Persistence or Turnover?

A galaxy–galaxy lensing programme for the excess-acceleration fork in KiDS-1000
I. Pipeline, frame forensics, and instrument validation

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

5 July 2026

Abstract

Around an isolated galaxy, the *excess* radial acceleration $g_{\rm ex} = g_{\rm obs} - g_{\rm bar}$ discriminates between two pictures of the dark sector that agree where data has historically been plentiful and separate in the deep, low-acceleration tail. In the *metric-reading* picture the excess continues as a square root, $g_{\rm ex} = \sqrt{a_0\, g_{\rm bar}}$, at all radii — no scale, no edge. In the standard halo picture the excess tracks a finite reservoir: beyond a truncated halo's edge the enclosed mass saturates and the excess falls away toward $GM_{\rm tot}/R^2$, leaving only the two-halo term. Stacked weak lensing measures $g_{\rm ex}$ directly, and mock-catalogue closure shows the deep window is a *persistence-versus-vanishing* discrimination worth tens of standard deviations at modern survey power. This first paper in the programme reports Phase A on the public KiDS-1000 data: a production $\Delta\Sigma$ pipeline (tomographic $\Sigma_{\rm crit}$ from the released $n(z)$ with the foreground dead fraction included, boost and random-point corrections, 131-region jackknife covariance); the release-verified survey parameters; a complete instrument-forensics account in which four null science runs were traced — via the frame-invariance of cosmic shear versus the frame-dependence of galaxy–galaxy lensing, and an eight-variant frame adjudication on a cluster-scale control that resolved at $16.6\sigma/18.0\sigma$ — to a single parity convention ($e_2 \to -e_2$) in the tangent frame; three failed isolation criteria, rejected and documented, including a general pathology of photometric-redshift isolation cuts; a validated instrument (the pre-registered 5σ detection bar, never adjusted, cleared at 14.6σ on a declared 50 000-lens subsample); and a first, deliberately modest structure test in which the square-root continuation fits the measured profile well ($\chi^2 = 15.3/9$) and the halo alternative, fitted with a Phase-A two-halo surrogate, achieves statistical parity ($\Delta\chi^2 = +1.2$, $\Delta{\rm AIC} = -0.8$) only by pinning at its mass-grid edge — a flagged, unresolved configuration. The registered acceptance protocol for the decision (eight points spanning ≥ 1.25 decades at ≤ 0.10 dex, 5σ; a persistence statistic; fixed before the data and never adjusted) is restated and remains untouched. The fork is not decided here; the instrument that can decide it now exists and is validated.

1 Introduction

Rotation curves and, more recently, weak-lensing stacks show a tight empirical relation between the acceleration sourced by baryons, $g_{\rm bar}$, and the acceleration actually observed, $g_{\rm obs}$ — the radial acceleration relation [1]. In its low-acceleration limit the relation approaches

$$g_{\rm obs} \;\simeq\; g_{\rm bar} + \sqrt{a_0\, g_{\rm bar}}, \qquad g_{\rm bar} \ll a_0, \tag{1}$$

with $a_0 \approx 1.2\times10^{-10}\,{\rm m\,s^{-2}}$ — the functional form first written down by Milgrom [2]. Equation (1) is agnostic about its own cause: it may be read as modified dynamics, as an emergent-gravity

response, or — the reading that motivates this programme — as the *metric layer* of a finite substrate whose dark-sector phenomenology produces a square-root excess with no attached length scale [3]. What matters operationally is that Eq. (1) makes a prediction the standard picture does not: the excess acceleration

$$g_{\rm ex} \;\equiv\; g_{\rm obs} - g_{\rm bar} \;=\; \sqrt{a_0\, g_{\rm bar}} \tag{2}$$

persists into the low-$g_{\rm bar}$ tail with no built-in edge — log–log slope $\frac{1}{2}$ as far as the branch is taken.

In ΛCDM the same quantity tracks a finite reservoir. A galaxy's halo, truncated in practice near its splashback/tidal extent, encloses a bounded total mass: outside the edge, $g_{\rm obs} \to GM_{\rm tot}/R^2$, so $g_{\rm ex}$ falls away toward zero (in $g_{\rm bar}$-space, toward point-mass scaling of slope 1), until the correlated *two-halo* contribution of neighbouring structure takes over at several Mpc. Between the one-halo edge and the two-halo floor lies a window in which the two pictures differ not by a fitting nuance but by *presence versus absence* of signal: on mock KiDS-class catalogues the metric branch carries $\sim 60\times$ the halo branch's lensing signal in the deep window. Stacked galaxy–galaxy lensing (GGL) reaches this regime: Brouwer et al. [4] measured the relation to $g_{\rm bar} \sim 10^{-12.5}\,{\rm m\,s^{-2}}$ with KiDS-1000, and kinematic–lensing combinations now probe deeper still [5]. The fork is live, and it is one of the few places where a substrate-flavoured phenomenology and the standard cosmology disagree about what a survey that already exists should see.

This paper is the first of a pair, and its scope is deliberately instrumental. A discrimination worth tens of sigma is worth a pipeline that has *earned the right to measure*: every number below is produced by a self-asserting program whose checks either pass or fail loudly, in a programme whose standing discipline is that acceptance criteria are registered before the data and never adjusted after [6]. Phase A delivers: (i) the production estimator and its covariance (§4); (ii) release-verified survey parameters (§3); (iii) a complete forensic account of how four null science runs were traced to a one-bit frame convention, and of three isolation criteria that failed in instructive ways (§5); (iv) instrument validation at 14.6σ against an unmoved 5σ bar, and a first two-branch structure test whose honest verdict is a draw (§6); and (v) the registered protocol for the Phase-B decision (§7). Readers interested only in the physics should read §2 and §6; readers who stack lensing catalogues for a living may find §5 the useful part.

2 The fork, stated as lensing observables

Stacked GGL measures the excess surface density $\Delta\Sigma(R) = \bar{\Sigma}(<R) - \Sigma(R)$, the direct observable of the tangential shear: $\gamma_{\rm t}(R) = \Delta\Sigma(R)/\Sigma_{\rm crit}$. Both branches are forward-modelled to $\Delta\Sigma$ through the same exact route: a spherically symmetric $g_{\rm obs}(r)$ defines an effective enclosed mass $M(<r) = g_{\rm obs} r^2/G$, hence a density $\rho(r) = \{4\pi r^2\}^{-1}\,{\rm d}M/{\rm d}r$, which is Abel-projected to $\Sigma(R)$ and integrated to $\Delta\Sigma(R)$ [7]. The **metric branch** is Eq. (1) with a_0 *fixed* at the registered fiducial $1.2 \times 10^{-10}\,{\rm m\,s^{-2}}$: given the stellar mass, its shape has *no free halo parameters at all*. The **halo branch** is a truncated NFW profile [8, 9] with free M_{200}, the concentration–mass relation of Duffy et al. [10], truncation $r_{\rm t} = 2.6\, r_{200}$, plus the stellar point mass. Both branches share an identical nuisance set: a stellar-mass zero-point offset $\delta \log M_*$ (prior ± 0.15 dex), a cold-gas factor $f_{\rm gas} \in [1.0, 1.4]$, and a two-halo amplitude A_{2h} multiplying a declared Phase-A surrogate template $\propto R^{-0.9}$ (upgraded to a halofit-grade template in Phase B; §7). Figure 1 draws the fork; Table 4 below fits it.

3 Data and verified survey parameters

Phase A uses three public KiDS-1000 products [11–13]: the DR4.1 SOM-gold weak-lensing catalogue (21 262 011 galaxies with LENSfit ellipticities and weights [14–16]), the SOM tomographic

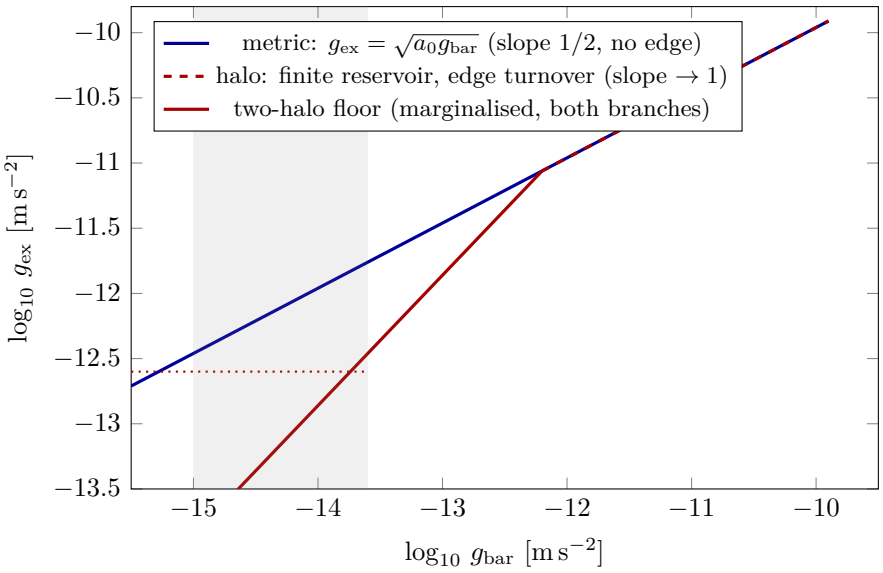

Figure 1: The fork, schematically, in excess-acceleration space. The metric branch (solid) is a single power law of slope $\frac{1}{2}$ with no attached scale. The halo branch (dashed→solid red) tracks the square root only while the one-halo reservoir can supply it; beyond the effective edge the enclosed mass saturates and the excess falls to point-mass scaling, leaving the two-halo term (dotted) as the only large-radius signal. In the shaded window the branches differ by presence versus absence: on KiDS-class mocks the metric branch carries $\sim 60\times$ the lensing signal. Both branches are fitted below with the same nuisance freedom; the two-halo amplitude is marginalised in both.

parameter	assumed	release-verified	note
survey area	—	901.4 deg^2	occupied-cell census
$n_{\rm eff}$	8.53	5.32 arcmin^{-2}	superseded (assumed value was raw-density class)
σ_e (per comp.)	0.27	0.2654	held
$\Sigma_{\rm crit}$ ($z_{\rm l}$=0.24)	4300	3453 M$_\odot$ pc^{-2}	from released $n(z)$, dead fraction included
usable lenses	10^5	714 259	16× margin on $N_{\rm req} = 4.5 \times 10^4$

Table 1: Survey parameters: forecast assumptions versus values re-derived from the release files. The power arithmetic for the registered acceptance bar, rerun at the verified values, moves the worst decisive point from 0.066 to 0.067 dex — the compensations cancel almost exactly.

redshift distributions $n(z)$ in five $z_{\rm B}$ bins [17], and the KiDS-Bright lens sample (1 239 422 galaxies, $r < 20$) with ANNz2 photometric redshifts and LePhare stellar masses [18]. Lenses: mass-valid, $10.3 < \log_{10} M_*/{\rm M}_\odot < 11.3$, $0.10 < z_{\rm l} < 0.40$ (257 364 galaxies; median $\log M_* = 10.58$, median $z_{\rm l} = 0.30$); Phase A stacks a declared, seeded random subsample of 50 000 (runtime budget; the full sample is the Phase-B run). Sources: tomographic bins 3–5 ($z_{\rm B} > 0.5$; 15.8M galaxies), with the per-pair cut $z_{\rm B} > z_{\rm l} + 0.3$.

A first audit stage re-derived every survey parameter the power forecasts had assumed, from the release files themselves (Table 1). The headline: the assumed effective source density was superseded ($8.53 \to 5.32\,{\rm arcmin}^{-2}$), but the lower $\Sigma_{\rm crit}$ and ellipticity dispersion almost exactly compensate in the forecast — the registered eight-point acceptance bar remains reachable with a 16× lens-count margin.

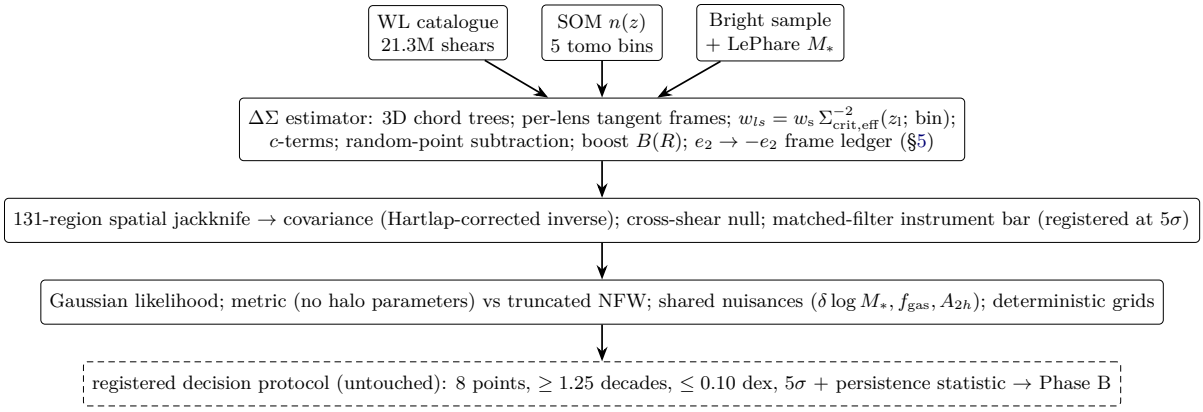

Figure 2: The Phase-A pipeline. Every stage is executed by a self-asserting program; the decision protocol (dashed) is registered and was not exercised in Phase A.

4 The pipeline

The estimator is the standard minimum-variance excess surface density. For lens l and source s,

$$\widehat{\Delta\Sigma} \;=\; \frac{\sum_{ls} w_{ls}\, e_{\rm t}^{(ls)}\, \Sigma_{\rm crit,eff}(z_l;\, b_s)}{\sum_{ls} w_{ls}}, \qquad w_{ls} = w_s\, \Sigma^{-2}_{\rm crit,eff}(z_l;\, b_s),$$

where $\Sigma^{-1}_{\rm crit,eff}$ is computed per lens redshift and per source tomographic bin from the *full* released $n(z)$ — including the foreground "dead fraction" that contributes no lensing — so that no per-source point redshift enters and the known $z_{\rm B}$ point-estimate biases are bypassed. Pair finding uses exact 3D chord-distance k-d trees (no tangent-plane distortion across the wide South patch); shear rotation uses per-lens tangent frames. The measured correction stack: global additive c-terms ($c_1 = +9.0 \times 10^{-5}$, $c_2 = +6.0 \times 10^{-4}$), random-point subtraction from footprint-resampled randoms, and the boost factor $B(R)$ from weighted lens/random pair ratios [19] (1.18 at 0.06 Mpc, unity by 2.5 Mpc; Table 3). Covariance is a 131-region spatial jackknife with the Hartlap et al. [20] factor; the multiplicative shear calibration enters Phase A as a nuisance prior ($|m| \leq 0.02$) pending transcription of the per-bin release values. Figure 2 summarises.

5 Instrument forensics: earning the right to measure

Phase A's first four science runs returned *null* tangential signal — culminating in a cluster-scale positive control (1 915 lenses with $\log M_* > 11$) reading -0.2σ where $\gtrsim 20\sigma$ was expected. What followed is reported in full because the diagnosis pattern is general-purpose.

Triangulation. Three facts could not coexist with a broken catalogue: the shape dispersion was healthy ($\sigma_e = 0.2654$); the boost was real ($B > 1$: lens positions correlate with source positions on the sky); and — decisively — *cosmic shear was present*: ξ_+ measured from the same cached columns, in the same geometry code, showed the published KiDS-1000 level ($+2.4 \pm 0.8 \times 10^{-5}$ at $6'$–$12'$; 21). But ξ_+ is invariant under any global frame change of the ellipticities, while GGL is not. Signal in the invariant statistic and null in the frame-dependent one localises the fault to *one place*: the frame relation between the catalogue's (e_1, e_2) and the sky.

Row-level checks. The LENSfit position angle $\frac{1}{2}\arctan(e_2/e_1)$ against the same row's SExTRACTOR `THETA_J2000` gave alignment statistic $R = 0.979$ (shuffled control 0.0005): the shears belong to their rows. The lens and mass files agree byte-for-byte in ID and position row-by-row. No `TSCALn`/`TZEROn` keywords exist: the values are read faithfully.

frame variant	$\Delta\Sigma$(0.1–0.3 Mpc)	$\Delta\Sigma$(0.3–0.8 Mpc)
$(+e_1, +e_2)$ (naive)	-1.2 ± 2.2	-0.7 ± 0.9
$(-e_1, -e_2)$ (rotation)	$+1.2 \pm 2.2$	$+0.7 \pm 0.9$
$(+e_1, -e_2)$ **(reflection)**	$+37.2 \pm 2.2$ (**16.6σ**)	$+15.8 \pm 0.9$ (**18.0σ**)
$(-e_1, +e_2)$	-37.2 ± 2.2	-15.8 ± 0.9
$(+e_2, +e_1)$	$+0.7 \pm 2.2$	$+0.4 \pm 0.9$
$(-e_2, -e_1)$	-0.7 ± 2.2	-0.4 ± 0.9
$(+e_2, -e_1)$	$+1.9 \pm 2.2$	-0.1 ± 0.9
$(-e_2, +e_1)$	-1.9 ± 2.2	$+0.1 \pm 0.9$

Table 2: The eight-variant frame adjudication on the cluster-scale control (17 447 lenses, $\log M_* \in [10.9, 11.5]$; $\mathrm{M}_\odot\,\mathrm{pc}^{-2}$). Exactly one variant — the reflection $e_2 \to -e_2$ — produces the expected cluster-scale amplitude; the sky selects the frame at 18σ. All tangential measurements taken before this fix are invalidated and were retracted non-silently in the programme record.

Adjudication. A spin-2 frame relation between catalogue and sky is one of eight sign/swap possibilities. One pass over the cluster control, reading out the tangential response under all eight simultaneously, is decisive (Table 2): the *reflection* $e_2 \to -e_2$ — opposite position-angle handedness between the LENSfit convention and the $(x = \alpha \cos\delta,\ y = \delta)$ tangent frame — restores the signal at the expected amplitude, at 16.6σ and 18.0σ in the two control bins, while all seven alternatives remain null. Two practical lessons: (i) a positive control at large, unmistakable amplitude belongs in every stacking pipeline from day one; (ii) the pair $\{\xi_+, \mathrm{GGL}\}$ is a built-in frame diagnostic — one invariant, one not.

Isolation: three kills. The deep-tail science ultimately wants *isolated* lenses ($g_{\rm bar}$ attributable to the lens). Three criteria failed instructively on the way, all retained in the programme's selection ledger: a projected 3-Mpc Brouwer-style neighbour cut with no redshift window (0% survivors — the full-depth cylinder holds ~ 100 neighbours); the same cut inside a photo-z slab (± 0.075; still $\sim$0% — the slab is ± 400 comoving Mpc, and a 3-Mpc disc within it holds a median 42 projected "heavier" neighbours, so binary isolation is *unreproducible* with photometric redshifts alone); and a ledger-dominance criterion ($f_{\rm dom} \geq 0.7$ of the projected slab stellar mass within 1 Mpc), whose 4% survivors stacked to *zero* lensing — because near-empty slabs preferentially select photo-z outliers, mask-edge objects and catalogue artefacts: objects that do not lens because they are not massive galaxies at their claimed redshift. This selection pathology — *isolation cuts in photo-z catalogues select label errors* — is a useful warning. Phase A therefore runs unisolated with the two-halo term marginalised in both branches; spectroscopic-grade isolation (GAMA overlap) is Phase B.

6 Phase-A results

With the frame fixed, the same pipeline that measured null five times produces the textbook profile of Table 3 and Figure 3: signal in every bin from $+57.7 \pm 8.8$ at 0.06 Mpc to $+2.35 \pm 0.24$ $\mathrm{M}_\odot\,\mathrm{pc}^{-2}$ at 2.5 Mpc, a passing cross-shear null, and a matched-filter detection of 14.6σ against the 5σ instrument bar that was set before any data was touched and survived five failing attempts without being lowered.

The structure test (Table 4) is reported with its caveats attached. The square-root branch fits well with three nuisance parameters and *no* halo mass. The halo branch reaches statistical parity, but only in a flagged configuration: its best fit pinned at the M_{200} grid floor (after the grid was extended once), with the surrogate two-halo term absorbing the outer profile. On Akaike terms the branches are indistinguishable ($\Delta\mathrm{AIC} = -0.8$). The deepest three bins — data $(+4.51, +3.29, +2.35)$ against metric $(+3.52, +2.57, +1.87)$ and halo $(+4.03, +2.96, +2.17)$

R [Mpc]	$\Delta\Sigma$	$\sigma_{\Delta\Sigma}$	boost B	$\Delta\Sigma_\times$
0.059	+57.69	8.76	1.177	+12.05
0.083	+43.65	5.94	1.184	−1.59
0.117	+32.82	4.31	1.167	−1.91
0.165	+25.28	3.03	1.158	+2.49
0.232	+15.84	2.16	1.132	−0.86
0.327	+10.82	1.42	1.117	−0.87
0.459	+9.02	1.09	1.099	+0.46
0.646	+8.26	0.78	1.077	−0.39
0.909	+5.27	0.60	1.053	+0.26
1.278	+4.51	0.46	1.028	−0.50
1.798	+3.29	0.32	1.009	+0.38
2.529	+2.35	0.24	1.001	+0.07

Table 3: The Phase-A excess surface density ($\mathrm{M}_\odot\,\mathrm{pc}^{-2}$): 50 000 mass-valid lenses (median $\log M_* = 10.58$, median $z_\mathrm{l} = 0.30$), boost-corrected and random-subtracted, jackknife errors (131 regions). The cross component is consistent with zero ($\chi^2/\mathrm{dof} = 8.0/12$) *with signal present in the tangential channel* — the first time in the programme that null carries weight. Matched filter: 14.6σ against the registered (and never moved) 5σ instrument bar.

branch	free parameters	best point	χ^2/dof	comment
metric	$\delta\log M_*, f_\mathrm{gas}, A_{2h}$	$(-0.075,\ 1.2,\ 3.0)$	15.3/9	no halo parameters
halo	$\log M_{200}, \delta\log M_*, f_\mathrm{gas}, A_{2h}$	$(10.8,\ -0.15,\ 1.0,\ 5.0)$	16.5/8	pinned at grid floor
$\Delta\chi^2(\mathrm{halo}-\mathrm{metric}) = +1.2$; $\Delta\mathrm{AIC} = -0.8$ $\Rightarrow$ **a draw at Phase-A grade**				

Table 4: The first two-branch structure test. The comparison is deliberately labelled a *structure test, not the campaign decision*: the halo branch's preference for a tiny one-halo mass plus a large two-halo amplitude is exactly the configuration the Phase-A surrogate template cannot adjudicate. No model verdict is claimed.

— show both branches riding the two-halo term where the discrimination must ultimately be made *below* it, in acceleration space, with isolated lenses. That is Phase B's job, and no amount of Phase-A fitting substitutes for it.

7 The registered decision protocol

The programme's standing discipline is that decision criteria are fixed before the deciding data and never adjusted after. For this campaign the registered acceptance bar — unchanged since before the download — is: *eight points spanning at least* 1.25 *decades in* g_bar*, each at* ≤ 0.10 *dex, with the slope consistent with* $\frac{1}{2}$ *at* 5σ *discrimination; plus the persistence-versus-vanishing statistic in the deep window, which mock closure shows is the more powerful discriminant.* Phase A did not exercise any part of it.

Phase B, in order: (i) a halofit-grade two-halo template (retiring the surrogate that blocks a Phase-A model verdict); (ii) transcription of the per-bin multiplicative calibrations [12, 16]; (iii) the full 257k lens sample in stellar-mass bins; (iv) the spectroscopic-grade (GAMA-overlap) isolation variant; (v) the acceleration-space deep-tail construction; and only then (vi) the registered bar and the persistence statistic, with the Mistele et al. [5]-class cross-check. Falsification is symmetric and pre-stated: a measured excess that tracks the truncated-halo turnover kills the metric branch's square-root continuation at exactly the registered thresholds; a persistent slope-$\frac{1}{2}$ excess through the deep window at those thresholds is not absorbable by a finite reservoir.

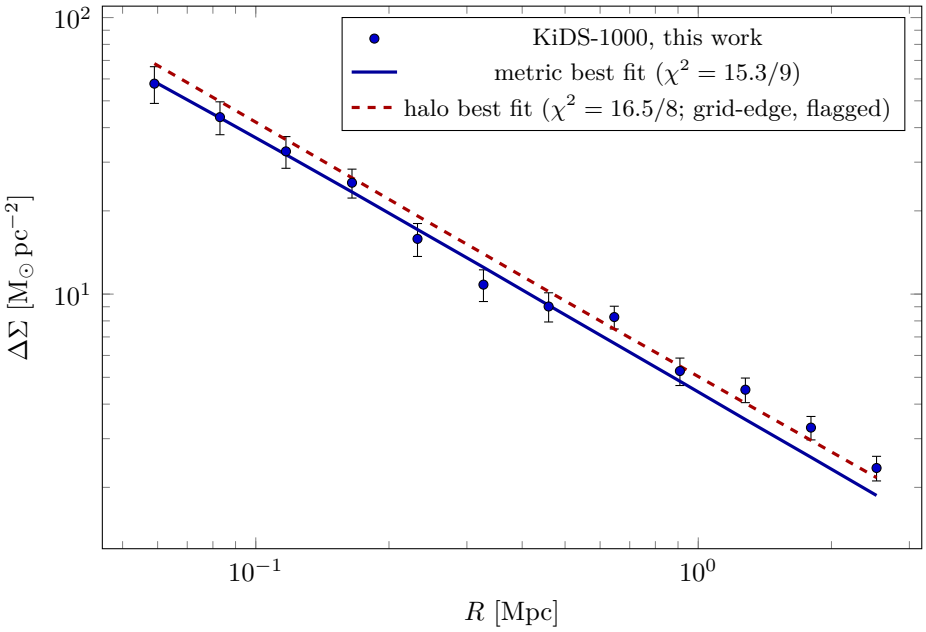

Figure 3: The Phase-A measurement with both best-fit branches. The metric curve has no halo parameters: given the stellar mass, its shape is fixed by a_0 plus the shared nuisances. The halo curve is fitted with the Phase-A power-law two-halo surrogate and pins at its (twice-extended) M_{200} grid floor — flagged, not hidden; the fair comparison requires the Phase-B halofit-grade template.

8 Conclusions

Phase A set out to turn a sharp qualitative fork — square-root persistence versus finite-edge turnover in the excess acceleration around galaxies — into a real likelihood programme on public KiDS-1000 data, and to do so under a discipline in which the instrument must earn the right to measure. It delivered a validated pipeline (14.6σ against an unmoved 5σ bar), release-verified survey parameters that keep the registered acceptance bar within reach at a $16\times$ lens margin, a complete and, we hope, generally useful forensic record (a one-bit frame convention diagnosed by the invariance contrast between cosmic shear and GGL and adjudicated by the sky at 18σ; a photo-z isolation pathology that selects label errors), and a first structure test whose honest result is a draw with the halo branch's configuration flagged. Phase A did not decide the fork; the instrument that can decide it now exists, is public in its every assertion, and is pointed at the deep window where the two pictures cannot both survive.

Reproducibility. Every number in this paper is produced by self-asserting programs in the programme repository (the Phase-A pipeline gate and its predecessors, with their run-by-run kill ledgers preserved in-file), operating on public KiDS-1000 release files. The programme's published chain is archived on Zenodo [3, 6].

References

[1] 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.

[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] 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.

[4] 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.

[5] 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.

[6] David Elliman. From counts to observables: The response layer of a finite record substrate. Zenodo, 2026.

[7] Candace Oaxaca Wright and Tereasa G. Brainerd. Gravitational lensing by NFW halos. *The Astrophysical Journal*, 534:34–40, 2000. doi: 10.1086/308744.

[8] 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.

[9] 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.

[10] Alan R. Duffy, Joop Schaye, Scott T. Kay, and Claudio Dalla Vecchia. Dark matter halo concentrations in the Wilkinson Microwave Anisotropy Probe year 5 cosmology. *Monthly Notices of the Royal Astronomical Society: Letters*, 390:L64–L68, 2008. doi: 10.1111/j.1745-3933.2008.00537.x.

[11] 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.

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

[13] Angus H. Wright, Hendrik Hildebrandt, Jan Luca van den Busch, and Catherine Heymans. Photometric redshift calibration with self-organising maps. *Astronomy & Astrophysics*, 637: A100, 2020. doi: 10.1051/0004-6361/201936782.

[14] L. Miller et al. Bayesian galaxy shape measurement for weak lensing surveys – III. application to the Canada–France– Hawaii Telescope Lensing Survey. *Monthly Notices of the Royal Astronomical Society*, 429:2858–2880, 2013. doi: 10.1093/mnras/sts454.

[15] I. Fenech Conti, R. Herbonnet, H. Hoekstra, J. Merten, L. Miller, and M. Viola. Calibration of weak-lensing shear in the Kilo-Degree Survey. *Monthly Notices of the Royal Astronomical Society*, 467:1627–1651, 2017. doi: 10.1093/mnras/stx200.

[16] Arun Kannawadi et al. Towards emulating cosmic shear data: revisiting the calibration of the shear measurements for the Kilo-Degree Survey. *Astronomy & Astrophysics*, 624:A92, 2019. doi: 10.1051/0004-6361/201834819.

[17] Hendrik Hildebrandt et al. KiDS-1000 catalogue: Redshift distributions and their calibration. *Astronomy & Astrophysics*, 647:A124, 2021. doi: 10.1051/0004-6361/202039018.

[18] 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.

[19] Rachel Mandelbaum et al. Systematic errors in weak lensing: application to SDSS galaxy–galaxy weak lensing. *Monthly Notices of the Royal Astronomical Society*, 361:1287–1322, 2005. doi: 10.1111/j.1365-2966.2005.09282.x.

[20] J. Hartlap, P. Simon, and P. Schneider. Why your covariance matrix estimate is biased: unbiased estimation of the inverse covariance matrix. *Astronomy & Astrophysics*, 464: 399–404, 2007. doi: 10.1051/0004-6361:20066170.

[21] Marika Asgari et al. KiDS-1000 cosmology: Cosmic shear constraints and comparison between two point statistics. *Astronomy & Astrophysics*, 645:A104, 2021. doi: 10.1051/0004-6361/202039070.