












Dynamical disequilibrium in dwarf galaxies: rethinking gas dynamics, rotation curves, and dark matter inference

Diego Dado ^{1,2★}, Kyle A. Oman ^{1,2}, Katherine E. Harborne ^{1,2}, Francesca Fragkoudi ¹,
Joop Schaye ³, Matthieu Schaller ^{3,4}, Alejandro Benítez-Llambay ⁵, Evgenii Chaikin ³,
Carlos S. Frenk ¹, Filip Huško ³, Sylvia Ploekinger ⁶ and Alexander J. Richings ^{7,8}

¹Department of Physics, Institute for Computational Cosmology, Durham University, South Road, Durham DH1 3LE, UK

²Department of Physics, Centre for Extragalactic Astronomy, Durham University, South Road, Durham DH1 3LE, UK

³Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands

⁴Lorentz Institute for Theoretical Physics, Leiden University, PO Box 9506, NL-2300 RA Leiden, the Netherlands

⁵Dipartimento di Fisica G. Occhialini, Università degli Studi di Milano Bicocca, Piazza della Scienza, 3 I-20126 Milano MI, Italy

⁶Department of Astrophysics, University of Vienna, Türkenschanzstrasse 17, A-1180 Vienna, Austria

⁷Centre for Data Science, Artificial Intelligence and Modelling, University of Hull, Cottingham Road, Hull HU6 7RX, UK

⁸E. A. Milne Centre for Astrophysics, University of Hull, Cottingham Road, Hull HU6 7RX, UK

Accepted 2026 February 12. Received 2026 February 3; in original form 2025 December 11

ABSTRACT

We quantify departures from hydrodynamical and centrifugal equilibrium in the gas discs of low-mass ($10^{10.75} < M_{200c}/M_{\odot} < 10^{11}$) galaxies from the COLIBRE cosmological hydrodynamical simulations. We evaluate the full Eulerian acceleration balance in the mid-plane and show that disequilibrium is widespread: equilibrium-based circular velocity estimates typically have errors of ≥ 10 per cent (≈ 75 per cent of mid-plane gas by mass). Disequilibrium is strongest and the largest associated errors occur in the inner few kiloparsecs that are crucial for constraining the dark matter density profile. Correcting the circular velocity to account for pressure and convective terms does not reliably improve its recovery in strongly perturbed systems where time-dependent forces dominate the residual acceleration budget. Stellar feedback, self-gravitating gas clumps, and active galactic nucleus energy injection account for most strong local perturbations, and large-scale gravitational asymmetries act as a scaffold for disequilibrium. We classify gas discs into coherent, perturbed, and slow/erratic rotators and show that this classification correlates with galaxy properties like mass, morphology, and tracers of recent feedback. A majority of galaxies in our sample would be unsuitable for standard rotation curve analysis. Much of the observed diversity in the shapes of dwarf galaxy rotation curves may stem from non-equilibrium gas motions rather than diversity in mass profiles – resolving the discrepancy is then first and foremost a problem in gas dynamics.

Key words: methods: numerical – galaxies: dwarf – galaxies: ISM – galaxies: kinematics and dynamics – dark matter.

1 INTRODUCTION

Despite the remarkable successes of the standard model of cosmology (Λ CDM; dark energy, Λ , and cold dark matter, CDM) in reproducing cosmic microwave background anisotropies (e.g. Planck Collaboration XIII 2016), the clustering signal in large-scale structure (e.g. D. J. Eisenstein et al. 2005), and gravitational lensing observations (e.g. C. Heymans et al. 2013), significant challenges remain on galactic scales – particularly for low-mass galaxies.

Dwarf galaxies – broadly defined as systems with stellar masses $M_{\star} < 10^9 M_{\odot}$ – fall in a crucial regime for testing both galaxy formation physics and the nature of dark matter. As the most dark-matter-dominated galaxies known, they (potentially) offer clean laboratories for studying gravitational dynamics, and

thus dark matter, in low baryon-fraction environments (see e.g. L. V. Sales, A. Wetzel & A. Fattahi 2022). At the same time, their shallow potential wells render them highly susceptible to feedback processes, introducing complex baryonic effects that complicate both simulations and observations (e.g. baryon-induced core creation, BICC; A. Pontzen & F. Governato 2012; A. Benítez-Llambay et al. 2019).

A long-standing issue in this context is the cusp–core problem: the discrepancy between the steep inner density profiles ($\rho \propto r^{-1}$) predicted by dark-matter-only N -body simulations (J. F. Navarro, C. S. Frenk & S. D. M. White 1997) and the flatter profiles inferred from observed galaxy rotation curves, once baryonic contributions are accounted for (R. A. Flores & J. R. Primack 1994; B. Moore 1994; W. J. G. de Blok 2010). This is typically evaluated by fitting a power law to the inferred dark matter density in the innermost resolved region, and comparing its slope to simulated haloes at similar radii (e.g. W. J. G. de Blok et al. 2001; W. J. G. de Blok & A. Bosma 2002; R. A. Swaters et al. 2003; S.-H. Oh et al.

* E-mail: diego.dado2@durham.ac.uk

2011a; V. H. Robles et al. 2017). A mismatch suggests tension with Λ CDM predictions.

Recent advances have shifted the focus from inner slopes to more global discrepancies. Cosmological parameters are now tightly constrained on large scales (e.g. Planck Collaboration XIII 2016), and dark matter halo scaling relations – such as the dependence of halo structural properties on cosmology – are well established (e.g. C. A. Correa et al. 2015; A. D. Ludlow et al. 2016; B. Diemer & M. Joyce 2019). This enables robust predictions of halo circular velocity profiles once a single parameter (e.g. the maximum circular velocity, V_{\max}) is known.

Following this logic, K. A. Oman et al. (2015) reformulated the cusp–core problem in terms of an inner mass deficit: a shortfall in the circular velocity at a fixed inner radius (e.g. 2 kpc, see also I. M. E. Santos-Santos et al. 2020, for a later, more flexible definition) relative to that expected from the dark matter halo alone. Some galaxies exhibit rotation velocities far below this baseline, even when sharing the same V_{\max} as systems without such deficits. The phenomenon spans a wide range of galaxy masses, from dwarfs to large discs, and varies substantially at fixed galaxy properties (e.g. the baryonic surface mass density Σ_{bar} , or the baryonic mass contribution in the inner regions $\eta_{\text{bar}} = (V_{\text{bar, fid}}/V_{\text{fid}})^2$).

In extreme cases, the inferred deficits exceed the total baryonic mass, challenging the hypothesis that they arise from BICC (K. A. Oman et al. 2015). While such processes – driven by bursty feedback and gravitational coupling (J. F. Navarro, V. R. Eke & C. S. Frenk 1996; S. Mashchenko, J. Wadsley & H. M. P. Couchman 2008; F. Governato et al. 2012; A. Pontzen & F. Governato 2012) – can flatten central density profiles, simulations cannot reproduce the most extreme cores inferred from observations and indicate that these baryonic processes are limited to a narrow range of halo masses (A. Di Cintio et al. 2014; T. K. Chan et al. 2015). Moreover, not all state-of-the-art simulations predict cores: cosmological hydrodynamical simulations like EAGLE (M. Schaller et al. 2015; J. Schaye et al. 2015, see also A. Benítez-Llambay et al. 2019) and Illustris (M. Vogelsberger et al. 2014; K. T. E. Chua et al. 2019) reproduce key galaxy statistics without featuring them.

In addition to a solution based on modifications due to baryonic feedback, alternative explanations fall into two main categories, as outlined by K. A. Oman et al. (2015): (i) dark matter is not cold and collisionless; or (ii) mass distributions are incorrectly inferred from observations.

Departures from the central assumption of the CDM framework – that dark matter behaves as a cold collisionless fluid – can alter the internal kinematics of dwarf galaxies. Self-interacting dark matter (SIDM) models have been proposed (D. N. Spergel & P. J. Steinhardt 2000), in which scattering interactions among DM particles are introduced. For interaction cross-sections per unit mass of order $\sigma/m = 1 \text{ cm}^2 \text{ g}^{-1}$, such interactions can thermalize the central regions of DM haloes, producing the constant-density cores observed in dwarfs while retaining the large-scale successes of CDM (see e.g. M. Rocha et al. 2013; O. D. Elbert et al. 2015; T. Ren et al. 2019). SIDM has also been invoked to explain the broad diversity of dwarf galaxy rotation curves; however, additional effects due to galaxy formation physics complicate this picture: introducing baryon–DM coupling and fine-tuning seems inevitable (P. Creasey et al. 2017; M. Kaplinghat, T. Ren & H.-B. Yu 2020; J. D. Burger et al. 2022; C. A. Correa et al. 2025).

Concerns about residual systematics in kinematic modelling and observational mass inference have also been a long-standing topic of discussion in the field. Emission line-widths can be

affected by ‘beam smearing’ (see R. A. Swaters et al. 2009, and references therein), or centring, alignment, and seeing in the case of $H\alpha$ slit spectroscopy (R. A. Swaters et al. 2003; K. Spekkens, R. Giovanelli & M. P. Haynes 2005). Mass models can be affected by ambiguities in the mass-to-light ratios of galaxies, or the rotation velocity-inclination degeneracy ($V_{\text{rot}} - i$). These effects can lead to biased inferences of the underlying dark matter distribution, creating the illusion of a core in intrinsically cuspy haloes (e.g. F. C. van den Bosch & R. A. Swaters 2001; A. A. Dutton et al. 2005). However, analyses based on high-resolution 21-cm radio observations of dwarf galaxies – particularly from the THINGS (F. Walter et al. 2008) and LITTLE THINGS (D. A. Hunter et al. 2012) surveys, and the SPARC compilation (F. Lelli, S. S. McGaugh & J. M. Schombert 2016) – continue to indicate the presence of some central cores, where concerns around many of these systematics are greatly alleviated. The mass-to-light ratios for galaxies in the THINGS and LITTLE THINGS surveys were estimated using empirical relations based on stellar optical colours (E. F. Bell & R. S. de Jong 2001; G. Bruzual & S. Charlot 2003), which have been extensively validated through independent methods (see e.g. S.-H. Oh et al. 2011b). In the SPARC compilation, stellar masses are derived from near-infrared luminosities at $3.6 \mu\text{m}$ – a wavelength widely considered one of the most reliable tracers of stellar mass (E. F. Bell & R. S. de Jong 2001), as the corresponding mass-to-light ratio depends only weakly on age and metallicity across a broad range of star formation histories.

Nevertheless, the possibility that the observed cores are merely artefacts of systematic uncertainties in observation or modelling has not been conclusively ruled out. Dark matter haloes have long been predicted to be triaxial in shape, as indicated by some of the earliest CDM simulations (M. Davis et al. 1985; C. S. Frenk et al. 1988). Such triaxiality gives rise to an aspherical gravitational potential, with the strongest deviations from spherical symmetry occurring in the central regions (E. Hayashi, J. F. Navarro & V. Springel 2007). A disc galaxy embedded within such a halo – typically aligned with the plane defined by the halo’s intermediate and major axes (E. Hayashi & J. F. Navarro 2006) – thus experiences a non-axisymmetric potential. This induces non-circular motions in the disc, as gas orbits become elongated, producing primarily bisymmetric ($m = 2$ harmonic) variations in azimuthal velocity as a function of projection angle (e.g. A. Marasco et al. 2018; K. A. Oman et al. 2019). Even mild asphericity in the potential can generate appreciable non-circular motions (e.g. E. Hayashi & J. F. Navarro 2006). Although harmonic decompositions of observed dwarf galaxy velocity fields generally indicate small non-circular motion amplitudes (C. Trachternach et al. 2008), these are likely underestimated in standard model fitting procedures (L. Chemin et al. 2020). In many cases, velocity variations arising from non-circular motions are absorbed by degenerate model parameters (e.g. F. A. Roper et al. 2023), potentially leading to substantial errors in rotation curve measurements, that is, significant discrepancies between the inferred and true rotational velocities at a given radius.

Despite the promising efforts to resolve the diversity of rotation curve shapes tension in studies targeting the effects of non-circular motions (A. Marasco et al. 2018; K. A. Oman et al. 2019; F. A. Roper et al. 2023) and/or pressure gradients (e.g. O. Valenzuela et al. 2007; L. Chemin et al. 2016; J. I. Read et al. 2016; P. E. Mancera Piña et al. 2025), they have yet to provide conclusive evidence that the tension has been fully resolved (I. M. E. Santos-Santos et al. 2020; see also L. V. Sales et al. 2022 for a recent review). This realization has motivated a closer examination of

the hierarchy of assumptions required to reduce the complex dynamics of gaseous discs to a one-dimensional rotation curve, i.e. an azimuthally averaged rotational velocity profile, $V_{\text{rot}}(r)$ – as a tracer of the underlying gravitational potential.

At the heart of this reduction lies a critical physical assumption: that the gas is in a quasi-steady state, or equivalently, that the system’s time-scale of intrinsic velocity evolution are long compared to dynamical time-scales – or, in general, $\partial_t \mathbf{v}$ is negligible compared to other terms. This (quasi-)equilibrium condition is invoked – either explicitly or implicitly – when equating the observed rotation velocity of the gas to the circular velocity induced by the gravitational potential. Mathematically, this equivalence follows from Euler’s equation for fluid motion. Hence, any correction that relies solely on time-independent terms, such as accounting for thermal or turbulent pressure gradients (e.g. asymmetric drift) or non-centrifugal convective motions (e.g. higher-order harmonics), remains inherently limited by neglecting the time-dependent evolution of the velocity field.

Recent simulation studies have shown that gas in dwarf galaxies often violates these equilibrium assumptions, owing to a range of perturbative processes. Bursty feedback, accretion flows, misaligned angular momentum, and environmental interactions can generate warps, asymmetries, and large-scale non-circular motions (e.g. K. El-Badry et al. 2018; E. R. Downing & K. A. Oman 2023; E. D. Jahn et al. 2023; M. P. Rey et al. 2024; I. S. Sands et al. 2024). These effects challenge the validity of conventional 3D or 2D velocity-field modelling pipelines, even when sophisticated corrections to the rotation curve are applied, as such methods implicitly assume steady-state dynamics.

Following this motivation, we present a novel, fully local, Eulerian analysis of mid-plane gas dynamics in a large sample of dwarf galaxies drawn from the high-resolution COLIBRE cosmological hydrodynamical simulations (E. Chaikin et al. 2025; J. Schaye et al. 2025). COLIBRE directly and self-consistently models the multiphase interstellar medium (ISM), including dust grains coupled to the chemical network. It reproduces key galaxy statistics at $z = 0$, such as the observed stellar mass function, sizes, star formation rates, HI and H₂ masses, gas and stellar metallicities, and dust abundances (J. Schaye et al. 2025). These successes, and in particular the self-consistent treatment of the cold and dense phases of the ISM and the supersampling of dark matter particles (hence enhanced resolution of the gravitational potential, mitigating spurious baryon–DM energy transfer A. D. Ludlow et al. 2019, 2021, 2023), make COLIBRE particularly well suited to identifying where steady-state assumptions break down, and to tracing the physical origins of such departures from equilibrium.

In this work, we quantify when and where traditional dynamical modelling assumptions fail by directly computing gravitational, pressure-gradient, and convective accelerations in the simulation mid-plane. Our framework also provides objective diagnostics for identifying dynamically perturbed systems and isolating the physical processes responsible for deviations from equilibrium.

In the sections that follow, we describe the simulation setup and analysis methodology (Section 2); present results on dynamical diagnostics and the accuracy of mass modelling across our sample (Section 3); identify the primary astrophysical processes driving these systems out of dynamical equilibrium (Section 4); and investigate some broad differences between perturbed and unperturbed systems (Section 5). We conclude by discussing the implications of our findings for

interpreting galaxy rotation curves and testing dark matter models (Section 6).

2 METHODS

2.1 COLIBRE simulations

The COLIBRE (COLd ISM and Better REsolution) simulation suite¹ consists of cosmological hydrodynamical runs performed with the public SWIFT code (M. Schaller et al. 2024), version 2025.04, extended with a bespoke subgrid physics implementation. The simulations adopt the flat Λ CDM cosmology from the ‘DES Y3 + All Ext.’ data set (T. M. C. Abbott et al. 2022), consistent with the fiducial FLAMINGO simulations (J. Schaye et al. 2023), with parameters: $h = 0.681$, $\Omega_{\text{m}} = 0.306$, $\Omega_{\text{b}} = 0.0486$, $\Omega_{\Lambda} = 0.6939$, $n_{\text{s}} = 0.967$, $\sigma_8 = 0.807$, $A_{\text{s}} = 2.099 \times 10^{-9}$, and $\sum m_{\nu} c^2 \approx 0.06$ eV.

COLIBRE uses the SPHENIX hydrodynamics scheme (J. Borrow et al. 2022), a smoothed particle hydrodynamics (SPH) formulation optimized for galaxy formation. This method employs a density–energy approach with artificial viscosity and conduction terms, using a quartic spline kernel with smoothing length $\eta = 1.2348$, corresponding to approximately 65 weighted neighbours.

Key subgrid physics modules in COLIBRE enhance the realism of the interstellar and circumgalactic media, allowing gas to cool to temperatures as low as ~ 10 K. Radiative cooling, heating, and species abundances (atoms, ions, and molecules) are computed using HYBRID-CHIMES (A. J. Richings, J. Schaye & B. D. Oppenheimer 2014a, b; S. Ploekinger et al. 2025), which evolves the non-equilibrium abundances of electrons and nine hydrogen and helium species on the fly, and treats an additional 147 metal species using quasi-equilibrium chemistry. Dust physics is modelled self-consistently for three grain species and two size bins, coupled to the chemistry and cooling network (J. W. Trayford et al. 2026). Star formation follows a Schmidt law based on gravitational instability, without imposing an effective ISM equation of state (F. S. J. Nobels et al. 2024).

Stellar mass-loss includes metal injection from asymptotic giant branch stars, core-collapse supernovae (CCSNe), and Type Ia supernovae (SNIa) with updated yields and a calibrated SNIa delay-time distribution (Correa et al., in preparation). Pre-supernova feedback from massive stars – including stellar winds, radiation pressure, and H II region photoionization – is explicitly modelled (A. Benítez-Llambay et al. 2026). CCSN and SNIa feedback follow a stochastic thermal scheme (C. Dalla Vecchia & J. Schaye 2012), supplemented by a lower energy kinetic mode that enhances turbulence in the case of CCSN (E. Chaikin et al. 2023). The simulations analysed here employ the fiducial purely thermal active galactic nucleus (AGN) feedback model (C. M. Booth & J. Schaye 2009), updated to improve feedback sampling for low-mass black holes. We have explicitly checked that similar runs using the hybrid AGN model of F. Huško et al. (2026), which tracks black hole spin and combines thermally driven winds with kinetic jet feedback, yield results consistent with the thermal AGN model in all parts of our analysis.

¹COLIBRE is a project of the Virgo Consortium for cosmological supercomputer simulations. Further information on the project, including publications, team members, and simulation visualizations, is available at <https://colibre.strw.leidenuniv.nl/>.

COLIBRE employs four times more dark matter than baryonic particles, improving sampling of the gravitational potential and suppressing spurious CDM–baryon energy transfer (A. D. Ludlow et al. 2019, 2021, 2023). The absence of an imposed ISM equation of state allows gas structure to develop on sub-kiloparsec scales, yielding a higher effective resolution than in previous simulations with comparable baryonic particle masses. Together, these features enable a fully multiphase ISM and improved modelling of the gravitational potential, representing a major resolution improvement over earlier galaxy formation simulations (see E. Chaikin et al. 2025; J. Schaye et al. 2025 for full model and calibration details, respectively).

Halo and galaxies are identified through a three-step process: (i) a friends-of-friends (FoF) algorithm applied to dark matter particles with a linking length of 0.2 times the mean interparticle spacing; (ii) the hydrodynamics-optimized HBT-HERONS subhalo finder (V. J. Forouhar Moreno et al. 2025), a descendant of HBT+ (J. Han et al. 2018), which identifies bound structures and constructs merger trees; and (iii) the SOAP tool (R. McGibbon et al. 2025), which computes galaxy and halo properties within fixed and overdensity-based apertures, adopting M_{200c} and R_{200c} as the halo mass and radius definitions.² Throughout the analysis, we refer to individual galaxies via their HBT-HERONS track ID, which uniquely links each $z = 0$ halo to its main progenitor across simulation snapshots and thus provides a stable and physically meaningful identifier for each system.

For this work, we use a COLIBRE run with a comoving box size of 25 cMpc at m5 resolution, corresponding to baryonic and dark matter particle masses of $\approx 2 \times 10^5 M_\odot$ and $\approx 3 \times 10^5 M_\odot$, respectively. The Plummer-equivalent gravitational softening length is 0.35 kpc (physical) for $z < 1.57$.

2.2 Gas discs of dwarf galaxies: sample and definitions

We analyse central haloes, defined as those with no more massive structure identified within their FoF group, with $10.75 \leq \log_{10}(M_{200c}/M_\odot) \leq 11$ at $z = 0$. This selection yields 122 galaxies. Their corresponding stellar masses, $M_\star \sim 10^8$ – $10^9 M_\odot$ (with only a few slight exceptions), place them in the regime of bright dwarf galaxies while ensuring sufficient particle resolution ($> 10^5$ total particles per system, and typically $\gtrsim 10^3$ stellar or gas particles). No pre-selection is applied based on gas content, morphology, or dynamical state to avoid biasing correlations with the (dis)equilibrium state of the gas.

Our sample spans diverse morphologies and gas fractions, with neutral hydrogen (HI) masses ranging from 10^7 to $10^9 M_\odot$ and HI fractions $f_{\text{HI}} = M_{\text{HI}}/M_\star$ between (roughly) 0.1 and 10. This diversity is essential for testing the robustness of dynamical modelling assumptions across different galaxy types.

2.2.1 Disc extent

The radial extent of each gas disc is defined using the cumulative HI mass profile,³ $M_{\text{HI}}(< r)$. We identify the radius $R_{\text{ext, HI}}$

²The subscript 200c indicates a quantity evaluated within a spherical aperture with a mean enclosed density of 200 times the critical density for closure, $\rho_c = 3H^2/8\pi G$.

³While not unique to COLIBRE, this approach is greatly facilitated by the self-consistent treatment of HI fractions directly within the simulation outputs, avoiding the need for post-processing schemes commonly required in other simulations (see e.g. the discussion in section 3.3 of K. A. Oman et al. 2019).

at which this profile flattens, i.e. where the logarithmic slope satisfies

$$\frac{d \log_{10}[M_{\text{HI}}(< r)/M_\odot]}{dr} = \lambda,$$

with $\lambda = 0.01 \text{ dex kpc}^{-1}$ chosen empirically to ensure robustness across the sample. For a rotationally supported disc this is usually similar to the radius enclosing 90 per cent of the HI mass in a galaxy ($R_{90, \text{HI}}$, see e.g. E. R. Downing & K. A. Oman 2023; I. S. Sands et al. 2024); but unlike an enclosed mass threshold, it is a more reliable tracer of the outer edge of a galaxy’s gas in disturbed or irregular morphologies.

2.2.2 Warped disc frame

To accommodate warped gas morphologies, we construct a coordinate system locally aligned with the radially varying angular momentum vector of the gas:

$$\hat{L}_{\text{gas}}(r) = \frac{\sum_i m_i \mathbf{x}_i \times \mathbf{v}_i}{\left| \sum_i m_i \mathbf{x}_i \times \mathbf{v}_i \right|},$$

where the sums run over gas particles in a spherical shell at radius r . Here, m_i is the mass of particle i , \mathbf{x}_i is its position vector relative to the coordinate origin, and \mathbf{v}_i is its peculiar velocity vector in the global frame. The profiles are evaluated in radial bins of fixed width $\Delta r = 0.2 \text{ kpc}$, which are identical to those used for the cumulative HI mass profiles described above. Interpolating $\hat{L}_{\text{gas}}(r)$ as a continuous function of r – in this work achieved via a broken third-order polynomial, ‘cubic spline’ – yields a smoothly varying local z -axis that defines the warped disc mid-plane.

The coordinate origin is fixed at the potential minimum as identified by HBT-HERONS. The velocity centroid \mathbf{v}_{cen} is computed as the mass-weighted average velocity of dark matter particles within a sphere of radius $r_{\text{vel cen}} = 2 r_{\text{conv}}$, where r_{conv} is the convergence radius following C. Power et al. (2003). This choice ensures the frame is centred on a physically meaningful dynamically stable origin and velocity frame, minimizing artificial harmonic modes (e.g. spurious $m = 1$ modes) in the potential and velocity fields caused by miscentring.

Sampling coordinates in this warped frame are

$$\{\mathbf{x}\}_{\text{local}} = \{(r_i, \phi_j, 0)\}_{\text{local}},$$

where the local spherical coordinates are radius r_i and azimuthal angle ϕ_j – we note that setting $\theta_{\text{local}} = 0$ for all sampling points is equivalent to mid-plane sampling, and that in this special case the spherical and cylindrical (r, ϕ) sets coincide. Specifically, r_i ranges from 0.1 kpc to $R_{\text{ext, HI}}$ in steps of $\Delta r = 0.2 \text{ kpc}$, and ϕ_j spans 0° to 358° in $\Delta\phi = 2^\circ$ increments. The azimuthal direction ϕ_{local} is defined such that the local tangential velocity satisfies $\langle v_\phi \rangle_\phi(r_i) \geq 0$. We use the symbol $\langle A \rangle_\phi$ to denote the azimuthal average of any quantity A .

2.3 Dynamics of gas discs: fluid models and circular velocity reconstruction

As introduced in Section 1, this work aims to critically assess the validity of commonly adopted assumptions when modelling the kinematics of galactic gas discs.

The evolution of the velocity field in a collisional inviscid fluid system – such as galactic gas – is governed by Euler’s acceleration equation:

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} + \frac{\nabla P_{\text{eff}}}{\rho} = -\nabla \Phi, \quad (1)$$

where \mathbf{v} is the fluid velocity, ρ the density, Φ the gravitational potential, and P_{eff} the effective pressure. The latter includes contributions from both the thermal pressure, $P_{\text{th}} = (\gamma - 1)u\rho$, where u is the internal energy per unit mass and γ the ratio of specific heats, and the turbulent pressure at the resolution limit, $P_{\text{turb}} = \rho\sigma_v^2/3 = \rho\sigma_{v,1D}^2$, where σ_v is the weighted velocity dispersion inside the SPH kernel. Equation (1) can be viewed as a Jeans-equation analogue for gaseous systems, relating observable gas kinematics to the underlying gravitational potential generated by baryonic and dark matter, and is always a statement of local force balance, regardless of whether the system is in equilibrium.

In observational analyses, the partial time-derivative term $\partial_t \mathbf{v}$ is often neglected under the assumption of quasi-steady dynamics. This omission is not justified by observational limitations (i.e. the fact that galaxies are observed at a single epoch, rendering the term unmeasurable) but rather by the implicit physical assertion that the system's time-scales of intrinsic velocity evolution are long compared to dynamical time-scales. When this assumption breaks down, the equilibrium approximation fails, and $\partial_t \mathbf{v}$ becomes physically significant. Measuring its contribution therefore provides a direct probe of the degree of disequilibrium in the gas, and quantifies the limitations inherent in steady-state kinematic and mass modelling pipelines.

Under the steady-state limit ($\partial_t \mathbf{v} \approx 0$), the radial component of equation (1) defines what we refer to as the ‘equilibrium model’⁴:

$$[(\mathbf{v} \cdot \nabla)\mathbf{v}]_r + \frac{1}{\rho} \frac{\partial P_{\text{eff}}}{\partial r} = -\frac{\partial \Phi}{\partial r} \equiv g_r. \quad (2)$$

Further simplifications arise by assuming an idealized axisymmetric disc with negligible radial and vertical motions, and where the pressure gradient is either small or corrected separately (e.g. via ‘asymmetric drift corrections’; see O. Valenzuela et al. 2007, appendix A, for a helpful discussion of this technique and related terminology). Under these conditions, the radial convective term reduces to the centrifugal acceleration, yielding the familiar relation

$$v_\phi = \sqrt{r \left| \frac{\partial \Phi}{\partial r} \right|} \equiv v_c, \quad (3)$$

where v_ϕ is the local azimuthal (rotational) velocity and v_c the corresponding local circular velocity.⁵

The evaluation of equations (1)–(3) requires a change of formalism for the hydrodynamics: in a particle-based (Lagrangian)

⁴In this framework, any residual difference between the left- and right-hand sides of equation (2) directly quantifies the neglected time-dependent term in equation (1):

$$\partial_t v_r = - \left\{ [(\mathbf{v} \cdot \nabla)\mathbf{v}]_r + \frac{1}{\rho} \frac{\partial P_{\text{eff}}}{\partial r} - g_r \right\},$$

which represents the local radial acceleration of the fluid at fixed spatial coordinates, and thus provides a physically meaningful measure of departures from equilibrium.

⁵The commonly used expression $v_c(r) = \sqrt{GM(< r)/r}$ only holds under spherical symmetry. In general, the fundamental dynamical variable is the radial gravitational acceleration $g_r(r, \phi)$, which is well defined for any potential, even in systems that do not admit closed circular orbits. The ‘circular velocity’ used throughout this work therefore refers to the local force-defined quantity $v_c(r, \phi) \equiv \sqrt{r |g_r(r, \phi)|}$, representing the centrifugal speed required to balance the instantaneous radial gravitational force in the disc mid-plane, and can vary as a function of ϕ .

representation, the intrinsic acceleration $\partial_t \mathbf{v}$ and the convective acceleration $(\mathbf{v} \cdot \nabla)\mathbf{v}$ are coupled within the total (material) derivative $d\mathbf{v}/dt$. In the following section (Section 2.4), we describe the SPH-based continuous-field reconstruction developed to evaluate all terms in equation (1) independently in Eulerian fashion, while ensuring full consistency with the simulation’s hydrodynamics conventions.

2.4 SPH-based continuous field sampling

Hydrodynamical quantities and their derivatives at warped disc grid points are computed using a post-processing scheme consistent with the SPHENIX SPH formalism, using the same smoothing kernel and neighbour criteria as in the COLIBRE simulation (see Section 2.1). Formally, our procedure involves the following:

(i) Smoothing length estimation: At each sampling point \mathbf{x} , we solve for the smoothing length $h(\mathbf{x})$ via the implicit SPH number density equation

$$n(\mathbf{x}) = \sum_i W(\mathbf{x} - \mathbf{x}_i, h(\mathbf{x})) = \left(\frac{\eta}{h(\mathbf{x})} \right)^3,$$

using Newton–Raphson iteration to relative tolerance 10^{-4} . Here, W is the smoothing kernel, η is a dimensionless kernel parameter fixing the weighted neighbour number (defined in Section 2.1), and the sum goes over all the particle neighbours of the point, i , with position vectors \mathbf{x}_i .

(ii) Density at sampling points and interpolation: The density at an arbitrary sampling point \mathbf{x} is computed via the traditional SPH equation,

$$\rho(\mathbf{x}) = \sum_i m_i W(\mathbf{x} - \mathbf{x}_i, h(\mathbf{x})),$$

where m_i is the mass of the neighbour particle i and $h(\mathbf{x})$ is the smoothing length evaluated at \mathbf{x} . Scalar, vector, and tensor fields are then interpolated at sampling points \mathbf{x} via

$$A^a(\mathbf{x}) = \frac{\sum_i m_i A_i^a W(\mathbf{x} - \mathbf{x}_i, h(\mathbf{x}))}{\rho(\mathbf{x})},$$

where A_i^a is the field value of particle i . The notation A^a indicates any component of an n -dimensional field.

(iii) First-order derivative operator and evaluation: The interpolation scheme described above allows a self-consistent first-order derivative operator for any n -dimensional field A . This operator is given by

$$\partial_b A^a(\mathbf{x}) = \frac{\sum_i m_i (A_i^a - A^a(\mathbf{x})) \partial_b W(\mathbf{x} - \mathbf{x}_i, h(\mathbf{x}))}{\rho(\mathbf{x})},$$

where ∂_b denotes the partial derivative of a quantity with respect to the b component of the sampling position vector, i.e. $\partial/\partial \mathbf{x}^b$.

Gravitational acceleration is computed as

$$\mathbf{g}(\mathbf{x}) \equiv -\nabla \Phi(\mathbf{x}) = \sum_j \frac{G m_j (\mathbf{x}_j - \mathbf{x})}{\left(\sqrt{|\mathbf{x}_j - \mathbf{x}|^2 + \epsilon_{\text{soft}}^2} \right)^3},$$

i.e. via direct N -body summation over all particles within R_{200c} (both baryonic and dark, indexed by j), employing a softened Newtonian potential consistent with the COLIBRE redshift-dependent softening length, ϵ_{soft} , and the adopted gravitational constant, G . Since we focus our analysis on the scale of the disc ($\approx 0.2R_{200c}$) this truncation of the particle distribution when evaluating $\mathbf{g}(\mathbf{x})$ is an acceptable approximation.

All calculations are performed in the global coordinate frame and projected into the warped local frame only at the final step, ensuring that radial, azimuthal, and vertical components of fields and their derivatives remain physically meaningful. For illustrative purposes, Appendix A presents a representative example from a galaxy in our sample, illustrating the warped disc frame construction alongside the output of our SPH-based continuous field sampling procedure. These visualizations provide context for the numerical procedures outlined here and in Section 2.2.

This continuous-field reconstruction extends particle-based analyses into smooth spatially resolved descriptions of the gas dynamics. Technically, it corresponds to a Shepard (D. Shepard 1968) SPH-gather scheme employing the same kernel definitions and neighbour criteria as the simulation itself, thereby ensuring consistency with its hydrodynamical framework. An important advantage of this scheme is that each sampling point attains an effective spatial resolution determined self-consistently by the local fluid density. This guarantees that reconstructed quantities vary smoothly across the disc, mitigating the shot noise typical of scatter-type interpolations, where individual particle contributions dominate the local estimate. The Shepard normalization further enforces that the effective kernel sum satisfies $\sum_i \tilde{W} = 1$, eliminating one of the main sources of bias in standard SPH interpolation – deviations from unity kernel normalization, which can otherwise lead to systematic underestimation or overestimation of field amplitudes in sparsely sampled regions.

A caveat of standard Shepard formalisms is that they are not strictly conservative in mass or momentum (S. Reinhardt et al. 2019). However, since the sampling points are not interpreted as finite-volume fluid cells, but as infinitesimal tracers defining a continuous warped mid-plane, the lack of strict conservation poses no conceptual inconsistency for the subsequent equilibrium analysis. Our reconstruction does not attempt to reproduce the exact numerical forces used during the simulation’s time integration, but rather to recover the instantaneous continuous field structure implied by the discrete SPH representation.

When evaluating the hydrodynamical equations (1)–(2), we adopt several deliberate departures from the native SPH formulation used during the simulation. In particular, we exclude (i) artificial viscosity terms, which are numerical constructs introduced for entropy dissipation and shock capturing (J. Borrow et al. 2022), and (ii) the contribution from the large-scale gravitational field of the full simulation box. Our focus is instead on the physical force balance of the gas under the assumption of an inviscid collisional fluid.

Artificial viscosity in SPH serves a necessary numerical role but introduces several well-known side effects: (i) spurious shear stresses, (ii) unphysical angular-momentum transport, and (iii) damping of turbulent motions at the resolution scale. In our reconstruction, the latter effect is at least partially accounted for by the inclusion of the turbulent pressure term, $P_{\text{turb}} = \rho \sigma_{v,1D}^2$, which captures the local subgrid velocity dispersion and its contribution to pressure support. Consequently, neglecting the explicit artificial-viscosity term does not significantly bias the inferred quasi-static balance of forces, since its dominant dynamical impact – turbulent damping – is represented implicitly through P_{turb} .

Framing the reconstruction in this way ensures that the recovered fields address the fundamental equilibrium problem central to rotation-curve analyses: whether pressure gradients and convective accelerations balance the internal gravitational potential, and where astrophysical processes such as feedback or accretion drive measurable departures from this balance.

2.5 Feedback tracing and filtering

Since stellar and AGN feedback are primary candidates for driving disequilibrium in the ISM, we construct a filtering scheme to identify regions of the disc that have been recently affected by these processes. This enables a direct spatial and temporal association between localized feedback sites and the gas dynamics in those sites.

For each warped mid-plane sampling point, we select gas and stellar particles within a spherical distance $r \leq 1$ kpc. This scale encompasses the local environment dynamically sensitive to recent feedback while minimizing contamination from unrelated structures. A subsequent temporal filter retains only:

- (i) gas particles flagged as having experienced feedback from an AGN or from a supernova (CCSN/SNIa) within a lookback time $t_{\text{fb}} \leq t_{\text{cutoff}}$; and
- (ii) young stars with age $\tau_{\star} \leq t_{\text{cutoff}}$.

We test multiple cutoff times (e.g. $t_{\text{cutoff}} = 10, 20, 50,$ and 100 Myr) to verify that the inferred spatial correlations and disequilibrium metrics are robust to this selection.

Feedback strength is quantified via energy-equivalent proxies:

- (i) Thermal feedback: for purely thermal energy injection (e.g. CCSN or SNIa events), the deposited energy is taken as proportional to the local temperature increase, $\Delta E \propto k_B \Delta T$, where k_B is Boltzmann’s constant.
- (ii) Kinetic feedback: for momentum-driven injection, we compute $\Delta E = \frac{1}{2} m (\Delta v)^2$, where m is the particle mass and Δv the velocity kick imparted at the last feedback event.
- (iii) AGN feedback: the simulation directly records the thermal or kinetic energy received per event, and this quantity is adopted without further conversion.
- (iv) Generic stellar feedback: additional stellar processes such as winds, radiation pressure, and ionizing radiation are traced via the mass of young stars formed. We have checked that the spatial distribution of active HII regions (tracked in the COLIBRE model) closely follows that of recently formed stars, justifying this proxy.

Feedback proxies are accumulated across the warped disc and, where indicated, normalized by the corresponding local disc area to yield surface energy-density measures of feedback impact. This filtering procedure isolates regions where the gas has been recently perturbed by feedback and thus provides a physically motivated framework for linking energetic injection to local deviations from force equilibrium. The motivation for, limitations of, and implications of this analysis are discussed in detail in Section 4.2.

3 DISEQUILIBRIUM IN DWARF GALAXIES

In rotation curve studies, it is standard practice to assume that the azimuthal velocity of gas directly traces the local gravitational potential. While this assumption is often motivated by the apparent rotational support, regularity of emission lines, and morphological coherence of gas discs, such observational signatures alone establish neither dynamical equilibrium, nor purely circular motion.

As outlined in Section 1, a range of effects can influence mass inferences from gas kinematics. Some stem from projection and observational artefacts – such as beam smearing, miscentring, or inclination uncertainties, while others arise from modelling

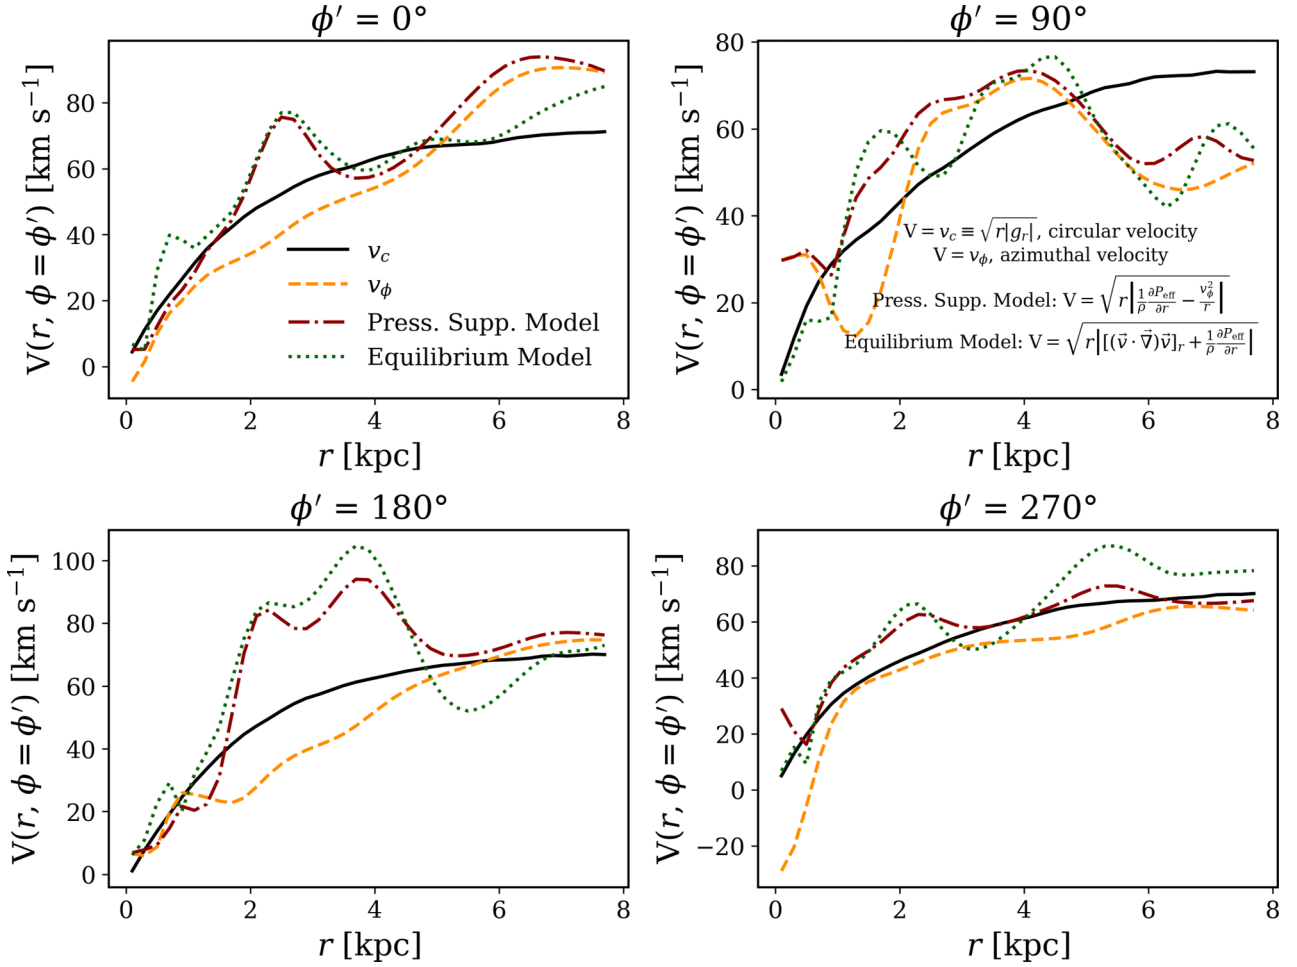


Figure 1. Comparison of different fluid dynamical models in recovering the true circular velocity profile for a representative dwarf galaxy. Each panel shows radial profiles of reconstructed circular velocity along a fixed azimuthal direction ($\phi' = 0, 90, 180, 270^\circ$). The black solid line shows the true v_c from the gravitational potential; lines that are coloured and non-solid show reconstructions from different models (legend). More physically complete models do not always outperform the naïve $v_\phi = v_c$ assumption.

degeneracies, particularly those between geometric assumptions and non-circular motions. Yet, almost all inference methods adopt, either explicitly or implicitly, the premise that the gas is in hydrodynamical equilibrium. This assumption is seldom tested in a rigorous, spatially resolved manner in either observations or simulations.

3.1 Local performance of fluid models

A central aim of this work is to test fluid models in a local physically self-consistent way. We evaluate whether the steady-state Euler equation (equation 2) is satisfied in simulated gas discs, and how well different models recover the true gravitational potential.

Fig. 1 presents a case study of one simulated galaxy (HBTHERONS track ID = 14514, stellar mass $M_\star \approx 2.80 \times 10^8 M_\odot$), showing reconstructed circular velocity profiles along four azimuthal slices. The black solid line is the true circular velocity from the potential (i.e. the local force-defined speed $v_c(r, \phi) = \sqrt{r |g_r(r, \phi)|}$); broken coloured curves correspond to different models as indicated in the legend. Two main points stand out.

First, all reconstructions exhibit significant azimuthal variability at fixed radius, indicating strong model-dependent sensitivity to local dynamics and asymmetries. Such non-axisymmetric

structures are well documented in empirical gas discs’ hydrodynamical fields (e.g. harmonic decomposition of moment maps revealing spiral patterns and/or n-symmetric velocity fields; see R. H. M. Schoenmakers, M. Franx & P. T. de Zeeuw 1997; C. Trachternach et al. 2008; J. A. Sellwood & R. Z. Sánchez 2010). However, the behaviours observed here underscore a deeper implication: azimuthal averaging, which is commonly used in rotation curve analyses (both observational and theoretical), can obscure critical localized discrepancies. These deviations are physical signatures of disequilibrium physics, driven by ongoing astrophysical processes (as we will discuss in Section 4).

Secondly, and perhaps more strikingly, the ‘Equilibrium model’ (dark green/dotted) – despite being the most comprehensive physical model – is frequently outperformed by the simplistic $v_\phi = v_c$ assumption (orange/dashed); see, for example, the increased absolute deviation of the equilibrium reconstruction from the true circular velocity in the $\phi' = 180^\circ$ slice. The simpler ‘Pressure support model’ (brown/dash-dotted), which attempts a correction to v_ϕ under the assumption that hydrostatic pressure gradients are the dominant non-equilibrium effect, fares similarly – in this galaxy this assumption is justified across most of the disc, although this does not hold for all galaxies. This counterintuitive outcome echoes the findings of I. S. Sands et al. (2024) who

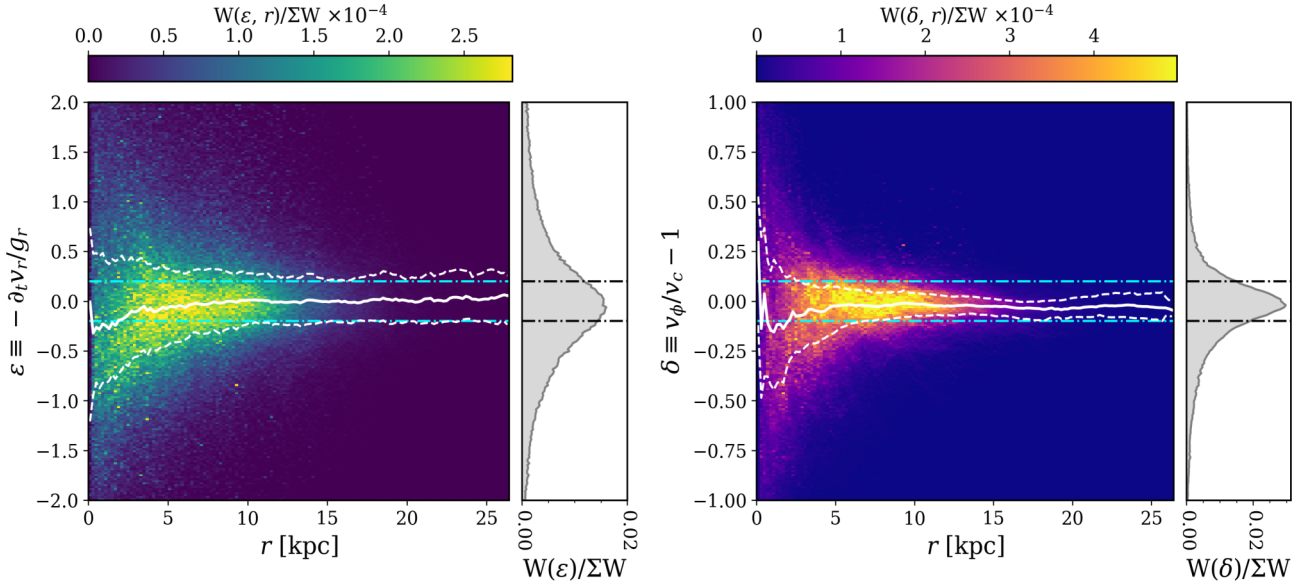


Figure 2. Radial distributions of the $\varepsilon(r, \phi)$ and $\delta(r, \phi)$ parameters (equation 4) weighted by gas mass across the full galaxy sample, respectively, $W(\varepsilon(r, \phi), r)$ and $W(\delta(r, \phi), r)$. Each panel displays both the joint and marginal distributions of the respective parameter. *Left panel:* stacked 2D histogram of ε versus galactocentric radius, r , where the colour scale indicates the joint weighted density, $W(\varepsilon, r)/\Sigma W$, summed across all galaxies. The vertical subpanel to the right shows the normalized marginal (1D) distribution, $W(\varepsilon)/\Sigma W$. *Right panel:* same as left panel, but for the δ parameter. The cyan dash-dotted lines in the main panel and black dash-dotted lines in the side panel denote the parameter range associated with ≤ 10 per cent error in circular velocity inference based on the corresponding dynamical model. In the main subpanel, the white solid and white dashed lines show the median and central 50 per cent envelopes of the respective conditional distributions, $W(\varepsilon|r)$ and $W(\delta|r)$.

studied disequilibrium in gas discs from the FIRE-3 simulation suite (P. F. Hopkins et al. 2023). They argue that the poor performance of the equilibrium model can be traced to the unbalanced nature of Eulerian terms. Specifically, radial components of the convective derivative, $(\mathbf{v} \cdot \nabla)\mathbf{v}$, often arise from the same non-equilibrium processes that generate time variations in the velocity field, $\partial_t \mathbf{v}$. This coupling undermines the validity of equilibrium reconstructions. That we find similar trends in COLIBRE reinforces a key point: adding correction terms for non-circular motions or pressure support does not guarantee more accurate mass inferences. The limiting factor is whether the gas is actually in hydrodynamical equilibrium.

In summary, the standard equality $v_\phi = v_c$ breaks down in a systematic and physically meaningful way. Even in the most complete time-independent fluid models, disequilibrium produces frequent significant errors in circular velocity estimates. This challenges the foundational assumptions of rotation curve analysis and suggests that purely analytic corrections are limited by unobservable time-dependent processes.

3.2 Population-wide radial trends in disequilibrium physics

To test whether the results from the galaxy shown in the previous section are representative of the broader dwarf galaxy population, we extend our analysis to the full simulation sample. We introduce two diagnostic parameters that quantify deviations from hydrodynamical equilibrium and centrifugal balance, respectively. These metrics also correspond to the fractional error in inferred circular velocity when adopting either the full equilibrium model (equation 2) or the naïve assumption that the rotational velocity exactly traces the mass distribution (equation 3).

The choice to express both discrepancies in fractional terms arises from the large variety of galaxies in the sample, which span a wide range of circular velocity profiles (i.e. rising at different rates and up to slightly different V_{\max}). Hence, we define

$$\varepsilon \equiv \left(\frac{[(\mathbf{v} \cdot \nabla)\mathbf{v}]_r + \frac{1}{\rho} \frac{\partial P_{\text{eff}}}{\partial r}}{g_r} \right) - 1 = -\frac{\partial_t v_r}{g_r}. \quad (4a)$$

$$\delta \equiv \frac{v_\phi}{v_c} - 1. \quad (4b)$$

While δ has a familiar physical interpretation – indicating the extent to which rotation supports a gas parcel against gravity – ε is better interpreted as a diagnostic of net radial imbalance. It compares the total time-independent accelerations experienced by a gas parcel to the local gravitational acceleration. As such, $\varepsilon \neq 0$ implies that the local fluid is undersupported ($\varepsilon < 0$) or oversupported ($\varepsilon > 0$), and would develop a net radial inflow or outflow if evolved in an Eulerian framework (assuming an initial $v_r \approx 0$).

Fig. 2 shows the radial distribution of both diagnostic parameters across the galaxy sample. The left panel corresponds to the mass-weighted density of $\varepsilon(r, \phi)$, i.e. $W(\varepsilon, r)/\Sigma W$, while the right panel shows the same for $\delta(r, \phi)$. In each case, the main subpanel presents a stacked 2D histogram of parameter values as a function of galactocentric radius, with colour indicating the joint distribution normalized by the total counts. Each galaxy here contributes 180 evenly spaced azimuthal measurements per radial bin up to the radius marking the extent of its gaseous disc – $R_{\text{ext, HI}}$, as defined in Section 2.2, enabling robust population-level statistics. The vertical subpanels display the normalized 1D distribution of the corresponding parameter. The cyan (main panels) and black (side panels) dash-dotted lines indicate the range of parameter values associated with ≤ 10 per cent error in circular velocity inference under the relevant fluid model (i.e. $|\varepsilon| \lesssim 0.2$,

$|\delta| \leq 0.1$). In the main subpanel, the white solid and white dashed lines indicate the median and central 50th percentile of the respective conditional distributions, $W(\varepsilon|r)$ and $W(\delta|r)$. Several results emerge.

First, both the ε and δ exhibit broad distributions at all radii, demonstrating that disequilibrium is a pervasive feature of dwarf galaxy gas discs. Crucially, this is not a subtle effect – typical deviations of the order of 10 per cent from radial force balance and centrifugal support are common, with a substantial fraction of measurements showing even larger discrepancies. This is more clearly seen in the marginalized 1D distributions shown in the side panels, where the parameter ranges corresponding to ≤ 10 per cent circular velocity inference error approximately coincide with the full width at half-maximum of the distribution functions. This overlap indicates that a significant portion of the gas disc lies outside the regime where circular velocity estimates are accurate to better than 10 per cent.

As evident in the main panels, the vast majority of the samples contributing to these broad distributions originate from the inner disc regions ($r \lesssim 5$ kpc). This finding runs counter to expectations based on dynamical time-scales, which predict more rapid relaxation in the central regions; instead, we find clear evidence for persistent perturbations. This is particularly relevant for core-cusp studies, since this radial range provides the strongest leverage for constraining dark matter density profiles. This is not the only hint that the centres of dwarf galaxies are continuously perturbed out of equilibrium. In fact, a recent investigation of radial trends in the radial acceleration scaling relation (S. S. McGaugh, F. Lelli & J. M. Schombert 2016) using the SPARC sample (F. Lelli et al. 2016) shows that the inner regions – defined as $r < 2R_d$, where R_d is the stellar disc scale radius – are much less tight (in terms of intrinsic scatter) than the outer disc (see fig. S1 in T. Ren et al. 2019). This trend can be interpreted as empirical evidence that inner regions of galaxies are indeed the most dynamically disturbed.

Secondly, the equilibrium-based model (quantified by ε) does not consistently outperform the naïve centrifugal model (quantified by δ). Across the population, many regions – particularly in the inner disc – fall outside the parameter space where equilibrium-based corrections yield more accurate circular velocities than the azimuthal velocity, reinforcing the discussion in Section 3.1. This is reflected in Fig. 2, where the mass-weighted peak of the ε distribution reaches a value $\max(W(\varepsilon)/\Sigma W) \approx 0.015$, compared to the peak of the δ distribution $\max(W(\delta)/\Sigma W) \approx 0.030$.

Moreover, the δ parameter reveals a population-wide tendency to underestimate the true circular velocity in the inner disc ($\delta < 0$), implying a systematic bias towards inferring smaller inner velocities, hence ‘cored’ profiles. As the white solid and dashed lines indicate, this bias arises preferentially within $r \lesssim 5$ kpc, where the median radial trend of δ departs from zero and the conditional distributions become increasingly skewed towards negative values. Although the median fractional offsets remain relatively modest – of order 10 per cent at $r \approx 2$ kpc – the growing asymmetry of the distributions implies that a substantial fraction of regions systematically underestimate the true circular velocity; crucially, this behaviour is mirrored, not mitigated, when adopting the equilibrium-based model. However, the δ distribution is broad in the inner kiloparsecs, with a minority of regions exhibiting the opposite trend ($\delta > 0$). This bidirectional scatter suggests that some of the observed diversity in dwarf galaxy rotation curves could arise from disequilibrium physics

rather than intrinsic variations in the underlying dark matter distributions. In other words, the rotational structure in the inner regions of gas discs is a biased and high-variance tracer of the underlying gravitational potential even before projection effects, and the limited ability of kinematic models to accurately recover azimuthal speeds, come into play.

Together, these results point to a deeper limitation of analytic fluid models in the context of mass inference for dwarf galaxies. While efforts to correct for non-circular motions or pressure support – such as asymmetric drift prescriptions – have become increasingly sophisticated, our analysis suggests that these models are fundamentally constrained by the intrinsic disequilibrium of the gas. The inner kiloparsecs – the most critical regime for dark matter constraints – are also the least likely to satisfy steady-state assumptions.

If gas kinematics in dwarf galaxies is not a reliable tracer of the gravitational potential, mismatches between observed rotation curves and Λ CDM predictions may reflect the breakdown of the equilibrium assumption, not a failure of the cosmological model itself. This can reframe the long-standing core-cusp and diversity of rotation curves tensions into a mismatch between observed quantities and the theoretical constructs that they are used to infer.

4 ORIGINS OF DISEQUILIBRIUM

Section 3 showed that hydrodynamical disequilibrium is a widespread and persistent feature of the gas discs in our low-mass galaxy sample, spanning a diverse set of simulated systems. The amplitude and ubiquity of these deviations from steady-state dynamics have immediate implications for dynamical mass inference. They challenge the robustness of classical rotation curve analyses and offer a plausible route to explain the diversity of dwarf galaxy rotation curves.

This naturally prompts a central question: ‘If disequilibrium is so prevalent and has such strong implications, what causes it?’ Implicit in this is a methodological concern: are these departures genuine astrophysical phenomena, driven by internal feedback, gravitational instabilities, etc., or are they artefacts of numerical limitations or rare stochastic events?

In this section we dissect the physical origins of disequilibrium in our simulations, using spatially resolved diagnostics to connect local equilibrium violations with specific dynamical drivers. Our aim is to separate plausible physical mechanisms from numerical or incidental causes, while quantifying and characterizing their respective dynamical signatures.

4.1 Disequilibrium drivers

We consider four broad classes of mechanisms that are expected to disrupt Eulerian steady-state velocity fields: (i) impulsive momentum and energy injection from stellar feedback; (ii) over-dense, self-gravitating clumps; (iii) energy injection from AGN activity; and (iv) disc-scale gravitational asymmetries.

We stress that this taxonomy is not derived from prior studies of hydrodynamical disequilibrium per se, but rather from physical intuition and analytic expectations about the stability of galactic discs: any process capable of driving non-circular motions, injecting turbulence, or introducing rapidly evolving unrelaxed components should, in principle, generate measurable deviations from equilibrium. While the first three processes are often causally and temporally linked, e.g. feedback can seed clump formation, which in turn can trigger gas inflow, which then may feed the AGN, they

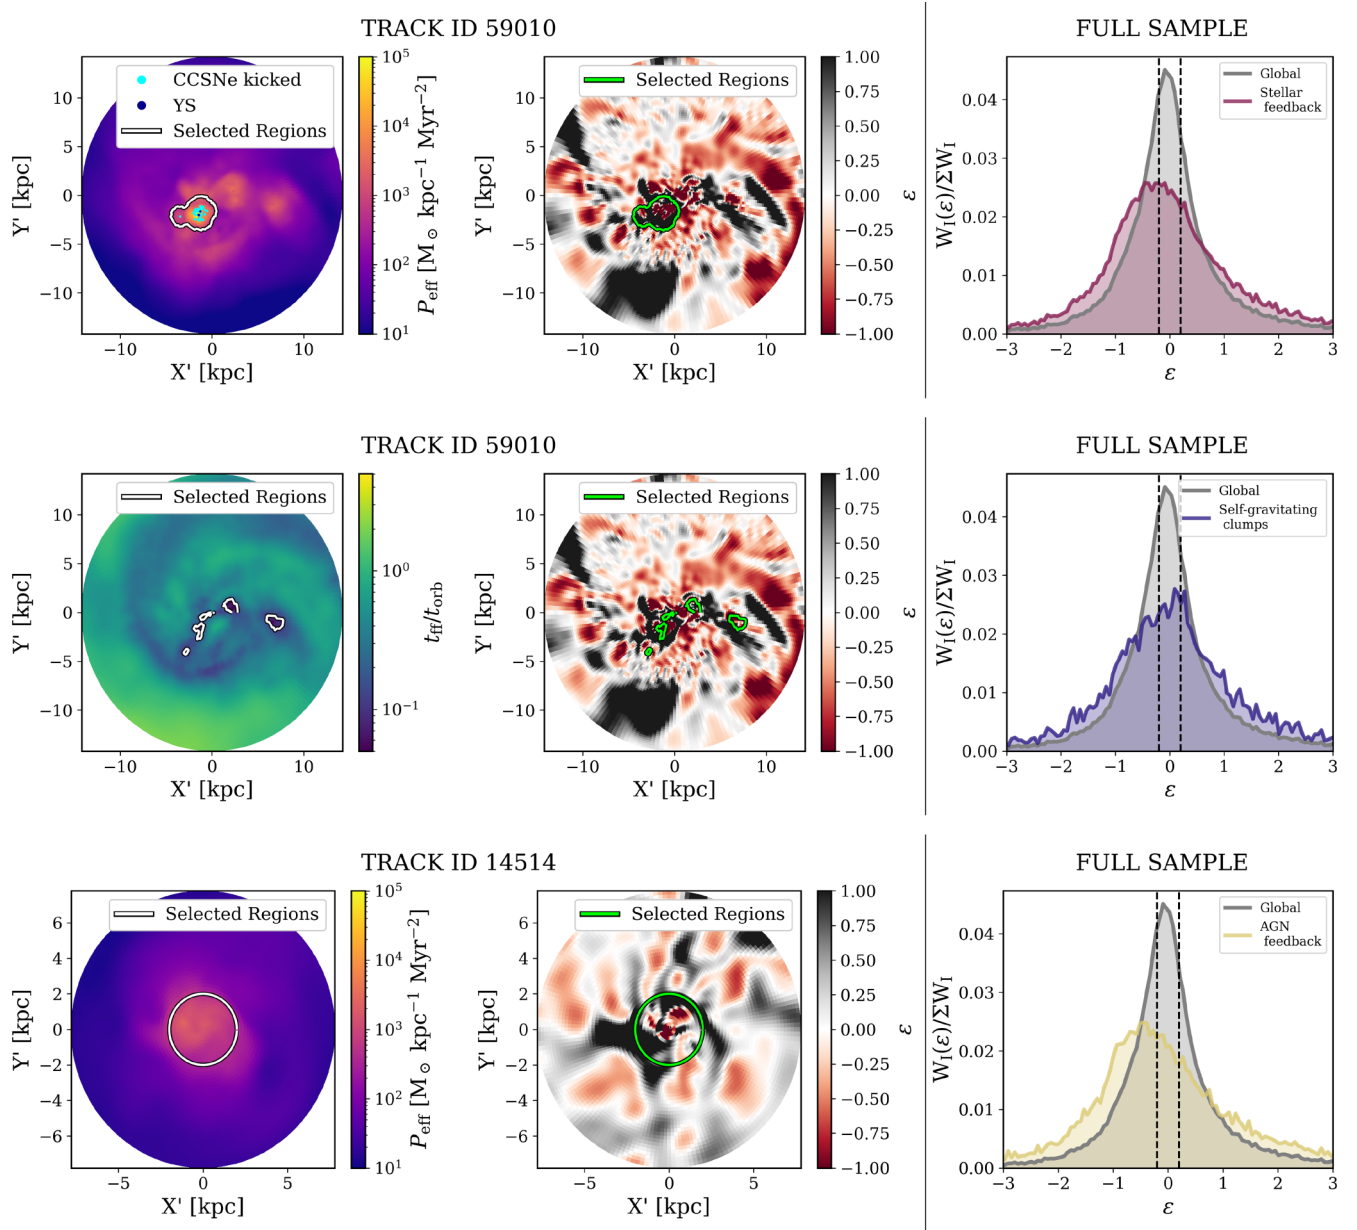


Figure 3. A visual summary of three different mechanisms inducing local disequilibrium in the ISM considered in this study, along with their associated tracers. The figure is organized in rows, each corresponding to one disequilibrium-driving mechanism: stellar feedback (top), self-gravitating gas clumps (middle), and AGN activity (bottom). For the first two rows (stellar feedback and clumps), the panels in the left and middle columns show the same representative galaxy, as indicated by the annotated track IDs, whereas the bottom row illustrates a different system; the right column in each row always shows the corresponding sample-wide ε distribution. *Left column:* 2D Cartesian maps of a representative deprojected warped disc mid-plane (X' , Y'), colour-coded by the associated tracer field – effective pressure for both stellar and AGN feedback, and the free-fall to orbital time ratio for self-gravitating clumps. The blue and cyan dots in the upper left panel indicate the positions of young stars (ages $\tau_* \leq 100$ Myr) and gas particles recently affected by CCSNe ($t_{\text{fb}} \leq 50$ Myr), respectively. The solid white contours enclose regions exceeding tracer thresholds, or satisfying proximity criteria. *Middle column:* the same maps, colour-coded by the disequilibrium parameter ε , with the same contours as in the left panels (now in lime) overlaid to highlight the spatial correlation between tracers and disequilibrium. *Right column:* normalized 1D distributions of ε , weighted by gas mass, computed across the full dwarf galaxy sample for each process (coloured), compared to the global distribution (grey). The vertical dashed lines indicate the quasi-equilibrium range defined by $|\varepsilon| \lesssim 0.20$. The spatial alignment of tracer-selected regions with extreme ε values, together with their skewed and/or broadened ε distributions relative to the global one, supports their role as drivers of hydrodynamical disequilibrium.

are none the less localized in space and time, and can therefore be identified through spatial tracer fields. By contrast, gravitational asymmetries are inherently non-local in their dynamical influence and are treated separately in Section 4.1.4. We also stress that this is not a comprehensive list of all possible drivers, but a selection of the most important (as we will discuss in Section 4.2).

Fig. 3 provides a compact summary of the spatial and statistical signatures of the first three mechanisms. Each row corresponds to one driver and one representative galaxy (unless otherwise noted), with the first column showing a tracer-field map of a representative warped and deprojected disc mid-plane, (X' , Y'); second column the corresponding disequilibrium parameter (ε)

map; and the third column showing the gas-mass-weighted ε distribution for tracer-selected regions (coloured), compared to the global distribution (grey), both normalized by their total gas-mass-weight and computed across the full dwarf galaxy sample. The tracer definitions are rooted in theoretical expectations: feedback regions are identified by effective (both thermal and turbulent) overpressure near young stars and recent CCSNe, or around an active AGN; self-gravitating clumps are traced by regions with short free-fall times relative to the local orbital period. In all three cases, the white (or lime) contours delineate regions exceeding tracer thresholds for overdense gas clumps, or proximity criteria for feedback. The visual and statistical correspondence between tracer-selected zones and high- $|\varepsilon|$ gas suggests that these mechanisms imprint distinct separable signatures on the ISM kinematics. We examine each in turn, beginning with stellar feedback.

4.1.1 Stellar feedback

Stellar feedback is a well-motivated source of ISM disequilibrium, particularly in shallow gravitational potentials where momentum and energy injection from massive stars can rival the disc binding energy (e.g. A. Dekel & J. Silk 1986; G. Stinson et al. 2006; P. F. Hopkins, E. Quataert & N. Murray 2012; C.-G. Kim & E. C. Ostriker 2017). Energy deposition from SNe, stellar winds, and radiation pressure drives expanding shells, turbulence, and large-scale pressure gradients, each of which can displace gas from local force balance for time-scales exceeding the local dynamical time (D. Martizzi, C.-A. Faucher-Giguère & E. Quataert 2015; M. E. Orr et al. 2020). Observations support this picture: nearby dwarfs often exhibit asymmetric H I profiles (e.g. F. Walter et al. 2008; A. M. Stilp et al. 2013) and vertically extended warm gas layers (R. Boomsma et al. 2008; C. Bacchini et al. 2020), with significant dispersive motions consistent with feedback-driven disturbances. A striking observational example of extreme disruption is the supergiant H I shell, or ‘superbubble’, in IC2574 (E. Brinks, F. Walter & J. Kerp 2003), which also coincides with one of the most pronounced apparent dark matter cores known in a dwarf galaxy.

In our analysis, feedback-affected gas is defined as material lying within a projected radius of 1 kpc in the warped disc mid-plane around either young stars (ages $\tau_* \lesssim 100$ Myr) or gas particles that have recently received kinetic injection from CCSNe ($t_{\text{fb}} \leq 50$ Myr). Particle positions are first assigned to their nearest point on the deprojected disc surface, and the selection is then made using a two-dimensional circular mask of radius 1 kpc (see Section 2.5 for details). This choice reflects both the geometry of the discs and the physical dispersion of feedback-driven material: although CCSNe impart an isotropic target kick of $v_{\text{kick}} = 50 \text{ km s}^{-1}$ in the rest frame of the star particle (E. Chaikin et al. 2023; J. Schaye et al. 2025), the residual velocity of affected gas in the rotating disc frame is typically smaller ($\sim 20 \text{ km s}^{-1}$), implying characteristic displacements of ~ 1 kpc over ~ 50 Myr (ignoring other forces). Furthermore, H I-dominated dwarf galaxy discs have expected vertical scale heights of ~ 1 kpc (A. Banerjee et al. 2011), making this spatial cut a physically motivated and conservative choice for identifying feedback-influenced regions in a population-wide analysis.

We focus on kinetically kicked gas because thermally heated ejecta leave the disc too rapidly to provide meaningful information: once heated, this material rises rapidly out of the mid-plane (typically to several kiloparsecs) and ceases to intersect the regions where we evaluate ε . It therefore leaves no coherent or spatially localizable imprint within the disc, whereas the kinetic

CCSNe mode produces gas that remains dynamically coupled to the disc for tens of megayears and can be robustly identified in our spatially resolved population-wide analysis.

The top row of Fig. 3 shows, via a representative system (HBTHERONS track ID = 59010, stellar mass $M_* \approx 1.00 \times 10^8 M_\odot$, star formation rate $\text{SFR} \approx 0.013 M_\odot \text{yr}^{-1}$), that stellar feedback-selected regions (white contours, left panel) coincide with zones of enhanced effective pressure, often exceeding the disc median by factors of several and frequently associated with coherent superbubble structures. The corresponding ε map (middle panel) indicates that these regions host some of the strongest disequilibrium in the disc, with $|\varepsilon|$ commonly exceeding unity.⁶

Statistically (right panel), the stellar feedback-affected gas exhibits a strongly skewed ε distribution with a pronounced tail towards large negative values, indicating a time-dependent Eulerian radially inward acceleration that offsets the excess outward force imparted by feedback. Physical intuition for this behaviour can be developed by first considering the rest frame of the young star launching the feedback. In this frame, a supernova injects an approximately isotropic impulse, producing a transient increase in the local radial velocity that decays as the blast wave evolves (L. I. Sedov 1946; G. Taylor 1950; J. von Neumann & A. H. Taub 1964), i.e. $\partial_t v_r < 0$, where v_r is the radial component measured away from the star. When transformed to the galactocentric frame, one might expect the inward- and outward-directed components of this symmetric impulse to contribute equally to $\partial_t v_r$ at a fixed location in the disc. However, the strong radial and vertical density gradients in galactic gas discs make it easier for feedback to displace gas outwards or out of the plane of the disc than inwards. As a result, the Eulerian $\partial_t v_r$ in the galactocentric frame is systematically biased negative when averaged across galaxies and mass-weighted, despite the underlying isotropy of the impulse.

Moreover, the stellar feedback-selected distribution exhibits a steep rise towards the peak starting at $\varepsilon \sim -1$, which aligns with a secondary bump in the global (full sample, disc-wide) distribution, suggesting that stellar feedback makes a large fractional contribution to the mass of gas in strong disequilibrium. A quantitative breakdown of the (dis-)equilibrium budget, together with a discussion of the limitations of our tracer-based method, is deferred to Section 4.2.

Elevated ε is not confined to stellar feedback-affected gas. Even in actively star-forming systems – such as the representative galaxy in the top row of Fig. 3 – substantial disequilibrium is present well outside the feedback contours, indicating the action of additional drivers. In the following subsections, we examine the role of self-gravitating clumps, AGN activity, and large-scale gravitational asymmetries in seeding and sustaining non-equilibrium dynamics.

4.1.2 Overdense and self-gravitating gas clumps

Local gravitational collapse within discs provides a natural route to hydrodynamical disequilibrium. Dense self-gravitating clumps – often arising from gravitational instability in the cold ISM (e.g.

⁶Since ε measures the ratio of the neglected time-dependent acceleration to the radial gravitational force, $\varepsilon = -\partial_t v_r / g_r$, values with $|\varepsilon| > 1$ imply that the intrinsic local velocity-evolution time-scale is comparable to or shorter than the dynamical time-scale set by gravity. In such regimes the dynamics cannot be treated as small perturbations about equilibrium, rather they signify strong departures from the assumptions underlying standard disc models.

M. Noguchi 1999; A. Dekel, R. Sari & D. Ceverino 2009; R. Genzel et al. 2011; N. J. Goldbaum, M. R. Krumholz & J. C. Forbes 2015) – develop internal dynamics that strongly deviate from the background shear flow.

We identify these regions using a time-scale diagnostic: the ratio of the local free-fall time to the orbital time, $t_{\text{ff}}/t_{\text{orb}}$. When this ratio falls below ~ 0.1 (a calibration value discussed in Section 4.2), collapse proceeds much more rapidly than the local orbital response time, signaling that gas self-gravity dominates over the ambient potential. This criterion, by definition, isolates dense knots and filaments whose dynamics are internally driven, rather than passively tracing the global disc rotation.

The middle row of Fig. 3 illustrates their spatial and dynamical imprint. In the left panel, the clump tracer field delineates spatially localized knots and filaments, frequently embedded within spiral arms or ring-like overdensities. The corresponding ε map (middle) reveals that these same sites are often among the most extreme disequilibrium regions, with ε commonly positive and exceeding unity. The right panel quantifies this: the clump-tagged distribution is markedly broader than the global disc-wide one, with a mild shift of the peak towards $\varepsilon \gtrsim 0$ and extended wings on both sides, reflecting a mix of accelerated collapse ($\varepsilon < 0$) and migration ($\varepsilon > 0$). This contrasts with the distribution characteristic of feedback zones, typically biased more strongly on the negative side (Section 4.1.1). In idealized scenarios, a clump drives a locally symmetric, converging flow in its own rest frame, such that $\partial_r v_r < 0$, where v_r is the radial component measured away from the clump (e.g. M. V. Penston 1969). In the Eulerian galactocentric frame, because the ambient gas density decreases with radius, flux from the denser inward side exceed those from the outward side, biasing $\partial_r v_r$ towards the positive end. This naturally produces the mild positive shift in the clump-selected distribution. At the same time, many dense clumps either form within, or evolve into, regions where young stars subsequently inject feedback; sampling these mixed environments reintroduces the strongly negative tail characteristic of feedback-driven gas. The joint effect is a broad wing-dominated distribution with a slight offset towards $\varepsilon > 0$, as seen in Fig. 3.

While some clumps form or reside within feedback-active environments, a substantial fraction lie in quiescent or outer disc regions lacking obvious stellar-driven perturbations. This is evident in Fig. 3: because the same representative galaxy appears in the first two rows, the stellar-feedback (top-row) and clump-selected (middle-row) contours can be directly compared. Many overdense regions fall well outside the feedback-affected zones, yet still exhibit extreme $|\varepsilon|$, demonstrating that clump-driven dynamics generate strong disequilibria independently of stellar feedback. Observational analogues – such as the massive clumps in $z \sim 2$ discs (e.g. D. M. Elmegreen et al. 2007; R. Genzel et al. 2011), or in the nearby WLM dwarf galaxy (S. M. Albers et al. 2019; C. Mondal, A. Subramaniam & K. George 2021) – likewise show that such structures can persist without contemporaneous feedback. We therefore treat self-gravitating clumps as a separate dynamically distinct driver of local disequilibrium.

4.1.3 AGN feedback

Energy and momentum injection from accreting massive black holes provide a third physical channel capable of generating local departures from hydrodynamical equilibrium. Although AGNs in dwarf galaxies have historically been difficult to identify –

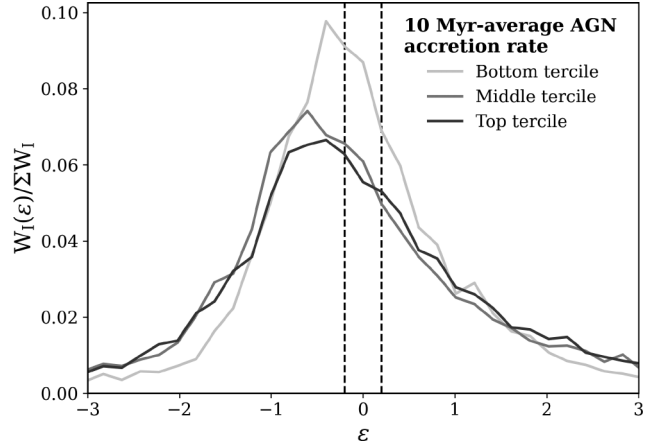


Figure 4. Comparison of the gaseous disequilibrium structure for galaxies binned by their 10 Myr-averaged AGN accretion rate. The three curves show the self-normalized mass-weighted distribution of the disequilibrium parameter, ε , i.e. $W_1(\varepsilon)/\Sigma W_1$, for systems in the bottom (light grey), middle (grey), and top (dark grey) terciles of the accretion-rate distribution. The vertical dashed lines mark $|\varepsilon| = 0.2$, as in Fig. 3. The systematic shift of the distributions – with higher accretion systems favouring slightly larger ε magnitudes – under identical geometric selection (a $r = 2$ kpc mask in the deprojected mid-plane) supports that AGN activity is an independent direct driver of part of the observed non-equilibrium behaviour.

owing to weak emission signatures, uncertain bolometric corrections, and confusion with stellar processes (e.g. A. E. Reines, J. E. Greene & M. Geha 2013; V. F. Baldassare et al. 2017) – both theory and observations now indicate that low-mass systems can indeed host active nuclei, and that even modest AGN episodes can imprint measurable disturbances on the surrounding ISM (e.g. C. M. Manzano-King & G. Canalizo 2020; S. Koudmani et al. 2025) affecting galaxy evolution (e.g. S. Koudmani, D. Sijacki & M. C. Smith 2022; E. Arjona-Gálvez, A. Di Cintio & R. J. J. Grand 2024).

To understand how AGN activity correlates with disequilibrium in the gas surrounding it, we examine the distribution of the disequilibrium parameter within fixed central apertures for galaxies binned by their 10 Myr-averaged AGN accretion rate (Fig. 4). Across the bottom, middle, and top terciles, the self-normalized mass-weighted distributions of ε exhibit a clear monotonic shift: galaxies with more strongly accreting AGNs show a pronounced enhancement in the mass fraction of strongly non-equilibrium gas, with substantially broader high- $|\varepsilon|$ tails. The middle tercile already has much less gas in the low- $|\varepsilon|$ regime (especially for $|\varepsilon| \lesssim 0.5$) than the bottom tercile, suggesting a threshold-like response rather than a smooth continuum. The top tercile roughly follows the middle-tercile distribution.

This behaviour motivates our working definition of AGN-affected regions: the central gas within a fixed $r = 2$ kpc aperture for galaxies in the middle and top accretion rate terciles, while the bottom tercile provides the appropriate passive comparison set under identical geometric selection. The inner-kiloparsec scale is motivated by both the feedback implementation and expectations for low-mass black holes: in the COLIBRE ‘Thermal’ AGN model energy is deposited only into the nearest SPH neighbours, producing a compact overpressurized region whose expansion is limited by the relatively modest energetics of dwarf galaxy AGN. As a result, the driven pressure perturbations are generally confined to the central region of the disc, and the 2 kpc aperture comfortably

encompasses their expected extent in the mid-plane. As for the other channels, discussion of the arbitrariness of these selection choices is deferred to Section 4.2.

With this selection in place, we can characterize the spatial and dynamical impact of AGN activity using the representative example in Fig. 3. The tracer-field map (left) reveals a compact centrally concentrated region of elevated effective pressure – often bubble-like in appearance, though generally less overpressurized than typical stellar feedback-driven bubbles. The corresponding ε map (middle) indicates that these AGN-adjacent zones host strong disequilibria, with $|\varepsilon|$ frequently exceeding unity. Statistically (right), AGN-selected gas displays slightly heavier high- $|\varepsilon|$ tails than either stellar feedback- or clump-selected gas, and a distinctly skewed distribution with a sharp rise towards the peak at negative values ($\varepsilon_{\text{peak}} \approx -0.6$). This reflects the centrally concentrated and impulsive nature of AGN energy deposition.

Two subtleties are worth noting. First, AGN-selected regions (i.e. $r \leq 2$ kpc) can overlap with stellar-feedback-selected zones in systems experiencing central star formation, which partly explains the qualitative similarity of their distributions and why the bottom tercile by 10 Myr-averaged AGN accretion rate still includes a substantial amount of non-equilibrium gas. Secondly, strong AGN events are rare in dwarfs, and thus they contribute a smaller total fraction of the disequilibrium budget than stellar feedback (we explore this further in Section 4.2). Nevertheless, AGN-driven disturbances form a distinct and robust source of disequilibrium, and therefore merit treatment as an independent driver. This claim is further reinforced by the representative system shown in Fig. 3 (HBT-HERONS track ID 14514; $M_\star \approx 2.8 \times 10^8 M_\odot$; $\text{SFR} = 0.00 M_\odot \text{yr}^{-1}$), in which AGN activity is the only local source of perturbation that we identify: its disequilibrium structure cannot be attributed to stellar processes or self-gravitating gas.

4.1.4 Gravitational asymmetries

The three mechanisms discussed so far – stellar feedback, self-gravitating clumps, and AGN activity – often operate within a broader structural context: large-scale gravitational asymmetries in the underlying potential. These asymmetries can arise from a variety of sources, including triaxial dark matter haloes (e.g. E. Hayashi & J. F. Navarro 2006), lopsided stellar mass distributions (e.g. H.-W. Rix & D. Zaritsky 1995; D. Zaritsky et al. 2013), perturbations from orbiting subhaloes (S. Kazantzidis et al. 2008), massive gas clumps, or more violent events such as minor mergers (e.g. N. Peschken, E. L. Łokas & E. Athanassoula 2020). In practice, gravitational asymmetries are difficult to disentangle from the other drivers because they can be both cause and consequence: a clump may deepen the potential in an off-centre fashion, generating streaming motions and torques; a minor merger can simultaneously deposit fresh gas, trigger a starburst, feed the AGN and distort the potential.

To probe these asymmetries directly, we examine the non-axisymmetric component of the local circular velocity field, defined as

$$v_{c, \text{asymm}}(r, \phi) \equiv v_c(r, \phi) - \langle v_c \rangle_\phi(r). \quad (5)$$

We then compare these asymmetries to the dimensional equilibrium deviation parameter, $\tilde{\varepsilon}$, defined by re-writing the steady-state Euler equation (equation 2), as the difference between the true local circular velocity and the velocity reconstructed from the

left-hand side, i.e.

$$\tilde{\varepsilon} \equiv \sqrt{r \left| [(\mathbf{v} \cdot \nabla)\mathbf{v}]_r + \frac{1}{\rho} \frac{\partial P_{\text{eff}}}{\partial r} \right|} - v_c. \quad (6)$$

This quantity can be interpreted as the offset between the actual and equilibrium-predicted orbital speeds, e.g. the deviation between the dotted dark green and solid black curves in Fig. 1.

Fig. 5 shows the conditional distribution of $\tilde{\varepsilon}$ at fixed $v_{c, \text{asymm}}$. The colour scale indicates the fraction of mass in each $\tilde{\varepsilon}$ bin relative to all gas at the same asymmetry amplitude. A solid coloured (light-red) line shows the mass-weighted median of $\tilde{\varepsilon}$ in each slice, with the shaded region marking the 50th percentile range.

The trend is clear: as $|v_{c, \text{asymm}}|$ increases from a few to roughly 20 km s^{-1} , the $\tilde{\varepsilon}$ distribution systematically shifts away from zero, with a slightly broadening scatter – as indicated by the ‘bend’ of the solid coloured (light-red) line and the widening shaded areas. Large positive $\tilde{\varepsilon}$ values are common for $|v_{c, \text{asymm}}| \gtrsim 5 \text{ km s}^{-1}$, indicating oversupport relative to equilibrium. At low asymmetry amplitudes ($\lesssim 5 \text{ km s}^{-1}$), the distribution flattens, consistent with a threshold below which the disc is largely unresponsive to the perturbation – either because velocity dispersion exceeds the asymmetry (i.e. a hot disc) or because other processes dominate the gas dynamics, effectively hiding or breaking any underlying correlation. In fact, although not visible in the figure, in this regime we find substantial low-weight scatter, largely attributable to the mechanisms discussed in the previous subsections. Furthermore, the increasing scatter at high $|v_{c, \text{asymm}}|$ likely reflects the combined influence of secondary processes such as triggered star formation, inflows, or shock-driven turbulence.

In this sense, gravitational asymmetries are not a sharply localized driver like feedback or clumps; rather, they are a persistent background field that both modulates and is modulated by other processes. They provide a structural scaffold for disequilibrium, shaping the global kinematic state of the disc while enabling local perturbations to couple more effectively to the gas.

4.2 Accounting for the disequilibrium budget

Throughout Section 4.1, we have examined how three localized physical mechanisms correlate with hydrodynamical disequilibrium. We also emphasized the structural role of gravitational asymmetries, which act as a persistent non-local driver modulating these localized processes. We now turn to a more quantitative question. How much of the total disequilibrium budget across our sample can be attributed to the specific mechanisms identified in Section 4.1?

Fig. 6 addresses this by decomposing the global ε distribution into the contributions from the three drivers. In the top panel, the coloured shaded distributions are stacked cumulatively, so that the total vertical extent at any given ε shows the fraction of gas mass linked to at least one driver (without double counting). Hatched and colour-blended regions mark overlaps between two drivers, making their joint contribution immediately visible. For clarity, the two rarest overlaps (together $\ll 1$ per cent of the total mass) are omitted. The solid black and grey distributions indicate, respectively, the sum of the drivers (i.e. including the omitted overlaps) and the global background. The bottom panel re-expresses the same data as fractions of the global distribution in each ε bin, making the relative importance of different drivers in each regime much easier to read at a glance.




Figure 5. Correlation between gravitational asymmetries and departures from hydrodynamical equilibrium. The horizontal axis shows the non-axisymmetric component of the local circular velocity, $v_{c, \text{asymm}}$, defined as the local deviation from the azimuthally averaged circular speed (equation 5). The vertical axis shows the ‘dimensional equilibrium deviation’ parameter, ξ , which measures the offset (in km s^{-1}) between the actual orbital speed and that predicted from the steady-state Euler balance (equation 6). Colours give the mass fraction in each ξ bin, normalized within each column (i.e. at fixed $v_{c, \text{asymm}}$), so each vertical slice represents a probability distribution. A solid and coloured (light-red) line indicates the mass-weighted median within each slice, with the shaded region showing the 50th percentile range, highlighting the trend and scatter. A clear pattern emerges: stronger gravitational asymmetries are associated with systematically larger equilibrium deviations, with increased scatter at high $|v_{c, \text{asymm}}|$.

We stress that the coloured curves are not mutually exclusive: a given gas element may be tagged by multiple criteria. This makes the overlap regions physically informative, as they highlight gas where multiple mechanisms act in concert – for example, the stellar/AGN feedback combination, which traces powerful central outflows, and the stellar feedback/self-gravitating clumps overlap, which often marks dense star-forming regions undergoing active disruption.

Looking first at the top panel, most of the mass not tagged by any driver lies in the low- $|\xi|$ regime, particularly for $|\xi| < 0.5$. In the bottom panel, this pattern becomes striking: the combined three-driver contribution falls from a near-saturation of ≈ 85 per cent for $|\xi| > 1$ to only ≈ 30 per cent at $\xi \approx 0$. This behaviour provides both a sanity check and strong evidence that our mechanisms are genuinely associated with disequilibrium. In an idealized scenario, one might expect a clean step from ~ 100 per cent recovery for $|\xi| \sim 1$ to 0 per cent in quasi-equilibrium, $|\xi| \sim 0.1$. The persistence of a small contribution near $\xi = 0$ and lacking some mass for $|\xi| > 0.2$ is unsurprising: the simple apertures that we use to select gas around disequilibrium drivers cannot perfectly capture the irregular morphology of the out-of-equilibrium gas. The contribution due to large-scale gravitational asymmetries is also not shown because labelling gas as affected by such asymmetries is ill defined.

The bottom panel also reveals a mild asymmetry in driver dominance. On the strongly negative side ($\xi < -0.8$), feedback dominates the budget, accounting – with overlaps – for a steady 80–85 per cent contribution. This reflects the prevalence of expanding SNe shells and wind-driven flows in regions with large negative $\partial_r v_r$. Interestingly, the majority of the mass in this extremely negative regime comes from the AGN/feedback overlap,

indicating that when the central regions host both young stars and accreting supermassive black holes (SMBHs), they are very strongly out of equilibrium. While both AGNs and self-gravitating clumps – when ignoring overlaps – are small contributors to the disequilibrium budget across the entire range, typically of order 1–5 per cent in total driver-tagged mass, on the positive end of the distribution (i.e. $\xi \gtrsim 0.2$), self-gravitating clumps have an elevated contribution, explaining ≈ 25 per cent of the mass when overlaps are included, and a steady ≈ 10 per cent contribution individually. Overall, the impression is that stellar feedback is the dominant perturbation mechanism driving dwarf galaxies out of steady-state evolution, as expected via analytic and energetic arguments. However, the large overlaps in Fig. 6 and, crucially, the local structure of disequilibrium in individual galaxies (Section 4.1) signify that disequilibrium in dwarf galaxies cannot be regarded as a consequence of just stellar processes.

A key caveat is the sensitivity of our results to the exact tagging criteria. Our fiducial thresholds are physically motivated: the stellar and AGN feedback mask sizes, maximum time since energy injection, and star particle age cap reflect typical SN bubble expansion time-scales and the duration of generic feedback episodes; the threshold for self-gravitating clumps are anchored to clearly outlying values in their respective tracer field. From these starting points, we iteratively tune the cuts: tightening a cut such that it reduces the mass tagged but leaves the ξ distribution unchanged suggests the cut is too strict, excluding relevant disequilibrium; loosening a cut such that abruptly increases the mass fraction at low- $|\xi|$ indicates overinclusiveness, contaminating the sample with recovered or ambient gas. The final selections are therefore both physically grounded and empirically calibrated.

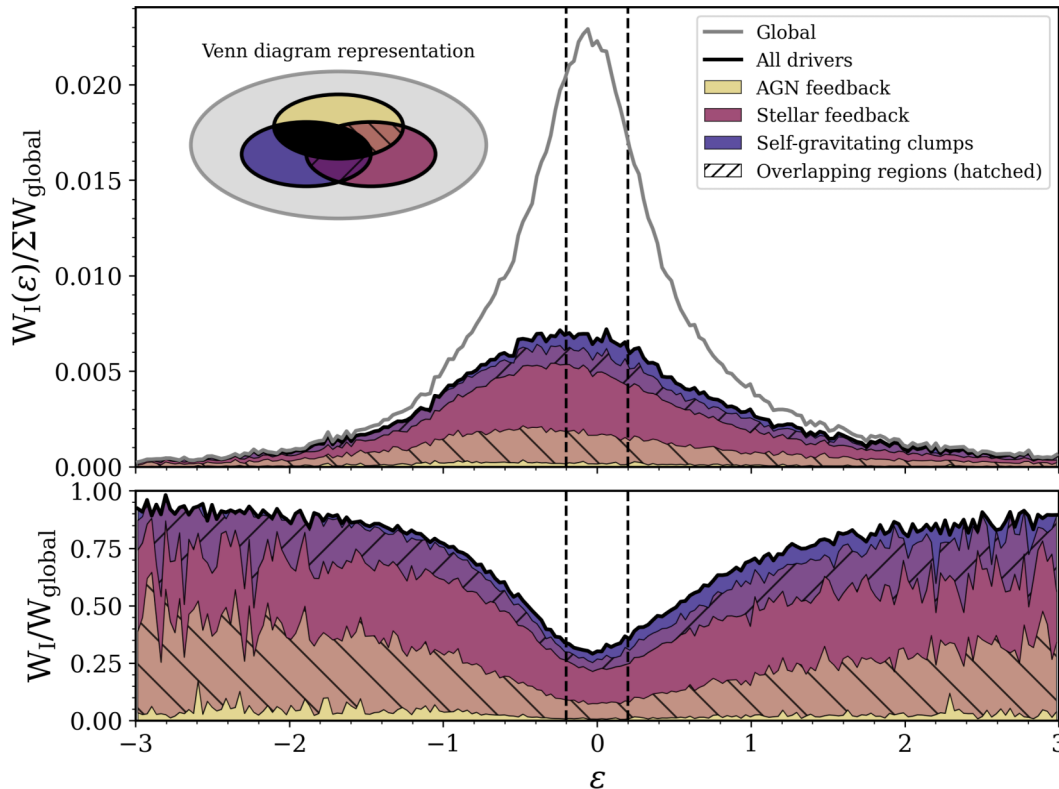


Figure 6. Cumulative accounting of the disequilibrium budget in simulated dwarf galaxy discs. *Top panel:* gas-mass-weighted distributions of ε for the three physical drivers considered – AGN energy injection (yellow), stellar feedback (magenta), and self-gravitating clumps (purple) – plotted cumulatively on top of one another. The grey solid line shows the global distribution for the full sample. Overlaps between drivers are explicitly visualized: regions where AGN and stellar feedback coincide are hatched with ‘\’ lines; overlaps between stellar feedback and self-gravitating clumps are hatched with ‘/’ lines. Colours in these overlaps follow a blending of the corresponding driver colours. For clarity, we omit the two remaining possible overlaps (AGN + clumps, and all three drivers simultaneously) as they contribute \ll 1 per cent of the total mass-weight. *Bottom panel:* fractional contribution of each process to the global mass-weight in each ε bin (same colour and hatch coding as above). This panel quantifies how each mechanism dominates, shares with, or yields to others across the disequilibrium spectrum. The vertical dashed lines in both panels mark the quasi-equilibrium range, $|\varepsilon| \leq 0.20$. *Inset:* the Venn diagram provides a schematic representation of how gas parcels tagged by the different physical drivers are subsets of the global gas population, illustrating their partial overlap rather than implying distinct or self-normalized distributions. To recover the distribution associated with a given driver, one should consider only the coloured regions uniquely associated with that driver and its direct combinations, rather than the cumulative stacked profile.

In integrated terms, roughly 75 per cent of the total gas mass in our sample lies outside the quasi-equilibrium range (i.e. $|\varepsilon| > 0.2$). Of this ‘non-equilibrium’ component, the three drivers recover ≈ 85 per cent for $|\varepsilon| > 1$, but only ≈ 30 per cent in the near-equilibrium regime ($|\varepsilon| \leq 0.2$) – a scale-dependent recovery fraction that matches the visual impression from Fig. 6. The remaining untagged mass, particularly at $0.2 < |\varepsilon| < 1$, most likely reflects transient processes and our omission in Fig. 6 of the asymmetric potential as a driver (Section 4.1.4), as well as mild turbulent, thermal, or gravitational fluctuations below our selection thresholds. We quantify the sensitivity of these recovery fractions to the adopted tagging apertures and thresholds in Appendix B, where we show that the qualitative trends and relative driver contributions are robust to factor-of-2 variations in the selection criteria.

To conclude, the cumulative-overlap decomposition shows that the disequilibrium budget is dominated by identifiable, physically interpretable mechanisms, often acting in concert and modulated by large-scale gravitational asymmetries. The small unaccounted fraction implies that any additional processes are either subdominant, short-lived, or subtle. By linking distribution shape and overlap structure with physical interpretation, Fig. 6 offers a framework for both simulation diagnostics and direct

observational comparison, particularly for radio interferometric and IFU surveys, and gas kinematics studies in low-mass galaxies. Moreover, the physically grounded origins of disequilibrium in our simulated sample reinforce the strong implications that non-equilibrium physics carries for rotation curve-based mass inference.

5 CLASSIFICATION OF DISEQUILIBRIUM GALAXIES

Determining which galaxies yield rotation curves suitable for dynamical mass modelling remains a long-standing problem. Common observational quality cuts – symmetric large-scale velocity gradients, high ordered-to-random motion ratios (V_{rot}/σ), or thin, axisymmetric disc morphologies – are widely employed (e.g. C. Trachternach et al. 2008; S.-H. Oh et al. 2015; P. E. Mancera Piña et al. 2025), yet none guarantees hydrodynamical equilibrium, nor do they constitute a uniform, physically grounded standard across environments.

In low-mass systems, the prevalence of disequilibrium (Section 4), its disproportionate impact on the inner disc, and the inability of more complete fluid models to outperform the naïve centrifugal balance assumption (Section 3) motivate a classification

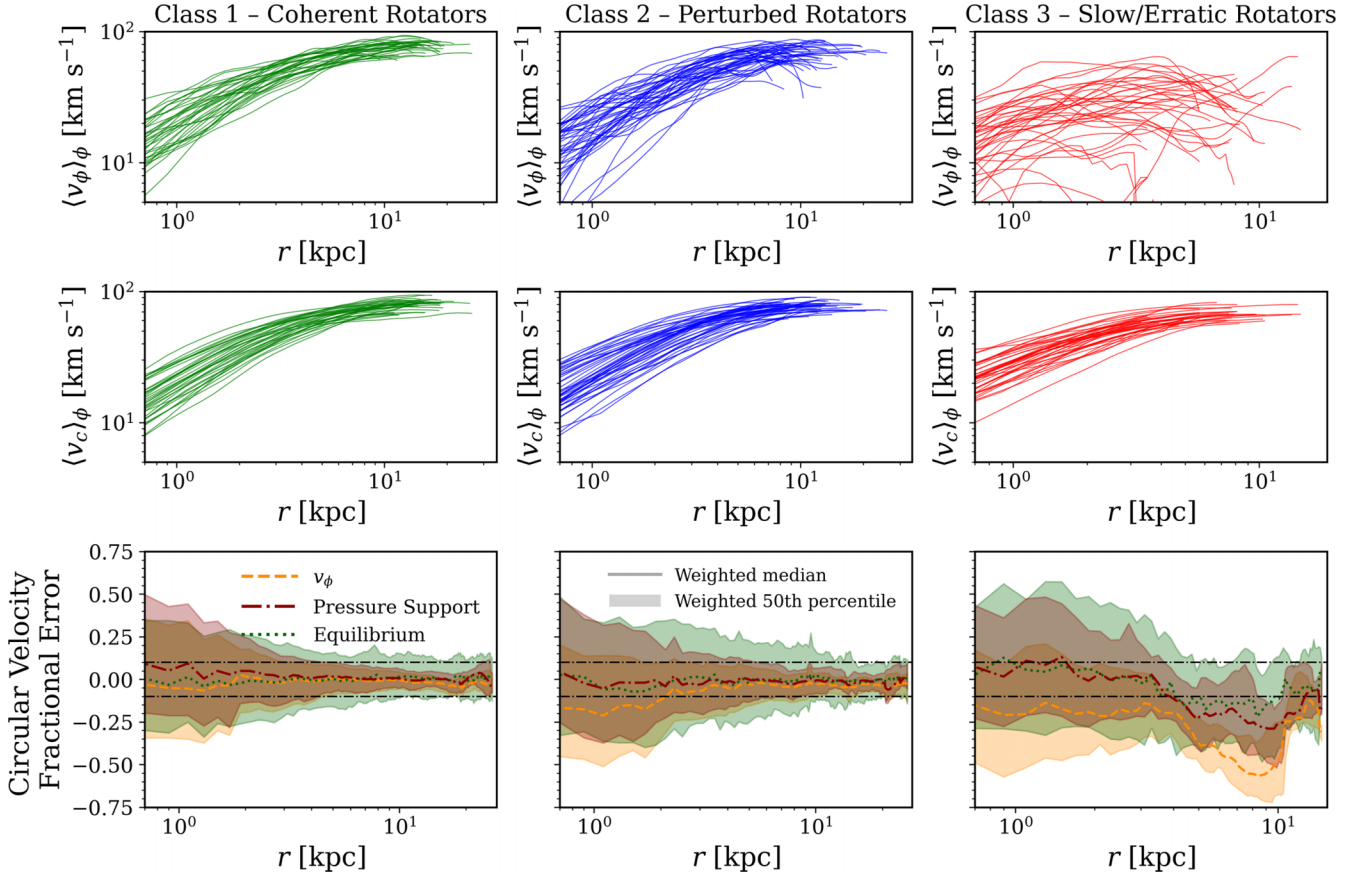


Figure 7. Visual summary of dynamical state classification of simulated gaseous discs. Each column corresponds to one of the three k -means classes (Appendix C), colour-coded consistently across rows: Class 1 (green), Class 2 (blue), and Class 3 (red). *Top row*: azimuthally averaged rotational velocity profiles, $\langle v_\phi \rangle_\phi(r)$, for all galaxies in each class. *Middle row*: corresponding azimuthally averaged circular velocity profiles, $\langle v_c \rangle_\phi(r)$. *Bottom row*: fractional error between the model-dependent reconstructed circular velocity, $v_{c,\text{model}}(r, \phi)$, and the true circular velocity from the potential $v_c(r, \phi)$, shown as class-stacked medians with 50th-percentile envelopes (both weighted) as a function of radius. The lines correspond to the centrifugal model (dark orange/dashed), the pressure-supported model (dark red/dash-dotted), and the equilibrium model (dark green/dotted), i.e. the same colour and line-style schemes as in Fig. 1. Despite similar $\langle v_c \rangle_\phi$ shapes across classes, the residuals with respect to $\langle v_\phi \rangle_\phi$ reveal distinct dynamical states, which are then reflected in the structure of the fractional error distributions across different models, including the hydrodynamical equilibrium-based, reinforcing the link between centrifugal balance and hydrodynamical equilibrium.

scheme tied directly to the underlying dynamics. Our starting point is to identify, within the simulation, those systems whose rotation curves trace the circular velocity with minimal bias, i.e. systems close to centrifugal balance.

5.1 The centrifugal–hydrodynamical equilibrium link

The condition of centrifugal balance (equation 3) directly follows from the hydrodynamical equilibrium condition (equation 2), under the assumptions of axisymmetry and negligible pressure gradients. The converse, however, is not strictly true: a galaxy may satisfy $v_\phi \approx v_c$ despite the presence of significant non-circular motions or pressure support. In practice, sustained departures from equilibrium are expected to produce systematic residuals between the azimuthally averaged rotational velocity, $\langle v_\phi \rangle_\phi$, and azimuthally averaged circular velocity, $\langle v_c \rangle_\phi$, which are then reflected in the radial behaviour of ε . The expectation that the morphology and amplitude of these residuals provide a useful first-order diagnostic of dynamical state is motivated by the quasi-spherical nature of CDM haloes dominating the potentials of low-mass systems; such geometry naturally biases the acceleration

field towards centrifugal dominance through the approximate conservation of angular momentum.

We quantify this connection by characterizing each galaxy with three complementary residual-based metrics computed from $\langle v_\phi \rangle_\phi$ and $\langle v_c \rangle_\phi$: a global integrated deviation (L_1 norm), a fractional pointwise deviation in amplitude (NRMSE-Amplitude), and a fractional pointwise deviation in slope (NRMSE-Slope). Applying k -means clustering in this three-dimensional metric space yields three statistically well-separated groups, interpreted as dynamical classes. For clarity, Appendix C presents two case-study galaxies – a stable and a disturbed disc – showing the resulting metrics, together with a visualization of the 3D parameter space used by our k -means clustering algorithm to identify the three dynamical classes and further details.

Fig. 7 provides a consolidated view of the dynamical classification. The diversity of azimuthally averaged rotational velocities, $\langle v_\phi \rangle_\phi$, is striking (top row): Class 1 systems show clean monotonic rises with minimal scatter; Class 2 systems exhibit coherent, wave-like modulations; and Class 3 systems rotate slowly, irregularly, or both. In contrast, the circular velocity profiles $\langle v_c \rangle_\phi$ are comparatively uniform across classes, demonstrating that differences in the gravitational potential – whether mildly cored or

cuspy – are not the primary drivers of the observed rotational diversity. Instead, the variation originates almost entirely from gas dynamics: non-circular motions, pressure forces, and departures from steady-state flow.

The bottom row makes this explicit by directly comparing the reconstructed circular velocity to the true value via the mass-weighted fractional difference profiles for the three dynamical models introduced in Section 3.1: (i) the centrifugal model (orange, dashed),

$$v_{c,\text{model}}(r, \phi) = v_{\phi}(r, \phi); \quad (7)$$

(ii) the pressure-supported model (red, dash-dotted),

$$v_{c,\text{model}}(r, \phi) = \sqrt{r \left| \frac{1}{\rho} \frac{\partial P_{\text{eff}}}{\partial r} - v_{\phi}^2/r \right|}(r, \phi); \quad (8)$$

(iii) the full hydrodynamical-equilibrium model (green, dotted),

$$v_{c,\text{model}}(r, \phi) = \sqrt{r \left| \frac{1}{\rho} \frac{\partial P_{\text{eff}}}{\partial r} + [(\mathbf{v} \cdot \nabla)\mathbf{v}]_r \right|}(r, \phi). \quad (9)$$

For each class, the shaded regions denote the weighted 50th-percentile interval, while the lines show the weighted median.

5.1.1 Class 1 – coherent rotators, 29 per cent of the sample (35/122)

These galaxies exhibit near-perfect alignment between $\langle v_{\phi} \rangle_{\phi}$ and $\langle v_c \rangle_{\phi}$, with median residuals ≈ 0 per cent across most radii and only mild (≈ -5 per cent) underrotation beyond $r \gtrsim 12$ kpc. This bias largely vanishes when pressure or advective terms are included. Inside ≈ 4 kpc – and especially within ≈ 2 kpc – all models show significantly broadened error distributions, with 50th-percentile ranges extending beyond the ± 10 per cent threshold, which might be critical to discuss local features (or bumps) in the rotation curve and determine whether a halo is cuspy or cored (e.g. I. S. Sands et al. 2024). Even for dynamically coherent discs therefore the innermost kiloparsec remains a somewhat challenging regime for reliable mass modelling.

5.1.2 Class 2 – perturbed rotators, 37 per cent of the sample (46/122)

These systems display a baseline underrotation of ≈ -5 per cent at $r \gtrsim 4$ kpc, steepening to ≈ -10 per cent at $r \approx 2$ kpc and ≈ -20 per cent at $r \approx 1$ kpc, with considerable scatter skewed further negative. For this class, including pressure support or the full equilibrium terms substantially improves the median recovery, yielding values near 0 per cent across most radii. However, the scatter remains large – with 50th-percentile widths exceeding ± 10 per cent for $r \lesssim 6$ kpc for both corrected models, and for $r \lesssim 10$ kpc in the case of the equilibrium model. Thus, even though the median can be corrected, the intrinsic variation remains quite large for robust inference in the inner disc.

5.1.3 Class 3 – slow/erratic rotators, 34 per cent of the sample (41/122)

This class yields the most biased and dispersed estimates. The centrifugal model never enters the ± 10 per cent accuracy band at any radius, with median errors reaching ~ -60 per cent. Pressure support reduces the bias, especially within the inner 5 kpc, but

still leaves systematic offsets of ≈ -20 to -30 per cent at larger radii. The equilibrium model achieves the largest improvement, but remains significantly biased at all radii and exhibits the widest dispersion of any model (typically ± 20 – 30 per cent). For these galaxies, no estimator provides a reliable reconstruction of the true circular velocity.

Taken together, these results clarify the dynamical meaning of the three classes. Coherent rotators lie close to both centrifugal and hydrodynamical equilibrium; perturbed rotators contain substantial but structured departures that can be partially corrected; and slow/erratic rotators are dominated by strong disequilibrium that no steady-state model can adequately capture. The close correspondence between the morphology of the $\langle v_{\phi} \rangle_{\phi}$ – $\langle v_c \rangle_{\phi}$ residuals and the radial structure of the model-dependent fractional errors demonstrates that centrifugal balance provides a practical, physically motivated proxy for hydrodynamical equilibrium. Crucially, only a minority of galaxies – less than one-third centrals in the halo mass range $10.75 \leq \log_{10}(M_{200c}/M_{\odot}) \leq 11.00$ – reside in a regime where circular velocity recovery is unbiased and tightly constrained. This underscores the importance of dynamical classification when interpreting rotation curves of low-mass systems and motivates the use of coherent rotators as the subset best suited for robust dark-matter mass distribution inference.

5.2 Galaxy-wide correlations with dynamical class

Having established the existence of three distinct dynamical classes, we now examine whether they correlate with broader galactic properties. While many of the quantities considered here are not directly measurable observationally, they serve two purposes: (i) to identify the dominant physical processes driving disc-wide perturbations, and (ii) to guide the development of observational pre-selection criteria for identifying systems dynamically suitable for rotation curve modelling and dark matter inference.

Fig. 8 shows the distributions of selected global and disc-scale properties for coherent rotators (Class 1, green), perturbed rotators (Class 2, blue), and slow rotators (Class 3, red). Each panel presents a violin plot for one property, allowing visual comparison of trends and scatter between classes. Our aim is to link a galaxy’s perturbation state – quantified via residual kinematics – to structural and feedback-related properties that may drive or sustain dynamical disequilibrium.

A clear trend is seen in total HI mass, i.e. $M_{\text{HI}} = M_{\text{HI}}(r \leq R_{\text{ext,HI}})$. Coherent rotators have the highest median M_{HI} ; this is consistent with the physical picture of more inertia yielding enhanced resistance to perturbative forces, and the maintenance of well-ordered discs. Perturbed rotators occupy an intermediate regime with broader scatter, indicating that moderately gas-rich systems can retain some rotational support despite disturbances. Slow rotators show systematically lower M_{HI} , suggesting that reduced HI content is strongly associated with disrupted kinematics. This correlation is evident in our narrow halo-mass-selected sample (M_{200c}) but may be obscured in $V_{\text{max,rot}}$ -selected samples, where lower rotation speeds in disturbed systems bias them out of the selection, i.e. the observationally accessible $V_{\text{max,rot}} \equiv \max(\langle v_{\phi} \rangle_{\phi})$ will be a biased estimate of $V_{\text{max}} \equiv \max(\langle v_c \rangle_{\phi})$ for this specific class. Similar trends are seen in the NIHAO simulations (A. A. Dutton, A. Obreja & A. V. Macciò 2019) and in the FIRE/FIRE-3 simulations (K. El-Badry et al. 2018; I. S. Sands et al. 2024), which find that lower mass HI-poor galaxies are less rotationally supported and more strongly affected by non-circular

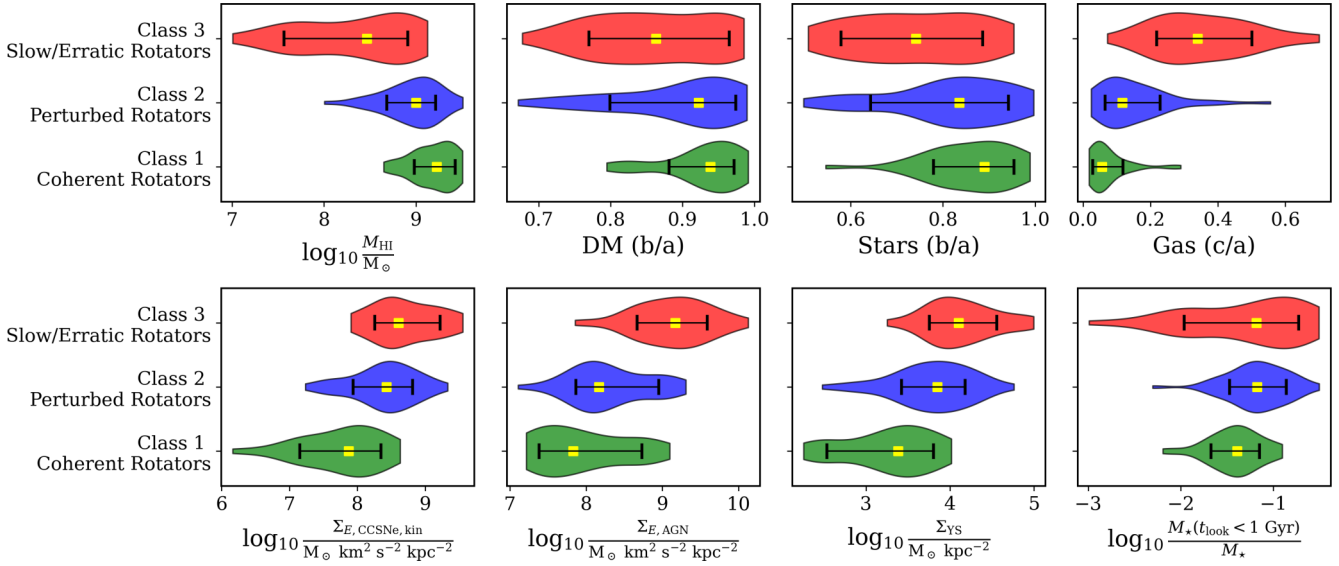


Figure 8. Distributions of global galaxy and disc properties across the three classes: coherent rotators, perturbed rotators, and slow rotators. Each panel shows a violin plot with class on the y-axis and a given physical property on the x-axis. The yellow squares indicate the class median; black horizontal lines span the 68th percentile range. *Top row, left to right:* total neutral hydrogen mass (M_{HI}), DM halo intermediate-to-major axis ratio (b/a), stellar distribution intermediate-to-major axis ratio (b/a), and gas distribution minor-to-major axis ratio (c/a). *Bottom row, left to right:* relative strength of kinetic feedback from CCSNe, AGN feedback, and generic stellar feedback from young stars (YS), as quantified in Section 2.5 and discussed in the text. The rightmost panel instead shows a proxy for recent star formation: the fraction of stellar mass formed within the last gigayear relative to the total stellar mass at $z = 0$. Systematic trends across classes offer insight into how global physical conditions and feedback processes relate to gas disc perturbation states.

motions. As noted in A. A. Dutton et al. (2019), this is confirmed observationally from the skewness of linewidths in bins of HI mass in both ALFALFA (M. P. Haynes et al. 2018) and HIPASS (B. S. Koribalski et al. 2004) observations.

We now examine the thickness and triaxiality of the gas, stellar, and dark matter components. To ensure that the shape parameters of the distributions are both comparable across galaxies and tied to the spatial region where the gas kinematics are classified, we measure their principal axes within a spherical aperture defined by the disc extent radius, $R_{\text{ext, HI}}$.

We quantify the gas disc thickness with the minor-to-major axis ratio (c/a). Thin discs are notoriously difficult to capture accurately in cosmological hydrodynamical simulations (e.g. O. Agertz, R. Teyssier & B. Moore 2011; S. Tacchella et al. 2019; A. D. Ludlow et al. 2023; J. A. Benavides et al. 2025). COLIBRE seems to fare better in this regard than past efforts: our sample includes 39/122 galaxies with stellar discs thinner than $c/a = 0.3$, and seven galaxies with $c/a \leq 0.2$. This still falls short of the frequency of (very) thin stellar discs in dwarf galaxies inferred from observed shape distributions, but is a substantial improvement over other models including TNG50, FIREbox, and ROMULUS25 (J. A. Benavides et al. 2025). This is encouraging: thin gas discs are a prerequisite to the formation of thin stellar discs, and the gas c/a distributions for our three dynamical classes show a monotonic progression that parallels the trend with total HI mass, indicating a tight connection between gas disc morphology and dynamical state. In vertical hydrostatic equilibrium, greater turbulence implies a thicker disc, and observationally V_{rot}/σ correlates tightly with disc aspect ratio (see section 5.3.1 of P. E. Mancera Piña et al. 2025). Here, coherent rotators are thin (median $c/a \sim 0.05$) with minimal scatter; perturbed rotators are thicker (median ~ 0.15) with a broader range; slow rotators are systematically thick (median ~ 0.3 , none below ~ 0.15). These results suggest that c/a is

a strong morphological indicator of disequilibrium in this halo-mass regime, and that its diagnostic potential in rotation curve reliability assessments should be considered of high fidelity.

For stellar and dark matter shapes, the distributions overlap between classes but show a bias towards greater prolateness (lower intermediate-to-major axis ratio b/a) in perturbed and slow rotators. This trend is consistent with findings from E. R. Downing & K. A. Oman (2023), who identified similar correlations between halo shape and dynamical perturbations in high-resolution APOSTLE dwarf galaxies. Theoretically, this connection is well motivated: a galaxy embedded in a prolate halo often resides in a non-axisymmetric potential, which naturally excites strong non-circular motions in the gas kinematics (E. Hayashi & J. F. Navarro 2006; E. Hayashi et al. 2007; A. Marasco et al. 2018; K. A. Oman et al. 2019) – a behaviour our classification is explicitly sensitive to. Moreover, when the gas disc is significantly misaligned with the plane defined by the halo’s intermediate and major axes, no stable configuration can be sustained (see e.g. the supplementary videos in E. R. Downing & K. A. Oman 2023), further reinforcing the link between halo shape and kinematic disequilibrium.

The bottom row of panels in Fig. 8 quantifies the recent baryonic feedback impact from three channels: CCSNe kinetic feedback, AGN thermal feedback, and a proxy for generic stellar feedback traced by the mass of young stars. We measure these by selecting gas (or young stars for the latter) within a spherical aperture of $R_{\text{ext, HI}}$ and with the most recent event within 100 Myr (or age for young stars), integrating the relevant energy proxy and normalizing by disc area. This choice ensures fair comparisons across discs of different sizes; full details are given in Section 2.5. All three feedback modes show systematic increases from coherent to perturbed to slow rotators, with median values rising by more than an order of magnitude in several cases, supporting a picture in which recent, concentrated feedback stirs turbulence,

thickens the disc, and drives sustained departures from hydrodynamical equilibrium. Our SNe proxy tracks only the kinetic channel of core-collapse explosions; although this constitutes roughly 10 per cent of the total SNe energy budget, the remaining thermal component scales proportionally (when averaged over the relevant time-scales) owing to the numerical implementation of the feedback model. As a result, including the thermal contribution would shift the absolute values but would not alter the relative trends across the dynamical classes.

In particular, coherent rotators exhibit uniformly low levels of feedback impact, consistent with their dynamically stable thin discs and quiescent gas kinematics. Perturbed rotators show moderate enhancement in all feedback measures, reflecting a plausible scenario where recent or ongoing feedback contributes to the development of wave-like residuals and asymmetric distortions in the gas velocity field. Finally, slow rotators display the highest overall feedback levels, especially for AGNs and generic stellar feedback channels, aligning with their strongly disrupted, thickened, and erratic gas discs.

These trends are both intuitive and well aligned with existing literature. For instance, A. Pontzen & F. Governato (2012) demonstrated, albeit under different subgrid prescriptions, that repeated momentum injection from supernova explosions can cause localized gas parcels to migrate substantially within the gravitational potential. This process induces rapid orbital fluctuations and delays the re-establishment of equilibrium in the disc; moreover, it aligns nicely with what we have shown and discussed in Section 4.1.1 and Section 4.2.

Furthermore, while AGN feedback has often been regarded as energetically unimportant in low-mass galaxies, recent work suggests that even modest AGN episodes can induce significant dynamical disturbances – see e.g. the analysis in S. Koudmani et al. (2025) and the detailed discussion in their sec. 4.2. High-resolution simulations show that intermittent AGN outbursts can drive cycles of expansion and re-accretion in the central gas, generating potential fluctuations and delaying the re-establishment of equilibrium (e.g. D. Martizzi, R. Teyssier & B. Moore 2013; S. Peirani et al. 2017), i.e. a mechanism similar to what described in A. Pontzen & F. Governato (2012) regarding supernova explosions.⁷ Observational evidence reinforces this picture: C. M. Manzano-King & G. Canalizo (2020) reported a strong association between AGN activity and disturbed ionized-gas kinematics in dwarf galaxies, suggesting that AGNs may play a more substantial role in shaping disc structure and rotation-curve morphology than previously assumed. This interpretation accords with our results, where slow rotators exhibit the highest AGN feedback signatures – with median energy proxies more than 1 dex above those of perturbed and coherent rotators – as well as with the broader trends discussed in Sections 4.1.3 and 4.2.

In summary, these correlations point to a coherent physical picture in which halo shape, gas disc structure, and recent feedback collectively determine a galaxy’s dynamical state. This not only clarifies the mechanisms that drive and sustain disequilibrium, but also begins to suggest concrete physically motivated diag-

nostics for identifying systems whose rotation curves are likely unreliable for dynamical mass modelling in observations.

6 CONCLUSIONS

We have used the currently available highest resolution largest boxsize COLIBRE cosmological hydrodynamical simulations (i.e. a 25 cMpc box at $\sim 10^5 M_\odot$ particle resolution) to quantify departures from hydrodynamical and centrifugal equilibrium in the gas discs of dwarf galaxies, defined as systems with virial mass in the range $10^{10.75} < M_{200c}/M_\odot < 10^{11}$. By evaluating the full Eulerian acceleration balance locally in warped mid-plane frames, we have assessed – for the first time on a large simulated sample – the spatial and radial prevalence of disequilibrium physics, and its implications for rotation-curve-based mass inference. We carried out our entire analysis for both the thermal and hybrid AGN feedback models and found no substantial differences affecting our conclusions; for clarity, we have presented only results from the fiducial thermal runs throughout. Our main conclusions are as follows:

(i) Disequilibrium is ubiquitous. Across our sample, both the full equilibrium model (equation 2) and the naïve centrifugal balance assumption (equation 3) frequently fail to recover the true circular velocity, with typical fractional deviations of ≥ 10 per cent. The inner few kiloparsecs, which provide the greatest leverage for dark matter profile constraints, are also the most strongly perturbed with fractional deviations typically much greater than 10 per cent. In particular, we find a population-wide tendency for the rotational velocity to systematically underestimate the true circular velocity in the inner disc, while a minority of regions exhibit the opposite bias; this bidirectional scatter can naturally result in a wide diversity of inferred inner rotation-curve shapes, including both cored and centrally concentrated profiles.

(ii) Related to the first point, more complete fluid models do not guarantee improved mass recovery. Even when including pressure-gradient and convective terms, equilibrium-based reconstructions often perform worse than the simple assumption of $v_\phi = v_c$. This reflects the fact that unobservable time-dependent terms limit the validity of analytic corrections. Understanding the dynamical state of the system should have higher priority than attempting corrections to the rotation curve.

(iii) Three physical drivers account for most strong disequilibria. Stellar feedback, self-gravitating gas clumps, and AGN energy injection collectively explain ≈ 85 per cent of the gas mass in extreme disequilibrium ($|\varepsilon| > 1$), often acting in concert. Large-scale gravitational asymmetries provide a persistent structural background that modulates and sustains these perturbations, and might also explain some residual mass unaccounted for by these three processes, especially in the $0.2 < |\varepsilon| < 1$ regime (mild disequilibrium).

(iv) A physically grounded classification of rotation-curve reliability is possible. Using residual-based metrics linking centrifugal and hydrodynamical equilibrium, we identify three dynamical classes – coherent rotators, perturbed rotators, and slow/erratic rotators, which differ primarily in gas dynamics, not underlying potential shapes. Population trends at a class level in structural- and feedback-related properties reveal that recent feedback activity, gas vertical thickness, and total HI mass are powerful indicators of the galaxy’s dynamical state, albeit for a sample ‘unrealistically’ selected in a narrow halo mass range.

⁷As noted by D. Martizzi et al. (2013), this mechanism is effective only when the AGN-driven cycle of gas expansion and subsequent re-accretion occurs on time-scales comparable to the local dynamical time. In Section 4.1.3 we indeed find $|\partial_r v_r| \sim |g_r|$ (i.e. $|\varepsilon| \sim 1$) in regions adjacent to an active AGN, indicating that the relevant gas motions occur on precisely the time-scales required for this mechanism to operate efficiently.

The quantitative fractions that we report – such as the ≈ 75 per cent of gas mass in non-equilibrium ($|\varepsilon| > 0.2$) and the ≈ 70 per cent of systems at $\log_{10}(M_{200c}/M_{\odot}) = 10.75\text{--}11$ unsuitable for rotation-curve analysis – should be interpreted as indicators of prevalence rather than universal numbers. Our 25 cMpc box is the largest volume that simultaneously reaches the gas-mass resolution required to resolve dwarf discs with $\gtrsim 10^3$ gas particles, enabling spatially resolved mid-plane kinematics. Larger volume lower resolution simulations would inevitably wash out the disequilibrium features that we measure and inject numerical noise, while higher resolution runs in even larger boxes remain computationally prohibitive.

A second limitation is the numerical implementation of feedback. While COLIBRE incorporates a self-consistent treatment of the cold and dense phases of the ISM, and an enhanced sampling of the gravitational potential, mitigating spurious baryon–DM energy transfer, hence making it an ideal choice for analyses of gas dynamics, our results will inevitably depend on the strength and coupling of feedback, hence remain model-dependent. Our cross-checks with the hybrid AGN feedback dual versions confirm that our main results are robust to these numerical choices: despite different calibration values and injection schemes, the nature and spatial pattern of disequilibrium are qualitatively unchanged. Ultimately, verifying the generality of our findings will require similarly resolved analyses in other simulation suites. Comparisons with recent FIRE-3 work using similar diagnostics (I. S. Sands et al. 2024) further suggest that our qualitative conclusions are not simulation specific. Related evidence for systematic departures from equilibrium and biased circular-velocity recovery has also been reported by E. D. Jahn et al. (2023) within idealized simulations using the SMUGGLE model, indicating that these effects arise across a range of numerical implementations and modelling strategies, even outside fully cosmological contexts. Within the EDGE simulation suite, M. P. Rey et al. (2024) found that inner circular velocities are recovered only in a rare subset of quiescent dwarf galaxies that temporarily host well-ordered HI discs. In particular, their results reinforce our conclusions that rotation-curve-based mass estimates critically depend on the dynamical state of the gas, which they also found tightly linked to morphology and star formation history, and that corrections for pressure gradients do not systematically improve recovery of the underlying gravitational potential.

To summarize, once hydrodynamical equilibrium is broken, the foundational assumption required to interpret the rotation curve as a circular velocity – or, equivalently, enclosed mass – tracer is itself invalid, irrespective of any detailed analytic correction. Given the prevalence of out-of-equilibrium gas (≈ 75 per cent of the mid-plane gas mass out of equilibrium), the disproportionate impact in the innermost regions of galaxies, and the ≈ 70 per cent of systems at $\log_{10}(M_{200c}/M_{\odot}) = 10.75\text{--}11$ that we find to be ill-suited for rotation-curve analysis (with this fraction expected to rise at lower halo mass), it is unavoidable that at least part of the observed diversity in dwarf galaxy rotation curves reflects mistaking non-equilibrium gas dynamics for variation in dark matter structure. This reframes long-standing issues such as the core–cusp and rotation-curve diversity problems primarily as matters of tracer reliability, implying that steady-state mass modelling requires either stringent dynamical pre-selection or methods that explicitly incorporate time-dependent gas physics.

Crucially, the near-symmetry of the ε distribution (Fig. 6) does not imply that disequilibrium effects cancel in the mean. The observationally inferred rotation curve will inevitably depend

sensitively on the detailed local structure of each individual galaxy and on how its disequilibrium features project on to the line of sight, so ensemble averages cannot be assumed to recover the true underlying mass – similar arguments apply for the pressure-supported and purely centrifugal models shown in the bottom row of Fig. 7. Also note that, by construction of the warped mid-plane geometry, we are evaluating the degree of disequilibrium in the local plane of maximal rotation; upper and lower vertical layers likely exhibit even stronger effects, enhancing these issues in thicker discs. This highlights the need to treat disequilibrium not as a nuisance that cancels in the mean, but as a first-order systematic shaping the kinematic signal itself.

Quantifying precisely the contribution of disequilibrium to rotation-curve diversity remains an open question. The next step is to confront our theoretical diagnostics and predictions directly with mock and real observations. By generating synthetic radio (21 cm) and IFU emission-line data from our simulations (e.g. respectively MARTINI and SIMSPIN, K. Harborne 2019; K. Oman 2024), we will assess the extent to which the disequilibrium drivers and classification criteria identified here can be recognized in practice, and how they propagate into dark matter inferences. Such work will ultimately determine whether the disequilibrium physics identified in simulations is a major contributor to small-scale tensions in Λ CDM, or a challenge peculiar to certain galaxy formation models.

ACKNOWLEDGEMENTS

The authors thank the researchers working within the broader COLIBRE collaboration for helpful discussions, and providing insights that helped refining the final version of the manuscript.

DD, KAO, and KEH acknowledge support by the Royal Society through a Dorothy Hodgkin Fellowship (DHF/R1/231105) held by KAO. FF is supported by a UKRI Future Leaders Fellowship (grant no. MR/X033740/1). ABL acknowledges support by the Italian Ministry for Universities (MUR) program ‘Dipartimenti di Eccellenza 2023–2027’ within the Centro Bicocca di Cosmologia Quantitativa (BiCoQ), and support by UNIMIB’s Fondo Di Ateneo Quota Competitiva (project 2024-ATEQC-0050). EC and FH acknowledge funding from the Netherlands Organization for Scientific Research (NWO) through research programme Athena 184.034.002. CSF was supported by the European Research Council (ERC) Advanced Investigator grant DMIDAS (GA 786910) and the STFC Consolidated Grant ST/T000244/1. SP acknowledges support by the Austrian Science Fund (FWF) through grant-DOI: 10.55776/V982.

This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University, and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure. This work has made use of NASA’s Astrophysics Data System Bibliographic Services.

For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

Author contributions: DD led the analysis, interpretation, and presentation of all results under the supervision of KAO, KEH, and FF. All other authors contributed to creating the COLIBRE

simulations, and provided input on methodology, interpretation of results, and draft versions of the manuscript.

SOFTWARE

The following software packages were used in this work: NUMPY (C. R. Harris et al. 2020), SCIPY⁸, SCIKIT-LEARN (F. Pedregosa et al. 2011), NUMBA (S. K. Lam, A. Pitrou & S. Seibert 2015), ASTROPY (Astropy Collaboration 2022), H5PY (A. Collette et al. 2022), HBT-HERONS (V. J. Forouhar Moreno et al. 2025), SOAP (R. McGibbon et al. 2025), SWIFTSIMIO (J. Borrow & A. Borrisov 2020), SWIFTGALAXY (K. Oman 2025).

DATA AVAILABILITY

The data supporting the plots within this article are available from the corresponding author upon request. The COLIBRE simulation suite will eventually be released publicly as part of the collaboration's planned data-release programme, although its multipetabyte volume will likely prevent the raw data from being hosted directly online. Until then, researchers wishing to use the simulations may contact the COLIBRE team. Further information on the project is available at <https://colibre.strw.leidenuniv.nl/>.

REFERENCES

- Abbott T. M. C. et al., 2022, *Phys. Rev. D*, 105, 023520
 Agertz O., Teyssier R., Moore B., 2011, *MNRAS*, 410, 1391
 Albers S. M. et al., 2019, *MNRAS*, 490, 5538
 Arjona-Gálvez E., Di Cintio A., Grand R. J. J., 2024, *A&A*, 690, A286
 Astropy Collaboration, 2022, *ApJ*, 935, 167
 Bacchini C., Fraternali F., Iorio G., Pezzulli G., Marasco A., Nipoti C., 2020, *A&A*, 641, A70
 Baldassare V. F., Reines A. E., Gallo E., Greene J. E., 2017, *ApJ*, 836, 20
 Banerjee A., Jog C. J., Brinks E., Bagetakos I., 2011, *MNRAS*, 415, 687
 Bell E. F., de Jong R. S., 2001, *ApJ*, 550, 212
 Benavides J. A., Sales L. V., Navarro J. F., White S. D. M., Oman C. S., Oman K. A., Cole S., 2025, (preprint [arXiv:2512.11035](https://arxiv.org/abs/2512.11035))
 Benítez-Llambay A., Frenk C. S., Ludlow A. D., Navarro J. F., 2019, *MNRAS*, 488, 2387
 Benítez-Llambay A. et al., 2026, *MNRAS*, 546, stag268
 Boomsma R., Oosterloo T. A., Fraternali F., van der Hulst J. M., Sancisi R., 2008, *A&A*, 490, 555
 Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
 Borrow J., Borrisov A., 2020, *J. Open Source Softw.*, 5, 2430
 Borrow J., Schaller M., Bower R. G., Schaye J., 2022, *MNRAS*, 511, 2367
 Brinks E., Walter F., Kerp J., 2003, *Ap&SS*, 284, 627
 Bruzual G., Charlot S., 2003, *MNRAS*, 344, 1000
 Burger J. D., Zavala J., Sales L. V., Vogelsberger M., Marinacci F., Torrey P., 2022, *MNRAS*, 513, 3458
 Chaikin E., Schaye J., Schaller M., Benítez-Llambay A., Nobels F. S. J., Ploekinger S., 2023, *MNRAS*, 523, 3709
 Chaikin E. et al., 2025, preprint ([arXiv:2509.04067](https://arxiv.org/abs/2509.04067))
 Chan T. K., Kereš D., Oñorbe J., Hopkins P. F., Muratov A. L., Faucher-Giguère C.-A., Quataert E., 2015, *MNRAS*, 454, 2981
 Chemin L., Huré J.-M., Soubiran C., Zibetti S., Charlot S., Kawata D., 2016, *A&A*, 588, A48
 Chemin L., Amram P., Carignan C., Balkowski C., Épinat B., 2020, *Boletín de la Asociación Argentina de Astronomía La Plata Argentina*, 61C, 45
 Chua K. T. E., Pillepich A., Vogelsberger M., Hernquist L., 2019, *MNRAS*, 484, 476
 Collette A. et al., 2022, *h5py/h5py: 3.7.0*, Zenodo
 Correa C. A., Wyithe J. S. B., Schaye J., Duffy A. R., 2015, *MNRAS*, 452, 1217
 Correa C. A., Schaller M., Schaye J., Ploekinger S., Borrow J., Bahé Y., 2025, *MNRAS*, 536, 3338
 Creasey P., Sameie O., Sales L. V., Yu H.-B., Vogelsberger M., Zavala J., 2017, *MNRAS*, 468, 2283
 Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 426, 140
 Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
 de Blok W. J. G., 2010, *Adv. Astron.*, 2010, 789293
 de Blok W. J. G., Bosma A., 2002, *A&A*, 385, 816
 de Blok W. J. G., McGaugh S. S., Bosma A., Rubin V. C., 2001, *ApJ*, 552, L23
 Dekel A., Silk J., 1986, *ApJ*, 303, 39
 Dekel A., Sari R., Ceverino D., 2009, *ApJ*, 703, 785
 Di Cintio A., Brook C. B., Macciò A. V., Stinson G. S., Knebe A., Dutton A. A., Wadsley J., 2014, *MNRAS*, 437, 415
 Diemer B., Joyce M., 2019, *ApJ*, 871, 168
 Downing E. R., Oman K. A., 2023, *MNRAS*, 522, 3318
 Dutton A. A., Courteau S., de Jong R., Carignan C., 2005, *ApJ*, 619, 218
 Dutton A. A., Obreja A., Macciò A. V., 2019, *MNRAS*, 482, 5606
 Eisenstein D. J. et al., 2005, *ApJ*, 633, 560
 El-Badry K. et al., 2018, *MNRAS*, 473, 1930
 Elbert O. D., Bullock J. S., Garrison-Kimmel S., Rocha M., Oñorbe J., Peter A. H. G., 2015, *MNRAS*, 453, 29
 Elmegreen D. M., Elmegreen B. G., Ravindranath S., Coe D. A., 2007, *ApJ*, 658, 763
 Flores R. A., Primack J. R., 1994, *ApJ*, 427, L1
 Forouhar Moreno V. J., Helly J., McGibbon R., Schaye J., Schaller M., Han J., Kugel R., Bahé Y. M., 2025, *MNRAS*, 543, 1339
 Frenk C. S., White S. D. M., Davis M., Efstathiou G., 1988, *ApJ*, 327, 507
 Genzel R. et al., 2011, *ApJ*, 733, 101
 Goldbaum N. J., Krumholz M. R., Forbes J. C., 2015, *ApJ*, 814, 131
 Governato F. et al., 2012, *MNRAS*, 422, 1231
 Han J., Cole S., Frenk C. S., Benítez-Llambay A., Helly J., 2018, *MNRAS*, 474, 604
 Harborne K., 2019, *Astrophysics Source Code Library*, record ascl:1903.006
 Harris C. R. et al., 2020, *Nature*, 585, 357
 Hayashi E., Navarro J. F., 2006, *MNRAS*, 373, 1117
 Hayashi E., Navarro J. F., Springel V., 2007, *MNRAS*, 377, 50
 Haynes M. P. et al., 2018, *ApJ*, 861, 49
 Heymans C. et al., 2013, *MNRAS*, 432, 2433
 Hopkins P. F., Quataert E., Murray N., 2012, *MNRAS*, 421, 3522
 Hopkins P. F. et al., 2023, *MNRAS*, 519, 3154
 Hunter D. A. et al., 2012, *AJ*, 144, 134
 Huško F. et al., 2026, *MNRAS*, in press
 Jahn E. D. et al., 2023, *MNRAS*, 520, 461
 Kaplinghat M., Ren T., Yu H.-B., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 027
 Kazantzidis S., Bullock J. S., Zentner A. R., Kravtsov A. V., Moustakas L. A., 2008, *ApJ*, 688, 254
 Kim C.-G., Ostriker E. C., 2017, *ApJ*, 846, 133
 Koribalski B. S. et al., 2004, *AJ*, 128, 16
 Koudmani S., Sijacki D., Smith M. C., 2022, *MNRAS*, 516, 2112
 Koudmani S., Rennehan D., Somerville R. S., Hayward C. C., Anglés-Alcázar D., Orr M. E., Sands I. S., Wellons S., 2025, *MNRAS*, 540, 1928
 Lam S. K., Pitrou A., Seibert S., 2015, in *Proc. Second Workshop on the LLVM Compiler Infrastructure in HPC*. ACM, p. 1
 Lelli F., McGaugh S. S., Schombert J. M., 2016, *AJ*, 152, 157
 Ludlow A. D., Bose S., Angulo R. E., Wang L., Hellwing W. A., Navarro J. F., Cole S., Frenk C. S., 2016, *MNRAS*, 460, 1214
 Ludlow A. D., Schaye J., Schaller M., Richings J., 2019, *MNRAS*, 488, L123
 Ludlow A. D., Fall S. M., Schaye J., Obreschkow D., 2021, *MNRAS*, 508, 5114

⁸docs.scipy.org/doc/

- Ludlow A. D., Fall S. M., Wilkinson M. J., Schaye J., Obreschkow D., 2023, *MNRAS*, 525, 5614
- Mancera Piña P. E., Read J. I., Kim S., Marasco A., Benavides J. A., Glowacki M., Pezzulli G., Lagos C. D. P., 2025, *A&A*, 699, A311
- Manzano-King C. M., Canalizo G., 2020, *MNRAS*, 498, 4562
- Marasco A., Oman K. A., Navarro J. F., Frenk C. S., Oosterloo T., 2018, *MNRAS*, 476, 2168
- Martizzi D., Teyssier R., Moore B., 2013, *MNRAS*, 432, 1947
- Martizzi D., Faucher-Giguère C.-A., Quataert E., 2015, *MNRAS*, 450, 504
- Mashchenko S., Wadsley J., Couchman H. M. P., 2008, *Science*, 319, 174
- McGaugh S. S., Lelli F., Schombert J. M., 2016, *Phys. Rev. Lett.*, 117, 201101
- McGibbon R., Helly J., Schaye J., Schaller M., Vandenbroucke B., 2025, *J. Open Source Softw.*, 10, 8252
- Mondal C., Subramaniam A., George K., 2021, *JA&A*, 42, 50
- Moore B., 1994, *Nature*, 370, 629
- Navarro J. F., Eke V. R., Frenk C. S., 1996, *MNRAS*, 283, L72
- Navarro J. F., Frenk C. S., White S. D. M., 1997, *ApJ*, 490, 493
- Nobels F. S. J., Schaye J., Schaller M., Ploekinger S., Chaikin E., Richings A. J., 2024, *MNRAS*, 532, 3299
- Noguchi M., 1999, *ApJ*, 514, 77
- Oh S.-H., de Blok W. J. G., Brinks E., Walter F., Kennicutt R. C. Jr, 2011a, *AJ*, 141, 193
- Oh S.-H., Brook C., Governato F., Brinks E., Mayer L., de Blok W. J. G., Brooks A., Walter F., 2011b, *AJ*, 142, 24
- Oh S.-H. et al., 2015, *AJ*, 149, 180
- Oman K., 2024, *J. Open Source Softw.*, 9, 6860
- Oman K., 2025, *J. Open Source Softw.*, 10, 9278
- Oman K. A. et al., 2015, *MNRAS*, 452, 3650
- Oman K. A., Marasco A., Navarro J. F., Frenk C. S., Schaye J., Benítez-Llambay A., 2019, *MNRAS*, 482, 821
- Orr M. E. et al., 2020, *MNRAS*, 496, 1620
- Pedregosa F. et al., 2011, *J. Mach. Learn. Res.*, 12, 2825
- Peirani S. et al., 2017, *MNRAS*, 472, 2153
- Penston M. V., 1969, *MNRAS*, 145, 457
- Peschken N., Lokas E. L., Athanassoula E., 2020, *MNRAS*, 493, 1375
- Planck Collaboration XIII, 2016, *A&A*, 594, A13
- Ploekinger S., Richings A. J., Schaye J., Trayford J. W., Schaller M., Chaikin E., 2025, *MNRAS*, 543, 891
- Pontzen A., Governato F., 2012, *MNRAS*, 421, 3464
- Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T., 2003, *MNRAS*, 338, 14
- Read J. I., Iorio G., Agertz O., Fraternali F., 2016, *MNRAS*, 462, 3628
- Reines A. E., Greene J. E., Geha M., 2013, *ApJ*, 775, 116
- Reinhardt S., Krake T., Eberhardt B., Weiskopf D., 2019, *ACM Trans. Graph.*, 38, 189
- Ren T., Kwa A., Kaplinghat M., Yu H.-B., 2019, *Phys. Rev. X*, 9, 031020
- Rey M. P. et al., 2024, *MNRAS*, 529, 2379
- Richings A. J., Schaye J., Oppenheimer B. D., 2014a, *MNRAS*, 440, 3349
- Richings A. J., Schaye J., Oppenheimer B. D., 2014b, *MNRAS*, 442, 2780
- Rix H.-W., Zaritsky D., 1995, *ApJ*, 447, 82
- Robles V. H. et al., 2017, *MNRAS*, 472, 2945
- Rocha M., Peter A. H. G., Bullock J. S., Kaplinghat M., Garrison-Kimmel S., Oñorbe J., Moustakas L. A., 2013, *MNRAS*, 430, 81
- Roper F. A., Oman K. A., Frenk C. S., Benítez-Llambay A., Navarro J. F., Santos-Santos I. M. E., 2023, *MNRAS*, 521, 1316
- Sales L. V., Wetzel A., Fattahi A., 2022, *Nat. Astron.*, 6, 897
- Sands I. S. et al., 2024, preprint (arXiv:2404.16247)
- Santos-Santos I. M. E. et al., 2020, *MNRAS*, 495, 58
- Schaller M. et al., 2015, *MNRAS*, 451, 1247
- Schaller M. et al., 2024, *MNRAS*, 530, 2378
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Schaye J. et al., 2023, *MNRAS*, 526, 4978
- Schaye J. et al., 2025, preprint (arXiv:2508.21126)
- Schoenmakers R. H. M., Franx M., de Zeeuw P. T., 1997, *MNRAS*, 292, 349
- Sedov L. I., 1946, *J. Appl. Math. Mech.*, 10, 241
- Sellwood J. A., Sánchez R. Z., 2010, *MNRAS*, 404, 1733
- Shepard D., 1968, in Proc. 23rd ACM National Conference. ACM, p. 517
- Spekkens K., Giovanelli R., Haynes M. P., 2005, *AJ*, 129, 2119
- Spergel D. N., Steinhardt P. J., 2000, *Phys. Rev. Lett.*, 84, 3760
- Stilp A. M., Dalcanton J. J., Skillman E., Warren S. R., Ott J., Koribalski B., 2013, *ApJ*, 773, 88
- Stinson G., Seth A., Katz N., Wadsley J., Governato F., Quinn T., 2006, *MNRAS*, 373, 1074
- Swaters R. A., Madore B. F., van den Bosch F. C., Balcells M., 2003, *ApJ*, 583, 732
- Swaters R. A., Sancisi R., van Albada T. S., van der Hulst J. M., 2009, *A&A*, 493, 871
- Tacchella S. et al., 2019, *MNRAS*, 487, 5416
- Taylor G., 1950, *Proc. R. Soc. A*, 201, 159
- Trachternach C., de Blok W. J. G., Walter F., Brinks E., Kennicutt R. C. Jr, 2008, *AJ*, 136, 2720
- Trayford J. W. et al., 2026, *MNRAS*, 545, staf2040
- Valenzuela O., Rhee G., Klypin A., Governato F., Stinson G., Quinn T., Wadsley J., 2007, *ApJ*, 657, 773
- van den Bosch F. C., Swaters R. A., 2001, *MNRAS*, 325, 1017
- Vogelsberger M. et al., 2014, *MNRAS*, 444, 1518
- von Neumann J., Taub A. H., 1964, *J. Am. Stat. Assoc.*, 59, 981
- Walter F., Brinks E., de Blok W. J. G., Bigiel F., Kennicutt R. C. Jr, Thornley M. D., Leroy A., 2008, *AJ*, 136, 2563
- Zaritsky D. et al., 2013, *ApJ*, 772, 135

APPENDIX A: SPH FIELD SAMPLING IN A WARPED GAS DISC: A REPRESENTATIVE EXAMPLE

To complement the methodological description in Sections 2.2 and 2.4, here we show an example of how the warped disc mid-plane and associated SPH-sampled fields are constructed and interpreted, for a representative galaxy in our sample.

Fig. A1 shows the cumulative HI mass profile, $M_{\text{HI}}(< r)$, and the orientation of the gas angular momentum vector as a function of radius for a representative galaxy in our sample (HBT-HERONS track ID = 14481). The disc extent radius, $R_{\text{ext,HI}}$, is identified from the flattening of the mass profile, using the slope criterion outlined in Section 2.2, and is marked in both panels. The lower panel illustrates the radial evolution of the gas angular momentum direction in the global frame, revealing a modest but coherent warp between the half mass radius, $R_{\text{half,HI}}$, and the disc extent radius, $R_{\text{ext,HI}}$, and a significantly stronger twist at radii $r < 5$ kpc.

This radial variation – in some cases significantly more pronounced – in the orientation of the gaseous plane of rotation motivates the construction of a warped radius-dependent local frame. Crucially, our definition of $R_{\text{ext,HI}}$ – based on the point where the cumulative mass profile flattens – ensures that the number of gas particles per radial bin (or shell) remains sufficiently high within the disc. This mitigates noise in the angular momentum calculation, which could otherwise be dominated by a few outlier particles in low-density outer shells, thus improving the stability and accuracy of the local disc frame construction. Nevertheless, an interpolation of both the enclosed HI mass and gas angular momentum is preferable when evaluating slopes and constructing the geometry of the mid-plane.

Fig. A2 provides a visual illustration of the warped disc frame construction motivated by the angular momentum profiles in Fig. A1; the same representative galaxy is shown in both figures to clarify the transition from a radially varying angular momentum direction to a warped mid-plane. The left panel shows the warped disc mid-plane and sampling points in the global coordinate frame (X, Y, Z), with points coloured by the local gas




Figure A1. Example of radial profiles for a representative halo in our sample. *Top panel:* enclosed HI mass, $M_{\text{HI}}(< r)$, as a function of radius (solid light green line). *Bottom panels:* components of the normalized gas angular momentum vector in the global x , y , and z directions (left to right), as functions of radius. The solid light green line show the profiles computed directly from the simulation particles (see Section 2.2); solid light red lines show smoothed versions obtained via a cubic spline interpolation. The dash-dotted and solid vertical grey lines shown in both panels indicate the HI half-mass radius, $R_{\text{half, HI}}$, and the gas disc extent, $R_{\text{ext, HI}}$, respectively.

density, ρ . The *right panel* shows the same points deprojected into the local warped disc frame (X' , Y'), where by construction the mid-plane lies at $Z' = 0$. This highlights how the local coordinate frame smoothly follows the radial variation in the gas angular momentum direction, effectively ‘unwarping’ the disc morphology.

Importantly, here it is also possible to grasp the utility of our SPH-based continuous field sampling procedure: it enables evaluation of hydrodynamical quantities, such as the gas volume density, directly on the warped disc mid-plane. This would not be possible using only the raw gas particle data, which are

irregularly distributed and not confined to the warped plane. Thus, our post-processing approach provides a physically meaningful high-resolution sampling of continuous fields aligned with the true gas disc geometry. Note that in strongly disturbed systems the ‘mid-plane’ may still be ill-defined, as large warps or oscillations in the angular momentum profile prevent the existence of a singular stable configuration – indeed the entire ‘tilted ring’ approximation may break down. In such cases, our method still recovers the locally defined surface of maximum rotation, which represents the most meaningful reference for kinematic analysis.

TRACK ID 14481

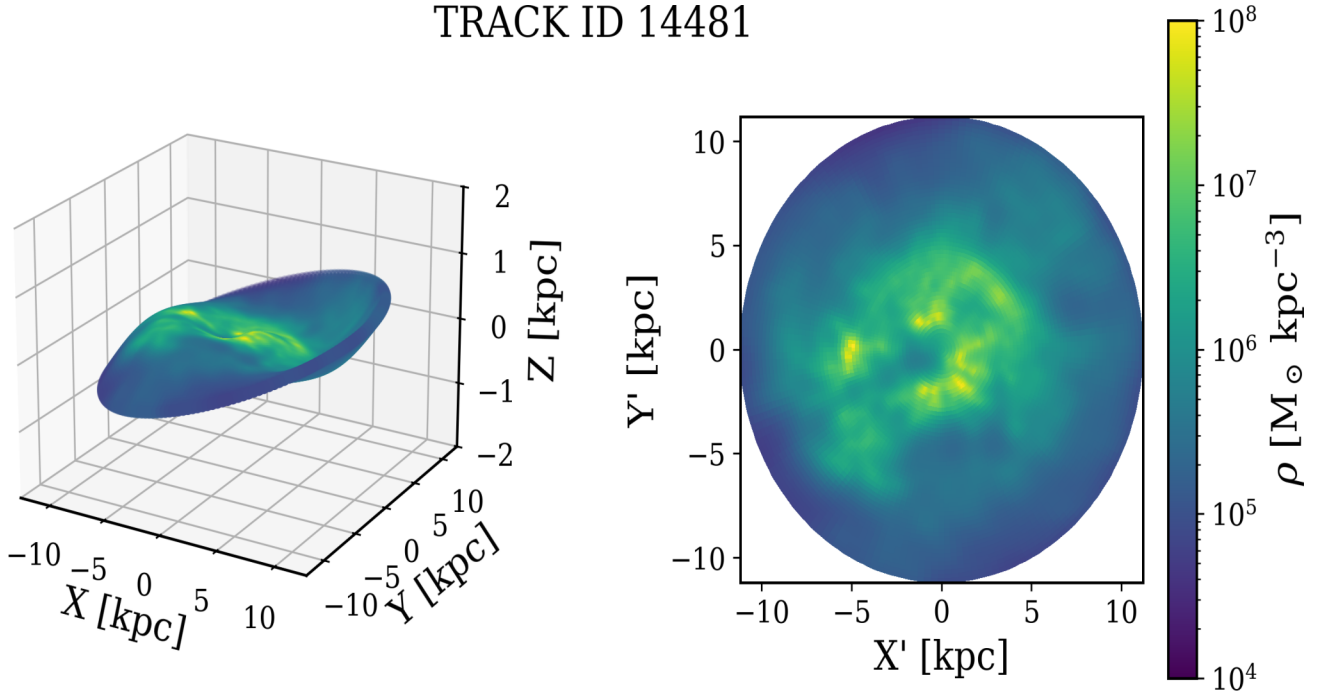


Figure A2. Example of a warped disc mid-plane and the corresponding gas volume density field at the sampling points. *Left panel:* 3D visualization of the warped disc mid-plane in the global coordinate frame (X, Y, Z), with sampling points coloured by the gas density, ρ , computed via the SPH-based continuous field sampling method described in Section 2.4. *Right panel:* the same sampling points and gas density field, deprojected on to the local warped disc frame (X', Y'), where $Z' = 0$ by construction. The same representative galaxy of Fig. A1 is shown here, clarifying the angular momentum-based frame construction.

APPENDIX B: ACCOUNTING FOR DISEQUILIBRIUM: SENSITIVITY TO THRESHOLD VARIATIONS

As discussed in Section 4.2, the quantitative contribution of individual physical drivers to the disequilibrium budget depends to some extent on the adopted criteria used to associate gas elements as affected by stellar feedback, AGN activity, and self-gravitating clumps. While our fiducial selections are physically motivated and empirically calibrated, it is important to assess the robustness of the resulting trends to reasonable variations in these somewhat arbitrary choices.

To this end, we repeat the full analysis using two alternative sets of tagging criteria. In a *permissive* configuration, we loosen the spatial apertures and selection thresholds such that the total disc area associated with each driver is increased by approximately a factor of 2 relative to the fiducial case. Conversely, in a *restrictive* configuration, we tighten the criteria so that the tagged area is reduced by a comparable factor. In all cases, the same physical definitions are retained, and only the extent of the tagged regions is modified.

Fig. B1 shows the resulting mass-weighted distributions of the disequilibrium parameter ε for each driver under the three configurations, compared to the global distribution. As expected, loosening the cuts increases the overlap with the near-equilibrium regime, while tightening them preferentially selects gas at larger $|\varepsilon|$. Importantly, however, the overall shapes of the distributions are preserved to a good degree. Self-gravitating clumps preferentially populate the positive side of the distribution, while AGN feedback remains strongly associated with large negative ε .

Stellar feedback exhibits the largest shape variation, particularly in the transition from the fiducial to the permissive configuration: in the latter case, the distribution becomes increasingly weighted towards the near-equilibrium regime, with a reduced relative contribution at negative ε . This behaviour indicates that overly permissive spatial selections (e.g. circular masks of order ≈ 2.5 kpc around young stellar particles or feedback-affected gas) begin to include substantial amounts of ambient or recovering gas. This effect naturally drives the mild differences in recovery fractions observed across configurations and directly motivates the calibration procedure adopted for the fiducial cuts.

The impact of these variations on the integrated disequilibrium budget is summarized in Fig. B2, which shows the fractional contribution of each driver – including overlaps – to the total gas mass tagged by the union of the three mechanisms. As anticipated, the relative ranking of the drivers is robust to changes in the tagging criteria: stellar feedback consistently dominates the budget, while AGN feedback and self-gravitating clumps contribute at a secondary but non-negligible level. The substantial overlap fractions are likewise largely preserved across configurations, reinforcing the conclusion that disequilibrium is frequently driven by multiple mechanisms acting simultaneously.

The only clear systematic difference induced by varying the thresholds is a redistribution of mass between the individual stellar feedback and AGN feedback components and their mutual overlap. This behaviour is expected, as the tagging criteria for both mechanisms are purely geometrical: increasing the characteristic radii naturally enlarges the region of intersection, reducing the fraction uniquely attributed to each driver (and vice versa for more restrictive cuts). By contrast, analogous trends are much weaker for self-gravitating clumps. This reflects the




Figure B1. Self-normalized mass-weighted distributions of the disequilibrium parameter ε , $W_I(\varepsilon)/\Sigma W_I$, for gas associated with each physical driver under different tagging criteria (summarized in Table B1). Each panel shows the distribution for one driver (stellar feedback, self-gravitating clumps, AGN feedback), comparing the fiducial selection (black) to permissive (orange) and restrictive (blue) variants. The grey shaded curve indicates the global gas distribution. The vertical dashed lines mark the quasi-equilibrium range, $|\varepsilon| \leq 0.2$. While the absolute normalization and degree of overlap with the low- $|\varepsilon|$ regime vary with the tagging criteria, the qualitative association between each driver and characteristic disequilibrium regimes is robust.

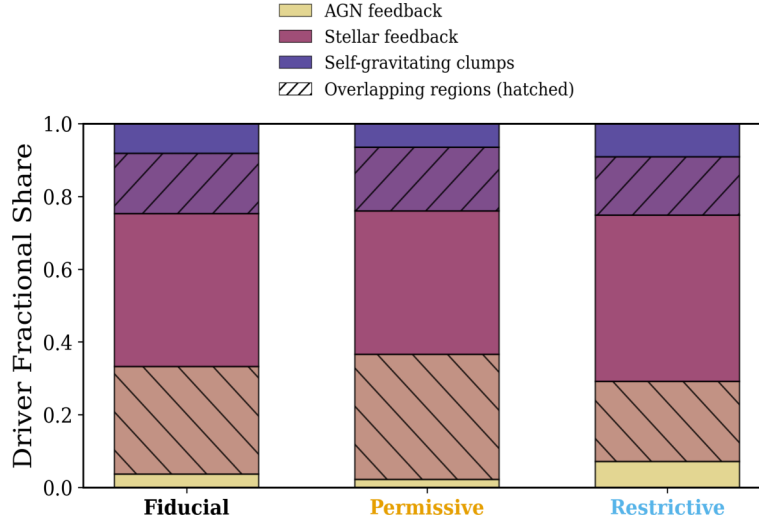


Figure B2. Fractional contribution of each physical driver to the total disequilibrium budget under fiducial, permissive, and restrictive tagging criteria (see Table B1). The coloured regions indicate the contributions from stellar feedback, AGN feedback, and self-gravitating clumps, while hatched regions mark overlaps between drivers as indicated in the legend. Although the absolute fractions vary with the adopted cuts in different ε regimes, the hierarchy of drivers and of overlap regions remain stable across all configurations.

fundamentally different nature of their selection: varying the threshold in $t_{\text{ff}}/t_{\text{orb}}$ does not simply expand or contract contiguous regions, but can instead select spatially disconnected structures. As a result, changes in the overlap between clumps and the other drivers are less straightforward and do not scale monotonically with the strictness of the cut (in fact, they are close to independent of them). Notably, the differences discussed here remain limited to the ≈ 5 –10 per cent level across configurations, indicating that they do not qualitatively affect the inferred disequilibrium budget.

For completeness, and in reference to the integrated percentages discussed in the final paragraphs of Section 4.2, the restrictive and permissive configurations, respectively, account for ≈ 20 and 50 per cent of the gas mass in the near-equilibrium regime ($|\varepsilon| \leq 0.2$), and ≈ 70 and 90 per cent in the extreme

disequilibrium regime ($|\varepsilon| \geq 1.0$). In the intermediate regime ($0.2 < |\varepsilon| < 1.0$), they encompass ≈ 35 and 65 per cent of the global gas mass. The systematic and differential response across $|\varepsilon|$ regimes further motivates the fiducial configuration as a balanced and optimized choice.

Taken together, these tests demonstrate that while the absolute normalization of the tagged mass varies with the adopted thresholds, the qualitative association between physical drivers and characteristic disequilibrium regimes – and their relative partitioning – remains remarkably stable. This robustness lends confidence to the fiducial tagging strategy and supports the broader conclusion that significant departures from equilibrium in the disc arise from the combined action of stellar feedback, AGN activity, and gravitational instability, rather than from any single dominant mechanism.

Table B1. Summary of the tagging criteria used to associate gas with the different disequilibrium drivers in the fiducial analysis and in the permissive and restrictive variants explored in Appendix B. The permissive and restrictive configurations are constructed to roughly double or halve, respectively, the total disc area tagged by each mechanism relative to the fiducial case, while preserving the same physical definitions.

Mechanism	Parameter description	Symbol [unit]	Restrictive	Fiducial	Permissive
Stellar feedback	Maximum distance from YS or from CCSNe kicked gas	$R_{\text{SF}}[\text{kpc}]$	≈ 0.423	$= 1.000$	≈ 2.450
Self-gravitating clumps	Free-fall to orbital time ratio	$t_{\text{ff}}/t_{\text{orb}}[\text{dimensionless}]$	≈ 0.094	≈ 0.111	≈ 0.131
AGN feedback	Maximum distance from active SMBH	$R_{\text{AGN}}[\text{kpc}]$	≈ 1.414	$= 2.000$	≈ 2.828

APPENDIX C: CLASSIFICATION ALGORITHM

As outlined in Section 5, our goal is to classify galaxies based on systematic deviations between the azimuthally averaged rotational velocity profile, $\langle v_\phi \rangle_\phi(r)$, and the corresponding circular velocity profile, $\langle v_c \rangle_\phi(r)$. These deviations are used as proxies for assessing the dynamical equilibrium state of each system.

To quantify such deviations in a statistically meaningful and comparative manner, we compute the following three metrics for each galaxy in our sample:

(i) L_1 norm: a global absolute measure of the discrepancy between the rotational and circular velocity profiles across the radial extent of the gas disc,

$$L_1 = \int_{r=r_0}^{r=R_{\text{ext,HI}}} |\langle v_c \rangle_\phi(r) - \langle v_\phi \rangle_\phi(r)| dr,$$

where r_0 is a small inner cutoff to exclude unresolved central regions.

(ii) Normalized Root Mean Square Error in Amplitude (NRMSE-Amplitude): a pointwise metric that emphasizes the fractional error between the two profiles,

$$\text{NRMSE} - \text{Amplitude} = \sum_i \sqrt{\frac{[\langle v_\phi \rangle_\phi(r_i) - \langle v_c \rangle_\phi(r_i)]^2}{\langle v_c \rangle_\phi(r_i)^2}},$$

where the sum runs over discrete radial bins r_i spanning the disc from r_0 to $R_{\text{ext,HI}}$ with $\Delta r = 0.2$ kpc spacing.

(iii) Normalized Root Mean Square Error in Slope (NRMSE-Slope): a scale-invariant measure comparing the local logarithmic slopes of the rotational and circular velocity profiles,

$$\text{NRMSE} - \text{Slope} = \sum_i \sqrt{\frac{[\alpha_{\langle v_\phi \rangle_\phi}(r_i) - \alpha_{\langle v_c \rangle_\phi}(r_i)]^2}{[1 + \alpha_{\langle v_c \rangle_\phi}(r_i)]^2}},$$

where

$$\alpha_A(r_i) = \left. \frac{d \ln A(r)}{d \ln r} \right|_{r=r_i}$$

denotes the logarithmic derivative of the quantity A evaluated at r_i . The normalization by $1 + \alpha_{\langle v_c \rangle_\phi}(r_i)$ ensures stability and prevents spurious amplification in regions where the slope becomes

small or negative. Again, the sum runs over discrete radial bins r_i spanning the disc from r_0 to $R_{\text{ext,HI}}$ with $\Delta r = 0.2$ kpc spacing.

These metrics are designed to capture both amplitude and shape mismatches between the rotational and circular velocity profiles, and they serve as quantitative inputs for the final step of our classification pipeline. Once all three parameters are computed for every galaxy in the sample, we apply a k -means clustering algorithm to the resulting feature space. This unsupervised learning method partitions the data set into distinct clusters – which we interpret as dynamical classes – based on similarities in the deviation metrics. The number of clusters is determined using standard internal validation criteria (i.e. the silhouette score), ensuring that the resulting classification reflects genuine structure in the data rather than arbitrary thresholds. These classes form the basis of the perturbed/equilibrium taxonomy discussed in the main text.

Fig. C1 provides both qualitative and quantitative insight into our classification methodology. The top panels contrast the kinematic behaviour of a stable disc (HBT-HERONS track ID = 6611) with that of a disturbed system (HBT-HERONS track ID = 58791). In the former, the rotational and circular velocity curves track each other almost exactly, indicative of a system close to centrifugal equilibrium. In the latter, systematic deviations in both amplitude and shape highlight significant dynamical disequilibrium. The bottom panel shows the full distribution of galaxies in the space defined by our three residual-based metrics.

The k -means algorithm successfully identifies well-separated groups corresponding to dynamically distinct populations. These groups, or clusters, serve as the basis for our perturbation classification framework. It is worth emphasizing that this scheme is designed to formalize the visual impression of how rotation curves depart from their expectation: the clustering provides a reproducible way to carve an otherwise continuous distribution of residual behaviours into classes that clarify the transition from ordered to disordered discs. In this context, limitations such as sample size or the use of a comparatively simple clustering method are not restrictive, since the purpose is organizational rather than predictive – any reasonable algorithm that captures the same large-scale structure in residual space would yield a broadly equivalent taxonomy.

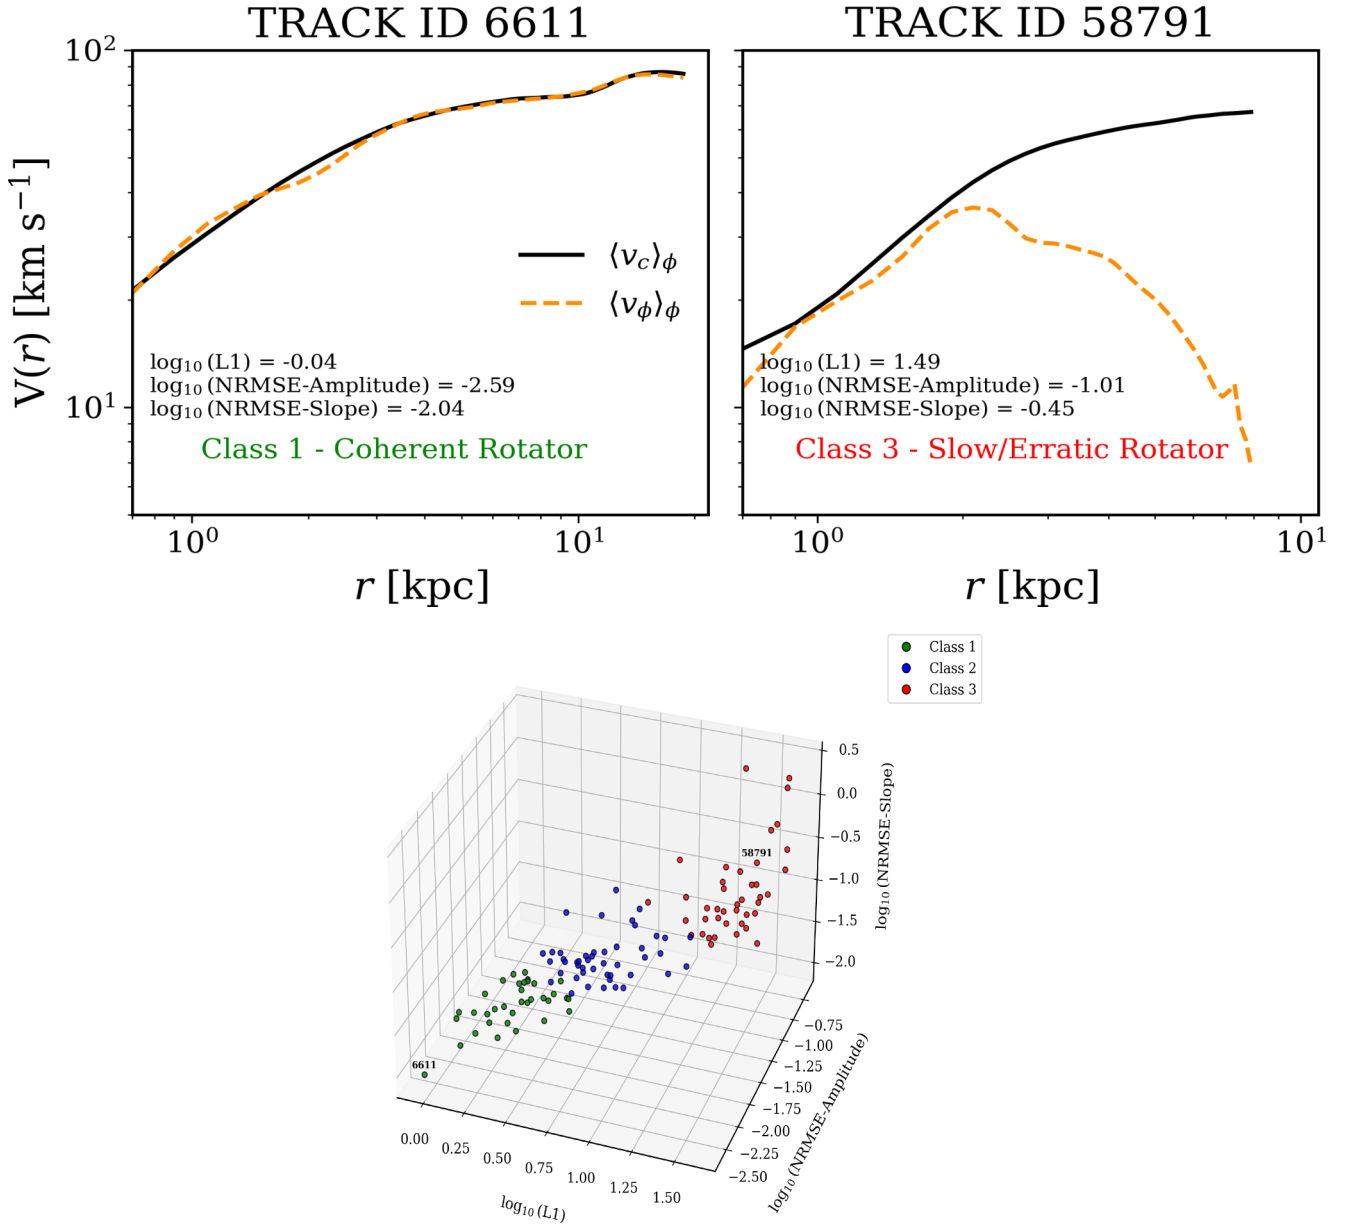


Figure C1. Illustrative examples and clustering result from our dynamical classification algorithm. *Top panels:* azimuthally averaged rotational velocity (dark orange/dashed lines) and circular velocity (black/solid lines) profiles as a function of radius for two representative galaxies in the sample. *Left:* a dynamically well-behaved system exhibiting excellent agreement between the two profiles across the entire disc extent. *Right:* a strongly perturbed system showing large and systematic residuals. The HBT-HERONS track IDs of these galaxies, and their associated values for the three classification metrics (L_1 , NRMSE-Amplitude, and NRMSE-Slope) are indicated in the legends and annotated text. *Bottom panel:* distribution of the full sample in the three-dimensional metric space, colour-coded by cluster assignment as determined via the k -means clustering algorithm. The resulting groups define the dynamical classes discussed in Section 5.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.