# Appendix A – Quantitative simulation benches

Frozen-hierarchy tests with explicit proxy or semi-blind numerics

**Scope.** Appendix A collects the benches for which the present Part II closure already supports an explicit quantitative proxy, audit scan, or semi-blind simulation. These are not claimed as final precision solutions. Rather, they show where the frozen hierarchy already produces numerically meaningful results, where the present closure still leaves structured residuals, and where further work is required. Some of those next steps belong naturally to Part III, especially when first-principles closure, history kernels, or matching conditions are involved. Others lie beyond the immediate Part II/III split and should be read as part of the longer programme of external population studies, higher-resolution orbital libraries, or self-consistent hydro/N-body modelling.

## 1 Ultra-diffuse galaxies: DF2, DF4, and Dragonfly 44

The ultra-diffuse-galaxy (UDG) benches provide one of the cleanest regime tests presently available. The key distinction is no longer a binary "dark-matter-rich vs. dark-matter-poor" reading, but a hierarchy-based separation between *ambient-suppressed* and *upper-regime / high-tail* systems. Within the current frozen hierarchy, NGC1052-DF2 remains robustly ambient-suppressed across plausible distance and velocity-dispersion solutions, whereas Dragonfly 44 remains in a much higher-tail regime even after revised kinematic inputs are considered. The point of the bench is therefore not a single tuned fit, but a stable regime separation under parameter perturbation.

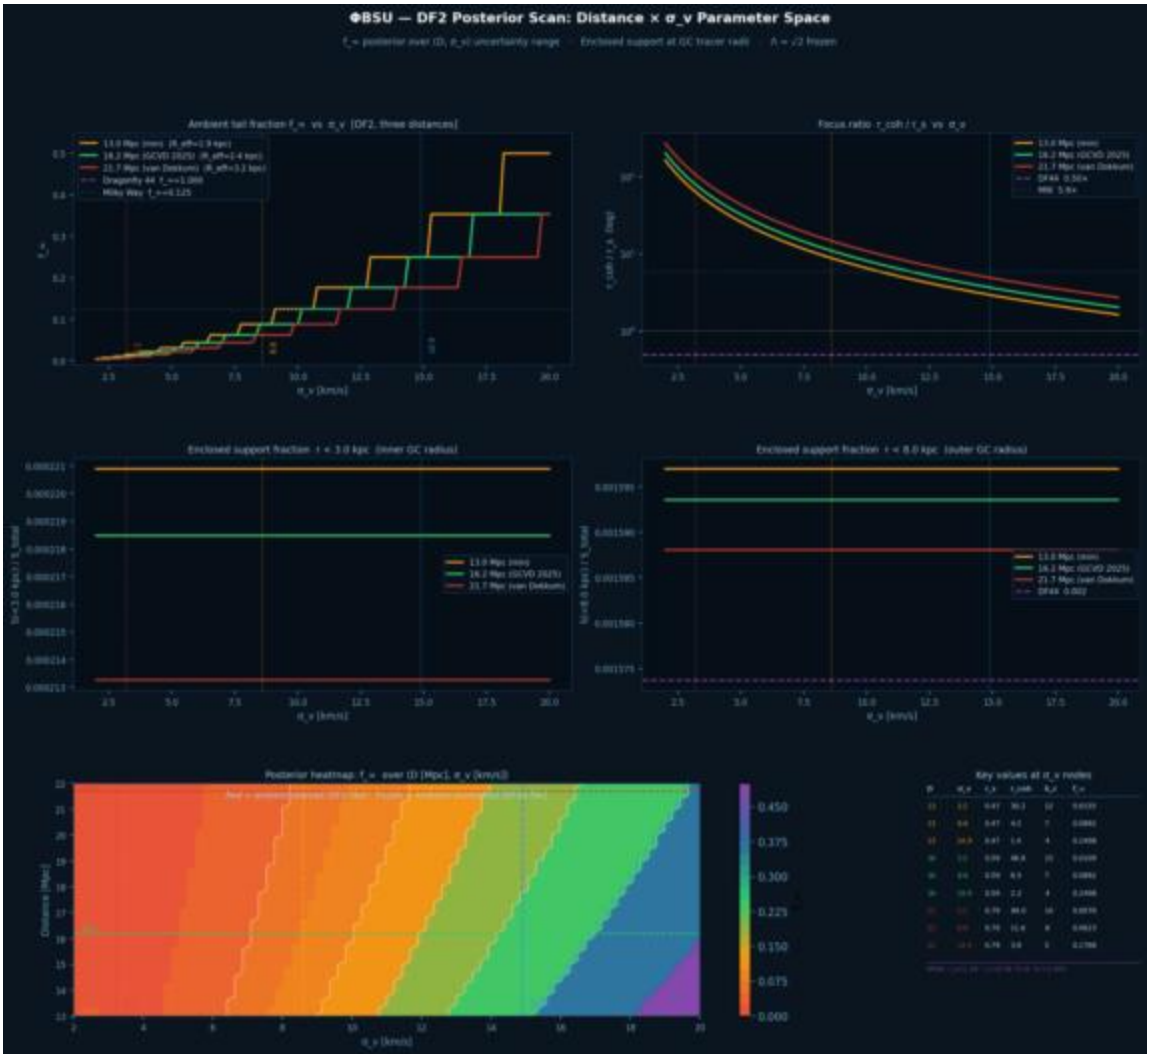


**Figure 1:** DF2 posterior scan over distance and line-of-sight dispersion. The ambient-tail fraction $f_\infty$ changes in discrete steps because the coherence index is discrete. The principal result is regime stability: DF2 remains ambient-suppressed throughout the plausible posterior region, while Dragonfly 44 sits far above it in the high-tail regime.

The semi-blind Jeans step is especially informative because it removes the observed velocity dispersion from the input side. Starting from baryons plus the frozen hierarchy, the model predicts intrinsically low DF2 dispersions. The prediction is therefore in the correct qualitative direction even when it remains somewhat below the currently preferred observational interval. This is one of the clearest examples in the present paper where the model gives a physically meaningful result without being fully closed.

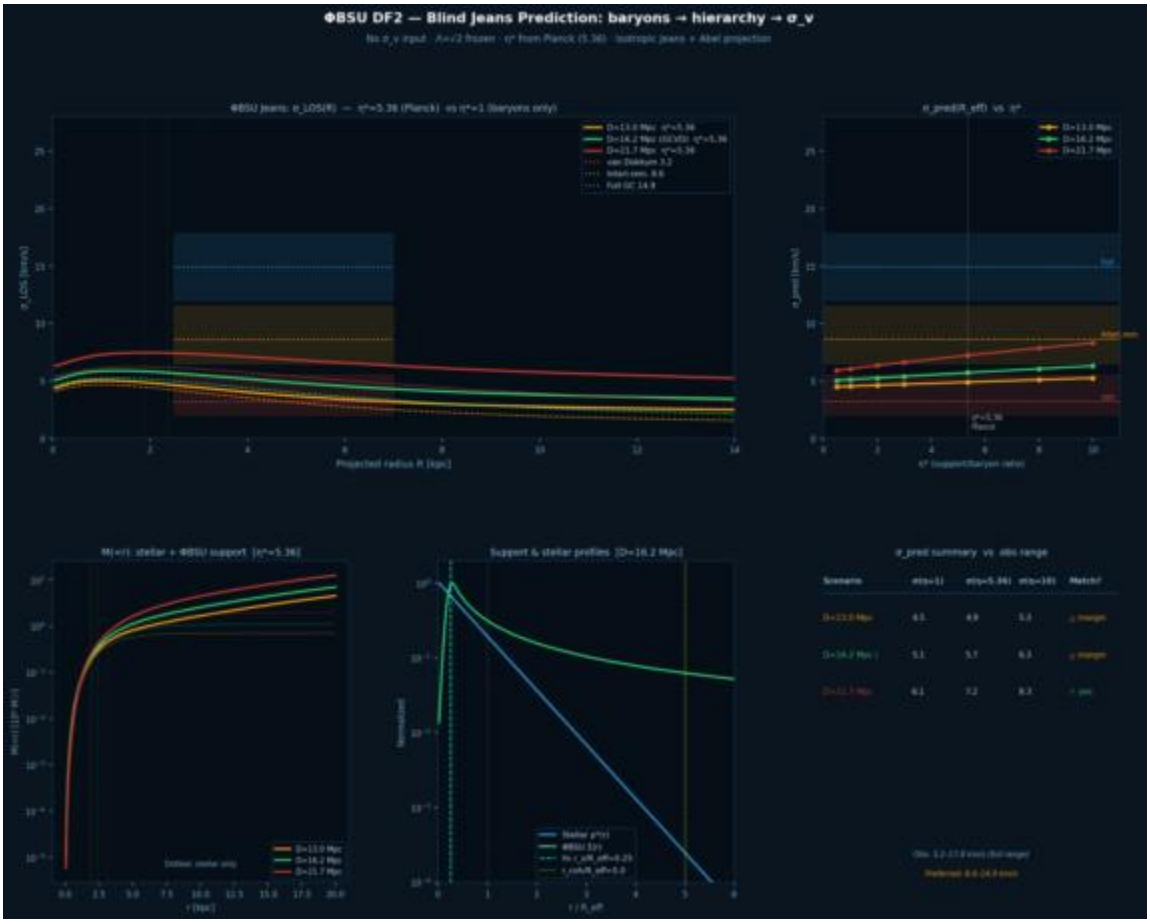


**Figure 2:** Semi-blind Jeans prediction for DF2. The frozen hierarchy places the system naturally in the low-dispersion regime, though typically somewhat below the currently preferred observational interval. This should be read as a constructive stress test rather than a final solved fit.

The revised Dragonfly 44 audit is valuable because it prevents the UDG section from resting on obsolete benchmark values. Updated kinematics soften the original “purely ambient-dominated” reading to a boundary or transitional case, but do not erase the DF2/DF44 separation. In that sense, the revised DF44 panel functions as an acid test on the regime claim rather than as a retreat from it.

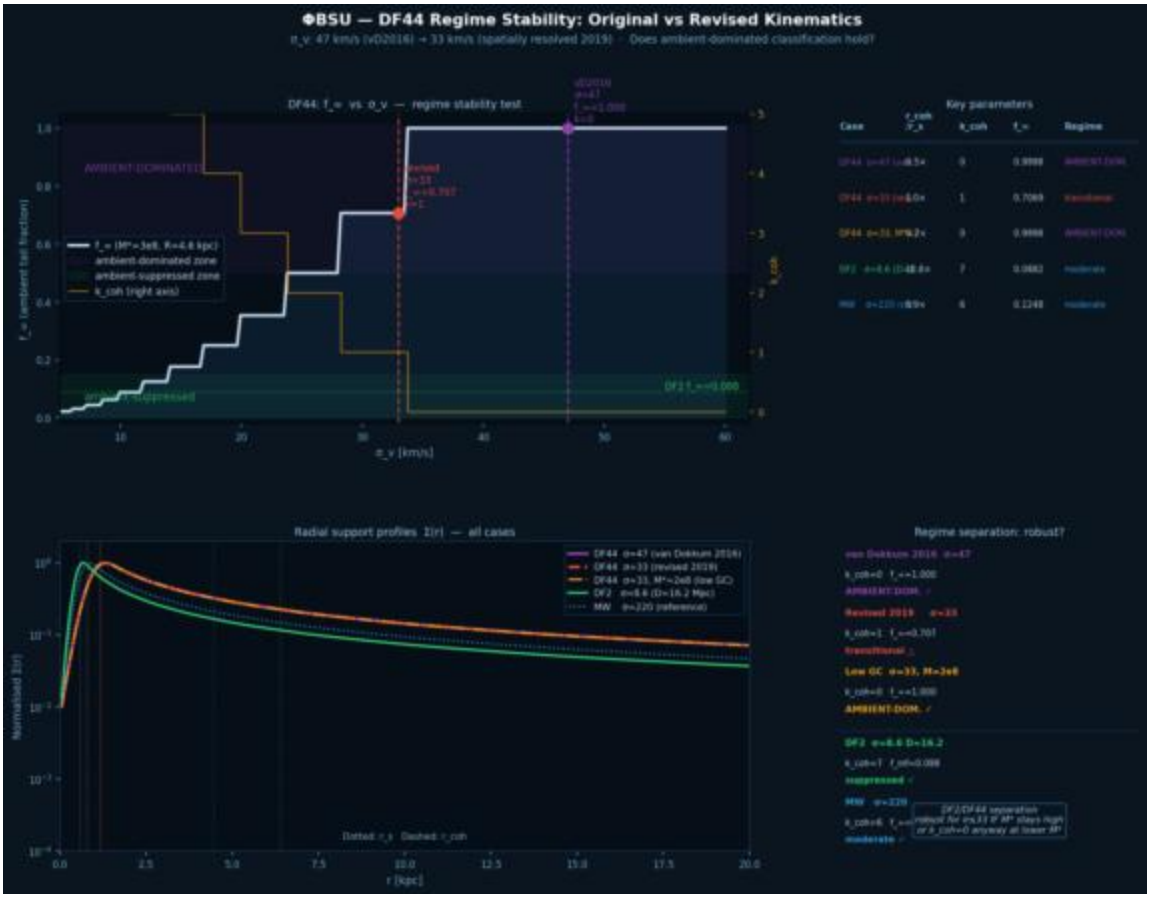


**Figure 3:** DF44 regime-stability audit under revised kinematics. The original van Dokkum benchmark remains useful as a hard upper-regime anchor, but the revised dispersion shows that DF44 should now be read as a boundary/high-tail case rather than an immutable point estimate. The DF2/DF44 separation nevertheless survives.

**Part II claim.** The frozen hierarchy already supports a robust UDG regime separation: DF2 is ambient-suppressed, DF44 is upper-regime/high-tail, and the semi-blind Jeans calculation naturally places DF2 in a low-dispersion domain.

**Part III interface.** A full closure requires tracer anisotropy, richer orbital libraries, and a principled relation between support partition and the measured velocity field. These are natural Part III interfaces.

**Further programme.** External UDG samples, broader posterior studies, and more realistic tracer-population modelling extend beyond the immediate Part II/III handoff and belong to the longer observational programme.

## 2 Galaxy-cluster bridges: MACS J0416 and the Bullet Cluster

The MACS J0416 and Bullet Cluster benches remain the clearest interacting-system tests of the deterministic overlap operator. The important quantitative result is not simply that a bridge-like term can be drawn, but that the active overlap window is selected by the same frozen first-contact logic in both systems, and that the cross-term improves bridge-region morphology without requiring a separate particle halo ontology.

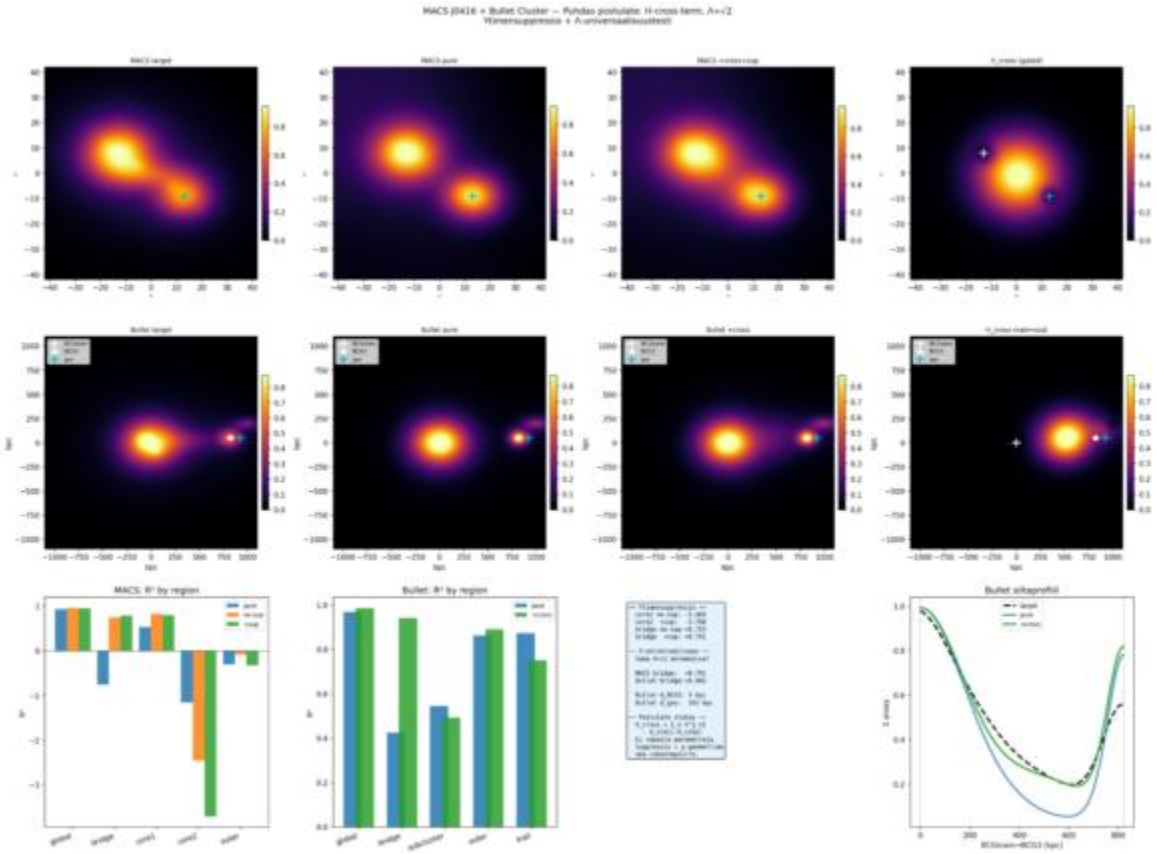

**Figure 4:** Cross-term geometry for MACS J0416 and the Bullet Cluster. The overlap contribution is most naturally read as a redistribution term within the hierarchy rather than as a new ontological sector. The quantitative Part II result is that the bridge morphology improves in the correct regions when the deterministic overlap closure is turned on.

Residual maps are especially useful here, because they prevent the bridge result from being read as mere visual smoothing. The gain is local and interpretable, not a global cosmetisation. At the same time, the residuals make it obvious that the present closure is not yet the final one: trail structure, asymmetries, and merger-history sharpening remain incompletely modelled.

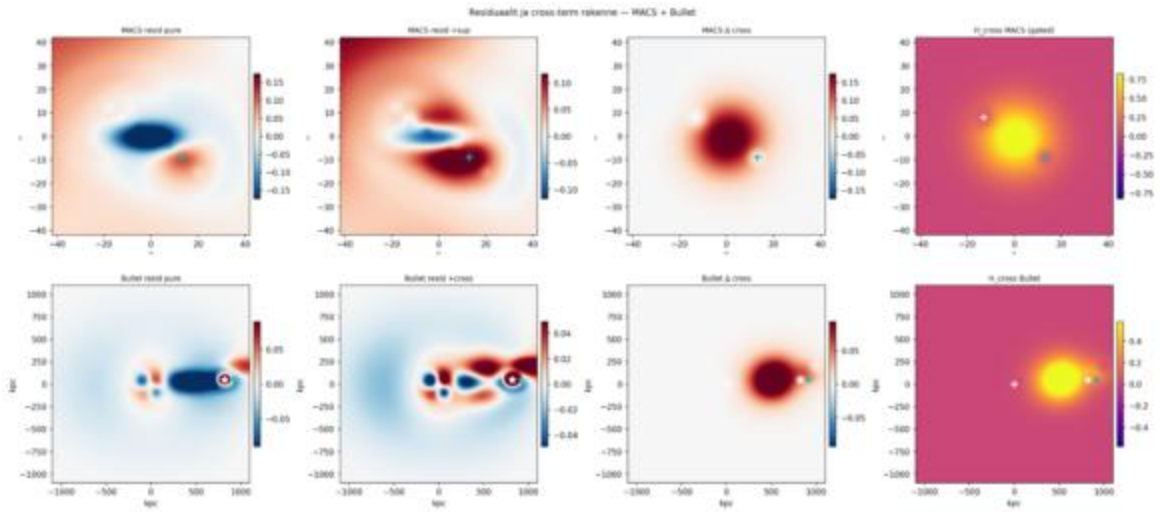

**Figure 5:** Residual structure for MACS and Bullet. The present overlap closure improves the bridge domain while leaving recognisable merger-history residuals elsewhere. This is the correct Part II outcome: the right regime concept appears numerically, but the full dynamical closure is not yet exhausted.

**Part II claim.** At cluster scale, the deterministic overlap operator already functions quantitatively as a bridge redistributor and scale selector. The bridge result is therefore not merely qualitative.

**Part III interface.** The remaining asymmetries, trail sharpening, and fully time-dependent merger geometry require history kernels and encounter dynamics beyond the frozen closure.

**Further programme.** A larger interacting-cluster sample and explicit image-plane forward modelling belong to the later programme once the Part II closure is stabilised.

# 3 Rotation-support benches: Milky Way and M31

The galactic rotation-support tests occupy an intermediate position between sourcewise structure and cosmological closure. For the Milky Way, the frozen hierarchy plus hidden-zone forecast already produces a coherent qualitative and semi-quantitative picture of the inner disk/bar support. For M31, the present benchmark remains morphology-first and sourcewise, but it is sufficient to show that the same hierarchy used in cluster and UDG tests also produces plausible galaxy-scale support distributions.

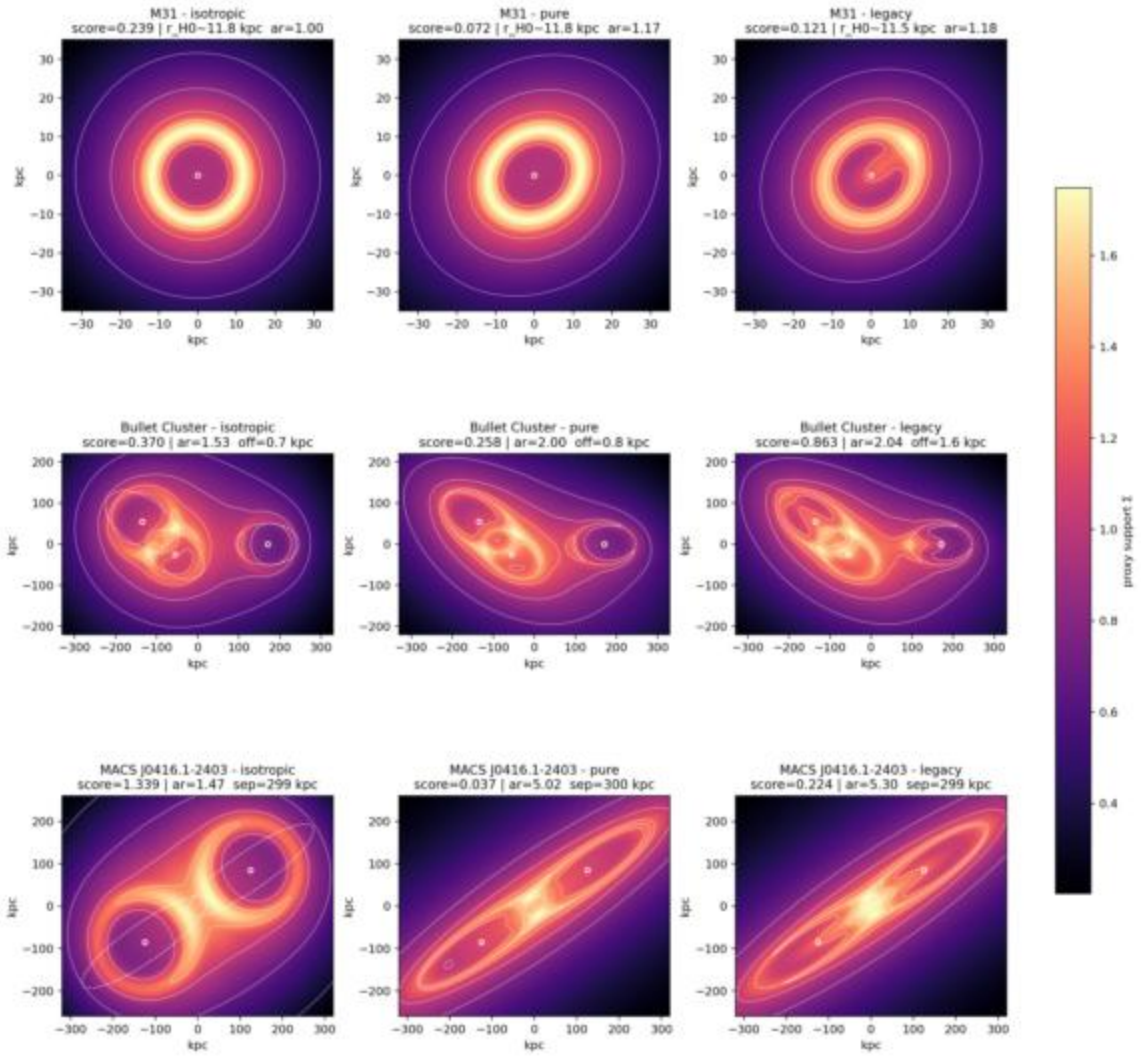


**Figure 6:** Common benchmark comparison across the current frozen-hierarchy tests. M31 appears here as a sourcewise galaxy-scale anchor, alongside the already discussed cluster benches. The point is not that all systems are fit to equal precision, but that the same hierarchy law continues to organise them without re-tuning the scale ladder.

For the Milky Way, the most useful current output is the forecast map itself. It already makes visible where support excess is expected in observationally difficult regions, and the same map provides the hidden-zone correction later used in the qualitative satellite-plane discussion. The present Part II result is therefore twofold: the frozen hierarchy yields a natural low-latitude support asymmetry, and the resulting rotation-support field is not arbitrary but tied to the same disk/bar baryonic seed.

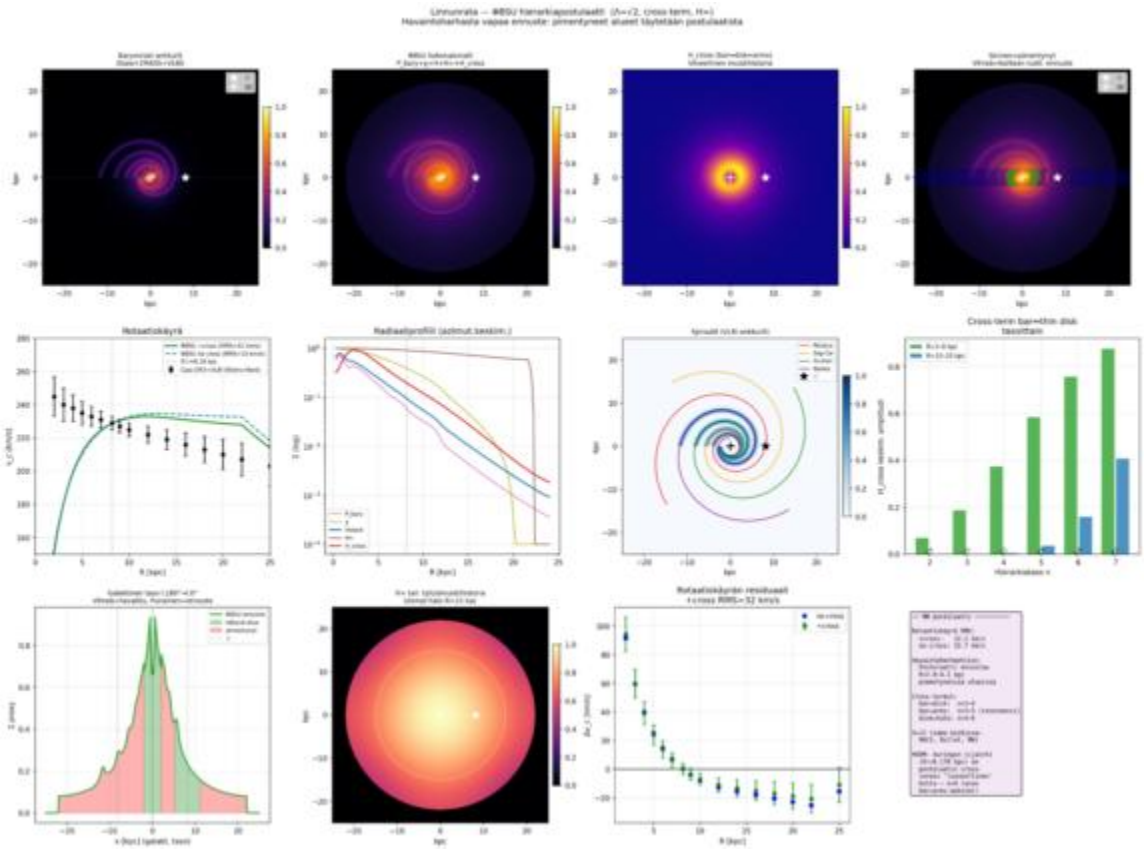

**Figure 7:** Milky Way forecast map and support decomposition. The hidden-zone excess and bar/disk anisotropy are not independent add-ons but consequences of the sourcewise baryonic hierarchy. This forecast later feeds into the qualitative satellite-plane discriminator discussed in Appendix B.

**Part II claim.** The Milky Way and M31 benches show that the frozen hierarchy remains numerically usable on galaxy scales: Milky Way inner support asymmetry and M31 morphology-first support are both organised by the same sourcewise law.

**Part III interface.** A sharper treatment of cylindrical mass integration, outer-halo closure, and the relation between support and tracer anisotropy belongs partly to Part III and partly to later dedicated galactic-dynamics work.

**Further programme.** Large external galaxy samples and stronger benchmark suites are required before the galaxy-scale closure can be treated as mature in the same sense as the best UDG and cluster tests.

# 4 Fornax dSph: globular-cluster survival

The Fornax dwarf spheroidal remains one of the best focused benches for a smooth-support interpretation. The decisive comparison is not between two arbitrary halo fits, but between a smooth support field and a particle-scattering halo analogue. In the present semianalytic implementation, this difference is enough to change the fate of the inner globular clusters.


**Figure 8:** Baseline Fornax globular-cluster survival test. In the smooth-support implementation, all five clusters remain outside the core over 12 Gyr, whereas the inner clusters sink rapidly in the particle-scattering halo analogue. The strongest discriminatory content lies in the survival of the inner clusters.

The extended mass-sensitivity and orbital-family scan strengthens the result. The advantage of the smooth-support field is largest in the observed globular-cluster mass range and for low-to-moderate eccentricities, i.e. precisely where the real Fornax clusters are thought to live. This is not yet a proof against all wave-like dark-sector alternatives, but it is a strong quantitative argument against reading the support field as a conventional scattering halo.

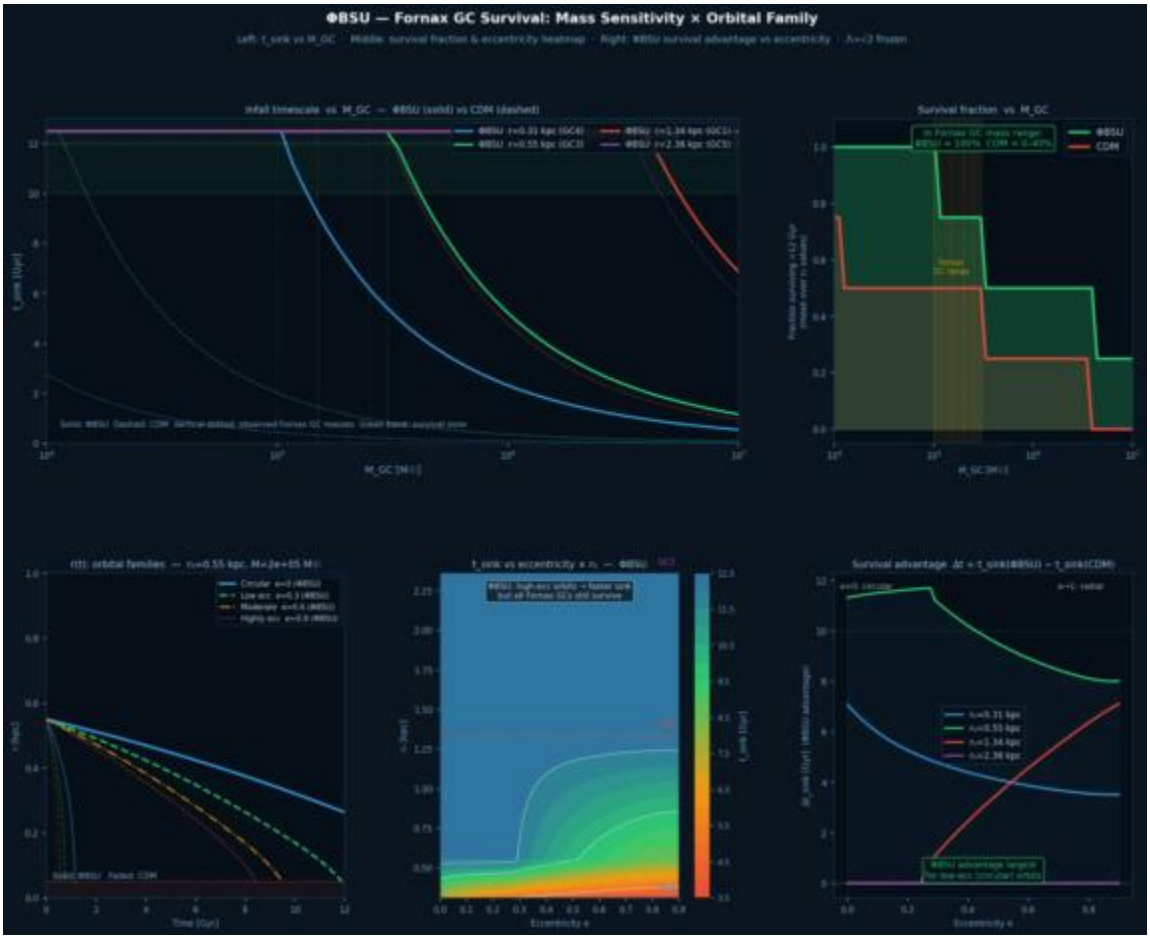


**Figure 9:** Mass sensitivity and orbital-family extension of the Fornax bench. The current result is strongest precisely in the observed cluster-mass range and for the more circular orbits favoured observationally. This makes the survival advantage physically meaningful rather than purely cosmetic.

**Part II claim.** The Fornax bench quantitatively favours a smooth support field over a particle-scattering halo analogue, chiefly because the inner clusters survive.

**Part III interface.** A deeper derivation of the friction law, together with possible connections to wave-like dark-sector alternatives and microphysical interior structure, lies at the Part II/Part III boundary.

**Further programme.** Broader cluster-mass scans, external dwarf samples, and fully orbit-averaged modelling extend beyond the present paper and should be treated as part of the future programme rather than deferred mechanically to Part III alone.

# 5 Solar-system and trans-Neptunian audit

The Solar-system branch is already quantitative in the sense that it produces concrete observability and falsification statements. What remains open is not whether the hierarchy can be parameterised, but how the outer annulus or band should be interpreted geometrically: as a torus, a shell with ecliptic enhancement, or a Solar×Galactic matching zone. The present bench is therefore a good falsifier and a good observability audit even before the deeper matching problem is solved.

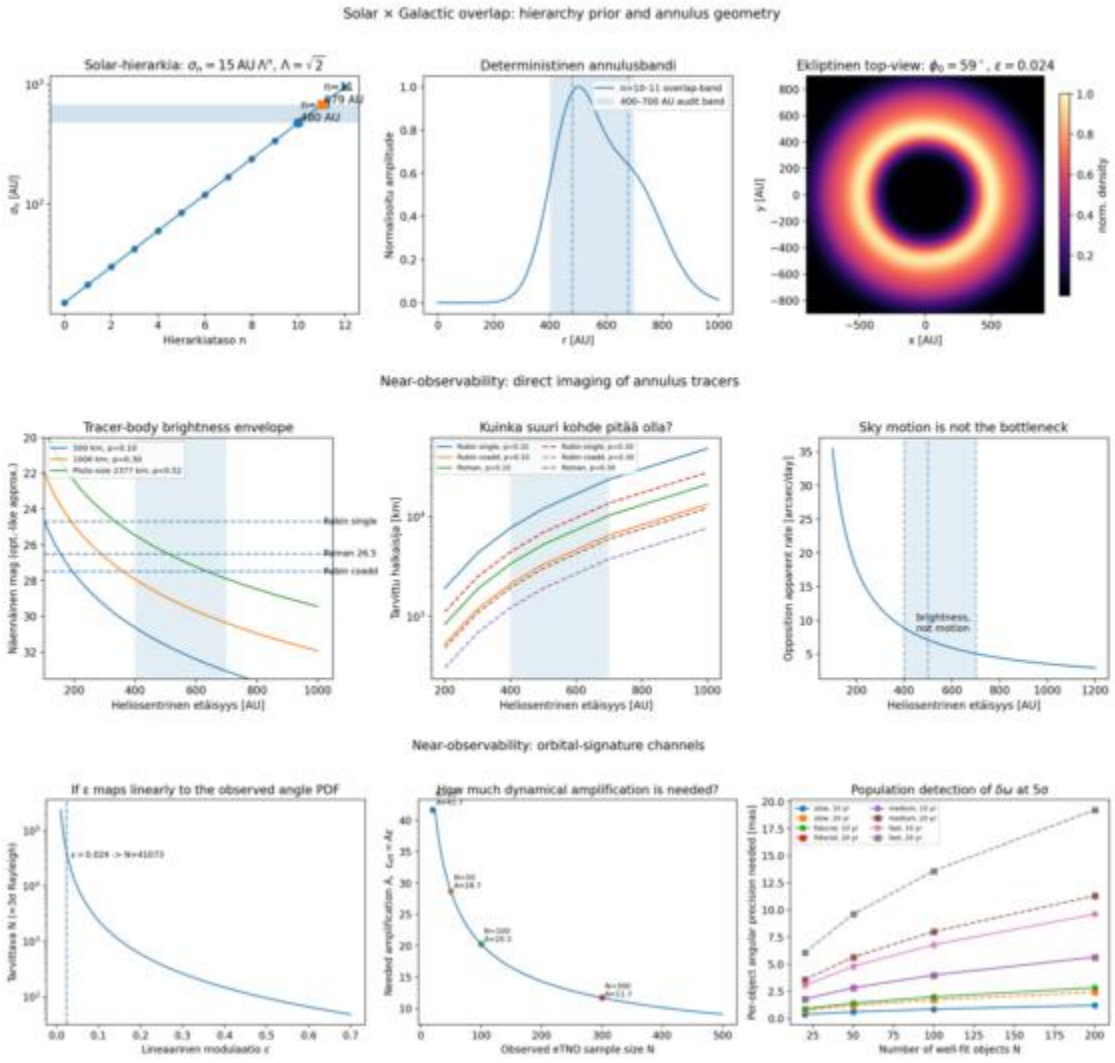


**Figure 10:** Solar-system overlap dashboard. The hierarchy supports explicit predictions for annulus-like outer support, broad clustering signatures, and small but testable precession effects while remaining consistent with inner-system ephemeris constraints at the present audit level.

The detectability outputs are especially useful because they force the discussion away from mere morphology and into real observational timelines. At the same time, they show the limits of the current closure: direct imaging of a smooth support annulus is not the near-term observable; tracer dynamics and population-level angular structure are more realistic targets.

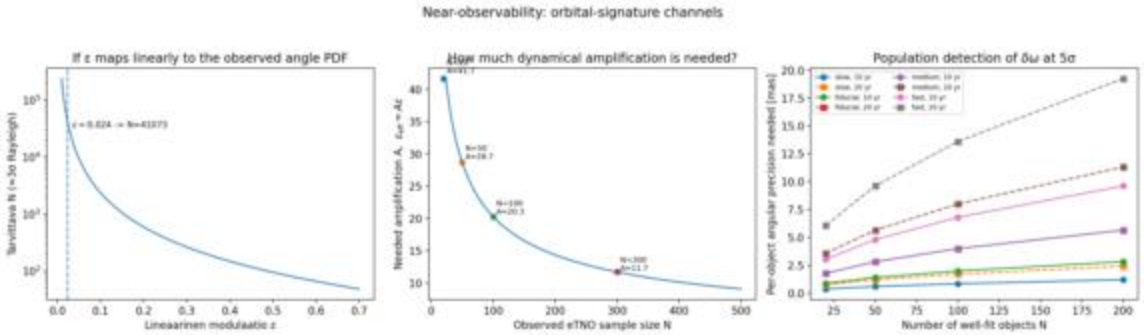


**Figure 11:** Detectability summary for the trans-Neptunian / Solar×Galactic overlap audit. The strongest near-term observable is not a direct "ring image" but a tracer-population signature, together with ephemeris-safe constraints and longer-baseline precession diagnostics.

**Part II claim.** The Solar-system branch is already quantitative as a falsifier and observability audit: the frozen hierarchy yields explicit outer-support scales, directional signatures, and ephemeris-safe bounds.

**Part III interface.** The geometric nature of the annulus or band, and its matching to the Galactic field, remains a Part III interface. The current Part II status is therefore predictive but not yet ontologically final.

**Further programme.** Selection-aware discovery simulations, external orbital-database reanalyses, and future Rubin/Roman inference pipelines belong to the longer observational programme even after Part III appears.

# 6 Working summary for Appendix A

Appendix A gathers the parts of the programme where the frozen hierarchy already supports explicit numerics. The UDG tests show a robust regime separation, the cluster bridges demonstrate quantitative overlap redistribution, the Milky Way and M31 provide galaxy-scale support anchors, the Fornax bench discriminates strongly against a particle-scattering reading, and the Solar-system branch functions as a hard falsifier and observability audit. None of these benches is claimed to be the final closure. Their purpose is instead to identify where the current Part II framework already produces structured numerical content, and where that content points beyond itself.

**Practical reading rule.** Appendix A should be cited as the quantitative face of the frozen hierarchy: not a catalogue of perfect fits, but a set of explicit numerical benches showing where the Part II closure already works, where it still leaves residual structure, and where further work belongs either to Part III or to the longer-term simulation and observation programme.