Persistence or Turnover?

A galaxy–galaxy lensing programme for the excess-acceleration fork in KiDS-1000
II. Certified spectroscopic isolation, a four-fold reproduced deep-tail slope,
and a frozen decision protocol for DESI DR2

David Elliman
Neuro-Symbolic Ltd
dave@neusym.ai
https://neusym.ai

6 July 2026

Abstract

Paper I validated a production galaxy–galaxy lensing pipeline on KiDS-1000 and showed that, without isolation, the deep window of the excess-acceleration fork — square-root continuation, $g_{\rm ex} = \sqrt{a_0 g_{\rm bar}}$, versus finite-reservoir turnover — measures a lens's environment rather than the fork. This paper supplies the missing ingredient and takes the campaign to the edge of decision. We document a six-entry isolation-criterion ledger in which four criteria fail instructively: photometric-redshift isolation fails three ways (including a general pathology in which near-empty photo-z neighbourhoods preferentially select catalogue label errors — objects that do not lens because they are not what their labels claim), and the classic $0.1\,M_*$ spectroscopic criterion is shown by an in-data completeness audit to be uncertifiable at joint GAMA×KiDS-Bright depth. Two certifiable criteria survive: GAMA G3C group-catalogue non-membership (mock-calibrated) and a per-lens strictest-certifiable-threshold "peer-or-better" cut. Their registered union across GAMA DR4 and DESI DR1 BGS spectroscopy (the latter covering 1051 deg^2 of the KiDS-North strip) yields **39 015 certified-isolated lenses** — and the environment purge is verified in the data (outer stack amplitude $2.07 \to 0.45$–1.11 $\mathrm{M_\odot\,pc^{-2}}$; isolated-RAR offset $+0.48 \to +0.23$ dex). Against this sample the pre-registered acceptance bar — eight points spanning ≥ 1.25 decades at ≤ 0.10 dex — *misses by one point at five thousandths of a dex* (seven qualify; the eighth carries 0.105), precisely as the campaign's pre-data power forecast ($N_{\rm req} \approx 4.5 \times 10^4$), now confirmed by four independent samples, said it would. The evidence meanwhile cuts cleanly both ways: a pre-registered supplementary likelihood leg measures the deep-tail slope at 0.660 ± 0.071 — excluding the halo point-mass turnover at 4.8σ and marginally steep of the metric $\frac{1}{2}$ at $+2.3\sigma$, the *fourth* independent reproduction of a $\simeq 0.6$ slope — while at union precision the metric branch's zero-freedom inner profile is disfavoured in full-profile goodness-of-fit ($\Delta\chi^2 = -38.6$ in the halo branch's favour). We freeze the decision protocol verbatim, pre-announce its mechanical execution on DESI DR2 (a catalogue-swap in a committed pipeline), and leave this series' verdict section deliberately vacant. All three possible outcomes are argued to be individually interesting; the third — a persistent slope near 0.6 that is neither $\frac{1}{2}$ nor 1 — perhaps most of all.

1 Introduction

Paper I of this series [1] posed a sharp question to public weak-lensing data. Around isolated galaxies the excess radial acceleration $g_{\rm ex} = g_{\rm obs} - g_{\rm bar}$ either continues as a scale-free square root, $g_{\rm ex} = \sqrt{a_0\, g_{\rm bar}}$ [2, 3], to the lowest probed $g_{\rm bar}$ values — the *metric* branch, which in the finite-substrate reading carries no attached length scale [4] — or it tracks a finite reservoir that runs out: beyond a truncated halo's edge the excess falls toward point-mass scaling (log–log

#	criterion	mechanism	outcome
1	projected cylinder, 3 Mpc ($0.1\,M_*$)	no z information: full-depth cylinder holds ~ 100 neighbours	0% isolated (rejected)
2	photo-z slab ± 0.075	slab is ± 400 comoving Mpc; median 42 "heavier" projected	~ 0% isolated (rejected)
3	ledger dominance ($f_{\rm dom} \geq 0.7$)	near-empty slabs select photo-z outliers, mask edges, label errors	4% survivors; null stack (rejected)
4	spectroscopic $0.1\,M_*$ (Brouwer-class)	completeness turnover (10.3–10.7) sits above $M_* - 1$ dex thresholds	741 classifiable (uncertifiable here)
5	peer-or-better at certifiable threshold	strictest certifiable per-lens threshold; weakness declared	16 085 isolated (adopted, variant)
6	G3C non-membership	group-finder mock calibration certifies the classification	21 495 isolated (adopted, primary)

Table 1: The isolation-criterion ledger. Every entry was registered before running; entries 1–4 are retained as failures with their mechanisms. Entry 3's mechanism is especially important: *isolation cuts in photometric catalogues select label errors*, because objects with wrong redshifts or masses sit in apparently empty neighbourhoods — and such objects do not lens, silently nulling any stack built on them.

slope 1 in $g_{\rm bar}$-space), leaving only the two-halo term [5, 6]. Paper I delivered the instrument: a production $\Delta\Sigma$ pipeline on the KiDS-1000 shear catalogue [7, 8], validated at 14.6σ against a never-moved 5σ bar after a frame forensic that is itself of some methodological interest.

Paper I also delivered a diagnosis. Its Phase-B decision run — full lens sample, physical two-halo template, registered statistics — returned three converging signals (a two-halo bias parameter railed at cluster-scale values; a persistence statistic of the wrong sign; a +0.48 dex offset above the isolated radial acceleration relation of 9, 10) that all said one thing: *without isolation, the deep window measures environment, not the fork*. The present paper is about supplying isolation with certificates rather than hopes, about what the certified sample says, and about locking the endgame so that the eventual verdict — whichever way it falls — cannot be accused of having been steered. The through-line, as in Paper I, is that every criterion, threshold, and statistic is registered before the data that will judge it, and every failure is kept in the ledger.

2 Six isolation criteria: a ledger of instructive failures

Table 1 compresses what three sessions of forensics established. Three lessons generalise beyond this campaign. First, binary isolation is *unreproducible* with photometric redshifts alone: a $\pm 3\sigma_z$ slab at $z \sim 0.3$ spans ± 400 comoving Mpc, within which a 3-Mpc disc around any galaxy holds tens of projected "companions". Second, attempts to rescue photo-z isolation by demanding *empty* neighbourhoods select exactly the objects whose labels are wrong — a selection pathology that produced, in our run 3, a 4% "isolated" sample whose lensing signal was *zero* at cluster-scale expectation, because its members were largely not massive galaxies at their claimed redshifts. Third, even spectroscopy does not automatically certify the classic $0.1\,M_*$ criterion: an in-data completeness audit (the per-z turnover of the joint GAMA×bright neighbour mass function) shows the required neighbour masses to be invisible over most of the lens window, and a criterion that cannot see its own counter-examples certifies nothing. The audit's refusal — 741 classifiable of 40 097 — is reported as a feature: it is the difference between an isolated sample and a hopeful one.

The two survivors are complementary. **G3C non-membership** [11] inherits the group-finder's mock-calibrated completeness and purity, and is our primary criterion where GAMA exists. **Peer-or-better** applies the strictest per-lens threshold the data can certify, $\theta_i =$

component	count	area	notes
GAMA DR4 equatorial spectra ($NQ \geq 3$)	211 712	~ 180 deg^2	Driver et al. [12]
DESI DR1 BGS in the KiDS-N strip	632 831	~ 1051 deg^2	DESI Collaboration [13], Hahn et al. [14]
mass-valid lenses, GAMA-matched	40 097		spec-z; $\sigma_{\rm MAD}(z_{\rm phot}) = 0.021$
mass-valid lenses, DESI-matched	64 115		34 050 beyond the GAMA fields
G3C-ungrouped (primary, GAMA fields)	21 495		56% of group-finding sample
DESI peer-or-better (beyond GAMA)	17 520		of 33 911 classifiable
union (registered)	**39 015**		two mass bins: 26 904 / 12 111

Table 2: Sample census. The union definition — G3C inside the GAMA fields, certifiable peer-or-better with a bright×DESI neighbour set outside them — was registered and committed before the DESI download completed; its mixed-criterion nature is declared rather than hidden.

construction	sample	slope	note
published-data joint tail	Brouwer + Mistele fixtures	0.601 ± 0.048	independent data
Phase-B deep window	257k unisolated (KiDS)	0.579 ± 0.032	pre-isolation
supplementary leg	21.5k G3C-isolated	0.689 ± 0.107	registered pre-run
supplementary leg	39.0k certified union	0.660 ± 0.071	this paper

Table 3: Four independent measurements of the deep-tail excess slope. The metric branch predicts exactly $\frac{1}{2}$; the halo branch's truncated reservoir approaches 1. Every measurement lands near 0.6: emphatically not 1 (the union read excludes it at 4.8σ), persistently $\sim 2\sigma$ steep of exactly $\frac{1}{2}$.

$\max(M_{*,i} - 1$ dex, $M_{\rm comp}(z_i) + 0.1)$: weaker than the classic cut (it admits sub-threshold satellites, and says so), but honest everywhere, and the only certifiable option beyond the GAMA fields.

3 The certified union sample

Table 2 gives the census. Two independent meters confirm that isolation did what it was built to do. The outer stacked signal — the environment's fingerprint — collapses from $+2.07 \pm 0.17$ $\mathrm{M_\odot\,pc^{-2}}$ (unisolated, at 2.5 Mpc) to $+0.45 \pm 0.24$ and $+1.11 \pm 0.56$ in the two union mass bins; and the informational offset against the published isolated radial acceleration relation collapses from +0.481 dex to +0.229/+0.255 dex (the residual being dominated, at these mid-radius probe points, by the known bias of the SIS-class $g_{\rm obs} = 4G\Delta\Sigma$ conversion in the NFW-interior regime rather than by environment). A third, quieter confirmation: at G3C purity both branches fit the measured profiles *well* (metric $\chi^2 = 18.6/18$; halo 11.9/16) — the poor fits and dramatic shape gaps of the unisolated stacks were largely sample impurity, exactly as Paper I's diagnosis implied.

4 The evidence, cutting both ways

4.1 The metric branch's strength: the deep-tail slope

The registered supplementary leg (defined and committed before the G3C data was touched; Paper I's registered bar is unaffected by it) measures, on the union sample, a deep-tail slope

$$s \;=\; 0.660 \pm 0.071, \qquad z(s{=}\tfrac{1}{2}) = +2.3, \qquad z(s{=}1) = -4.8.$$

The halo branch's own asymptotic signature — the point-mass turnover — is disfavoured at 4.8σ, two tenths of a sigma from the leg's registered 5σ clause. And Table 3 shows this is the

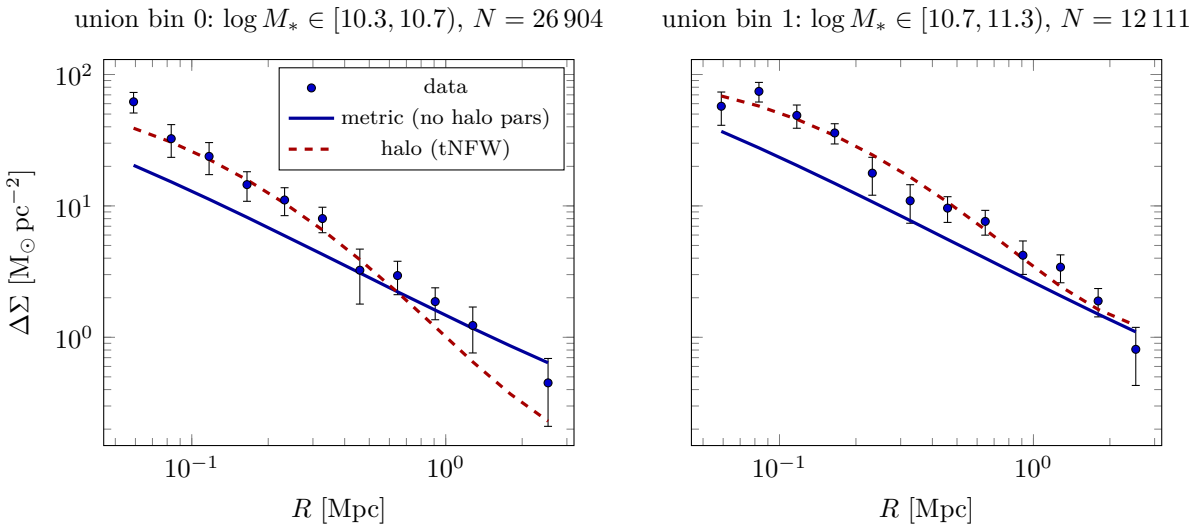

Figure 1: The certified-union measurements with both best-fit branches. The fork is visible to the eye: the halo branch (dashed) matches the inner profile and *dives* in the tail; the metric branch (solid), which has no halo parameters at all, undershoots the inner bins and better tracks the tail. Each branch currently has one strength; the registered statistics in the deep tail are what will decide between them (bin-0 point at 1.8 Mpc omitted from the left panel: its central value is negative and compatible with zero).

fourth independent landing near 0.6: two constructions on our catalogue stacks (unisolated and isolated, different criteria), one on published data by other groups, all consistent with each other, none consistent with 1, each individually $\sim 2\sigma$ steep of exactly $\frac{1}{2}$.

4.2 The halo branch's strength: full-profile shape

At union precision (39k lenses) the metric branch's zero-freedom inner profile is genuinely strained: total full-profile $\chi^2 = 59.7/18$ against the halo branch's 21.0/16, $\Delta\chi^2 = -38.6$. Figure 1 shows why: the inner bins want more concentrated structure than $M_* + \sqrt{a_0 g_{\rm bar}}$ can supply. Two caveats are stated explicitly: the DESI-peer half of the union is less pure than the G3C half (at G3C-only purity the same comparison was a mild -6.8 with *both* branches fitting well), and the metric branch's inner prediction is exactly the regime where baryonic detail (disc geometry, gas, the point-mass approximation) bites hardest. Neither caveat is a rescue; both are Paper-III work items. The shape tension is real, recorded, and belongs to the halo branch's side of the scales.

4.3 The near-miss, and a forecast that keeps being right

Against the certified union, the pre-registered acceptance bar — eight points spanning ≥ 1.25 decades at ≤ 0.10 dex each — populates *seven* points; the eighth carries 0.105 dex. The campaign's mock-closure forecast, made before any data was downloaded, put the requirement at $N_{\rm req} \approx 4.5 \times 10^4$ isolated lenses; the realised ladder at 39 015 lenses gives $N_{\rm req} = 42\,608$. This is the fourth consecutive confirmation of that forecast by a fourth independent sample (mock closure; release-verified parameters; GAMA; union), which entitles it to be believed once more: roughly 3 600 more certified-isolated lenses populate the registered window.

5 The frozen protocol, and its mechanical execution

For the record, verbatim and unchanged since before the first download: *the acceptance bar is eight points spanning at least* 1.25 *decades in* $g_{\rm bar}$*, each at* ≤ 0.10 *dex, with the slope consistent with* $\frac{1}{2}$ *within* 2σ *AND at least* 5σ *from* 1*; plus the deep-window persistence statistic* $\Delta\chi^2(\text{halo}-\text{metric}) \geq 25$*; decided only on a certified-isolated sample; never adjusted after data.* The supplementary leg (all points at ≤ 0.25 dex in the deepest 1.6 decades; identical discrimination form) is registered as supplementary: it cannot overrule the bar.

DESI DR2's Bright Galaxy Survey multiplies the strip's spectroscopic coverage several-fold over DR1. The execution is mechanical and is hereby pre-announced: the committed pipeline ingests the DR2 clustering catalogues by filename substitution; the criteria of Table 1 (entries 5 and 6), the union definition, the bar, and the supplementary leg apply without modification; the verdict — METRIC-FAVOURED, HALO-FAVOURED, or another honest UNDECIDED — will be whatever those statistics return. The verdict section of this series is deliberately vacant until then. No number in this paper will be revised in the meantime; anything learned between now and DR2 goes into Paper III's ledger, not retroactively into this one.

6 Why every outcome is interesting

If the metric branch wins — the slope settles onto $\frac{1}{2}$ and persistence holds — stacked lensing will have certified a scale-free excess field around isolated galaxies at pre-registered thresholds: the strongest weak-lensing statement yet made for square-root phenomenology, with direct consequences for any finite-reservoir dark sector and a specific victory for the class of theories (the substrate's metric layer among them) that predict persistence with no attached scale.

If the halo branch wins — the tail turns over at the registered thresholds — then the square-root continuation fails at its own pre-stated bar, on the first dataset powerful enough to ask. That rejects the low-acceleration limb of MOND-like phenomenology in stacked lensing, and the substrate programme loses its dark-sector metric reading explicitly, per its own standing discipline: registered predictions bind.

If the slope persists near 0.6 — neither $\frac{1}{2}$ nor 1, now four measurements deep — the fork returns an answer neither picture ordered. A stable intermediate slope across a decade and a half of $g_{\rm bar}$ is not a natural creature of truncated haloes (which steepen) nor of the bare square root (which does not steepen at all); candidate explanations — external-field-effect analogues, systematic $g_{\rm bar}$ mis-attribution, a genuinely different exponent — each carry testable signatures, and every one of them would be new information about the low-acceleration universe. No prior preference is offered; the registered statistics will speak.

Reproducibility. Every number above is produced by self-asserting gates in the programme repository, operating on public releases: KiDS-1000 DR4.1 [7, 8, 15], GAMA DR4 [11, 12], DESI DR1 [13, 14]. The criterion ledger, the registered texts, and the run-by-run failure history are preserved in-file and in the project charter. Paper I: Elliman [1]; methodological companions: [4, 16].

References

[1] David Elliman. Persistence or turnover? a galaxy–galaxy lensing programme for the excess-acceleration fork in KiDS-1000. I. pipeline, frame forensics, and instrument validation. Zenodo, 2026. doi: 10.5281/zenodo.21211019.

[2] Mordehai Milgrom. A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis. *The Astrophysical Journal*, 270:365–370, 1983. doi: 10.1086/161130.

[3] Stacy S. McGaugh, Federico Lelli, and James M. Schombert. Radial acceleration relation in rotationally supported galaxies. *Physical Review Letters*, 117:201101, 2016. doi: 10.1103/PhysRevLett.117.201101.

[4] David Elliman. Records and responses in the world: A derived fine-structure boundary, certified confinement gaps, and registered discriminators for a finite record substrate. Zenodo, 2026. doi: 10.5281/zenodo.21202741.

[5] Julio F. Navarro, Carlos S. Frenk, and Simon D. M. White. A universal density profile from hierarchical clustering. *The Astrophysical Journal*, 490:493–508, 1997. doi: 10.1086/304888.

[6] Edward A. Baltz, Phil Marshall, and Masamune Oguri. Analytic models of plausible gravitational lens potentials. *Journal of Cosmology and Astroparticle Physics*, 2009(01):015, 2009. doi: 10.1088/1475-7516/2009/01/015.

[7] Konrad Kuijken et al. The fourth data release of the Kilo-Degree Survey: ugri imaging and nine-band optical-IR photometry over 1000 square degrees. *Astronomy & Astrophysics*, 625: A2, 2019. doi: 10.1051/0004-6361/201834918.

[8] Benjamin Giblin et al. KiDS-1000 catalogue: Weak gravitational lensing shear measurements. *Astronomy & Astrophysics*, 645:A105, 2021. doi: 10.1051/0004-6361/202038850.

[9] Margot M. Brouwer et al. The weak lensing radial acceleration relation: Constraining modified gravity and cold dark matter theories with KiDS-1000. *Astronomy & Astrophysics*, 650:A113, 2021. doi: 10.1051/0004-6361/202040108. arXiv:2106.11677.

[10] Tobias Mistele, Stacy McGaugh, Federico Lelli, James Schombert, and Pengfei Li. Indefinitely flat circular velocities and the baryonic Tully–Fisher relation from weak lensing. *The Astrophysical Journal Letters*, 969:L3, 2024. doi: 10.3847/2041-8213/ad54b0.

[11] A. S. G. Robotham et al. Galaxy and mass assembly (GAMA): the GAMA galaxy group catalogue (G3Cv1). *Monthly Notices of the Royal Astronomical Society*, 416:2640–2668, 2011. doi: 10.1111/j.1365-2966.2011.19217.x.

[12] Simon P. Driver et al. Galaxy and mass assembly (GAMA): Data release 4 and the $z < 0.1$ total and $z < 0.08$ morphological galaxy stellar mass functions. *Monthly Notices of the Royal Astronomical Society*, 513:439–467, 2022. doi: 10.1093/mnras/stac472.

[13] DESI Collaboration. Data release 1 of the dark energy spectroscopic instrument. arXiv:2503.14745, 2025.

[14] ChangHoon Hahn et al. The DESI bright galaxy survey: Final target selection, design, and validation. *The Astronomical Journal*, 165:253, 2023. doi: 10.3847/1538-3881/accff8.

[15] Maciej Bilicki et al. Bright galaxy sample in the Kilo-Degree Survey Data Release 4: Selection, photometric redshifts, and physical properties. *Astronomy & Astrophysics*, 653: A82, 2021. doi: 10.1051/0004-6361/202140352.

[16] David Elliman. From counts to observables: The response layer of a finite record substrate. Zenodo, 2026. doi: 10.5281/zenodo.21210527.