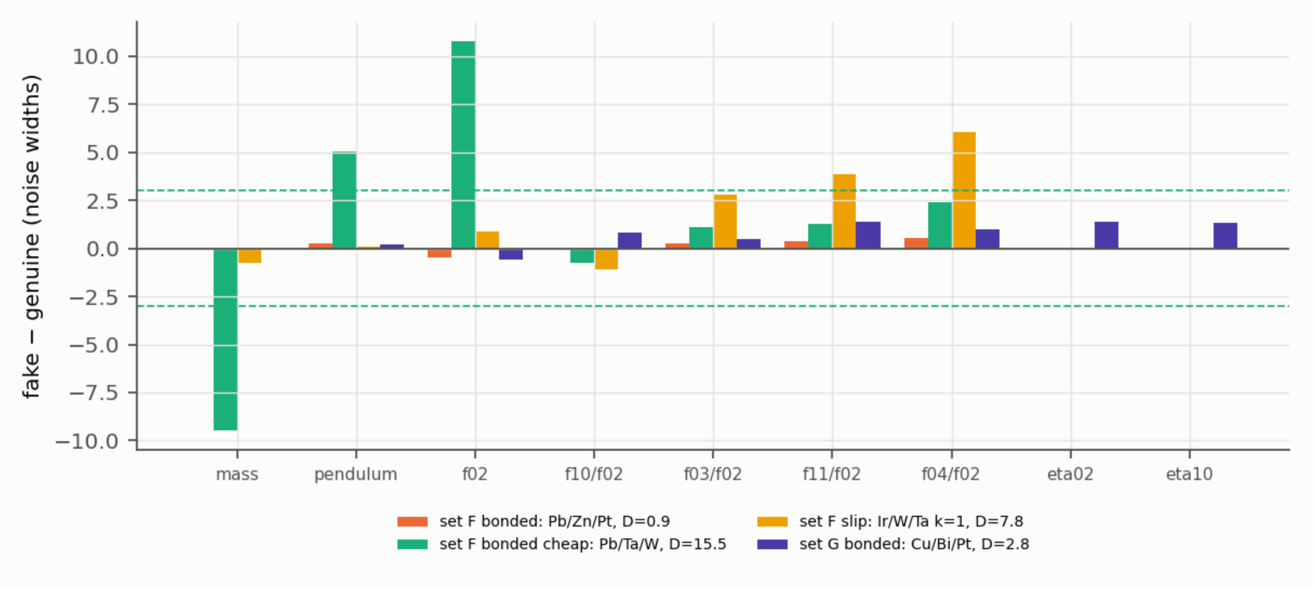

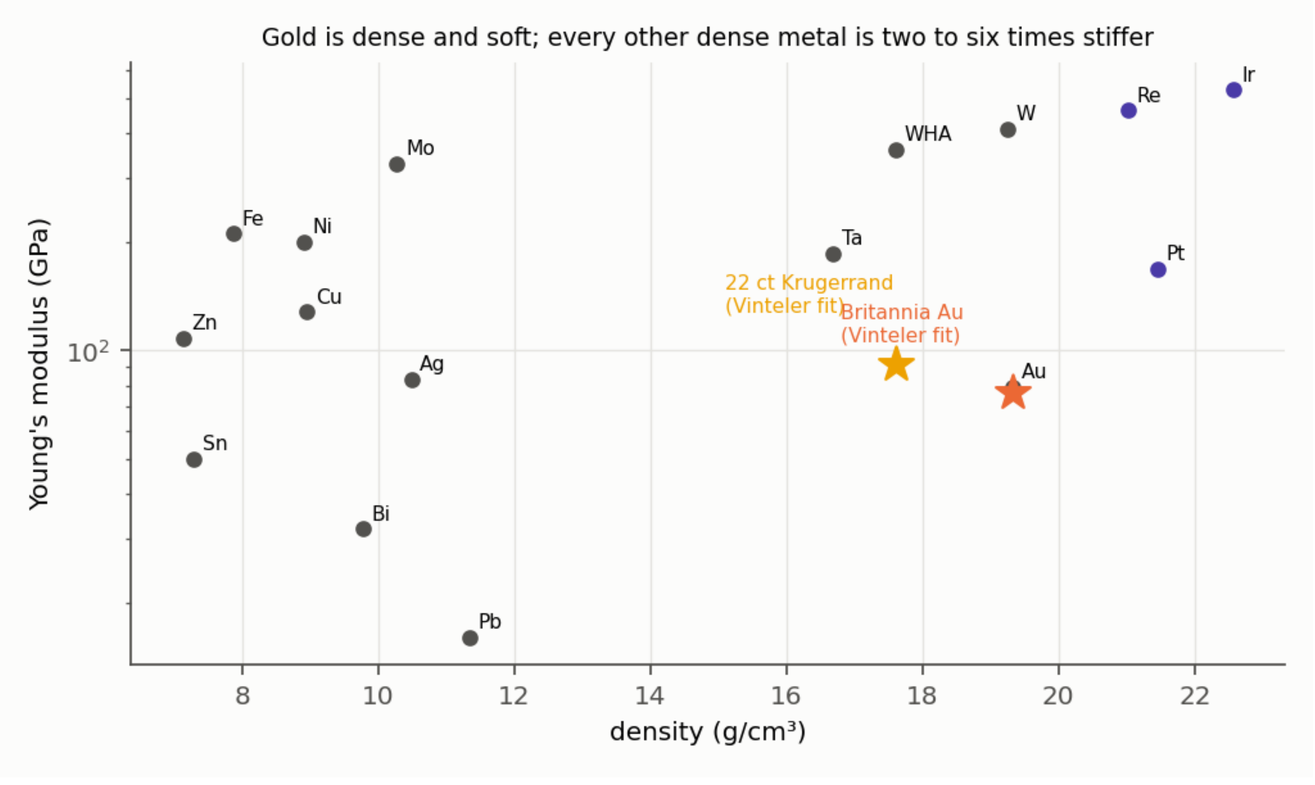

What this rests on.

- The noise widths for f02 and the ratios are assumptions. If the real coin-to-coin ratio spread is 1% rather than 0.5%, the foil-deck margin halves. The platinum stack stays invisible whatever the spread, because its largest deviation is 0.3%.
- FSDT with an energy-consistent shear factor is approximate for layers of very different stiffness (Pb against Pt). The *difference* between fake and genuine should be good to about 20–30% of itself; that is not checked.
- Handbook moduli [memory]; loss factors are orders of magnitude only.

- The plug's pendulum weighting is crude.
- Only symmetric stacks were searched. Asymmetric ones are caught by the pendulum flip test (v1).

3. Why it may be useful

- The February database and the site's ping gate should not be relied on. Most "verified" values are theory, and the ±3% gate fails genuine split coins. The audit says exactly which 14 values have independent support.
- The ping and the pendulum see different physics: bending stiffness per mass and ν_eff, against conductance. The scale fixes the mass. Because they measure different physics, a fake built to beat one pays on the other. The search shows that for every cheap construction, and it narrows the open case to one expensive, precise family: platinum-cored stacks.
- Everything that limits the design is now one unmeasured number, the coin-to-coin spread of mode ratios. It can be measured with coins already held.

4. What to do next

1. **Record genuine replicates (no purchase needed).** Every Krugerrand, Britannia, Maple and Eagle in stock: five tap-and-hold taps and one bounce each, raw 48 kHz, one phone, plus mass, diameter and rim thickness. Ten coins of one type give the coin-to-coin spread of f02 and the four ratios, which replaces every [assumed] noise width above.
2. **Record the plated-tungsten Eagle** the same way. The prediction is that it rings far higher than genuine, not lower (for a tungsten-cored Maple, 11.5 kHz against 4.8). That tests the physics on a real fake at no cost.
3. **Build a small offline analysis script** from `phone_sim.py` (FFT peaks, doublet centre, envelope Q) and run it on those WAV files. Do not reuse the site's AnalyserNode path.
4. **Retire from the site:** the YouTube-derived database, the ±3% gate, "match %" and counterfeit_signatures.json. Keep the 14 paper-sourced values as a check.
5. **For the platinum case,** the only fake that gets through, a thickness-resolving test is needed: an ultrasonic pulse-echo (layer interfaces reflect) or a density-by-depth method. That is the one gap the phone app cannot close. I recommend noting it rather than building for it until a Pt-cored fake is seen.

Design of the next version: one phone app, three tests

1. **Weight and dimensions.** A 0.01 g scale and callipers, typed in or read by Bluetooth. A phone photo of the diameter is worse than callipers (about 0.2 mm against 0.02 mm) [assumed]. These give mass and the outline; with the ping they give ν_eff and stiffness per mass.
2. **Ping.** Five tap-and-hold recordings and one bounce, analysed offline in the app (FFT peaks, envelope Q, doublet handling). Modes are assigned to the reference library's modes by nearest match within ±2.5%, not by "closest peak to one number".
3. **Pendulum video.** The phone-cradle design, with a reference-coin comparison.
4. **Verdict.** Each channel is compared with the same-type reference distribution (mean and spread from real genuine coins), and the channels are combined as a Mahalanobis distance with its covariance. The report shows the distance in noise widths, the channel that drives it, and, from the fusion search, which known constructions the result could *not* exclude. There is no single "match %", and no gate is set until replicate spreads exist.
5. **Data plan.**
    - **Genuine:** ten coins each of Krugerrand 1 oz, Britannia 1 oz, Maple 1 oz, Eagle 1 oz and Sovereign. Five taps and one bounce each, on two phones, with mass, diameter and rim thickness. Ten repeats on one coin for repeatability.

    - **Fakes:** the plated-tungsten Eagle already held; two or three bought replicas (brass or tungsten, stated), recorded, then sectioned so that the construction is known.
    - **Pendulum:** ten releases per coin, as in the pendulum v2 note.
    - **Storage:** raw WAV plus JSON metadata (coin id, phone, hold method, strike position, temperature).
6. **What not to carry over:** the YouTube-derived database, the ±3% gate, counterfeit_signatures.json, and the "match %" display.