Orbital Support and Evolution of Flat Profiles of Bars (Shoulders)

Leandro Beraldo e Silva[1,2] , Victor P. Debattista[2] , Stuart Robert Anderson[2] , Monica Valluri[1] , Peter Erwin[3] ,
Kathryne J. Daniel[4,5] , and Nathan Deg[6]
[1] Department of Astronomy, University of Michigan, 1085 S. University Avenue, Ann Arbor, MI 48109 USA; lberaldo@umich.edu, lberaldoesilva@gmail.com
[2] Jeremiah Horrocks Institute, University of Central Lancashire, Preston PR1 2HE, UK
[3] Max Planck Institut für Extraterrestrische Physik, Giessenbachstrasse, D-85748 Garching, Germany
[4] Department of Astronomy & Steward Observatory, University of Arizona, Tucson, AZ 85721, USA
[5] Center for Computational Astrophysics, Flatiron Institute, New York, NY 10010, USA
[6] Department of Physics, Engineering Physics, and Astronomy, Queen's University, Kingston, ON, K7L 3N6, Canada
Received 2023 March 7; revised 2023 June 27; accepted 2023 July 13; published 2023 September 14

Abstract

Many barred galaxies exhibit upturns (shoulders) in their bar-major-axis density profile. Simulation studies have suggested that shoulders are supported by looped x_1 orbits, occur in growing bars, and can appear after bar buckling. We investigate the orbital support and evolution of shoulders via frequency analyses of orbits in simulations. We confirm that looped orbits are shoulder-supporting, and can remain so, to a lesser extent, after being vertically thickened. We show that looped orbits appear at the resonance $(\Omega_\varphi - \Omega_P)/\Omega_R = 1/2$ (analogous to the classical inner Lindblad resonance, and here called ILR) with vertical-to-radial frequency ratios $1 \lesssim \Omega_z/\Omega_R \lesssim 3/2$ (vertically *warm* orbits). *Cool* orbits at the ILR (those with $\Omega_z/\Omega_R > 3/2$) are vertically thin and have no loops, contributing negligibly to shoulders. As bars slow and thicken, either secularly or by buckling, they populate warm orbits at the ILR. Further thickening carries these orbits toward crossing the vertical ILR [vILR, $(\Omega_\varphi - \Omega_P)/\Omega_z = 1/2$], where they convert in-plane motion to vertical motion, become chaotic, kinematically hotter, and less shoulder-supporting. Hence, persistent shoulders require bars to trap new stars, consistent with the need for a growing bar. Since buckling speeds up trapping on warm orbits at the ILR, it can be followed by shoulder formation, as seen in simulations. This sequence supports the recent observational finding that shoulders likely precede the emergence of BP-bulges. The python module for the frequency analysis, naif, is made available.

Unified Astronomy Thesaurus concepts: Galactic bar (2365); Galaxy dynamics (591); Orbits (1184); Galaxy bars (2364); N-body simulations (1083)

1. Introduction

Stellar bars are present in ≈60%–70% of nearby spiral galaxies (Eskridge et al. 2000; Menéndez-Delmestre et al. 2007; Sheth et al. 2008; Nair & Abraham 2010; Erwin 2018). Barred galaxies have surface-brightness profiles along the bar-major-axis, which are traditionally classified as either exponential or as having, on top of an exponential profile, a nearly flat part in the bar outskirts (on both sides)—hereafter the "shoulders" (e.g., Elmegreen & Elmegreen 1985; de Carvalho & da Costa 1987; Elmegreen et al. 1996; Prieto et al. 2001; Gadotti et al. 2007; but see Erwin et al. 2023 for an updated classification scheme).

Elmegreen et al. (1996), Regan & Elmegreen (1997), and Elmegreen et al. (2011) found that shoulders are generally present in profiles of early-type barred spiral galaxies, while late-type galaxies normally have bars with exponential profiles. In a sample of 144 barred galaxies from the Spitzer Survey of Stellar Structure in Galaxies (S⁴G; Sheth et al. 2010), Kim et al. (2015) found that massive and bulge-dominated galaxies tend to have flat bar profiles, while bulgeless galaxies tend to have exponential profiles, in accordance with early predictions based on N-body simulations (Combes & Elmegreen 1993). The same trends were found by Kruk et al. (2018) in a sample of 3461

barred galaxies from SDSS, while Erwin et al. (2023) found that the presence of shoulders is more strongly correlated with stellar mass. Shoulders in the bar-major-axis profile have also been observed in M31 (Athanassoula & Beaton 2006) and in several simulation studies (e.g., Athanassoula & Misiriotis 2002; Bureau & Athanassoula 2005; Anderson et al. 2022, hereafter Paper I).

Figure 1 shows examples of bar-major-axis profiles for the galaxies NGC 3681 and NGC 4340 from S⁴G (Sheth et al. 2010) and for the simulation SD1 (described in Section 2) at 7 Gyr (dotted black/red) and at 10 Gyr (dashed–dotted blue/orange). The bars are aligned with the x-axis and for the simulation we select all star particles within $|y|/a_{bar} < 0.145$, where a_{bar} is the bar length—see Section 2 for details. For the observed galaxies, we show 1 pixel wide profiles along the bar-major-axis from Spitzer 3.6 μm images. NGC 3681 has a pure-exponential profile, while NGC 4340 and the model SD1 have profiles with shoulders (red), as detected by Paper I.

The galaxy NGC 4340 was chosen as an example of a bar profile with particularly pronounced shoulders, and these look much more pronounced than the ones in the model SD1 at 7Gyr. The shoulder strength, estimated as the excess mass of the shoulder over a purely exponential profile, normalized by the total mass within the shoulder region, typically grows over time, as shown in Paper I. In fact, we see that the shoulders of SD1 get stronger at 10 Gyr (blue/orange dashed–dotted lines), although its length (in units of a_{bar}) is approximately constant and smaller than in the observed galaxy NGC 4340.
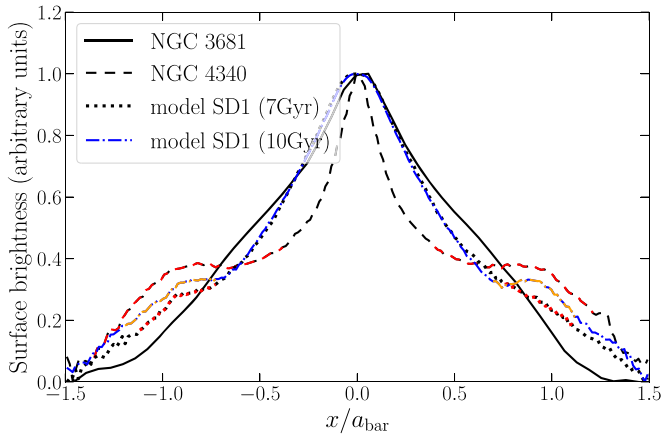
Figure 1. Normalized observed bar-major-axis brightness profiles for the Spitzer 3.6 μm band for NGC 3681 and NGC 4340, in arbitrary units (S^{4}G; Sheth et al. 2010). Also shown are the bar-major-axis surface density profiles for simulation SD1 at 7 Gyr (dotted black) and at 10 Gyr (dashed–dotted blue) —see Section 2 and Figure 3. NGC 3681 has a pure-exponential profile, while NGC 4340 and model SD1 have shoulders (red and orange), as detected in Paper I.

Table 1
Selection Times (t_s), Bar Speed Parameter $\mathcal{R} = R_{\mathrm{CR}}/a_{\mathrm{bar}}$, Bar Length a_{bar}, Number of Selected Star Particles and Whether the Model Has Gas or a Strongly Buckling Bar

Model	t_s	$\mathcal{R}$	a_{bar} kpc	N	Gas?	Buckling?
SD1	7 Gyr	1.99	6.58	106032	N	N
SD1S	7 Gyr	1.31	3.78	46973	N	N
HG1	8 Gyr	1.15	2.46	177426	Y	N
4	5 Gyr	1.60	6.85	136488	N	Y, 3.8 Gyr

Notes. Model SD1 and model 4 have reasonably strong shoulders. Model HG1 has weak shoulders, and SD1S has a barely evolving bar and no shoulders. The simulation naming convention is the same as in Paper I.

In principle, profiles with shoulders might simply signal the presence of a strong and long bar. However, exploring a diverse set of simulations, Paper I found only a weak dependence on bar strength or bar length, while Erwin et al. (2023) found no evidence of such correlation in their observational sample. Importantly, Paper I demonstrated that shoulders tend to be present in galaxies with growing bars, although their results suggest that this is a necessary but not a sufficient condition. Furthermore, Paper I found that in secularly evolving (nonbuckling) bars, once formed, shoulders tend to persist. On the other hand, in bars that buckle, shoulders tend to appear right after the buckling, and to be erased (at least temporarily) if a second buckling happens. Finally, Paper I suggested that shoulders are mostly supported by looped x_1 orbits.

In this paper, we further investigate the shoulders in bar-major-axis profiles of galaxies, focusing on their orbital support and the time evolution of these orbits. The main questions we aim to answer are: how and why do some bars develop shoulders? Which orbits support the shoulders? Why do bars need to grow in order to retain the shoulders once formed? What is behind the role of buckling events in creating or erasing shoulders?

We answer these questions by selecting particles from N-body simulations, integrating their orbits in the "frozen" potentials of different snapshots, and performing frequency analyses of these orbits. The N-body simulations are described in Section 2. In Section 3, we explain the procedures for orbit integration and frequency analyses. The resulting frequency maps and stacked density profiles for different orbital groups, as well as their time evolution, are presented in Section 4. We demonstrate the somewhat surprising result that, once shoulders have formed, vertical resonances act as the main culprit for erasing them, shedding light on the suggestion of Paper I that, in order for shoulders to be present for long times, the bar needs to be growing and trapping new stars. In Section 5 we discuss our results, and the conclusions are summarized in Section 6.

2. Simulations

Here we describe the N-body simulations explored in this paper. These simulations are chosen to represent a diverse subset of the models analyzed in Paper I. Briefly, one model (SD1, our fiducial model) has clear and persistent shoulders appearing around 4 Gyr; one model (SD1S) does not develop shoulders; model HG1 is a fully self-consistent star-forming N-body+SPH simulation with weak shoulders; and Model 4 has a bar that buckles around 3.8 Gyr, developing persistent shoulders after that. These simulations are described in detail in Paper I, but we summarize them here—see Table 1.

Model SD1 is a pure N-body model, built using a modified version of GALACTICS (Kuijken & Dubinski 1995; Widrow & Dubinski 2005) that generates disks using a Sérsic surface density profile:

$$\Sigma_d = \frac{M_d}{2\pi n R_d^2 \Gamma(2n)} e^{-(R/R_d)^{1/n}} \mathrm{sech}^2(z/z_d), \qquad (1)$$

where M_d is the disk mass, R_d is the scale length, n is the Sérsic index, and Γ is the gamma function. The Sérsic disk has $M_d = 5.74 \times 10^{10} \, M_\odot$, $R_d = 0.265$ kpc, $z_d = 0.25$ kpc and $n = 2.05$. The model has a Hernquist dark matter (DM) halo, defined with GALACTICS parameters $\sigma_h = 550$ km s^{-1}, $R_h = 30$ kpc, $\alpha = 1$, and $\beta = 4$. It has 5×10^6 DM particles and 4.4×10^6 disk particles. In this, and in the other pure N-body simulations described below, the softening length is $\epsilon = 50$ pc and $\epsilon = 100$ pc for star and DM particles, respectively.

Model SD1S suppresses most secular bar growth by setting the halo of model SD1 in full prograde rotation (Debattista & Sellwood 2000; Long et al. 2014; Collier et al. 2018). This results in a large halo spin ($\lambda \simeq 0.091$), rare in cosmological simulations (Bullock et al. 2001). Our goal is to contrast this model with its counterpart SD1 whose bar is growing, so we are not concerned by its fully rotating halo being unrealistic.

HG1 is a star-forming simulation described in Cole et al. (2014), Gardner et al. (2014), and Debattista et al. (2017). It starts out with a gas corona but no stars; hence, the evolution is totally self-consistent. It is evolved with the smooth particle hydrodynamics code GASOLINE (Wadsley et al. 2004). The simulation uses high resolution ($\epsilon = 50$ pc for the gas, and gas particles with initial mass $2.7 \times 10^4 \, M_\odot$) and after being evolved for 10 Gyr has $\sim 1.1 \times 10^7$ stellar particles, and a total stellar mass of $\sim 6.5 \times 10^{10} M_\odot$. It uses the gas cooling, star formation, and stellar feedback prescriptions of Stinson et al. (2006). The DM halo consists of 5×10^6 particles with $\epsilon = 103$ pc, where 4.5×10^6 of them have mass $8.5 \times 10^4 \, M_\odot$, and $1.7 \times 10^6 \, M_\odot$ for the remainder.

Finally, Model 4 is a disk N-body model (no gas or star formation)—see also Debattista et al. (2020). It is a baryon-dominated disk Milky Way–like model, again set up using
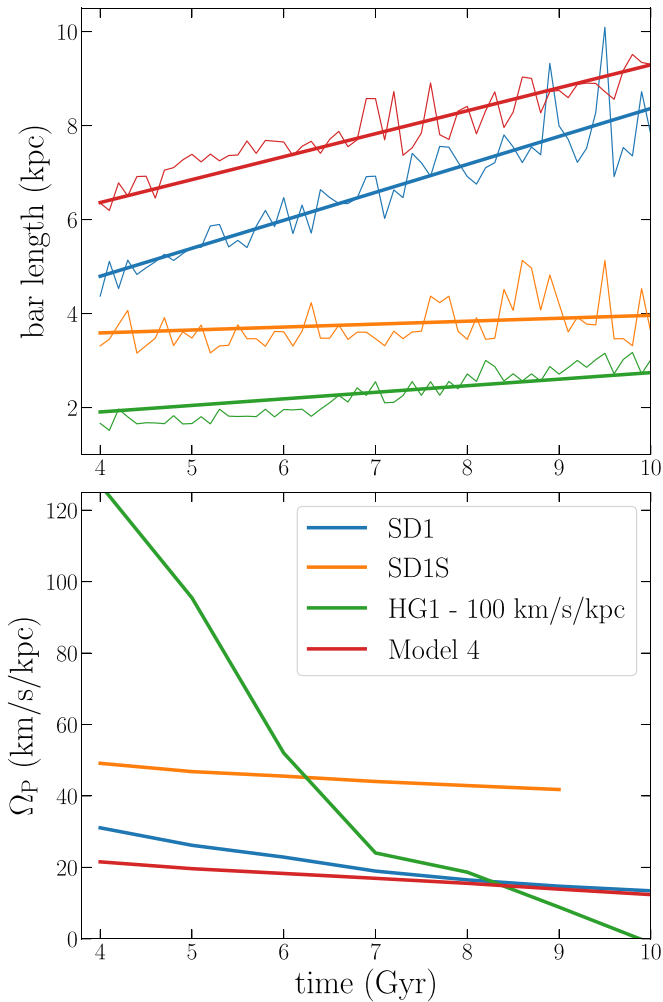
Figure 2. Evolution of bar lengths (top) and pattern speeds (bottom) for our simulations. In the upper panel, thick lines show the linear fit to the bar length evolution (thin), which provides $a_{\rm bar}$ used in our analysis. In all models, the bars slow down (but less significantly for SD1S and Model 4), and grow (but less significantly for models SD1S and HG1). For HG1, the values are shifted down by 100 km s^{-1} kpc^{-1}.

GALACTICS. It has a Navarro–Frenk–White DM halo (Navarro et al. 1996), and an exponential disk with an isothermal vertical profile. The disk has a central radial velocity dispersion of 128 km s^{-1}, a disk scale radius of 2.4 kpc, and a scale height of 150 pc. It has 6×10^6 particles in the disk and 4×10^6 in the halo.

All N-body models were evolved using PKDGRAV (Stadel 2001) for 10 Gyr. In what follows, we take model SD1 as our fiducial model, because of its clear trends illustrating the main results of this paper, although model HG1 may be more realistic.

2.1. Simulation Global Properties

We use PYNBODY (Pontzen et al. 2013) to read, center, and align the simulations such that, at each snapshot analyzed, the bar is initially along the x-axis. We estimate the gravitational potentials Φ and integrate orbits using AGAMA (Vasiliev 2019). Smooth potentials due to star and gas particles (when present) are separately modeled as triaxial distributions and estimated with the "cylspline" potential type, while the halo contribution is assumed to be axisymmetric and is estimated with the

"multipole" type. For each snapshot analyzed, we then estimate the total potential.

We compute the bar length (radius) $a_{\rm bar}$ as the mean of the cylindrical radius at which the amplitude of the $m = 2$ Fourier moment of the (x, y)-plane surface density distribution reaches half its maximum value after its peak, and the radius at which the phase of the $m = 2$ component deviates from a constant by more than $10°$—see Paper I for further details. We fit a straight line to the bar length as a function of time and use the fitted function to represent $a_{\rm bar}$ at each snapshot. Both estimated and fitted bar lengths are shown in the upper panel of Figure 2 as thin and thick lines, respectively, and we see that a straight line is a reasonable fit.

The bar pattern speeds $\Omega_{\rm P}$ are computed as the time derivative of the phase angle of the $m = 2$ Fourier mode of star particles, as done by Debattista & Sellwood (2000). We use 21 snapshots spaced by 5 Myr around the snapshots of interest. We then fit a straight line to the bar phase angle as a function of time in each of these time intervals, where the phase is assumed to increase monotonically with time. The pattern speeds are shown in Figure 2 (bottom panel) as functions of time. As often observed in simulations (e.g., Hernquist & Weinberg 1992; Debattista & Sellwood 1998; O'Neill & Dubinski 2003; Athanassoula 2003), the bars slow down significantly over the time span of the simulation—see also Fragkoudi et al. (2021) for a recent study of cosmological simulations and Hamilton et al. (2023) for a recent theoretical treatment.

The bar in the Model SD1 starts with a slightly higher $\Omega_{\rm P}$ in comparison to Model 4, but the pattern speeds in both models evolve to very similar, almost constant, values toward the end of the simulations. The model HG1 has a very centrally concentrated mass distribution, with a steep rotation curve in the central parts. Thus it has a very fast bar, which decelerates quickly without increasing in length significantly. However, these differences do not seem to alter the general trends of the shoulders we investigate.

Figure 3 shows edge-on (top row) and face-on (middle row) views of the four models at the times when particles are selected for orbit integration. The coordinates are normalized by the respective bar lengths. Colors represent surface densities in log-scale. The x–y (x–z) contours show nine (seven) equally spaced percentiles, from the 85th–99th (84th–99th). The edge-on view shows clear signs of a BP-bulge in all models, although the "X" shape is weaker in models SD1S and HG1.

In the face-on view, we see the bars along the x-axis, and hints of density excess at their outskirts (except for SD1S, and barely for HG1) shown by the contours. The shaded rectangles represent the shoulder region, defined as $|y|/a_{\rm bar} < 0.145$ and x-edges detected with the algorithm of Paper I. Model SD1S has no shoulders, and we define its shoulder region using the same fractions of the bar length as those of SD1. In all models with shoulders, these are detected in the transition region between the thick and thin parts of the bar. The bottom row shows the surface density along the stripe defined by $|y|/a_{\rm bar} < 0.145$ in each model, with red snippets showing where shoulders are detected.

Table 1 shows the number of star particles selected in the shoulder region, the bar length $a_{\rm bar}$ and bar speed parameter $\mathcal{R} = R_{\rm CR}/a_{\rm bar}$, where $R_{\rm CR}$ is the corotation radius (Debattista & Sellwood 2000) at the selection times, and whether the model has gas or a strongly buckling bar (and when it buckles).
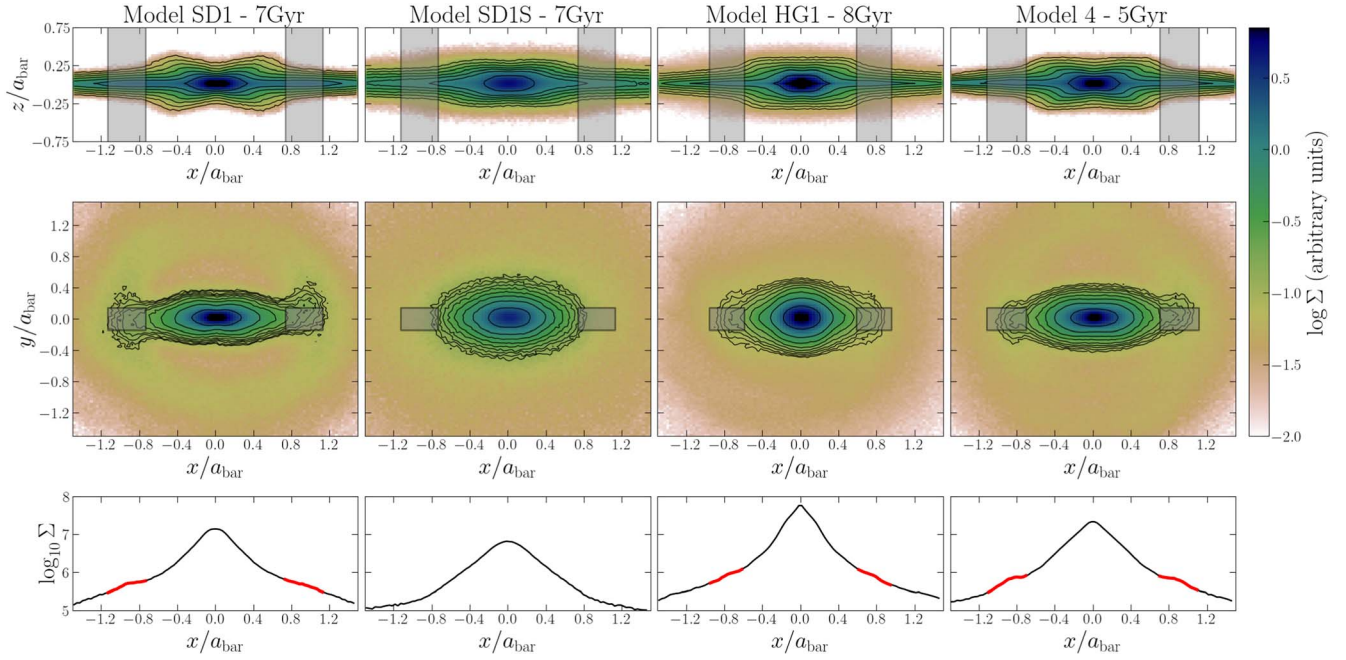
Figure 3. Edge-on (top row) and face-on (middle row) normalized surface density maps for the four simulations at the selection times. Shaded rectangles show the shoulder regions $|y|/a_{\rm bar} < 0.145$ and x-edges identified in Paper I (except for SD1S, which does not have detected shoulders), where particles are selected. The bottom row shows the surface density along the stripe $|y|/a_{\rm bar} < 0.145$, with the red snippets showing the shoulders (where detected).

Models SD1S and HG1 have fast bars ($\mathcal{R} \lesssim 1.4$) at these snapshots, while SD1 and Model 4 have slow bars ($\mathcal{R} \gtrsim 1.4$).

2.2. Epicyclic Frequencies, Resonance Names, and Selection Region

Figure 4 shows the rotation curve $\Omega(R)$ (solid black), and the curves $\Omega(R) - \kappa(R)/2$ (dashed) and $\Omega(R) - \kappa(R)/4$ (dotted), where

$$\Omega(R) \equiv \sqrt{\frac{1}{R}\frac{\partial \Phi}{\partial R}}, \tag{2}$$

and

$$\kappa(R) \equiv \sqrt{\frac{\partial^2 \Phi}{\partial R^2} + \frac{3}{R}\frac{\partial \Phi}{\partial R}} \tag{3}$$

is the radial epicyclic frequency (Binney & Tremaine 2008), evaluated along the bar-major-axis, at the times when star particles are selected. The times are chosen when clear and stable shoulders are present (except for SD1S, which does not have shoulders), and which are far from buckling events (in Model 4).

The bar pattern speeds at these times are shown as thin red lines, with the thick part representing the bar length. Gray rectangles show the area within the mean shoulder x-edges, as detected in Paper I, within which particles are selected. For simulation SD1S, we adopt the same fraction of the bar length as that of the shoulder extent of SD1 at the same time.

It is interesting to note that for all simulations with shoulders, these start outside the outer Inner Lindblad radius (where $\Omega - \kappa/2 = \Omega_P$). This is the region where the (main bar-supporting) x_1 orbits can start developing loops at their ends (Contopoulos & Papayannopoulos 1980; Binney & Tremaine 2008), and this is the first hint of the importance of

these orbits to shoulders, as we confirm below. On the other hand, this might suggest that the frequency ratio $(\Omega_\varphi - \Omega_{\rm P})/\Omega_R = 1/2$, where Ω_φ and Ω_R are the azimuthal and radial frequencies for each orbit, is not relevant for the shoulders. Conversely, Figure 4 shows that the shoulders typically appear near the ultraharmonic radius (where $\Omega - \kappa/4 = \Omega_P$), which might suggest an important role of the resonance $(\Omega_\varphi - \Omega_{\rm P})/\Omega_R = 1/4$. However, our frequency analysis (Section 4) reveals a fundamental role for the resonance $(\Omega_\varphi - \Omega_{\rm P})/\Omega_R = 1/2$ in supporting the shoulders, and no important role for the resonance $(\Omega_\varphi - \Omega_{\rm P})/\Omega_R = 1/4$.

The resonance $(\Omega_\varphi - \Omega_{\rm P})/\Omega_R = 1/2$ is the analog of the classical inner Lindblad resonance (ILR), but applies to orbits of arbitrary eccentricity. Keeping in mind that resonances are defined in the space of frequencies (rather than physical space), in this paper we will use actual frequency ratios to define the main resonances. For instance, the resonance $(\Omega_\varphi - \Omega_{\rm P})/\Omega_R = 1/2$ will be called the ILR, and the resonance $(\Omega_\varphi - \Omega_{\rm P})/\Omega_R = 1/4$ will be called ultraharmonic resonance (UHR). When identifying the approximate locations of these resonances, based on the epicyclic frequencies, we will refer to their radii, e.g., the inner Lindblad radius and the ultraharmonic radius—see Athanassoula (2003) for a discussion on the pros and cons of different ILR definitions.

With the bar aligned with the x-axis, we select star particles within $|y|/a_{\rm bar} < 0.145$ and in the x-intervals where shoulders are identified (with no cut in z)—see Figures 3 and 4. Since the detected shoulders are not exactly symmetric around $x = 0$, we define the "shoulder region" averaging their edges on the $x < 0$ and $x > 0$ sides. This selection is convenient for identifying orbits supporting the shoulder, but also because, being in the bar outskirts, this region is in the transition to the disk, and particles supporting both structures are selected. Table 1 shows the total number of star particles selected in each model.
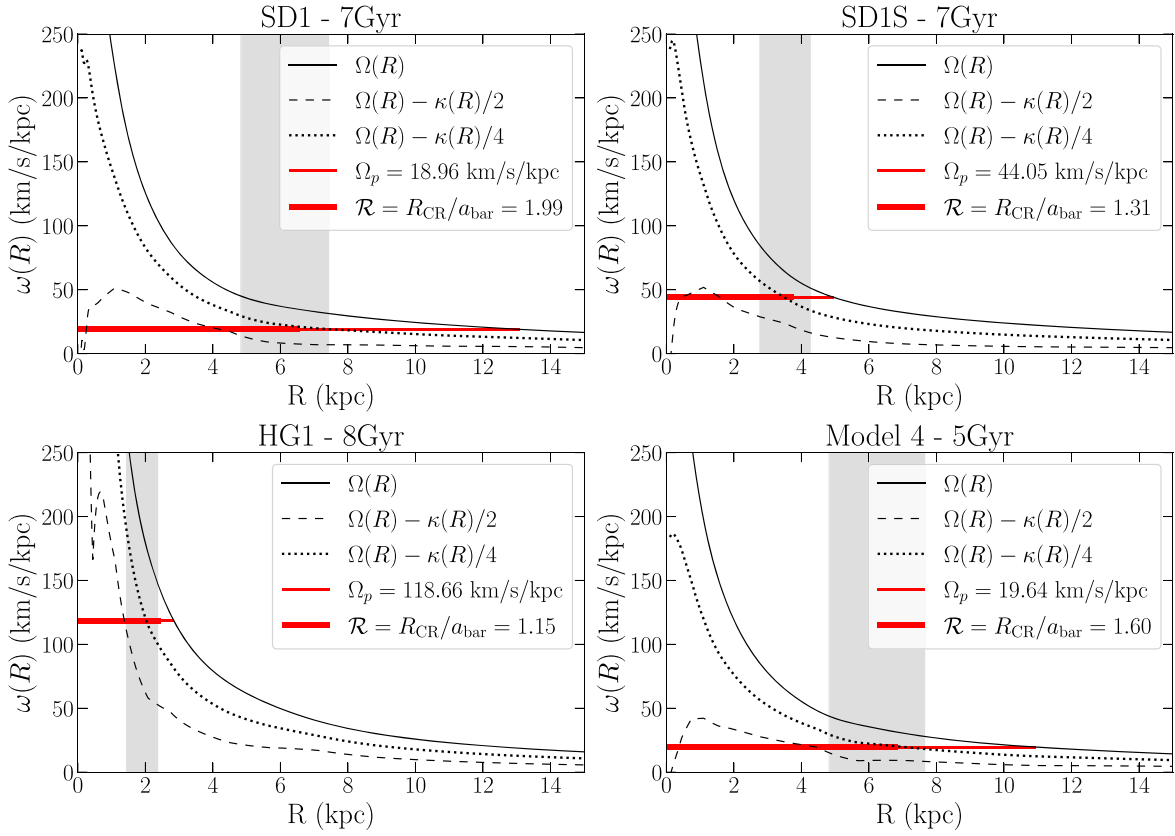
Figure 4. Rotation curves $\Omega(R)$ (solid black), $\Omega(R) - \kappa(R)/2$ (dashed), and $\Omega(R) - \kappa(R)/4$ (dotted) for the different simulations at the time when star particles are selected. Gray areas show the average radial interval where shoulders are detected (SD1S has no shoulders, and we define this region using the same fractions of the bar length as those of SD1), within which we select particles in the simulations. Horizontal red lines show the bar pattern speeds, with the thick part representing the bar lengths, at the respective times. Model SD1 and Model 4 have slow bars ($\mathcal{R} \gtrsim 1.4$), while models SD1S and HG1 have fast bars. The shoulder region is around the ultraharmonic radius where $\Omega - \kappa/4 = \Omega_P$.

Figure 5 illustrates the selection (shoulder) region in terms of the effective potential in the rotating frame,

$$\Phi_{\text{eff}} = \Phi - \frac{1}{2}\, \Omega_P^2 x^2, \qquad (4)$$

evaluated along the bar-major-axis x. The points show the Jacobi integral

$$E_J = E - \Omega_P L_z, \qquad (5)$$

which is conserved in the frame rotating with angular speed Ω_P, where E is the energy in the inertial frame and L_z is the angular momentum. These points are color-coded by the eccentricity $\text{ecc} \equiv (r_{\text{apo}} - r_{\text{per}})/(r_{\text{apo}} + r_{\text{per}})$, where r_{apo} and r_{per} are the apo- and pericenter radii. These quantities are calculated after the orbit integration explained in Section 3.1. The horizontal lines show the effective potential evaluated at the corotation radius, R_{RC} (solid thin), at the ultraharmonic radius (dotted), and at the outer inner Lindblad radius, R_{ILR} (dashed). In agreement with Figure 4, the selected particles have $E_J > \Phi_{\text{eff}}(R_{\text{ILR}})$. We note that for the simulations with shoulders, the most bound orbits have high eccentricities (thus small r_{per}) and these orbits extend right up to the ultraharmonic radius. Additionally, there are orbits with $E_J > \Phi_{\text{eff}}(R_{\text{CR}})$ and reasonably large eccentricities, which can visit both the disk and the bar regions—see Section 4.1.3.

3. Orbits

The coordinates of particles selected at each snapshot set the initial conditions for orbit integration in the frozen potential rotating with the corresponding pattern speed (see Athanassoula 2013). This means that in the rotating frame the potential is static for the whole integration time. We then estimate the frequencies of motion in the three directions (as detailed below), producing a frequency map for each snapshot analyzed.

3.1. Orbit Integration

For particles selected at a time t_s, we integrate orbits in the frozen potentials of five snapshots, $[t_s - 2, t_s - 1, t_s, t_s + 1, t_s + 2]$ Gyr, with a relative error tolerance of 10^{-15} for each coordinate, which is the best we could achieve for all particles at all snapshots. When plotting orbits, we use coordinates in the bar rotating frame, but the frequency analysis uses coordinates in the inertial frame, as detailed below.

We integrate orbits for a (pseudo)time of 5 Gyr (storing 10^5 points per orbit), which is a good compromise between two limiting time scales: a lower limit of, at least, tens of orbital periods for precise frequency estimates (Valluri & Merritt 1998), and an upper limit given by the predictability horizon for chaotic orbits with finite precision integration (e.g., Roy 1991). Note that this is not the real time of the simulation, but a parameter in the determination of instantaneous orbital properties in the frozen potential. This integration time represents at least $20 \times T_{\text{circ}}(E)$, where $T_{\text{circ}}(E)$ is the period
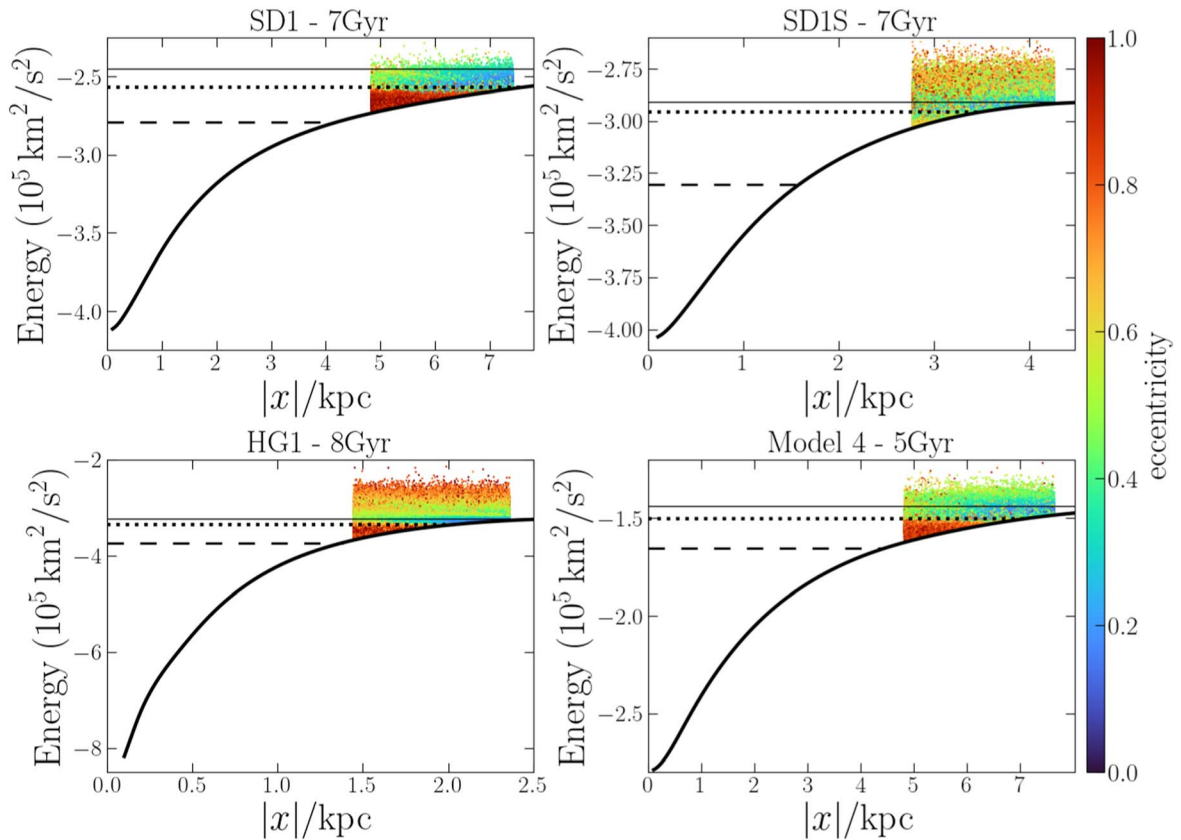
Figure 5. Effective potential along the bar-major-axis (thick black solid) and Jacobi integral (points) of the particles selected in the shoulder region from the different simulations, at the selection time. Also shown are the effective potential values at L1 (corotation, thin solid), at the ultraharmonic radius (dotted), and outer inner Lindblad radius (dashed).

of a circular orbit with energy E, for almost all particles, and this is typically the case for all selections. Furthermore, among the cylindrical coordinates (R, φ, z), the azimuthal period is typically the longest, so the integration time is also large enough for the other coordinates.

3.2. Frequency Analysis

We employ the Numerical Analysis of Fundamental Frequencies (NAFF) algorithm originally proposed by Laskar et al. (1992) and further developed by Valluri & Merritt (1998) and Valluri et al. (2010) to extract the leading frequencies of orbits. Regular orbits are quasiperiodic, so the time-series of each coordinate q can be approximated by

$$q(t) = \sum_{k=1}^{k_{max}} a_k e^{i\omega_k t}, \tag{6}$$

where a_k is a complex amplitude, and ω_k is an integer combination of the fundamental frequencies Ω_1, Ω_2, and Ω_3 (time-derivatives of the canonical angle variables). The power spectrum of a discrete Fourier transform (DFT) of a regular orbit has discrete peaks at frequencies ω_k (spectral lines; see Binney & Spergel 1982). Typically, although not always, the frequency with the largest amplitude (hereafter the *leading frequency*) corresponds to the fundamental frequency.

In nonintegrable potentials, quasiperiodicity cannot be assumed a priori, which might hinder the extraction of meaningful fundamental frequencies. However, most realistic galactic potentials, although nonintegrable, still host a large number of regular or weakly chaotic orbits, a picture (formally) supported by the KAM theorems (Kolmogorov 1954; Arnol'd 1963; Moser 1962). Additionally, in nonintegrable potentials, resonance trapping significantly affects the system's dynamical evolution. Once the fundamental frequencies are identified, resonances stand out in a plot showing ratios of the frequencies, i.e., a frequency map (Laskar 1990, 1999), such as the one shown in Figure 6.

We have written a new pure-python implementation of the NAFF algorithm, which we dub naif,[7] based on the FORTRAN implementation of Valluri & Merritt (1998). A summary of the algorithm is presented in Appendix A. Variants of this method have been developed and applied in different contexts (e.g., Binney & Spergel 1982; Carpintero & Aguilar 1998; Price-Whelan et al. 2016; Valluri et al. 2016; Beraldo e Silva et al. 2019; Dodd et al. 2022; Lucey et al. 2023). The analysis involves important decisions on implementation details, such as the coordinate system (e.g., Cartesian or cylindrical), the use of real or complex time-series, the reference frame, and the window function. In the study of bars (e.g., Portail et al. 2015; Valluri et al. 2016; Gajda et al. 2016; Parul et al. 2020; Smirnov et al. 2021; Sellwood & Gerhard 2020), it is common to adopt Cartesian coordinates (or mixed Cartesian and cylindrical) in the bar rotating frame, where the frequency of oscillation along the major axis, Ω_x, is used as a proxy for the azimuthal frequency $\Omega_\varphi - \Omega_P$. Furthermore, spatial coordinates as (real) time-series are often

---

[7] The documentation and installation instructions can be accessed at http://naif.readthedocs.io/en/latest/index.html.
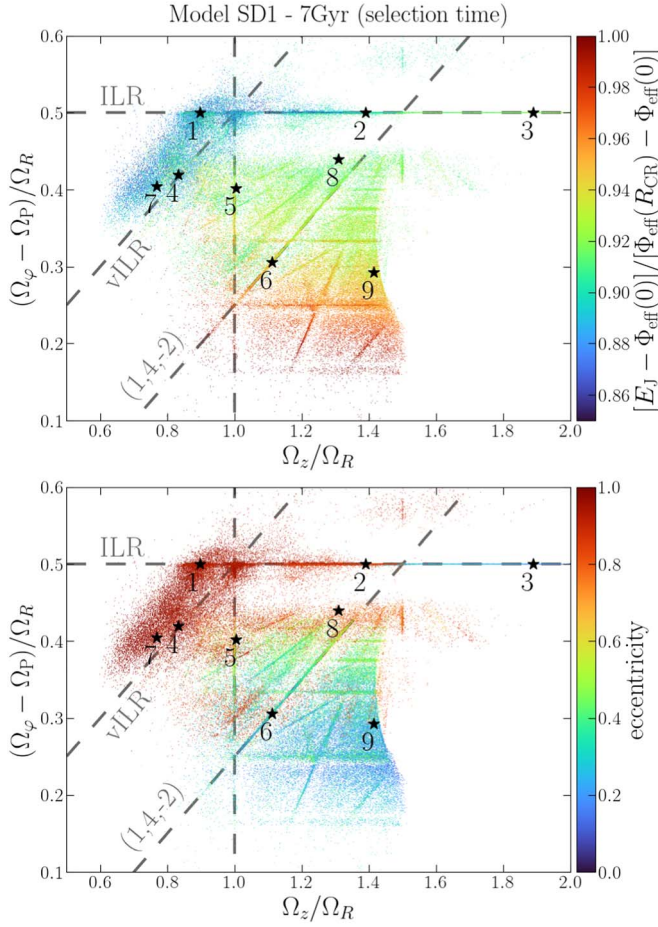
Figure 6. Frequency map for the shoulder region of model SD1 at 7 Gyr, with dashed lines showing prominent resonances. We indicate the ILR, the vertical ILR (vILR) and the 3D resonance $\Omega_R + 4(\Omega_\varphi - \Omega_P) - 2\Omega_z = 0$. The upper panel is color-coded by the normalized Jacobi integral, Equation (8). The upper-left corner is dominated by typical orbits in the central region, and moving toward the lower-right corner, orbits get more disk-like. The *vILR cloud* is the high-eccentricity region around the vILR. Numbered stars are the sample orbits shown in Figures 7, 9, and 10.

used. See Section 5.1 for a discussion on different choices and techniques.

To get some guidance in these decisions, and to test our algorithm, in Appendix B we compare our frequency estimates with analytical expressions for orbits in the isochrone potential. We compare the use of real and complex time-series (Appendix A), and show that the radial frequency Ω_r is estimated with relative errors $|\delta\Omega_r/\Omega_r| \approx 10^{-6}$ in both cases. For the azimuthal frequency Ω_φ, we show that the complex time-series $f_\varphi = \sqrt{2|L_z|}(\cos\varphi + i\sin\varphi)$ (Papaphilippou & Laskar 1996, 1998), where L_z is the z-component of the angular momentum, does offer advantages, compared with $f_\varphi = \varphi$: it improves the accuracy, and it does not require setting $\text{sign}(\Omega_\varphi) = \text{sign}(L_z)$, which is the case for the real time-series. After excluding the precessing frequency $\Omega_r - \Omega_\varphi$, which is the leading one for orbits with large apocenter, we estimate Ω_φ with relative errors $|\delta\Omega_\varphi/\Omega_\varphi| \lesssim 10^{-3}$ for almost all orbits.

All frequency estimates in this paper use the Hanning window, i.e., $p = 1$ in Equation (A2). In the following analysis, for an easier identification of important known resonances, we use cylindrical coordinates (R, φ, z) stored in the inertial frame of the galaxy, rather than the bar rotating frame. This allows for

identification of the corotation resonance, which would require an infinite integration time in the rotating frame—see Appendix C. Based on the discussions above, we use as input in the frequency analysis the complex time-series:

$$
\begin{aligned}
f_R &= R + iv_R \\
f_\varphi &= \sqrt{2|L_z|}(\cos\varphi + i\sin\varphi) \\
f_z &= z + iv_z,
\end{aligned}
\tag{7}
$$

where v_R and v_z are the radial and vertical velocities, respectively. For the radial and vertical components, we take the absolute value of the leading frequency, while we retain the original sign for Ω_φ. Since our cuts do not select particles with significantly large apocenter distances, we expect the number of orbits with misidentified fundamental frequencies to be negligible.

4. Results

4.1. Frequency Maps: Overview

Having demonstrated the accuracy of our numerical procedure, we move on to the analysis of orbits in the frozen potentials of our N-body simulations. Figure 6 shows the frequency map obtained for star particles selected in the shoulder region of model SD1 at 7 Gyr (see Figures 3 and 4). The upper panel is color-coded by the normalized Jacobi integral (see Figure 5)

$$
\frac{E_J - \Phi_{\text{eff}}(0)}{\Phi_{\text{eff}}(R_{\text{CR}}) - \Phi_{\text{eff}}(0)},
\tag{8}
$$

where Φ_{eff} and E_J are defined in Equations (4) and (5), respectively, and R_{CR} is the corotation radius. The bottom panel is color-coded by the eccentricity, and numbered stars indicate the example orbits presented in Figures 7, 9, and 10.

Given the typical shape of rotation curves (see Figure 4), we can anticipate that orbits with high $(\Omega_\varphi - \Omega_P)/\Omega_R$ in Figure 6 are typical of the galaxy central regions. Moving down in Figure 6, we go toward the end of the bar. For fast bars, this happens near the corotation, at $\Omega_\varphi - \Omega_P = 0$, and transitioning to the disk we have $\Omega_\varphi - \Omega_P \lesssim 0$, corresponding to retrograde stars in the bar frame. Evaluating the functions $\Omega(R)$ and $\kappa(R)$, Equations (2)–(3), for this model at the average outer shoulder edge (the outermost region of our selection), we obtain $(\Omega - \Omega_P)/\kappa \approx 0.25$, in broad agreement with the lower values (≈ 0.15) on the y-axis of Figure 6, and with the fact that the shoulder outer edge in this model matches the ultraharmonic radius—see Figure 5.

The ratio Ω_z/Ω_R is a proxy for the vertical thinness of the orbit: it is significantly larger than unity for typical orbits in the disk (in the solar neighborhood, $\Omega_z/\Omega_R \approx 2$; Binney & Tremaine 2008), and can be $\Omega_z/\Omega_R \lesssim 1$ in thicker or more spheroidal regions. Thus, the upper-left corner in Figure 6 is dominated by typical orbits that can visit the very central regions of the galaxy, and orbits get more disk-like when we move toward the lower-right corner, in agreement with the eccentricity trends seen in the bottom panel. With the radial $\kappa(R)$, Equation (3), and vertical $\nu(R) \equiv \sqrt{\partial^2\Phi/\partial z^2}$ epicyclic frequencies evaluated at the outer shoulder edge, we get $\nu/\kappa \approx 1.70$, also in broad agreement with the rightmost limit of Figure 6 (neglecting orbits at the ILR, which are well inside the
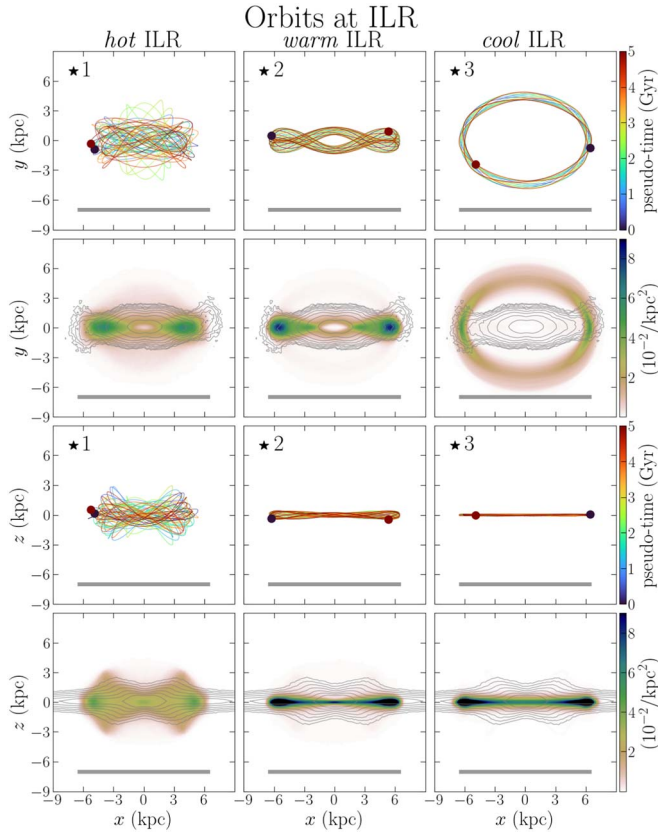
Figure 7. First and third rows: three typical orbits at the ILR selected at different Ω_z/Ω_R ranges from simulation SD1 at 7 Gyr. Orbits are color-coded by time, with dots showing the initial and final locations. Numbered stars identify these orbits in Figure 6. Second and fourth rows: stacked surface density maps of all orbits at the ILR in the respective Ω_z/Ω_R ranges. The orbits and stacked density maps are shown in the bar rotating frame, and the thick gray lines show the bar size. Thin gray lines are the same contours shown in Figure 3 for the total mass. The vertically warm orbits ($1 < \Omega_z/\Omega_R < 3/2$) produce a marked density peak in the outskirts of the bar. A similar, but diluted, shape is also produced by the hot orbits ($\Omega_z/\Omega_R < 1$).

bar region and thus less well approximated by the linear expressions).

We identify the main regions in the frequency map that we explore in detail below. The horizontal line at $(\Omega_\varphi - \Omega_P)/\Omega_R = 1/2$ characterizes the x_1 orbits, which are known to be the backbone orbits of bars (Contopoulos & Grosbol 1989; Skokos et al. 2002; Athanassoula 2003). In principle, x_2 orbits, which are aligned with the bar minor-axis, can also be present along this resonance. However, these orbits are only expected within the inner Lindblad radius (or between two inner Lindblad radii; see van Albada & Sanders 1982) and, as Figure 4 shows, we select particles outside this region.

As already mentioned in Section 2.1, we refer to the frequency ratio $(\Omega_\varphi - \Omega_P)/\Omega_R = 1/2$ as the ILR (see, e.g., Athanassoula 2003; Weinberg & Katz 2007). We emphasize that this differs from the classical definition of the ILR based on the epicyclic frequencies, Equations (2)–(3), or on the existence of x_2 orbits (van Albada & Sanders 1982), but can be seen as their generalizations, in the frequency space, to orbits of arbitrary eccentricities. The UHR, $(\Omega_\varphi - \Omega_P)/\Omega_R = 1/4$, is also well populated in Figure 6. Also highlighted are the vertical ILR (vILR), where $(\Omega_\varphi - \Omega_P)/\Omega_z = 1/2$; the resonance $\Omega_z/\Omega_R = 1$; and the 3D resonance $\Omega_R + 4(\Omega_\varphi - \Omega_P) - 2\Omega_z = 0$, denoted by (1,4,-2).

We call the cloud of highly eccentric orbits that are to the left and above the vILR the *vILR cloud*. We also note two other diffuse populations more clearly identified in the bottom panel: a large blue/green region with low eccentricities and permeated by several resonances (the region of orbits 6 and 9), and an orange stripe below the ILR, with $0.65 \lesssim \mathrm{ecc} \lesssim 0.85$ and $0.4 \lesssim (\Omega_\varphi - \Omega_P)/\Omega_R \lesssim 0.45$ (the region of orbit 8).

We now illustrate the orbital morphologies in these different regions, noting that Figure 6 does not represent the whole bar region, but only orbits that are selected at its outermost parts at a particular time. The whole menu of orbits in bars has been extensively studied in other works and is beyond the scope of this paper (see, e.g., Athanassoula et al. 1983; Sparke & Sellwood 1987; Combes et al. 1990; Patsis et al. 2002; Skokos et al. 2002; Valluri et al. 2016; Patsis & Athanassoula 2019). For a review, see Sellwood & Wilkinson (1993).

4.1.1. Orbital Shapes at the ILR

We first select orbits at the ILR by requiring $|(\Omega_\varphi - \Omega_P)/\Omega_R - 1/2| < 0.0025$. By inspection, this cut is a good choice to ensure minimal contamination from neighboring orbits. Unless otherwise stated, we use the same tolerance in selections of other resonances. Figure 6 suggests grouping the orbits at the ILR into three groups, which we will refer to as vertically *cool*, *warm*, and *hot* orbits at the ILR (in some instances we will simply refer to them as cool, warm, and hot orbits for conciseness). Vertically cool orbits are those at the ILR and with $\Omega_z/\Omega_R > 3/2$. They are the least bound and are generally less eccentric. A typical orbit from this group is shown in the right column (first and third rows) of Figure 7. It has an elliptical shape that is elongated in the direction of the major axis, and it is vertically very thin (see third row).

Transitioning to the middle section of the ILR, the vertically warm orbits have $1 < \Omega_z/\Omega_R < 3/2$. They are more bound (lower E_J) and eccentric (ecc ≈ 0.85). The upper-middle panel in Figure 7 shows a typical orbit in this region (orbit 2). These orbits still avoid the center and have characteristic loops at their ends, which have previously been identified in both analytical (e.g., Contopoulos 1978; Papayannopoulos & Petrou 1983) and numerical studies (e.g., Sparke & Sellwood 1987; Petersen et al. 2021). In particular, Contopoulos (1988) found that these orbits typically appear near the ultraharmonic radius, which corresponds to our selection region—see Figure 4. We refer to these as "looped" orbits.

Furthermore, for all simulations analyzed in this paper, we find that this transition from elliptical-like to this looped shape occurs at $\Omega_z/\Omega_R \approx 3/2$, suggesting some kind of influence of vertical resonances. In fact, we note that the crossing of the ILR with $\Omega_z/\Omega_R = 3/2$ is the convergence point of several 3D resonances in Figure 6. In the Sections that follow, we demonstrate the fundamental role of this looped morphology in building the shoulders in the bar-major-axis density profiles.

Finally, the vertically hot orbits are those at the ILR and with $\Omega_z/\Omega_R < 1$ (see Figure 6). They are slightly deeper in the potential, and thus have slightly smaller $r_{\rm apo}$. Despite this, they are significantly more eccentric (ecc ≈ 0.98; see also Figure 5), because they can get very close to the center, as illustrated in the upper-left panel in Figure 7 (orbit 1).

Figure 8 shows a normalized histogram of $\log_{10}(r_{\rm per}/{\rm kpc})$ for cool, warm, and hot orbits at the ILR of the model SD1 at 7 Gyr. It is clear that a significant fraction of the hot orbits can get as close as 10–100 pc to the center. Moreover, Figure 7
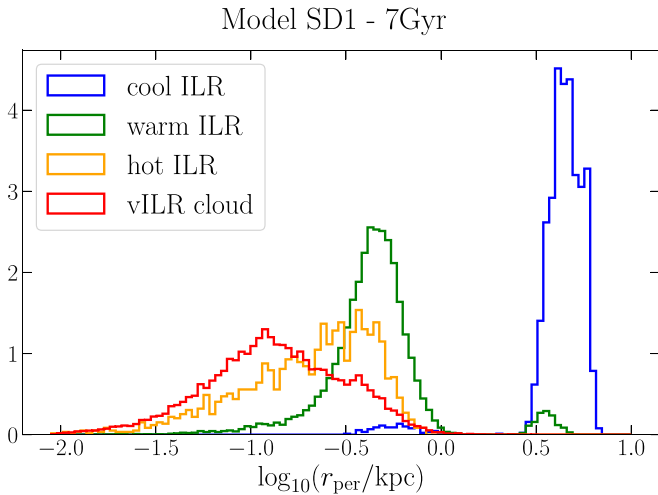
Figure 8. Normalized histogram of $\log_{10}(r_{\mathrm{per}}/\mathrm{kpc})$ for orbits along the ILR, in the different Ω_z/Ω_R intervals, and for orbits at the vILR cloud—see the main text. A large fraction of the vertically hot orbits at the ILR (and the vILR cloud in general) can get as close as $\sim$10–100 pc from the center. The small secondary peaks for cool and warm orbits at the ILR are due to cross-contamination of these orbits near the edge $\Omega_z/\Omega_R = 3/2$.

shows that typical orbits in this region do not have the coherent shapes observed before, scattering more in the x–y plane (a signature of their chaotic nature, as demonstrated in Section 4.2). Finally, their vertical thickness is significantly larger than in orbits 2 and 3.

The second and fourth rows in Figure 7 show the orbit-averaged surface densities of orbits at the ILR in each of the three groups in Ω_z/Ω_R. These plots are 2D histograms built by stacking all points of a single orbit (a time average) and all orbits (an ensemble average), normalized by the total number of orbits in each region to emphasize characteristic shapes rather than relative contributions of each group. We see that the sample orbits are representative of the populations.

The thin gray lines in the second and fourth rows are the same face-on and edge-on contours shown in Figure 3 for the whole star-particle distribution. We anticipate the role of the looped morphology of vertically warm orbits at the ILR to produce the density excess in the bar outskirts. A similar morphology is also present for the hot orbits at the ILR, but the density excess appears more diluted—see Sections 4.4 and 4.6. We also note the X-shape in the x–z plane of the orbit-averaged density of the hot orbits, suggesting their contribution to the BP-bulge visible in the contours.

4.1.2. Orbital Shapes at Vertical Resonances

We now select orbits at the three vertical resonances highlighted in Figure 6 and mentioned before: the vILR, the resonance $\Omega_z/\Omega_R = 1$, and the resonance $(1, 4, -2)$. In this selection we exclude the intersections with the ILR and orbits above it by requiring $(\Omega_\varphi - \Omega_{\mathrm{P}})/\Omega_R - 1/2 < -0.01$. Typical orbits (orbits 4, 5, and 6 in Figure 6) are shown in the first and third rows of Figure 9, while the panels below show the stacked density maps of all of the orbits selected along these resonances. From right to left, we observe a similar trend of orbits becoming more elongated, less regular, and vertically thicker, with orbit 4 reaching all the way to the center. This orbit, at the vILR, has the typical "banana" shape in the x–z plane (Pfenniger & Friedli 1991). Overall, orbits along these resonances imprint noticeable features in the x–z stacked
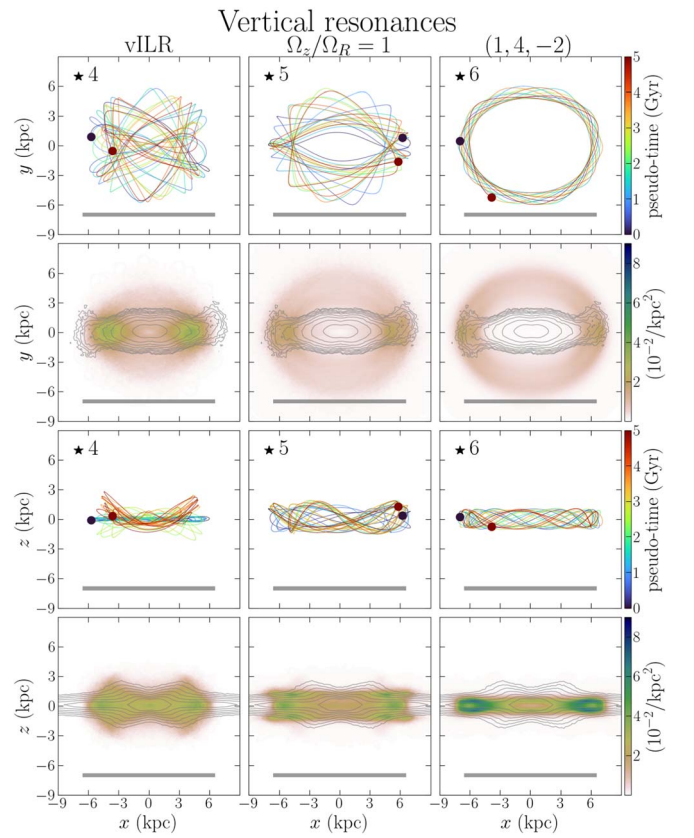

Figure 9. Similar to Figure 7, but for typical orbits selected at three vertical resonances (titles), and requiring $(\Omega_\varphi - \Omega_{\mathrm{P}})/\Omega_R < 1/2$. The numbered stars indicate these orbits in Figure 6.

density maps, but not as distinctly concentrated in the x–y plane as those shown in Figure 7.

4.1.3. Orbital Shapes at the vILR Cloud and Other Regions

Finally, we select orbits in three additional regions previously mentioned: the vILR cloud (with sample orbit 7); the stripe with intermediate eccentricities (with sample orbit 8), requiring $0.65 < \mathrm{ecc} < 0.85$ and not belonging to any of the previously selected regions; and the low-eccentricity region (with sample orbit 9), requiring $\mathrm{ecc} < 0.65$ and not belonging to any of the previously selected regions.

Figure 10 shows the sample orbits and the orbit-averaged surface density maps for these regions. From right to left, i.e., in order of decreasing Ω_z/Ω_R, we note again the tendency of orbits to become more elongated and less regular, finally approaching the center and becoming significantly thicker. Orbit 9 circulates a little bit farther out in the bar (in comparison to orbits 6 and 3). Its orbital group has $(\Omega_\varphi - \Omega_{\mathrm{P}})/\Omega_R \lesssim 0.4$, and these resemble disk orbits in the vicinity of the bar. Orbit 8 has a peculiar shape, with a strong resemblance to the looped orbits at the ILR (e.g., orbit 2), but oscillating between this and a close-to resonant triangular shape. Its orbital group has $0.4 \lesssim (\Omega_\varphi - \Omega_{\mathrm{P}})/\Omega_R \lesssim 0.45$, and it produces an excess in the outskirts of the bar similar to that produced by the looped orbits, but less prominent. These orbits are located near the convergence of several 3D resonances and near the ILR and, as will be demonstrated in Section 4.2, are chaotic. These orbits have $0.4 \lesssim (\Omega_\varphi - \Omega_{\mathrm{P}})/\Omega_R \lesssim 0.45$, similar to the location of a peak already observed by

Martinez-Valpuesta et al. (2006) and Smirnov et al. (2021). Some of these orbits (but not all) have $E_J > \Phi_{eff}(R_{RC})$ and reasonably large eccentricities (see Figure 5), allowing them to visit both the bar and the disk region, as those identified in N-body simulations by Sparke & Sellwood (1987) and Pfenniger & Friedli (1991). Figure 11 is similar to Figure 5, but only shows points for the orbits in this group. We see that in models SD1S and HG1, a large fraction of these orbits have $E_J > \Phi_{eff}(R_{RC})$.

An interesting aspect of Figures 7, 9, and 10 is the similarity of the orbit-averaged density maps in the left column of each figure, especially regarding the X-shape in the x–z projection supporting the BP-bulge (contours). All of these orbits are within the vILR cloud, now specifically defined as $(\Omega_\varphi - \Omega_P)/\Omega_z - 1/2 > -0.01$, i.e., including hot orbits at the ILR. Figure 8 shows that approximately half of these orbits can get as close as 10–100 pc to the galactic center.

Having shown the typical spatial distribution of orbits selected in the shoulder region, we now compare the frequency maps of the different simulations.

4.2. Frequency Maps: Comparing Simulations

Here we compare the frequency maps for orbits selected in the shoulder regions for the four simulations (see Section 2), at the snapshots where particles are selected (see Figure 4). Figure 12 shows the results for the model SD1, with the same data as in Figure 6, but with the upper panel color-coded by z_{max} and the bottom panel by the frequency drift $\log \Delta\Omega$, where

$$\Delta\Omega = \max_{i=1,2,3} \left| \frac{\Omega_i(T_2) - \Omega_i(T_1)}{\Omega_i(T_1)} \right|, \qquad (9)$$

and $\Omega_i(T_2)$ and $\Omega_i(T_1)$ are the frequencies estimated in two separate halves of the time-series (Valluri et al. 2010). Since regular orbits conserve frequencies, a large $\log \Delta\Omega$ indicates a chaotic orbit. The side-panels show histograms of the frequency ratios color-coded by the mean z_{max} (or $\log \Delta\Omega$) in bins of width 0.0025.

Figures 13–15 show the equivalent frequency maps for the models SD1S, HG1, and Model 4. The most noticeable feature in the four models is the well-populated ILR, producing a very prominent peak in the histogram to the right (note the logarithmic scale). This strong prominence of the ILR could be a selection effect since we chose particles that are predominantly part of the bar. However, in a random selection of 10^6 star particles in HG1 at 10 Gyr, we identify more than half of the particles at the ILR; such a strong ILR has also been observed in previous works (e.g., Athanassoula 2003). In our selection, within ± 0.0025 (± 0.01) from the ILR, we identify 19.8% (22.2%) of the orbits for model SD1, 4.5% (4.5%) for SD1S, 13.5% (18.3%) for HG1, and 15.1% (18.1%) for Model 4, meaning the ILR is well populated in all of the models, except for model SD1S.

In the models SD1S and HG1, we identify the corotation resonance at $(\Omega_\varphi - \Omega_P)/\Omega_R = 0$. This is in agreement with SD1S and HG1 having fast bars, and the selection region, always in the outermost parts of the bar, being close to corotation in these models–see Figure 4. We emphasize that identifying the corotation is only possible when using coordinates in the inertial frame—see Figure C1 for a comparison with the bar frame. We also note, in all simulations, big gaps around the corotation and ILR, testifying
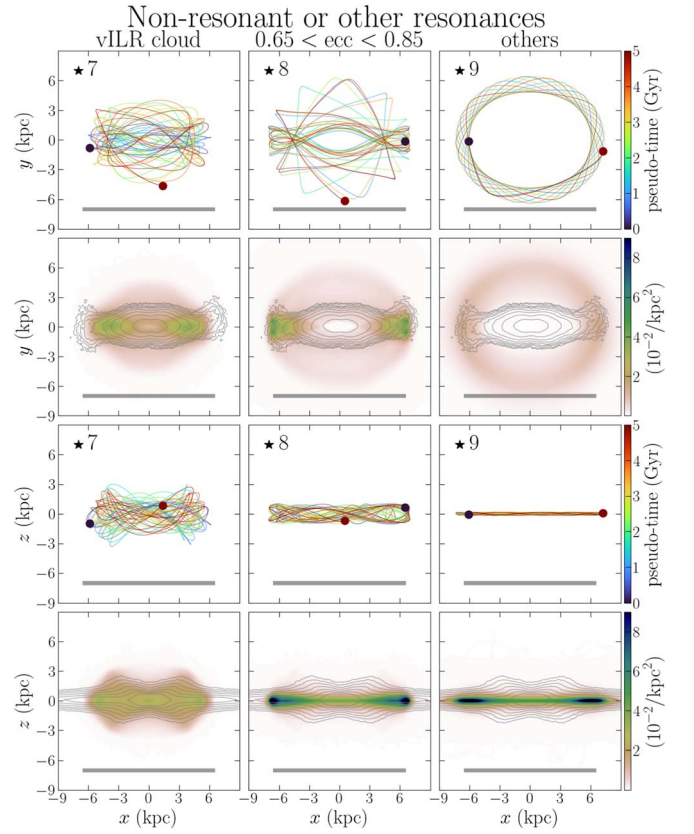

Figure 10. Similar to Figures 7–9, but for orbits in the vILR cloud and other locations in the frequency map (see the main text). The numbered stars indicate these orbits in Figure 6.

to their influence: orbits that get too close may be trapped or scattered by them.

Another remarkable feature in all models, except SD1S, is the vILR cloud, with significantly large z_{max} and $\log \Delta\Omega$. In HG1, which is the only model with gas, the z_{max} at the vILR cloud is smaller than in the other models, although $\log \Delta\Omega$ is still comparably large. In a certain sense, the vILR clearly separates the orbits that support the thin and thick parts of the bar. The looped orbits at the (vertically warm) ILR mostly support the thin part, and the orbits at the vILR cloud mostly support the thick part—see also Figure 16.

Finally, we see a plethora of additional resonances, particularly in HG1 (Figure 14). Orbits at the main resonances (and away from any crossing resonances) tend to have smaller $\log \Delta\Omega$, i.e., to be regular, as expected since these resonances are stable. On the other hand, orbits in areas with multiple strong resonances, such as the crossing of the ILR and the vILR, or on the floor below the ILR (the group of orbit 8 in Figure 6), tend to be chaotic, as is generally the case at resonance crossings (Chirikov 1979). These results illustrate the richness of the dynamics in the region between the bar and the disk (see, e.g., Contopoulos 1981), and also the power and precision of the frequency map technique we employ.

4.3. More on Morphologies along the ILR

We explore in greater depth the different morphologies for orbits along the ILR. Figure 16 shows, for the different models, the vertical excursion z_{max} (top panels) versus Ω_z/Ω_R for orbits at the ILR ($|(\Omega_\varphi - \Omega_P)/\Omega_R - 1/2| < 0.01$) at the selection
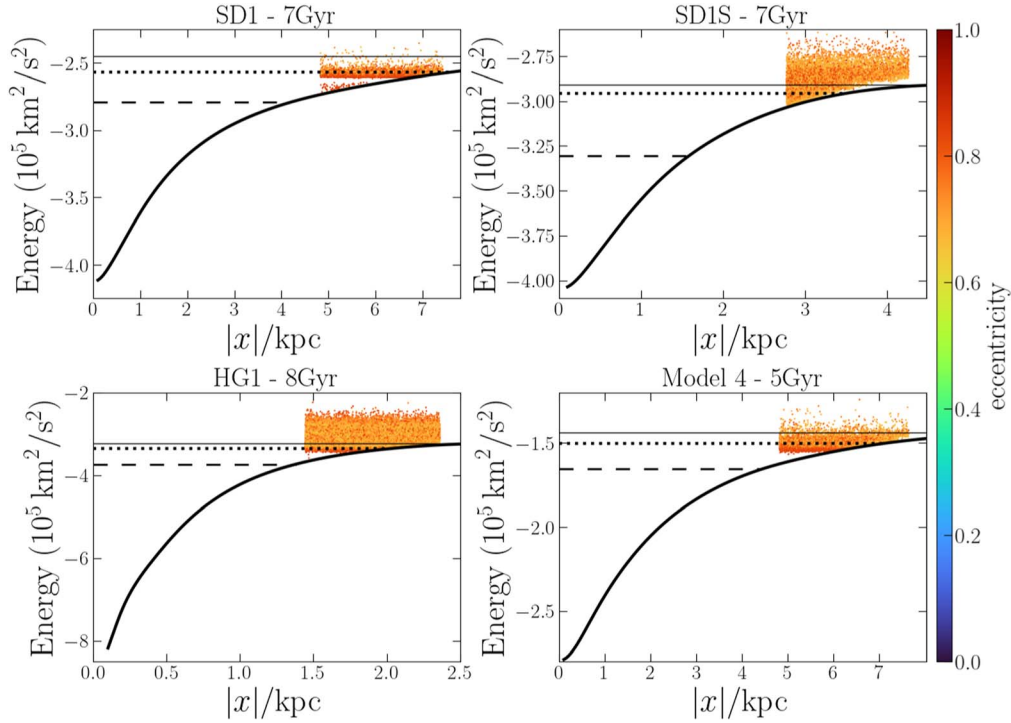
Figure 11. Similar to Figure 5, but only showing E_J (points) for the particles selected in the group of orbit 8 in Figure 6. These orbits typically have $0.4 \lesssim (\Omega_\varphi - \Omega_P)/\Omega_R \lesssim 0.45$ and in some simulations a significant fraction of these orbits have $E_J > \Phi_{\rm eff}(R_{\rm RC})$, i.e., can visit both the bar and the disk region, as those identified by Sparke & Sellwood (1987) and Pfenniger & Friedli (1991).

times, color-coded by the eccentricity. All models (except SD1S) show three distinct groups: the cool orbits ($\Omega_z/\Omega_R \gtrsim 3/2$) have low $z_{\max}$ and low-eccentricity (blue cloud)—the elliptical orbits in Section 4.1. The warm orbits ($1 \lesssim \Omega_z/\Omega_R \lesssim 3/2$) have slightly larger $z_{\max}$ and relatively high eccentricity (orange cloud)—the looped orbits. And the hot orbits ($\Omega_z/\Omega_R \lesssim 1$) have significantly larger $z_{\max}$ and eccentricity (red cloud)—these are chaotic orbits within the vILR cloud. We also see some hot orbits at the ILR (orbits with high $z_{\max}$) contaminating the Ω_z/Ω_R-interval of the warm orbits, which is not surprising given the crowd of orbits around the crossing of the ILR and the vILR—see, e.g., Figure 6.

As suggested by Figure 7 and demonstrated below, the loops at the ends of warm orbits at the ILR play an important role in producing the shoulders along the bar-major-axis density profiles. As a consequence of these loops, these orbits spend a considerable time in retrograde motion in the rotating frame of the bar. The bottom panels of Figure 16 show the fraction of time the orbits at the ILR are in retrograde motion, i.e., with $L_z < \Omega_P R^2$. We see that this fraction is zero for all cool orbits, except for some contamination of the warm orbits in model HG1. In the range $1 < \Omega_z/\Omega_R < 3/2$ of the warm orbits, the fraction of time in retrograde motion is typically high (≈ 0.2–0.55), a signature of the loops developed by these orbits. Finally, the vertically hot orbits also have large fractions of time in retrograde motion, but slightly smaller than the warm orbits.

Figure 16 demonstrates the common Ω_z/Ω_R-dependence of the different morphologies seen in Section 4.1 for the models with shoulders: vertically warm orbits at the ILR have loops and spend a considerable time in retrograde motion; and hot orbits at the ILR jump to large $z_{\max}$ and also spend a considerable time in retrograde motion, but less than warm orbits on average. The consequence of this for the shoulders is

discussed in Section 4.6. Moreover, in agreement with the importance of the looped morphology to the shoulders, as suggested in Paper I and confirmed in this work, we note that this morphology is absent in the shoulder-less model SD1S.

4.4. The Orbital Support of Shoulders

Here we quantify the contribution of the main orbital groups to the density profiles in the shoulder region. Since the ILR is the main bar-supporting resonance, and inspired by the morphological study in Sections 4.1 and 4.3, we investigate three groups: the vertically cool and warm orbits at the ILR, and the vILR cloud (which includes the hot orbits at the ILR).

Figure 17 (panel (a)) shows the N-body snapshot surface density profile for SD1 at 7 Gyr, for particles in the shoulder region. The density profile is calculated as the number of particles in $(x/a_{\rm bar})$−bins of width 0.015 and within $|y|/a_{\rm bar} < 0.145$, divided by $(\delta x/a_{\rm bar})(\Delta y/a_{\rm bar})N$, where $\Delta y = 2 \cdot 0.145$, and N is the total number of selected particles. Since these profiles are almost perfectly symmetric around $x = 0$, for better visualization, we present them averaged between the two sides $x < 0$ and $x > 0$. The gray areas show the average shoulder region, where particles are selected. The black solid curve shows the total density profile, and the colored solid curves show the relative contribution of particles in the different orbital groups, in the same colors as in the titles of panels (b)–(d).

In this model, the contributions from particles at the vILR cloud (red) and the warm ones at the ILR (green) are similarly large. The black dashed curve shows the total density profile excluding the contribution of the latter. Removing the warm orbits significantly reduces the density profile in the shoulder region, demonstrating their important role in the shoulders. The dotted curve shows the profile after also subtracting particles at
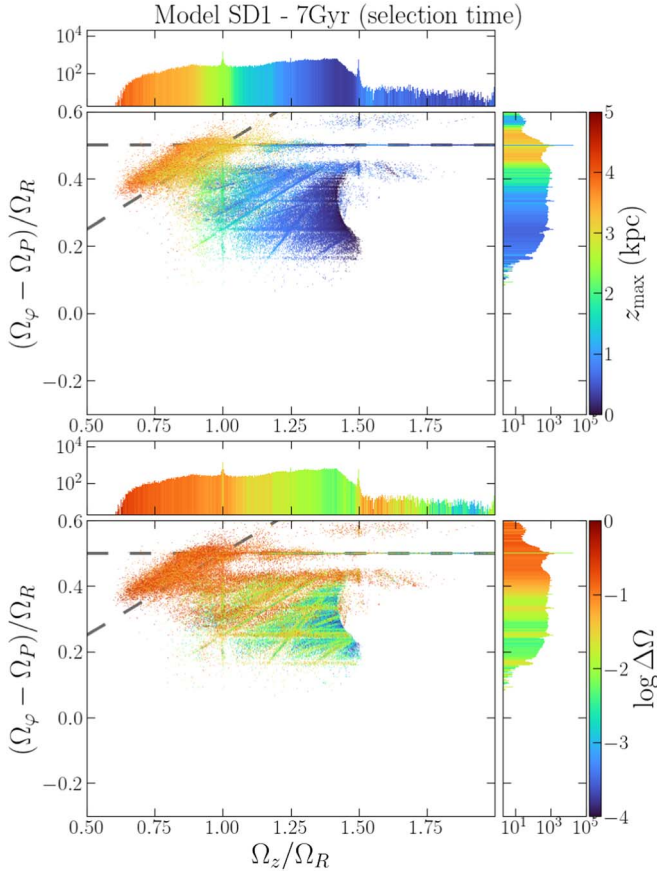
Figure 12. Frequency map for orbits in the shoulder region of simulation SD1 at 7 Gyr. Smaller panels show histograms (in log-scale) with bin-widths of 0.0025. The upper panel is color-coded by the vertical excursion, and the bottom panel is color-coded by the frequency drift (high values indicating chaotic orbits). The ILR, $(\Omega_{\varphi} - \Omega_{P})/\Omega_R = 1/2$, is the most prominent resonance. Orbits in the vILR cloud, to the left and above the vILR $(\Omega_{\varphi} - \Omega_{P})/\Omega_z = 1/2$, have significantly larger z_{max} and are strongly chaotic.

Figure 13. Similar to Figure 12, but for simulation SD1S (fast bar and no detected shoulders). The most prominent resonances are the ILR, the corotation, and $\Omega_z/\Omega_R = 1$. The ILR is not as strongly populated as in Figure 12, and there is no vILR cloud. Orbits with $\Omega_z/\Omega_R < 1$ still have larger z_{max} than the rest, but not anywhere as large as in Figure 12.

the vILR cloud, and we see that they mostly contribute to the innermost parts of the shoulders. Together, these two groups contribute so much to the density profile that their subtraction produces significant deficits in it.

The picture provided by panel (a) in Figure 17 is limited because it shows particles currently in the shoulder region. To evaluate the global profile produced by these particles, one can calculate the profile produced by them at a later time, allowing the particles to phase-mix, as done in Paper I. However, in the mean time, the bar pattern speed and the orbits themselves may have evolved into different ones. For a global *and instantaneous* view of the density profiles, we stack the orbits integrated for these particles in the frozen potential and plot the surface density in the *x–y* plane in Figure 17, panels (b), (c), and (d). To emphasize each group's relative contribution rather than the orbital shapes, we now normalize colors by the total number of particles in the shoulder region.

The orbit-averaged density of the orbits at the vILR cloud (panel (b)) resembles that of looped orbits, although it is more diffuse in the *x–y* plane, which arises because of the chaotic nature of orbits at the vILR cloud—see the bottom panels of Figures 12 and 14, and 15. Furthermore, its orbit-averaged density is less extended along the major axis, peaking outside of the shoulder region. Panel (c) shows the orbit-averaged density of the warm orbits at the ILR. We see that this profile is
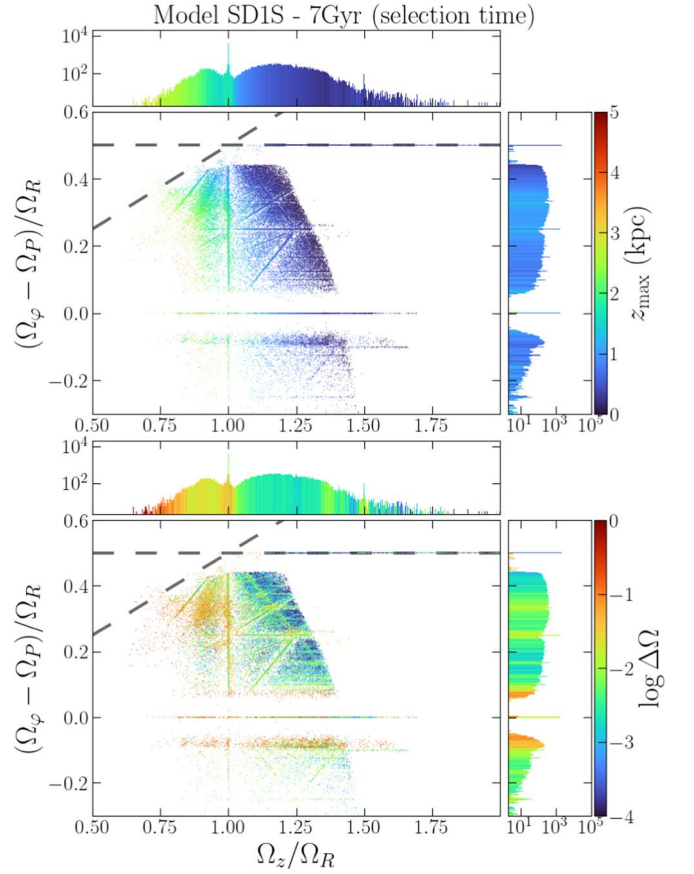
more elongated and has a very pronounced peak in the shoulder region. On the other hand, the cool orbits (panel (d)) do not represent any visible excess in the shoulder region.

Panel (e) shows the density profiles of these stacked orbits (for points within $|y|/a_{bar} < 0.145$) along the *x*-axis, analogous to the left panel, and using the same color/style convention. For a comparison with panel (a), these profiles are further normalized by the number of points per orbit. The smaller values, in comparison to panel (a), are due to the "dilution" of each particle along its orbit out of the selection region.

In these orbit-averaged profiles, the warm orbits at the ILR contribute much more significantly to the shoulder because: (i) they are more elongated than those at the vILR cloud; (ii) particles stay longer near apocenters; and (iii) the sharply defined loops at their ends make them spend even more time near apocenters. This confirms the important contribution of the looped morphology of the warm orbits at the ILR to the shoulders, but also demonstrates the importance of the vILR cloud (compare the green and red curves in panels (a) and (e)). Finally, the contribution of the cool orbits at the ILR is negligible also in the orbit-averaged profiles.

Figures 18–20 show the corresponding profiles for models SD1S, HG1, and Model 4, with the same color-normalization as in Figure 17. Recall that model SD1S has no shoulders, and that its "shoulder region" is set as the same fraction of the bar length as in SD1 at the equivalent snapshot. The warm orbits at the ILR (green) in this model do not develop loops, as already
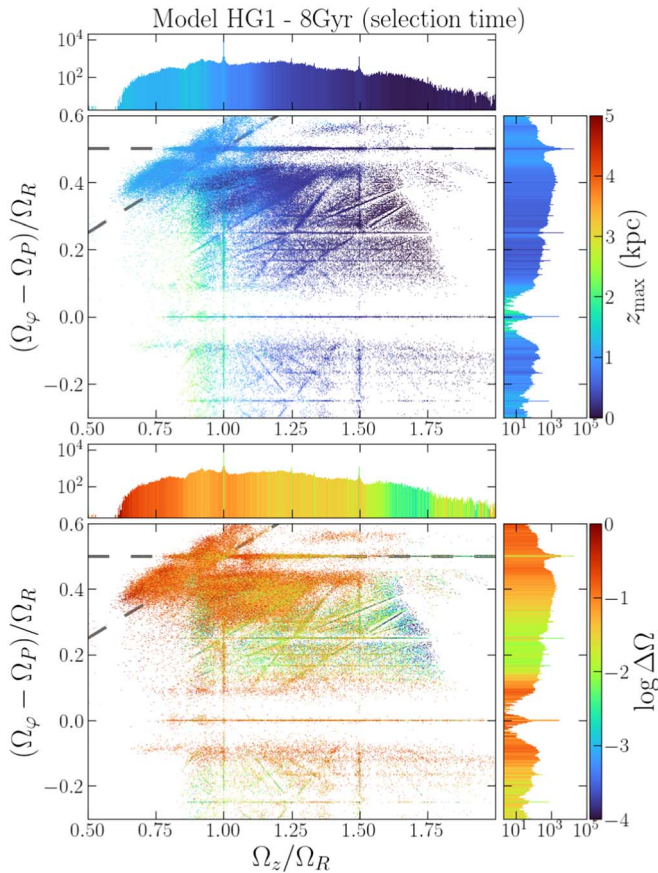
Figure 14. Similar to Figure 12, but for model HG1, which has gas and a fast bar. We see several well-populated resonances, particularly the ILR, the corotation and the vertical resonances $\Omega_z/\Omega_R = 1$ and $\Omega_z/\Omega_R = 3/2$. The vILR cloud is strongly populated with chaotic orbits and relatively high z_{max}, but not as high as in the model SD1, Figure 12.
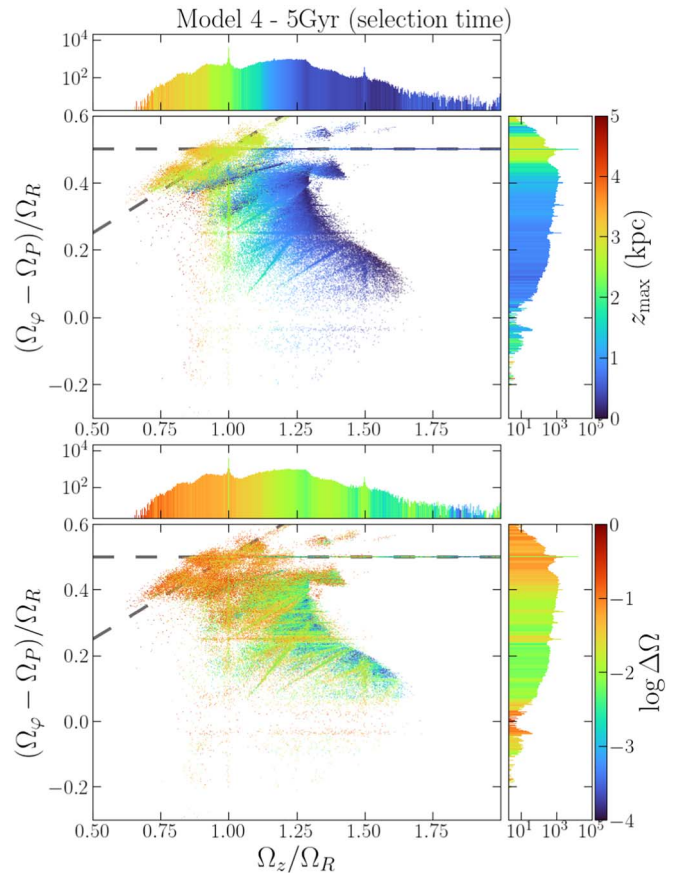
Figure 15. Similar to Figure 12, but for Model 4, whose bar buckles at 4 Gyr and is slow at 5 Gyr. This figure is similar to that of model SD1, with z_{max} values in the vILR cloud intermediate between those in SD1 and HG1, Figures 12 and 14.

shown in Figure 16, and these orbits' contribution to the density profile in the shoulder region is negligible, as is the contribution from the vILR cloud. In fact, just a handful of orbits populate the vILR cloud in this model (Figure 13).

Figure 19 shows analogous profiles for model HG1, which has weak shoulders—see Paper I, Figure 11. In this case, the contribution of the warm orbits is relatively small, and its subtraction produces a profile (dashed) that is close to the total one. On the other hand, in HG1 the vILR cloud orbits are very numerous and produce an overall shape in the x–y plane that resembles the looped morphology, contributing significantly to the density in the shoulder region. Figure 20 shows the equivalent profiles for Model 4. This model is intermediate between SD1 and HG1 in terms of bar length and pattern speed evolution (see Figure 2), and so are the relative contributions of warm orbits at the ILR and of the vILR cloud. We also observe significant contributions from the warm orbits and the vILR cloud, with the latter producing a more diluted looped-like shape in the x–y plane, in comparison with the warm orbits.

In summary, we have confirmed the important role of looped orbits to the shoulders, as previously suggested in Paper I, and have shown that this morphology is found in vertically warm orbits at the ILR. We also demonstrate that, although individual orbits in the vILR cloud are chaotic and do not have a well-defined looped shape, they still produce an overall shape in the x–y plane that can be seen as a diluted version of the shape produced by the looped orbits. We note that the morphology of

vILR cloud orbits in SD1 is the most diffuse in the x–y plane, and also has the largest z_{max}, among models with shoulders. Conversely, the vILR cloud of HG1 produces the least diffuse shape in the x–y plane, and the smallest z_{max}, while Model 4 is intermediate between the two. This connection between z_{max} and the looped shape is explored via the evolution of the frequency maps presented in Sections 4.5 and 4.6.

4.5. Evolution of the Frequency Maps

Having discussed the orbital support for shoulders, and having shown the frequency maps at the time of selection in Section 4.2, we now show their time evolution, shedding light on an important mechanism for the shoulders' evolution. We select the same particles from the model snapshots discussed above and integrate their orbits at four additional snapshots, with the corresponding frozen potentials and bar pattern speeds, and using the particles' coordinates at each snapshot as initial conditions (for HG1, which has star formation, we only use star particles born before the selection time). The results are shown in Figures 21–24, with the selection times shown in the central panels.

Figure 21 shows the evolution of the frequency maps for the fiducial model SD1. The main trends, which also occur in the other models with shoulders, are a global shift upward, i.e., to higher $(\Omega_\varphi - \Omega_P)/\Omega_R$, and to the left, i.e., to lower Ω_z/Ω_R. One obvious reason for the shift upward is the decline in Ω_P as the bar slows down (see Figure 2), with its main resonances
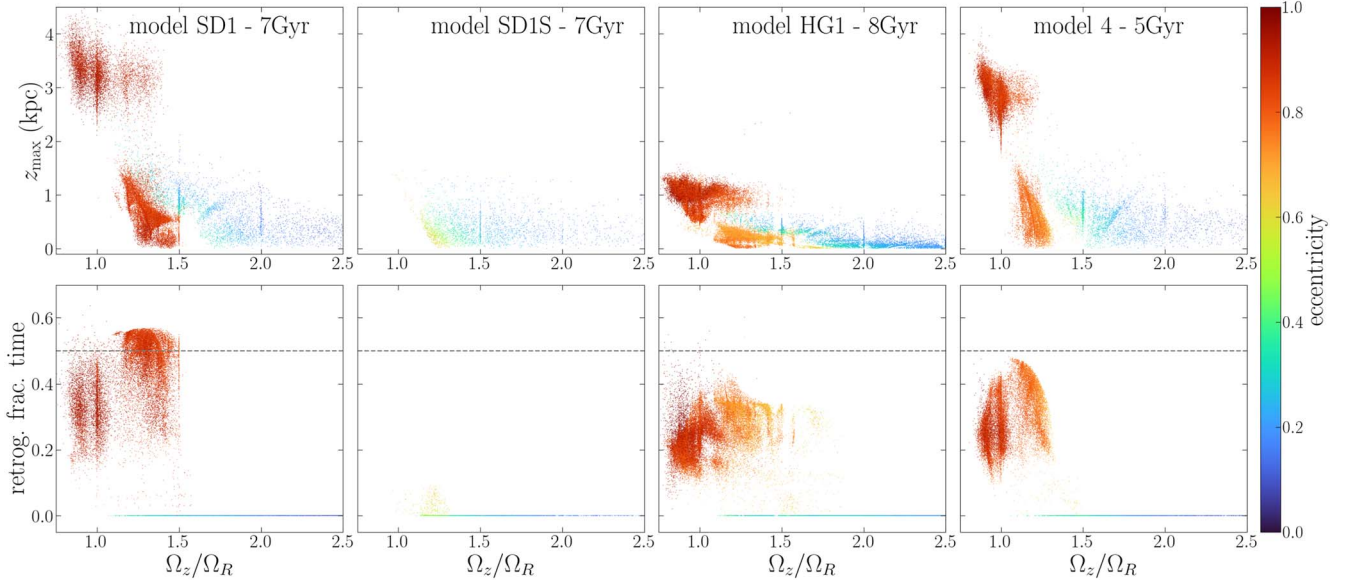
Orbits at the ILR

Figure 16. Top: vertical excursion vs. Ω_z/Ω_R for particles selected in the shoulder region and at the ILR for all models. Bottom: the fraction of time spent in retrograde motion (in the rotating frame). The horizontal dashed line is at 0.5. The different morphologies identified in Section 4.1 are clearly seen here (except for SD1S): cool orbits at the ILR have low z_{max}, low eccentricity, and are almost exclusively prograde; warm orbits at the ILR have slightly larger z_{max}, relatively high eccentricity, and spend a significant time in retrograde motion (a symptom of the strong loops at their ends); hot orbits at the ILR have significantly larger z_{max} and eccentricity, and moderate fractional time in retrograde motion. The looped morphology is not detected in the shoulder-less model SD1S. We also note the overall lower z_{max} for the SPH model HG1, in comparison to models SD1 and Model 4.
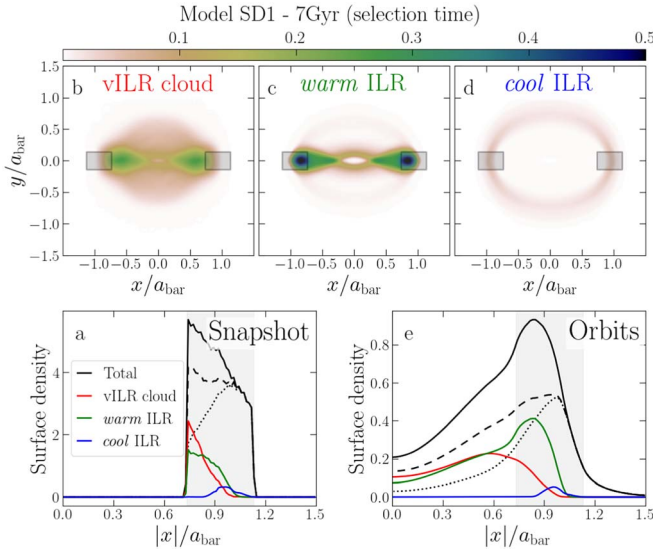

Figure 17. Surface density of particles in the shoulder region (gray shaded areas) of model SD1 at 7 Gyr, with coordinates normalized by the bar length at 7 Gyr. Panel (a): instantaneous density profile along the major axis ($|y|/a_{bar} < 0.145$) in the simulation. The solid black curve shows the total profile, and colored curves show contributions from the groups in panels (b)–(d) (see the main text). The dashed curve is the total profile after subtracting the warm orbits at the ILR ($1 < \Omega_z/\Omega_R < 3/2$), and the dotted one results from also subtracting orbits at the vILR cloud. Panels (b), (c), and (d) show the 2D orbit-averaged maps for the respective orbital groups, with colors normalized by the total number of particles in the shoulders. Panel (e) shows the 1D density profile, similar to panel (a), but for the orbits, and restricting to $|y|/a_{bar} < 0.145$. In this panel, the large contribution of warm orbits to the profile in the shoulder region becomes more evident, with a smaller but considerable contribution from the vILR cloud. The contribution from the cool orbits (blue) is negligible.

sweeping the disk outward and trapping new particles. Once trapped by the strong ILR, the orbits tend not to shift any longer upward in the frequency maps.
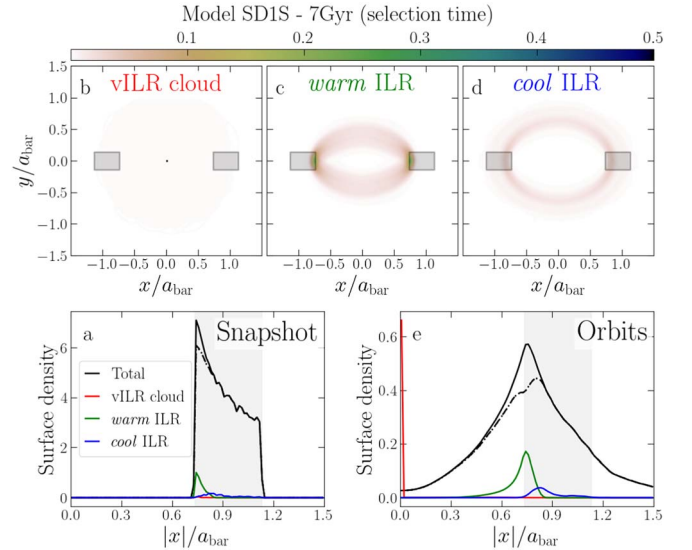
Figure 18. As in Figure 17, but for model SD1S, with no detected shoulders. In this case, warm orbits at the ILR do not have loops, and their contribution to the density profile (solid green) is much smaller than in the model SD1, Figure 17.

The ratio Ω_z/Ω_R is a proxy for the vertical thinness of the orbit (Section 4.1). Loosely speaking, the net shift to the left in the frequency maps indicates evolution to thicker orbits, which is confirmed by their large z_{max} (colors). While the thickness grows gradually for decreasing Ω_z/Ω_R, the vertical motion is amplified when orbits reach the vertical resonance $\Omega_z/\Omega_R = 1$, and a more dramatic amplification happens at the vILR. This is particularly important near the crossing point of these two resonances and the ILR. While large vertical excitation by the vILR for orbits close to the ILR [$(\Omega_\varphi - \Omega_P)/\Omega_R = 1/2$] has been predicted by Binney (1981) and confirmed by, e.g.,
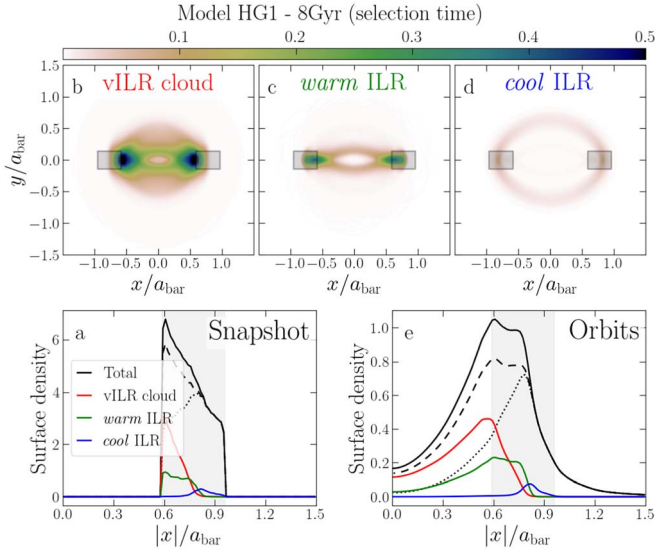
Figure 19. As in Figure 17, but for model HG1. The contribution from warm orbits at the ILR (solid green) to the density profile is much smaller than in SD1. However, in this case the number orbits at the vILR cloud is much larger, and they produce an orbit-averaged density in the x–y that resembles the one produced by looped orbits.
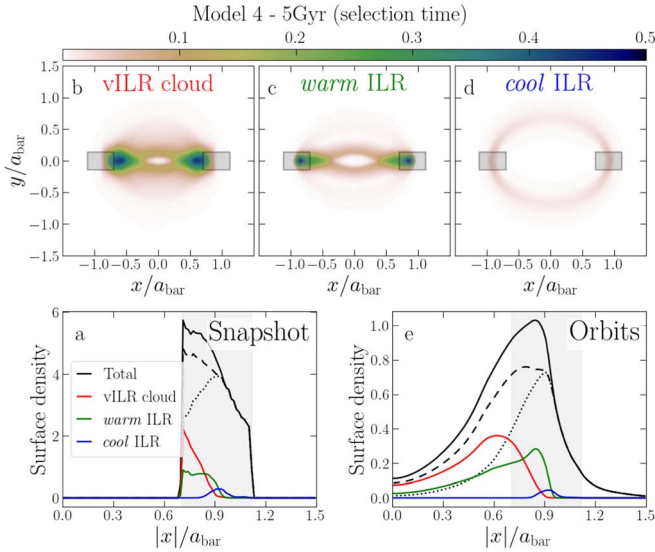

Figure 20. As in Figure 17, but for Model 4. The contributions from warm orbits at the ILR (green) and the vILR cloud (red) are comparable, and a looped-like morphology is produced by both groups, as in the model HG1, Figure 19.

Pfenniger & Friedli (1991), Figure 21 demonstrates this in a particularly transparent way. A comparison with Figure 12 also suggests that, once orbits cross the vILR, their quasiperiodicity is destroyed, in agreement with Valluri et al. (2016), who found that "banana" orbits are commonly chaotic—see also Manos et al. (2022).

Note that Figure 21 shows the evolution of a fixed set of particles selected in the shoulders at 7 Gyr, instead of the evolution of the shoulders themselves. In particular, starting from ≈4 Gyr, this simulation persistently manifests shoulders that move outward following the bar growth—see Figure 8 of Paper I.

Figure 22 shows the frequency map evolution for the model SD1S. In this case, a small overall shift upward and to the left

is barely noticeable, in agreement with the very slow bar evolution in this simulation (Figure 2). We also note a significant increase in z_{max} for orbits crossing to $\Omega_z/\Omega_R < 1$, but the vILR seems relevant only very close to the crossing point with the ILR in this model, and there is no vILR cloud.

For HG1, the evolution of the frequency map is shown in Figure 23. Once more, we note a shift upward and to the left. In this case, the resonance $\Omega_z/\Omega_R = 1$ also seems more effective in producing large vertical excursions. In the vILR cloud, z_{max} is not as significantly amplified as in the model SD1 (Figure 21). However, the vILR cloud is more populated in this model.

Finally, Figure 24 shows the frequency map evolution for Model 4. The bar in this model buckles slightly before 4 Gyr, and shoulders are observed right after that—see Section 2. The abrupt shift to lower Ω_z/Ω_R from 3 to 4 Gyr is a consequence of this buckling, which vertically thickens the bar region significantly. This interpretation and the frequency map at 4 Gyr itself must be considered with caution, since this is too close to the buckling, and the gravitational potential may lack the triaxial symmetry we assume in order to estimate it. On the other hand, the main features in the frequency map at 4 Gyr are still observed at 5 Gyr, by which point the potential is more symmetric, and therefore our fit is more reliable, with the same trend of a shift upward and to the left. This suggests that the frequencies at 4 Gyr are reasonably accurate, despite the disturbance from the buckling. In this model, the strong role of the vILR in exciting vertical motion is again observed, producing a well-populated vILR cloud with high z_{max}.

4.6. Evolution of Shoulder-supporting Orbits

In Section 4.5 we showed that orbits selected in the shoulder region were either already trapped by the ILR at previous snapshots or previously had smaller $(\Omega_\varphi - \Omega_P)/\Omega_R$, with the latter evolving toward larger $(\Omega_\varphi - \Omega_P)/\Omega_R$, until possibly being trapped by the ILR. The orbits also evolve toward smaller Ω_z/Ω_R, until possibly crossing the vILR, thereby being excited to high z_{max} and becoming chaotic (e.g., Figure 12). In Section 4.4 we demonstrated the important role of vertically warm orbits at the ILR, and of the vILR cloud, in building the shoulders. We now analyze the morphological evolution of these warm orbits.

Figure 25 shows, for the model SD1, the orbit-averaged surface density maps at different times for the same star particles selected to be warm orbits at the ILR at 7 Gyr, with the central column representing the selection time. The upper (middle) row shows the maps in the x–y (x–z) plane, with coordinates normalized by the bar length at the selection time. Here again, colors are normalized by the total number of star particles in the shoulder region. The shaded areas show the normalized shoulder region at the selection time. The bottom panels show the same orbits in frequency maps.

The upper row shows that the warm orbits at the ILR at 7 Gyr mostly come from less elongated orbits, which gradually become more elongated, until being trapped by the ILR. A fraction of these orbits had $\Omega_z/\Omega_R > 3/2$ before the selection time, and either before or after being trapped by the ILR. Before they cross the vILR ($t < 7$ Gyr), these trends are accompanied by negligible evolution in the x–z plane. Once they cross the vILR (7–8 Gyr), these orbits occupy the whole vILR cloud region, with large z_{max}. This is accompanied by a dissolution of the density excess in the x–y plane. Additionally, the orbits start to shrink along the major axis, and become confined to the inner parts at 9 Gyr, no
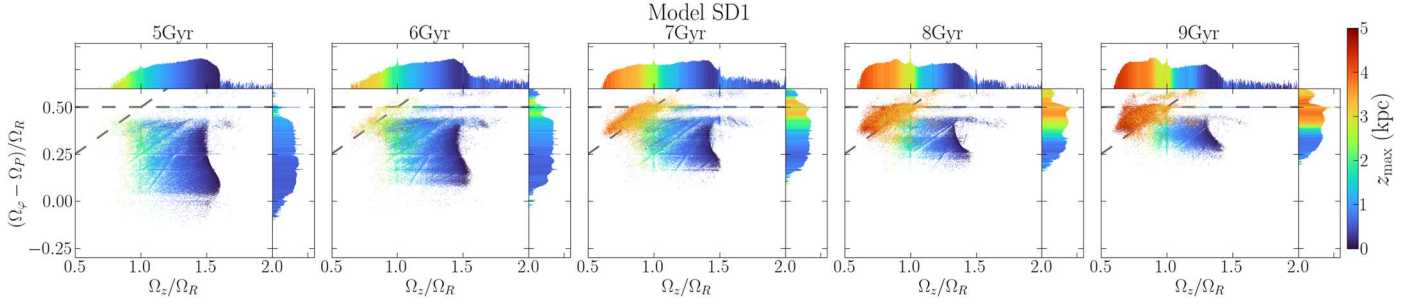
Figure 21. Time evolution of the frequency map for the same set of star particles selected at 7 Gyr in the shoulder region of model SD1. Points are color-coded by the vertical excursion z_{max}. We see an overall shift upward in this plot mainly due to the bar slowing down, i.e., decreasing Ω_P, and to the left, with orbits becoming vertically thicker over time. Orbits crossing the vILR from right to left have vertical motion dramatically amplified.
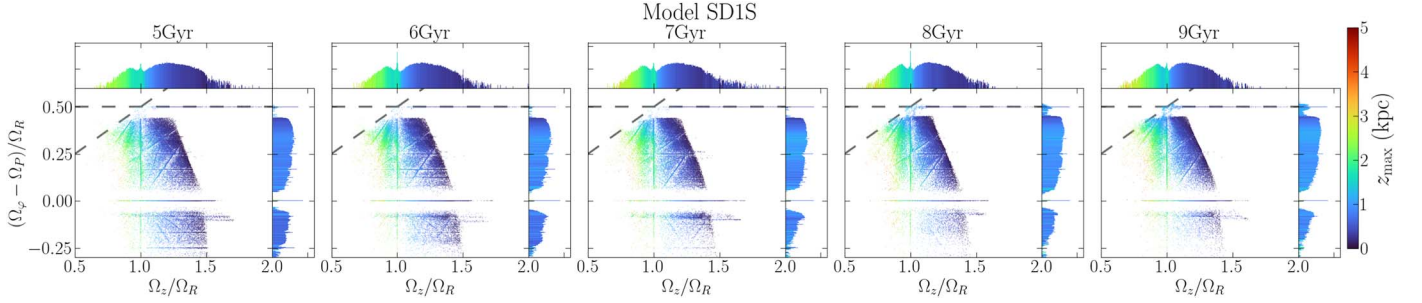

Figure 22. Similar to Figure 21, but for model SD1S. In this case, we do not observe a significant time evolution of the frequency map, in agreement with the very mild evolution of the bar parameters—see Figure 2.
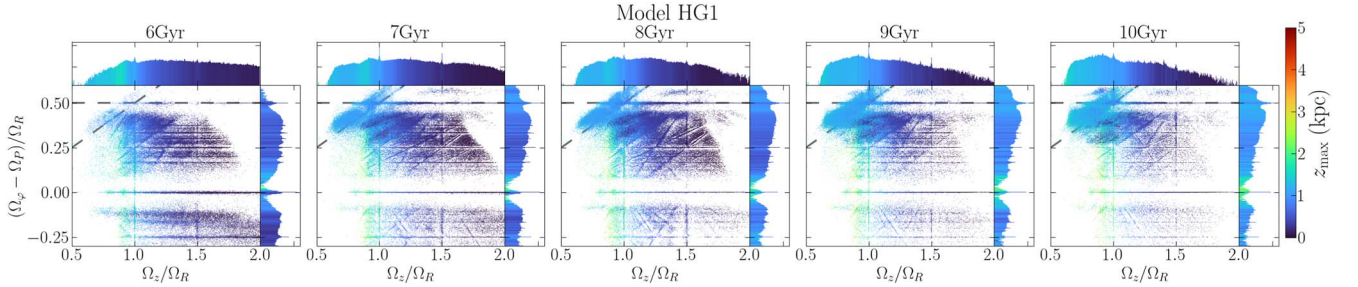

Figure 23. Similar to Figure 21, but for model HG1. We observe the same trends as in model SD1, i.e., a net shift upward and to the left. Orbits crossing the vertical resonances $\Omega_z/\Omega_R = 1$ and the vILR have their vertical motion significantly amplified.
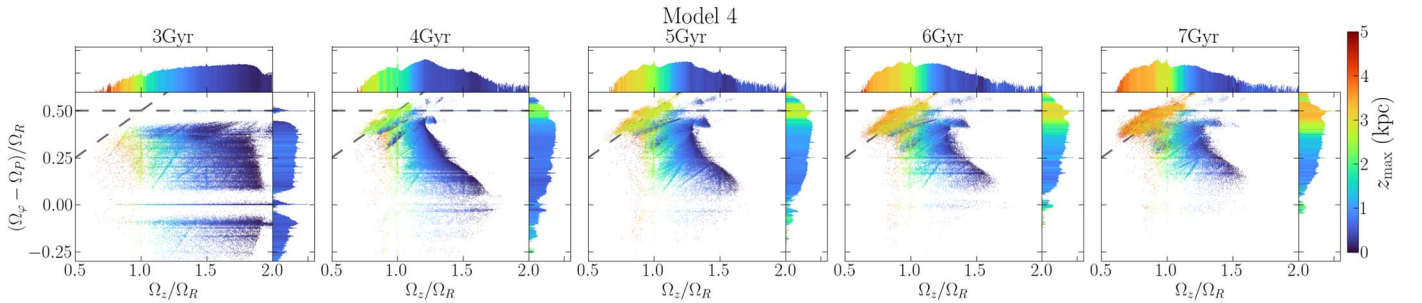

Figure 24. Similar to Figure 21, but for Model 4. In addition to the same trends observed for SD1, we note the abrupt shift from 3 to 4 Gyr toward smaller Ω_z/Ω_R. This is produced by the buckling bar, slightly before 4 Gyr, and after which shoulders are detected in the model.

longer falling in the original (7 Gyr) shoulder region. This is closely connected to the abrupt change in the x–z plane from 7 to 8 Gyr. As shown in Figure 12, this is associated with the chaotic nature of the orbits in the vILR cloud, which dilutes the density distribution in the x–y plane and weakens their contribution to the shoulders. But perhaps most importantly, we see that, to the

extent that the vILR increases the vertical excursion, it decreases the in-plane amplitude, converting in-plane motion to vertical motion, like a pole vaulter.

For completeness, Figure 26 shows the corresponding plot for model SD1S. However, since the bar barely evolves in this model (see Figure 2), we see no significant evolution in this plot.
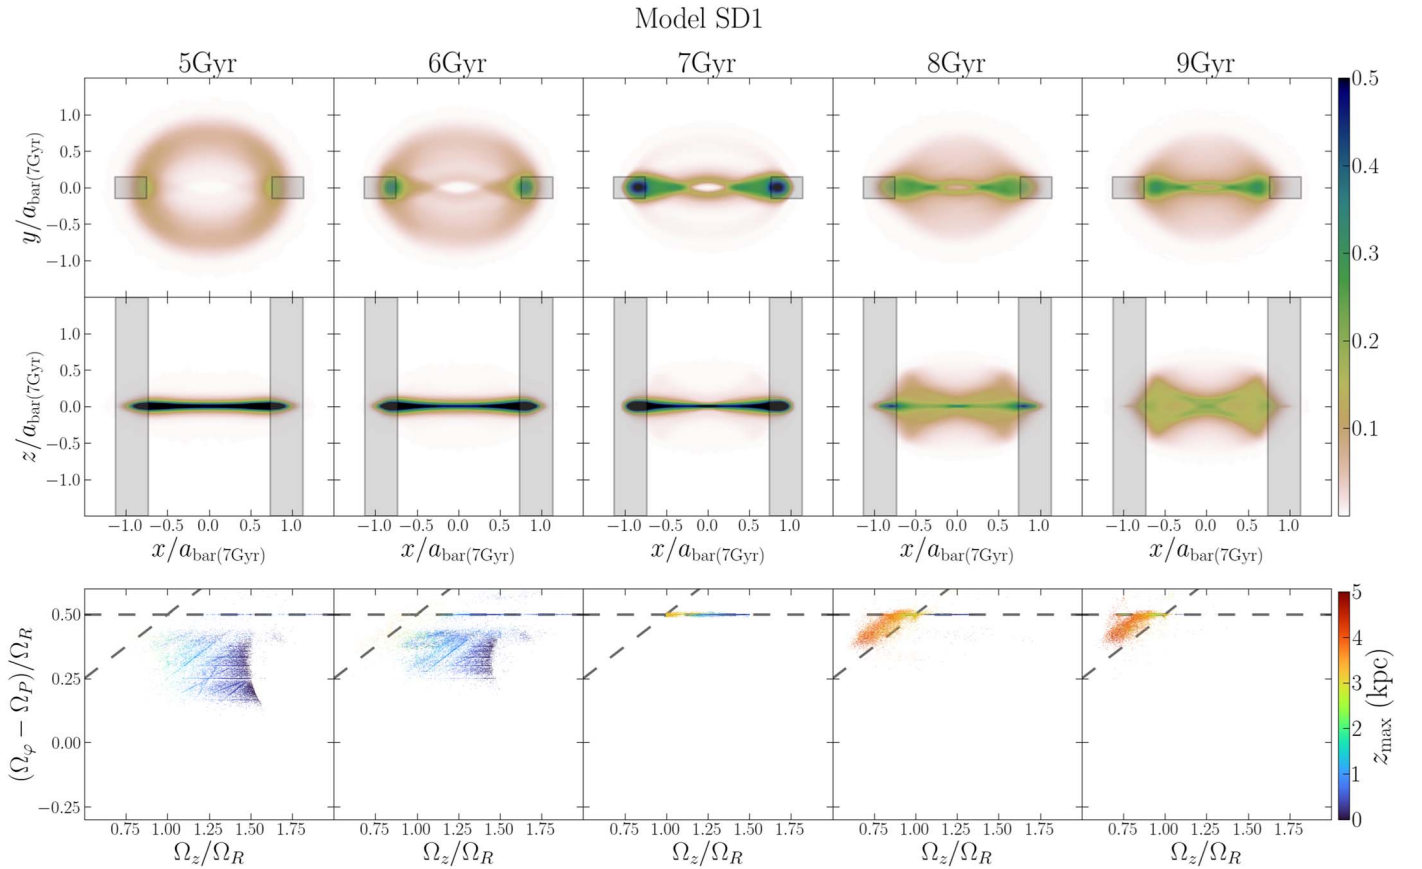
Figure 25. Evolution of the orbit-averaged density maps in the x–y (top) and x–z planes (middle) for star particles selected to be on warm orbits at the ILR at 7 Gyr in model SD1. The coordinates are normalized by the bar length at the selection time. Vertical shaded gray areas show the shoulder region at the selection time. The bottom row shows the same orbits in the frequency map. Those that at 7 Gyr have the loop-shape, and support the shoulder morphology, evolve from less elongated orbits and typically away from smaller Ω_z/Ω_R. Once they cross the vILR, they spread in the vILR cloud, with much larger z_{max}, diluting the overdensities in the x–y plane.

Figure 27 shows a corresponding plot for the model HG1, where we observe the same trends. However, in HG1, the vILR is less vigorous in exciting vertical motion and, accordingly, the orbit-averaged density in the x–y plane is less disturbed for particles crossing the vILR to the vILR cloud. Figure 28 shows the equivalent plots for Model 4, and we see the same connection between the dilution of the density excess in the x–y plane and the increase in the vertical motion after the orbits cross the vILR, populating the vILR cloud.

These results demonstrate the role of the vILR in perturbing the shoulder-supporting orbits: once these orbits cross this resonance, their in-plane motion is converted to vertical motion and they become chaotic, contributing to diluting the shoulders.

5. Discussion

5.1. Comparison with Different Techniques

We have performed frequency analyses for orbits integrated in frozen potentials of barred galaxy simulations. Frequency analysis of the coordinates in ongoing N-body simulations (without introducing frozen potentials) has also been proven possible (e.g., Ceverino & Klypin 2007; Gajda et al. 2016; Parul et al. 2020; Smirnov et al. 2021). While this has the apparent advantage of tracking the orbital evolution in a more self-consistent manner, the short time-series (typically ≈1 Gyr), the unavoidable changes in the potential in the considered time interval, and all sources of noise in N-body potentials introduce uncertainties in the estimated frequencies. On the other hand, orbit integration in frozen potentials provides *instantaneous* orbital properties, and it allows for the use of densely sampled time-series (allowing investigation of higher frequencies), with the narrowness of the main resonances in the frequency maps testifying to the high accuracy of our procedure.

The frequency analysis involves the choice of the coordinate system, the reference frame, and the use of real or complex time-series. Regarding the last point, Hunter (2002), using Cartesian coordinates for a few orbits, reported no improvements in the accuracy of estimated frequencies when using complex time-series. Here, we chose cylindrical coordinates for their easier interpretation, given the approximate symmetries of the potential, and complex time-series. In particular, we use the complex Poincaré polar coordinates for the azimuthal angle.

Difficulties in estimating the azimuthal frequency Ω_φ have been reported in several works (e.g., Athanassoula 2002, 2003; Ceverino & Klypin 2007), but often not in explicit terms. In Appendix B, exploring the isochrone model, we show how to overcome these difficulties, demonstrating that the use of the complex Poincaré polar coordinates improves the accuracy of the azimuthal frequency and facilitates its extraction. The accuracy of estimated frequencies in more generic potentials with (semi)analytic expressions for the fundamental frequencies deserves further investigation.

Unlike the common practice of performing frequency analysis in the rotating frame of the bar, we opted to store
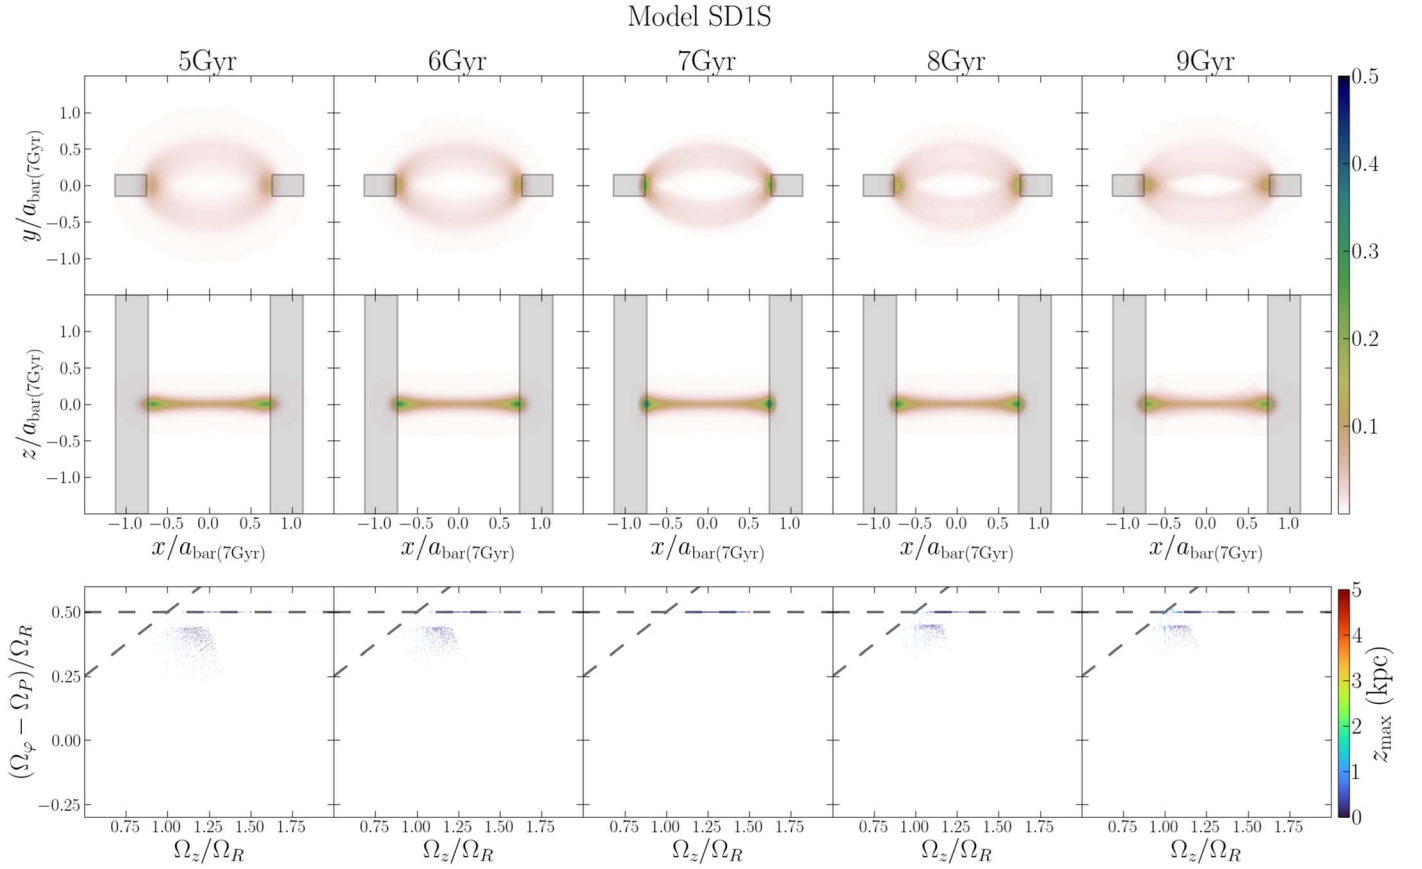
Figure 26. Similar to Figure 25, but for model SD1S. As the bar barely evolves in this model, the morphology of the warm orbits selected at 7 Gyr at ILR does not change significantly.

the coordinates in the inertial frame of the galaxy. As demonstrated in Appendix C, this choice allows for a precise identification of the corotation resonance, which is not possible, at least with the techniques employed in this work, if the rotating frame of the bar is adopted.

5.2. Implications for the Evolution of Shoulder Profiles

Paper I demonstrated that shoulders are present in the density profiles of growing bars. It also suggested the important contribution of looped orbits to the shoulders, which is confirmed with the detailed orbital study performed here (Section 4.4). We now return to the questions that motivated this work and that we posed in Section 1.

The evolution of the density profile is complex and must depend to some extent on all orbital groups, but we draw a simple picture that is based on the evolution and prevalence of looped orbits along the ILR. In Section 4.4 we demonstrated that the looped morphology in the x–y plane is very well defined for vertically warm orbits at the ILR (but not in the model SD1S, where orbits are not looped), while cool orbits at the ILR are more elliptical and contribute negligibly to the shoulders. Hot orbits at the ILR (and at the vILR cloud) are chaotic and do not share the same looped morphology of warm orbits, but overall still spend a considerable amount of time in retrograde motion and produce a similar, more diluted, looped-like shape in the x–y plane. Their contribution to the shoulders depends on the strength of the vILR in exciting vertical motion.

Thus, in order for shoulders to emerge in the density profiles, the ILR needs to trap a significant number of orbits as the bar

slows and its resonances sweep through the disk. Additionally, a significant fraction of these orbits need to develop loops at their ends, which typically happens in the warm orbits at the ILR ($1 \lesssim \Omega_z/\Omega_R \lesssim 3/2$). Since orbits in the disk and in the bar's thinnest parts typically have $\Omega_z/\Omega_R > 3/2$, to support shoulders these orbits need to decrease Ω_z/Ω_R, either before or after being trapped by the ILR, which occurs when the orbits thicken vertically. While the bar thickens, either via secular processes or via buckling, the shoulder-supporting warm orbits at the ILR can sooner or later be populated. This is compatible with the findings of Paper I that shoulders can appear in either buckling or nonbuckling bars and that, in buckling bars, shoulders commonly appear just after a buckling event—see their Figure 7.

Once shoulders have emerged in the density profile, their supporting orbits can depart from the looped morphology in the x–y plane if they move out of the range $1 < \Omega_z/\Omega_R < 3/2$. This typically happens with further bar thickening and orbits at the ILR crossing the vILR toward $\Omega_z/\Omega_R \lesssim 1$—see Figures 25–28. Once more, this thickening can happen either secularly or via a secondary buckling. Thus, if the bar is not replenished with looped orbits, the shoulders can eventually be erased. This explains why Paper I noted several instances of sudden erasure of shoulders after a secondary buckling. It is worth noting that the newly vertically excited orbits will contribute to a thicker potential, thus decreasing the ratio Ω_z/Ω_R of nearby orbits, bringing them toward the vILR, in a positive feedback process.

Thus, the common thickening of orbits in the bar region would make shoulders transient features, unless new orbits are trapped at the ILR with the appropriate Ω_z/Ω_R range and
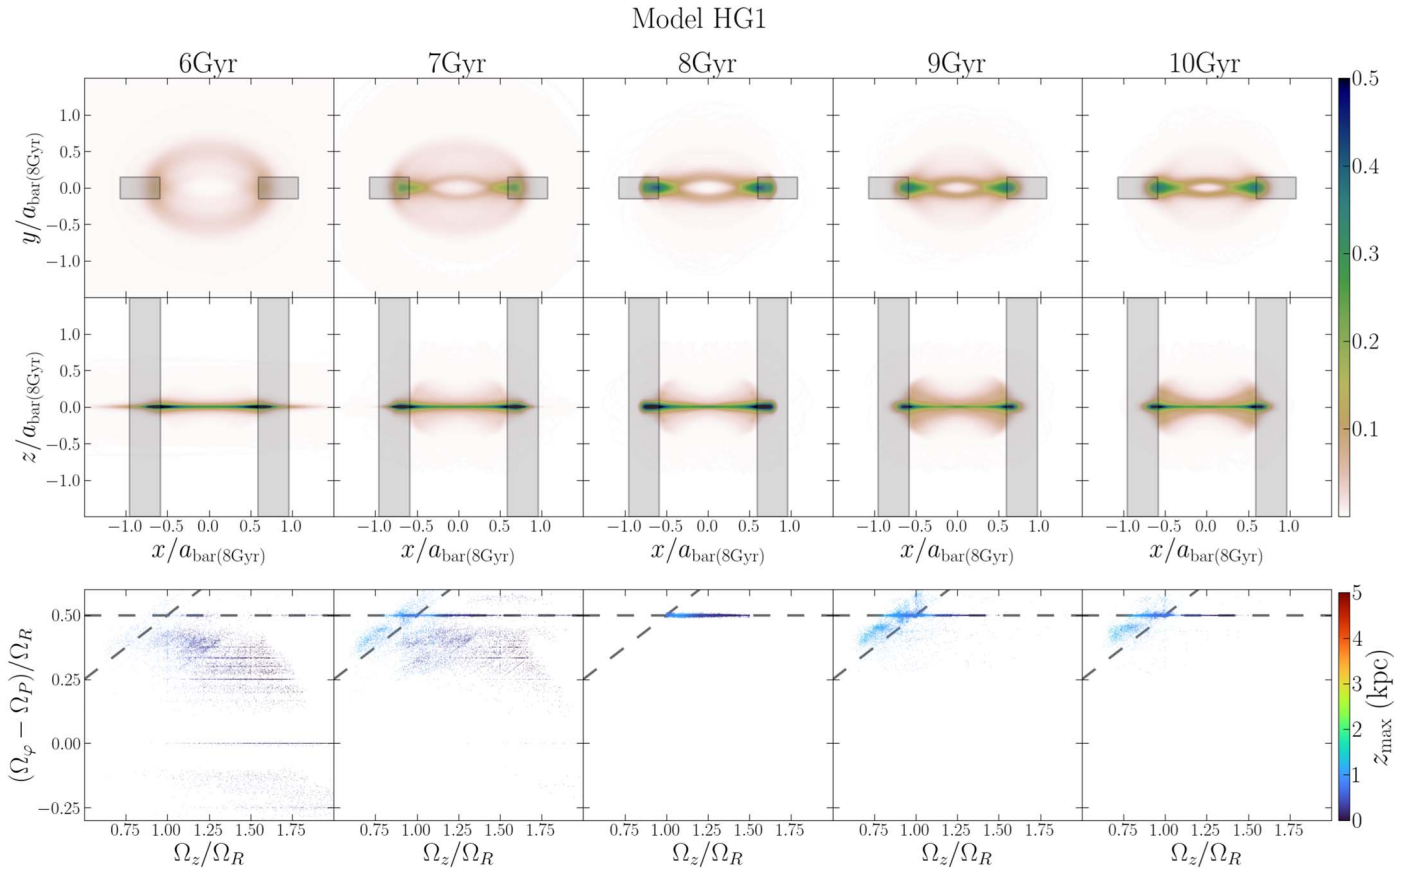
Figure 27. Similar to Figure 25, but for model HG1. For orbits crossing the vILR, the amplification of z_{max} and the dilution of the overdensities in the x–y plane are smaller than in SD1, but we still observe the vertical heating by the vILR.

develop the looped morphology, a process that seems to happen often for bars evolving secularly—see Figure 8 of Paper I.

The fact that the looped morphology tends to appear in orbits at the ILR and with $1 \lesssim \Omega_z/\Omega_R \lesssim 3/2$ suggests some connection with the vertical motion. This is intriguing since those orbits were first found in 2D bars (Contopoulos 1988). An interesting possibility is that this looped morphology might be somehow triggered by vertical resonances—for instance, comparing Figures 12–15, we see that the model SD1S, the only one without looped orbits or shoulders, is particularly unpopulated in the region near the crossing of the ILR and $\Omega_z/\Omega_R = 3/2$.

5.3. Connection with Bulges

Investigating mechanisms for bar thickening, Sellwood & Gerhard (2020) contrasted three scenarios: a bar that buckles, one in which orbits are vertically excited by staying trapped for significant times at the vILR (Quillen 2002), and another one in which orbits cross the vILR, get vertically heated, and remain that way after leaving the resonance (Quillen et al. 2014). Having found evidence for the long-lasting trapping mechanism only in an artificially vertically symmetrized model, Sellwood & Gerhard (2020) concluded that the third mechanism is probably more relevant in real galaxies without bar buckling.

Our results, in particular, e.g., Figures 21 and 25 (bottom row), confirm the rapid crossing of the vILR, as opposed to a long trapping time, and the rapid vertical excitation with a

long-lasting X-morphology. It is worth mentioning that the orbits at the vILR cloud, after crossing and being perturbed by the vILR, can get very close to the galactic center (Figure 8), possibly within the sphere of influence of a supermassive black hole (SMBH). This can potentially represent an important mechanism to persistently replenish the sphere of influence of an SMBH with new stars. This possibility is even more interesting if galactic bars can coevolve in harmony with black holes, as recently demonstrated by Wheeler et al. (2023).

In Section 4, we showed that orbits crossing the vILR populate what we call the vILR cloud, characterized by chaotic orbits with large z_{max} and X-shapes in the x–z plane. This morphology being produced by the vILR agrees with several previous works (e.g., Pfenniger & Friedli 1991), where the contribution of the "banana" orbits to BP-bulges has been emphasized. On the other hand, Portail et al. (2015) found that these orbits only contribute in the bulge's outermost regions, while orbits with $\Omega_x/\Omega_z \approx 3/5$ ("brezel" orbits) are more relevant for the bulge as a whole (in their study, Ω_x is the frequency of oscillation about the bar-major-axis and can be seen as a proxy for our $\Omega_\varphi - \Omega_P$)—see also Abbott et al. (2017).

Sections 4.3 and 4.6 demonstrate that after crossing the vILR, the banana orbits can also retain, to a certain extent, the loops near apocenters inherited from warm orbits at the ILR. In fact, these loops have already been found in orbits supporting either the thin (Patsis & Katsanikas 2014) or the thick (Combes et al. 1990) parts of bars.

Altogether, our results suggest that: (i) the banana orbits can simultaneously contribute to the outskirts of the bulge peanut
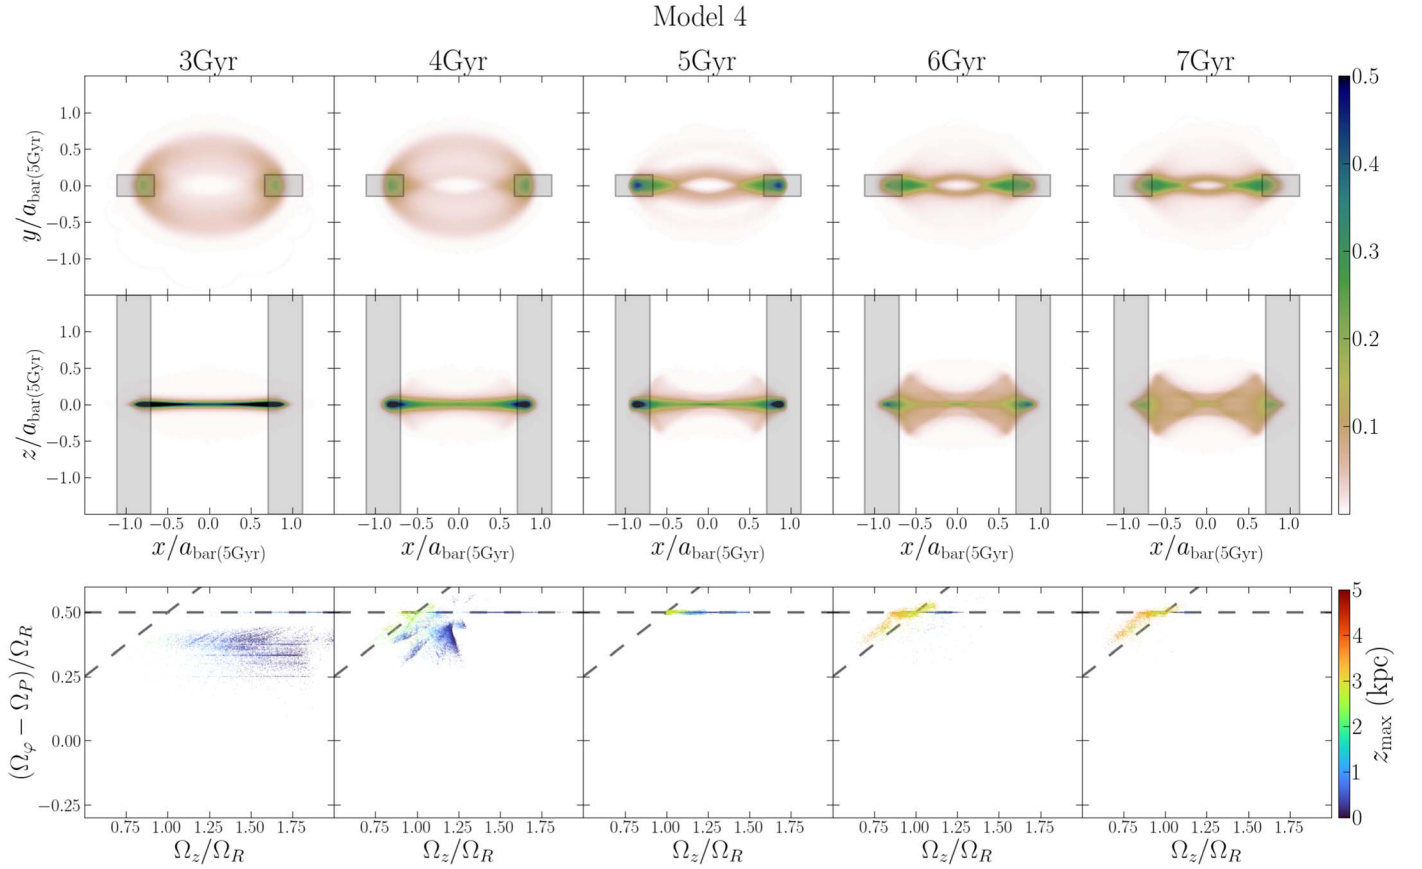
Figure 28. Similarly to Figure 25, but for Model 4. We see the same trends observed in the models SD1 and HG1, but with the abrupt change from 3 to 4 Gyr due to the buckling bar.

shape (in the *x–z* plane) and, depending on the strength of the vertical perturbation, to the shoulders (in the *x–y* plane)—see, e.g., Figure 27 (9 Gyr); (ii) in the study of the orbital support of box-peanut (BP) bulges, the importance of the vILR may be underestimated if we are restricted to orbits trapped at it and exclude the vILR cloud, since the latter is populated by orbits crossing the vILR and acquiring a long-term BP-morphology.

In a detailed observational study with the S[4]G survey (Sheth et al. 2010), Erwin et al. (2023) found some cases of galaxies with shoulders but no BP-bulge, and no clear case of bars with BP-bulges and no shoulders, suggesting that shoulders appear before BP-bulges. Given the role of the vILR in producing BP-bulge-supporting orbits and the evolution of shoulder-supporting orbits along the ILR, from the vertically warm ones to the vILR cloud after crossing the vILR, the suggestion of Erwin et al. (2023) emerges as a natural outcome.

The presence of shoulders has been found to correlate with massive (classical) bulges, or centrally concentrated halos (Kim et al. 2015; Kruk et al. 2018; Athanassoula & Misiriotis 2002). Furthermore, Athanassoula (1992) found that looped orbits tend to appear in galaxies with massive bulges (or high central concentration, in general), and with more elongated and slowly rotating bars (looped orbits in a slow bar have also been detected by Chaves-Velasquez et al. 2017). This is in agreement with the important contribution of looped orbits to shoulders, suggested by Paper I and confirmed here. Additionally, among our three models with shoulders, two (SD1 and Model 4) have slow bars and one (HG1, with weak shoulders) has a fast bar, suggesting that shoulders are not restricted to slow bars, but can also be present, with varying strength, in fast

bars. All models have clear signs of BP-bulges, although they are weaker for models SD1S and HG1, which suggests that the correlation shoulder-bulge is not restricted to either BP- or classical bulges.

A possible interpretation of the shoulder-massive bulges correlation might be that massive bulges induce thicker orbits in their outskirts, facilitating occupation of the shoulder-supporting warm orbits at the ILR, while bars in bulgeless galaxies would tend to trap stars in thinner orbits, i.e., with larger Ω_z/Ω_R such as the cool orbits at the ILR, with a morphology that does not support shoulders—see Figure 7. This line of thought would shift a massive bulge toward being the *cause* for shoulders, which would be in contradiction with the apparent precedence of shoulders with respect to BP-bulges suggested by Erwin et al. (2023).

On the other hand, warm orbits at the ILR crossing the vILR are vertically excited and change into chaotic orbits. As we have shown, this is accompanied by these orbits shrinking in the *x–y* plane and passing close to the center (see, e.g., Figure 17, panels (b) and (c)), overall increasing the density of the central regions. This may explain the sudden increase in the central surface density after buckling events observed in simulated galaxies (Raha et al. 1991; Debattista et al. 2006), since this brings orbits toward the vILR—see Figure 24. Therefore, this mass transfer to the center may contribute to the growth of massive bulges, which, instead of causing the appearance of shoulders, would be the *consequence* of the evolution of shoulder-supporting orbits toward the vILR.

Kim et al. (2015) suggested that shoulders are indicative of dynamically evolved bars, which should be formed with
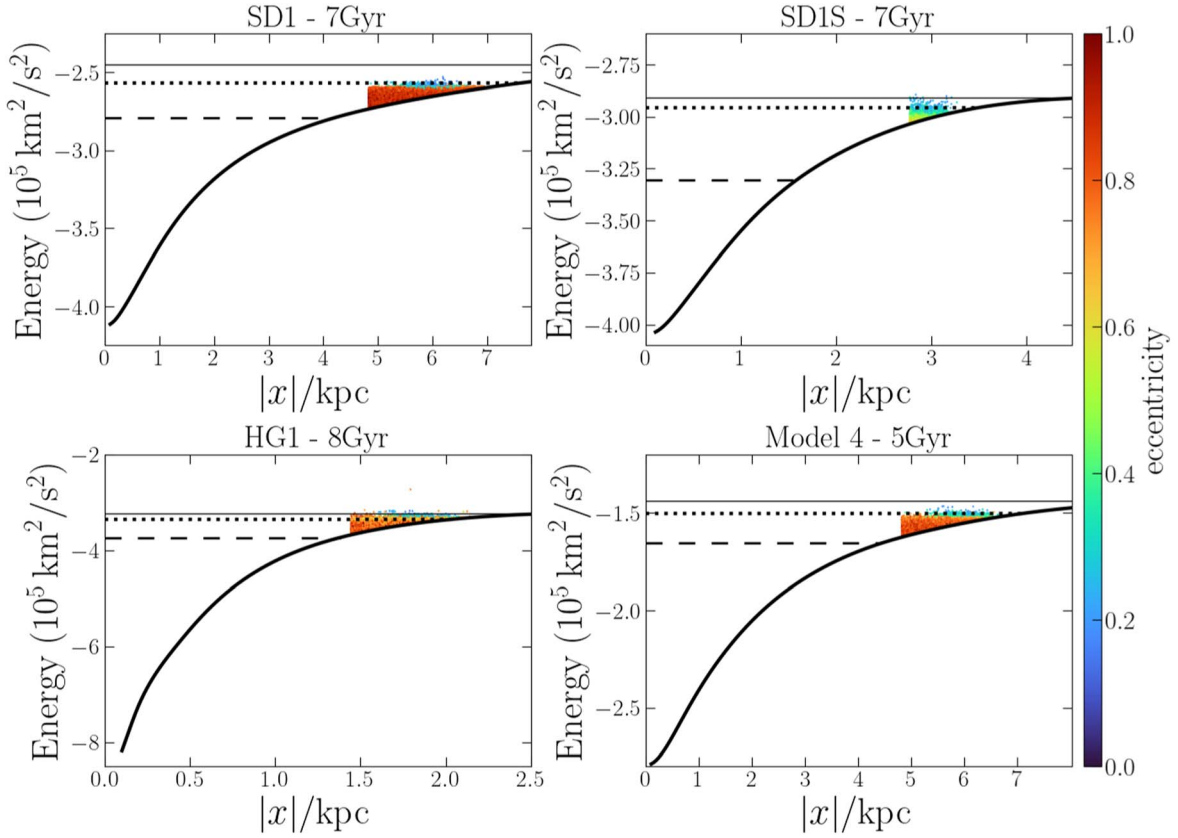
Figure 29. Similar to Figure 5, but only showing E_J (points) for warm orbits at the ILR. These orbits are confined to within the UHr.

exponential profiles inherited from the disk. They also suggest that this might explain the correlation with the galaxy's stellar mass and Hubble type, with more-massive early-type spirals forming bars earlier. Although the eventual erasure of shoulders found by Paper I complicates this narrative, this would agree with the strong dependence of shoulders on stellar mass found by Erwin et al. (2023). Our finding that the most shoulder-supporting morphology appears in orbits at the ILR with intermediate thickness (thus requiring a minimum thickness) is also in line with shoulders being associated with dynamically evolved bars. Additionally, the evolution of these orbits toward vertically thicker orbits supporting bulges also agrees with the suggestion of Erwin et al. (2023) that shoulders seem to precede BP-bulges.

5.4. Shoulder Definition and Possible Observational Implications

Our results show that the most relevant orbits to shoulders are the looped ones at the vertically warm ILR. In this way, one can expect shoulders to be restricted to the region where these orbits are present. Similarly to Figure 5, in Figure 29 we show the effective potential in the bar-major-axis, as well as the approximate locations of the corotation, UHR, and the ILR. We also show the Jacobi integral, but now restricting to warm orbits at the ILR.

It is clear that these orbits are restricted to the region inside the ultraharmonic radius. In Figure 5, we see that for the model SD1, the outer edge of the shoulder approximately coincides with the ultraharmonic radius. However, for the other simulations, the shoulder outer edge (as detected in Paper I) can be significantly outside this radius.

This observation is relevant for two reasons: on one hand, this suggests that identifying the outer edge of the shoulders based on a minimum of the curvature radius of the density profile (as implemented in Paper I) might be overestimating the shoulder extent, if this is to be identified with the region of looped orbits at the ILR. If that is the case, our current selection may be contaminated by a significant number of orbits that are not associated with the shoulders. Thus, the contribution of the looped orbits to shoulders we inferred in Section 4.4 represents a lower bound to their true importance to shoulders. On the other hand, if we can infer that the shoulders should end at the ultraharmonic radius, this can potentially be translated into constraints on the mass distribution and/or bar pattern speed of galaxies where shoulders are detected, or even in the Milky Way.

6. Summary

In this paper, we explored three pure *N*-body models and one *N*-body+SPH model, which form bars self-consistently. Three of these models, including the one with gas, develop shoulders in the bar-major-axis density profile. Among these, one has a bar that buckles. We select star particles in snapshots and regions where shoulders are identified, and integrate their orbits at different times in frozen potentials rotating with the corresponding bar pattern speeds. We perform frequency analysis of these orbits, using cylindrical coordinates in the inertial frame and selecting the leading frequency, presenting results in the frequency space ($\Omega_\varphi - \Omega_P)/\Omega_R$ versus Ω_z/Ω_R. We evaluate the contribution of different orbital types to the density in the shoulders region, and follow their evolution. Our main conclusions are as follows:

1. We confirm the suggestion of Paper I of the importance of orbits with loops at their ends (in the x–y plane) to the shoulders in bar density profiles.

2. We show that this orbital shape typically happens for orbits at the ILR $[(\Omega_\varphi - \Omega_P)/\Omega_R = 1/2]$ and with $1 \lesssim \Omega_z/\Omega_R \lesssim 3/2$ (vertically warm). Among the orbits at the ILR, those with $\Omega_z/\Omega_R \gtrsim 3/2$ (cool orbits) have a mildly elliptical shape and contribute negligibly to the shoulders.

3. As the bar slows down, the main resonances sweep the disk and, for a fixed set of particles, the ratio $(\Omega_\varphi - \Omega_P)/\Omega_R$ increases, until they are trapped by the ILR. While the bar thickens, the ratio Ω_z/Ω_R tends to decrease, causing a significant number of orbits to cross the vILR (Figures 21–24).

4. In agreement with Quillen et al. (2014) and Sellwood & Gerhard (2020), we find that the vILR does not trap stars for long periods of time. However, orbits crossing it are permanently vertically excited, developing an X-shape in the x–z plane—see also Wheeler et al. (2023).

5. Warm orbits at the ILR crossing the vILR toward $\Omega_z/\Omega_R \lesssim 1$ are vertically excited at the expense of in-plane motion, become chaotic, and populate a region around the vILR in the frequency map that we dub the *vILR cloud*. In comparison to warm orbits at the ILR, those at the vILR cloud are less extended along the x-axis and produce a more diffuse net figure in the x–y plane (Figures 25–28).

6. Orbits at the vILR cloud also contribute significantly to shoulders, but the larger their vertical excursions, the less shoulder-supporting they are.

7. As the bar thickens, the orbital support for shoulders can increase with orbits at the ILR entering the range of warm orbits, and potentially decrease once they move to $\Omega_z/\Omega_R \lesssim 1$ (where they cross the vILR). This explains why shoulders can develop right after a buckling event but can also be erased right after a second buckling (as noted in Paper I), since these can suddenly either populate or de-populate those orbits. This simple evolutionary sequence supports the suggestion of Erwin et al. (2023) that shoulders seem to appear before BP-bulges (with a possible co-formation of shoulders and BP-bulges in buckling bars).

8. An analysis with the isochrone model (Appendix B) shows that for orbits with large apocenters, the leading frequency in the azimuthal spectrum is the precession one, instead of Ω_φ. This has possible implications for the analysis of Milky Way halo stars and deserves further investigation. Moreover, frequency estimates with the time-series $f_\varphi = \sqrt{2|L_z|}(\cos \varphi + i \sin \varphi)$ provide higher precision (compared to $f_\varphi = \varphi$) and facilitate its extraction —see Figure B1.

In conclusion, bar thickening and vertical resonances can dilute the shoulders in the bar-major-axis density profiles, unless the bar keeps trapping new orbits in the appropriate frequency ratios, with the looped morphology. This sheds light on the conclusion of Paper I that long-term shoulders require a growing bar. Our results help to build a coherent picture tying together known analytical results with several recent observational and simulation-based results on the properties and evolution of shoulders in the density profiles of bars.

Appendix A
Frequency Analysis Algorithm

Here we present the algorithm used to extract frequencies of orbits, which is essentially the same as NAFF, proposed by Laskar et al. (1992) and further developed by Valluri & Merritt (1998)—see also Price-Whelan et al. (2016). We refer the reader to these papers for a complete description. We also comment on some details and improvements implemented in the python package used in this work, which we dub naif (Beraldo e Silva & Valluri 2023)—the documentation is available at http://naif.readthedocs.io/en/latest/index.html. This preliminary version extracts peak frequencies very accurately, as demonstrated in Appendix B, and will be improved in the future with several tools for frequency analysis.

Assume that an orbit is integrated for a total time T and let, for simplicity of notation, the time variable be defined in the symmetric interval $-T/2 \leqslant t \leqslant T/2$. In practice, t is an array with N elements t_n, and we have a (real or complex) time-series $f(t_n)$ associated with each coordinate. We start by performing a windowed discrete Fourier transform

$$F_j = \frac{1}{N}\sum_{n=0}^{N-1} f(t_n)\chi_p(t_n)e^{-2\pi i n j/N}, \qquad (A1)$$

where $j = -N/2,\ldots,0\ldots N/2$, and

$$\chi_p(t) = \frac{2^p(p!)^2}{(2p)!}\left[1 + \cos\left(\frac{2\pi t}{T}\right)\right]^p \qquad (A2)$$

is the window function, with $p \in \mathbb{N}$. In this work, we use $p = 1$, but other values might improve the estimate (Laskar 1999; Hunter 2002).

The spectrum is calculated at frequencies $\omega_j = 2\pi j/T$, and the minimum frequency that can be probed is $\omega_{\min} = 2\pi/T$. On the other hand, the maximum (Nyquist) frequency is $\Omega_{\mathrm{Ny}} = \Omega_{j=N/2} = N\pi/T$. If we conservatively require $\Omega_{\mathrm{Ny}} \approx 10^3\,\mathrm{Gyr}^{-1}$, this requires $N \gtrsim (10^3/\pi)(T/\,\mathrm{Gyr})$. For an orbit with a period of circular motion $T_{\mathrm{circ}} \approx 0.25\,\mathrm{Gyr}$, integrated for $T \approx 20 T_{\mathrm{circ}} \approx 5\,\mathrm{Gyr}$ ($T \approx 40 T_{\mathrm{circ}} \approx 10\,\mathrm{Gyr}$), we have $\omega_{\min} \approx 1.25\,\mathrm{Gyr}^{-1}$ ($\omega_{\min} \approx 0.63\,\mathrm{Gyr}^{-1}$), and we need $N \gtrsim 1500$ ($N \gtrsim 3000$) points. Using $N \sim 10^4$ points is thus safe for most applications.

A rough estimate of the leading frequency is the spectral line with maximum amplitude $|F_j|$, $\omega_{\max}$. Since the spectrum is defined in the discrete set $\omega_j = 2\pi j/T$, this estimate's precision is limited by the resolution $\Delta\omega = 2\pi/T$. This estimate is then refined by calculating the projection $\phi(\omega) = \langle f(t), e^{i\omega t}\rangle$ of the time-series $f(t)$ onto the "frequency vector" $e^{i\omega t}$, where

$$\langle f(t), g(t)\rangle \equiv \frac{1}{T}\int_{-T/2}^{T/2} f(t)g^*(t)\chi_p(t)\,\mathrm{d}t. \tag{A3}$$

The refined frequency is found by maximizing $|\phi(\omega)|$ around $\omega_{\max}$. The maximum is determined with the optimized Brent's algorithm, as implemented in scipy. If the maximum cannot be identified in this first shot, a local maximum is identified via a brute force search. The interval around $\omega_{\max}$ used in the original NAFF implementation is $\omega_{\max} \pm 2\pi/T$, but we have found that the interval $\omega_{\max} \pm 4\pi/T$ is more appropriate to guarantee that Brent's algorithm will find the maximum, and this is the interval used in naif. Having thus identified this first leading frequency ω_1, we calculate the complex amplitude a_1 associated with it, where

$$a_k = \langle f(t), u_k\rangle, \tag{A4}$$

and $u_k = e^{i\omega_k t}$. If we want to extract additional frequencies, we "subtract" this spectral line from the time-series, defining in general

$$f_k(t) = f_{k-1}(t) - a_k e_k, \tag{A5}$$

where $e_1 = u_1$, $f_k(t)$ is the time-series after extracting the kth frequency ω_k, and $f_0(t)$ is the original one. Here, e_k denotes an orthonormal basis, built from the u_k's via a Gram–Schmidt process:

$$e_k = u_k - \sum_{j=1}^{k-1} \langle u_k, e_j\rangle e_j. \tag{A6}$$

We then iterate through Equations (A1)–(A6) to extract as many frequencies as desired.

Compared to the original NAFF algorithm, a small (but relevant) improvement of naif is as follows: NAFF first applies a DFT without windowing to look for the rough estimate $\omega_{\max}$, and then uses the window function $\chi_p(t)$ in the refinement step when maximizing $|\phi(\omega)|$. This different procedure to look for the same frequency ends up changing the order of the highest amplitudes in some cases. This is later corrected by sorting the frequencies by amplitude and taking the leading one, but requires extraction of some frequencies even if we are only interested in the leading one. On the other hand, naif performs a windowed DFT, Equation (A1), already in the first search for $\omega_{\max}$, which ends up always being the leading frequency. This removes the need to extract additional frequencies if we are only interested in the leading one.

Finally, naif provides easy access to the spectrum before the extraction of each frequency and also to $|\phi(\omega)|$ around each extracted frequency, which can be useful for detailed investigations of individual orbits.

Appendix B
Numerical Frequencies for the Isochrone Model

We use AGAMA to generate a self-consistent sample of the isochrone model and integrate each orbit for ≈ 100 azimuthal periods. A self-consistent ensemble is not strictly necessary, since we simply compare frequencies of orbits individually and do not focus on their relative numbers. However, this guarantees a broad phase-space region for investigation. Since this is a spherical model, we only estimate the radial and azimuthal frequencies, Ω_r and Ω_φ, respectively (the transversal frequency $\Omega_\vartheta = \pm\Omega_\varphi$). For this model, we set $GM = 1$.

Figure B1 shows the relative errors of the estimated frequencies for 10^3 orbits, in comparison to the analytic expressions for the isochrone model (Binney & Tremaine 2008). Points are color-coded by the apocenter radii (in units of the scale radius). The left column shows results obtained using as input the real time-series

$$
\begin{aligned}
f_r &= r \\
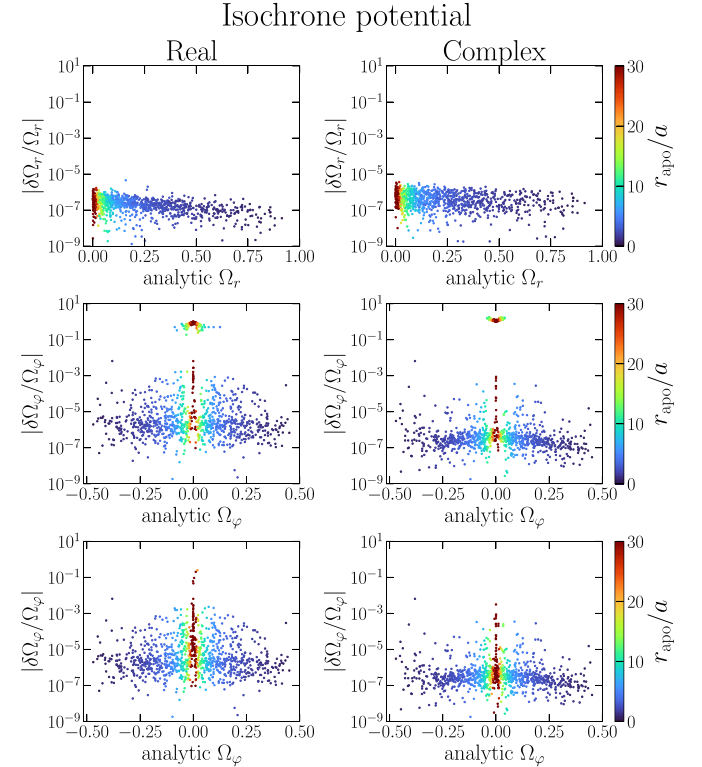f_\varphi &= \varphi,
\end{aligned}
\tag{B1}
$$

Figure B1. Relative errors in the fundamental frequencies for 10^3 orbits in the isochrone potential. Left column: real time-series, Equation (B1). Right column: complex time-series, Equation (B2). The radial frequency (top row) is estimated as the absolute value of the leading frequency. Real and complex time-series give similar errors. For Ω_φ (middle row), the real time-series requires setting $\mathrm{sign}(\Omega_\varphi) = \mathrm{sign}(L_z)$, while this is automatically set for the complex time-series. Requiring $\Omega_\varphi > \Omega_r/2$ (see the main text) and taking the leading frequency, we have $|\delta\Omega_\varphi/\Omega_\varphi| \lesssim 10^{-3}$ (bottom row). The complex time-series provides $\approx 10\times$ smaller errors for Ω_φ.

and the right column shows results for the complex time-series (Papaphilippou & Laskar 1996, 1998)

$$f_r = r + iv_r$$
$$f_\varphi = \sqrt{2|L_z|}\,(\cos\varphi + i\sin\varphi), \tag{B2}$$

where $v_r = dr/dt$ is the radial velocity, and L_z is the z-component of the angular momentum. For the radial coordinate (first row), the absolute value of the leading frequency has a relative error $|\delta\Omega_r/\Omega_r| \approx 10^{-6}$, using either the real or complex time-series.

For the azimuthal angle (second row), using the real time-series $f_\varphi = \varphi$, the leading frequencies provide errors $|\delta\Omega_\varphi/\Omega_\varphi| \lesssim 10^{-3}$ for a fraction of the orbits, but only after imposing $\mathrm{sign}(\Omega_\varphi) = \mathrm{sign}(L_z)$, while the complex time-series (Equation (B2)) does not require this intervention. Additionally, there are a few orbits with small analytic Ω_φ and $|\delta\Omega_\varphi/\Omega_\varphi| \approx 1$ for both real and complex time-series. These orbits have large apocenters ($r_{\rm apo}/a \gtrsim 15$), and spend most of their time far from the center, effectively behaving as precessing ellipses in a Kepler-like potential. As a consequence, the leading frequency in the spectrum is $\Omega_r - \Omega_\varphi$ (the precessing frequency), although Ω_φ is also present, with a smaller amplitude. In a Kepler-like potential, $\Omega_r \approx \Omega_\varphi$, thus the leading frequency is ≈ 0, which implies $|\delta\Omega_\varphi/\Omega_\varphi| \approx 1$. In realistic galactic potentials, typical stars have $\Omega \lesssim \kappa \lesssim 2\Omega$. Thus, if we heuristically require $|\Omega_\varphi| > |\Omega_r|/2$, and take the leading frequency respecting this restriction, the fundamental Ω_φ is recovered with good accuracy ($|\delta\Omega_\varphi/\Omega_\varphi| \lesssim 10^{-3}$) for almost all orbits (third row). Finally, the complex time-series provides a $\approx 10\times$ better precision (although the estimates are almost identical to those of the real series if orbits are integrated for only ≈ 50 circular periods).

Appendix C
Inertial versus Rotating Frame

Here we compare the frequencies obtained using coordinates in the inertial frame and in the bar rotating frame. We use the model SD1S, which has a fast bar at 7 Gyr, and the region where star particles are selected is near the corotation radius—see Figure 4. Thus, we can expect the corotation resonance to be populated in the frequency map, as is the case when we use the inertial frame (Figure 13). In this case, the frequency analysis outputs Ω_φ, and we literally subtract the bar pattern speed Ω_P before plotting the frequency map in Figure 13.

Figure C1 shows the frequency map for the same snapshot and the same orbits, but now using coordinates in the bar rotating frame. With Ω_φ still representing the azimuthal frequency in the inertial frame, the output of the frequency analysis is now $\Omega_\varphi - \Omega_P$.

Comparing Figures 13 and C1, we see that the two frequency maps are almost identical, except for the corotation ($\Omega_\varphi = \Omega_P$), which is sharply defined in the inertial frame, but is completely lost if we use coordinates in the rotating frame. This is probably due to the leading frequency in this case being the libration around the corotation, which might be relevant for detailed investigations of orbits trapped at corotation, but it is not of interest for this work.
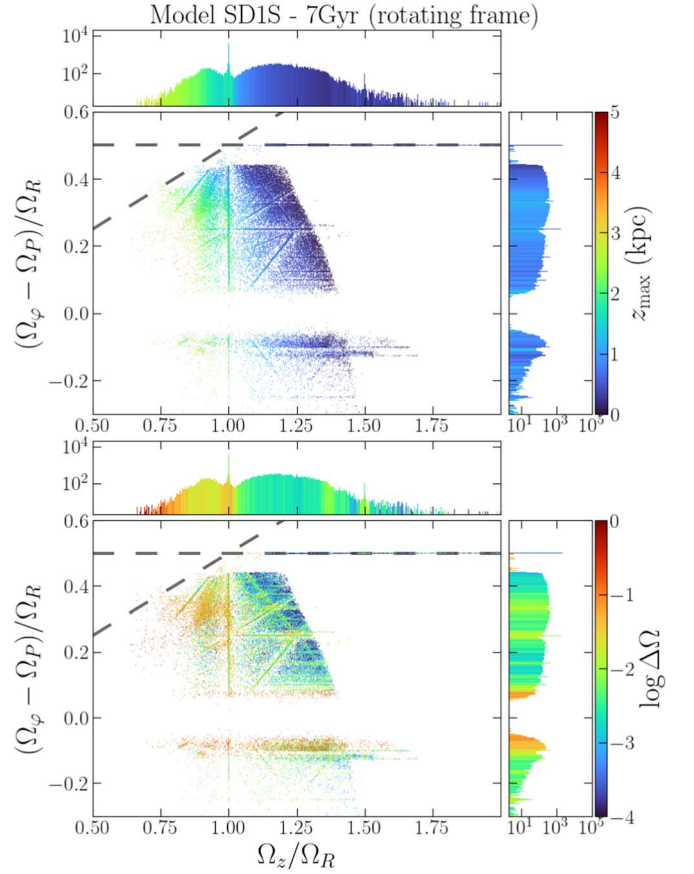

Figure C1. Frequency maps for the same orbits as in Figure 13, but here with frequencies estimated using coordinates in the rotating frame. The two frequency maps are almost identical, except for the corotation resonance, which is not correctly identified in this case.

ORCID iDs

Leandro Beraldo e Silva ⓘ https://orcid.org/0000-0002-0740-1507
Victor P. Debattista ⓘ https://orcid.org/0000-0001-7902-0116
Stuart Robert Anderson ⓘ https://orcid.org/0000-0002-6379-5053
Monica Valluri ⓘ https://orcid.org/0000-0002-6257-2341
Peter Erwin ⓘ https://orcid.org/0000-0003-4588-9555
Kathryne J. Daniel ⓘ https://orcid.org/0000-0003-2594-8052
Nathan Deg ⓘ https://orcid.org/0000-0003-3523-7633

References

Abbott, C. G., Valluri, M., Shen, J., & Debattista, V. P. 2017, MNRAS, 470, 1526
Anderson, S. R., Debattista, V. P., Erwin, P., et al. 2022, MNRAS, 513, 1642
Arnol'd, V. I. 1963, RuMaS, 18, 9
Athanassoula, E. 1992, MNRAS, 259, 328
Athanassoula, E. 2002, ApJL, 569, L83
Athanassoula, E. 2003, MNRAS, 341, 1179
Athanassoula, E. 2013, in Secular Evolution of Galaxies, ed. J. Falcón-Barroso & J. H. Knapen (Cambridge: Cambridge Univ. Press), 305
Athanassoula, E., & Beaton, R. L. 2006, MNRAS, 370, 1499
Athanassoula, E., Bienayme, O., Martinet, L., & Pfenniger, D. 1983, A&A, 127, 349
Athanassoula, E., & Misiriotis, A. 2002, MNRAS, 330, 35
Beraldo e Silva, L., de Siqueira Pedra, W., Valluri, M., Sodré, L., & Bru, J.-B. 2019, ApJ, 870, 128
Beraldo e Silva, L., & Valluri, M. 2023, naif: Frequency analysis package, Astrophysics Source Code Library, ascl:2303.004

Binney, J. 1981, MNRAS, 196, 455
Binney, J., & Spergel, D. 1982, ApJ, 252, 308
Binney, J., & Tremaine, S. 2008, Galactic Dynamics (2nd ed.; Princeton, NJ: Princeton Univ. Press)
Bullock, J. S., Dekel, A., Kolatt, T. S., et al. 2001, ApJ, 555, 240
Bureau, M., & Athanassoula, E. 2005, ApJ, 626, 159
Carpintero, D. D., & Aguilar, L. A. 1998, MNRAS, 298, 1
Ceverino, D., & Klypin, A. 2007, MNRAS, 379, 1155
Chaves-Velasquez, L., Patsis, P. A., Puerari, I., Skokos, C., & Manos, T. 2017, ApJ, 850, 145
Chirikov, B. V. 1979, PhR, 52, 263
Cole, D. R., Debattista, V. P., Erwin, P., Earp, S. W. F., & Roškar, R. 2014, MNRAS, 445, 3352
Collier, A., Shlosman, I., & Heller, C. 2018, MNRAS, 476, 1331
Combes, F., Debbasch, F., Friedli, D., & Pfenniger, D. 1990, A&A, 233, 82
Combes, F., & Elmegreen, B. G. 1993, A&A, 271, 391
Contopoulos, G. 1978, A&A, 64, 323
Contopoulos, G. 1981, A&A, 102, 265
Contopoulos, G. 1988, A&A, 201, 44
Contopoulos, G., & Grosbol, P. 1989, A&ARv, 1, 261
Contopoulos, G., & Papayannopoulos, T. 1980, A&A, 92, 33
de Carvalho, R. R., & da Costa, L. N. 1987, A&A, 171, 66
Debattista, V. P., Liddicott, D. J., Khachaturyants, T., & Beraldo e Silva, L. 2020, MNRAS, 498, 3334
Debattista, V. P., Mayer, L., Carollo, C. M., et al. 2006, ApJ, 645, 209
Debattista, V. P., Ness, M., Gonzalez, O. A., et al. 2017, MNRAS, 469, 1587
Debattista, V. P., & Sellwood, J. A. 1998, ApJL, 493, L5
Debattista, V. P., & Sellwood, J. A. 2000, ApJ, 543, 704
Dodd, E., Helmi, A., & Koppelman, H. H. 2022, A&A, 659, A61
Elmegreen, B. G., & Elmegreen, D. M. 1985, ApJ, 288, 438
Elmegreen, B. G., Elmegreen, D. M., Chromey, F. R., Hasselbacher, D. A., & Bissell, B. A. 1996, AJ, 111, 2233
Elmegreen, D. M., Elmegreen, B. G., Yau, A., et al. 2011, ApJ, 737, 32
Erwin, P. 2018, MNRAS, 474, 5372
Erwin, P., Debattista, V. P., & Anderson, S. R. 2023, MNRAS, 524, 3166
Eskridge, P. B., Frogel, J. A., Pogge, R. W., et al. 2000, AJ, 119, 536
Fragkoudi, F., Grand, R. J. J., Pakmor, R., et al. 2021, A&A, 650, L16
Gadotti, D. A., Athanassoula, E., Carrasco, L., et al. 2007, MNRAS, 381, 943
Gajda, G., Łokas, E. L., & Athanassoula, E. 2016, ApJ, 830, 108
Gardner, E., Debattista, V. P., Robin, A. C., Vásquez, S., & Zoccali, M. 2014, MNRAS, 438, 3275
Hamilton, C., Tolman, E. A., Arzamasskiy, L., & Duarte, V. N. 2023, ApJ, 954, 12
Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Natur, 585, 357
Hernquist, L., & Weinberg, M. D. 1992, ApJ, 400, 80
Hunter, C. 2002, SSRv, 102, 83
Kim, T., Sheth, K., Gadotti, D. A., et al. 2015, ApJ, 799, 99
Kolmogorov, A. 1954, Dokl. Akad. Nauk. SSR, 98, 2
Kruk, S. J., Lintott, C. J., Bamford, S. P., et al. 2018, MNRAS, 473, 4731
Kuijken, K., & Dubinski, J. 1995, MNRAS, 277, 1341
Laskar, J. 1990, Icar, 88, 266
Laskar, J. 1999, in Introduction to Frequency Map Analysis, ed. C. Simó (Dordrecht: Springer), 134
Laskar, J., Froeschlé, C., & Celletti, A. 1992, PhyD, 56, 253
Long, S., Shlosman, I., & Heller, C. 2014, ApJL, 783, L18

Lucey, M., Pearson, S., Hunt, J. A. S., et al. 2023, MNRAS, 520, 4779
Manos, T., Skokos, C., & Patsis, P. A. 2022, MNRAS, 509, 1995
Martinez-Valpuesta, I., Shlosman, I., & Heller, C. 2006, ApJ, 637, 214
Menéndez-Delmestre, K., Sheth, K., Schinnerer, E., Jarrett, T. H., & Scoville, N. Z. 2007, ApJ, 657, 790
Moser, J. 1962, Nachr. Akad. Wiss. Gött., II. Math.-Phys. Kl., 1962, 1
Nair, P. B., & Abraham, R. G. 2010, ApJL, 714, L260
Navarro, J. F., Frenk, C. S., & White, S. D. M. 1996, ApJ, 462, 563
O'Neill, J. K., & Dubinski, J. 2003, MNRAS, 346, 251
Papaphilippou, Y., & Laskar, J. 1996, A&A, 307, 427
Papaphilippou, Y., & Laskar, J. 1998, A&A, 329, 451
Papayannopoulos, T., & Petrou, M. 1983, A&A, 119, 21
Parul, H. D., Smirnov, A. A., & Sotnikova, N. Y. 2020, ApJ, 895, 12
Patsis, P. A., & Athanassoula, E. 2019, MNRAS, 490, 2740
Patsis, P. A., & Katsanikas, M. 2014, MNRAS, 445, 3546
Patsis, P. A., Skokos, C., & Athanassoula, E. 2002, MNRAS, 337, 578
Petersen, M. S., Weinberg, M. D., & Katz, N. 2021, MNRAS, 500, 838
Pfenniger, D., & Friedli, D. 1991, A&A, 252, 75
Pontzen, A., Roškar, R., Stinson, G., & Woods, R. 2013, pynbody: N-Body/SPH analysis for python, Astrophysics Source Code Library, ascl:1305.002
Portail, M., Wegg, C., & Gerhard, O. 2015, MNRAS Lett., 450, L66
Price-Whelan, A. M., Johnston, K. V., Valluri, M., et al. 2016, MNRAS, 455, 1079
Prieto, M., Aguerri, J. A. L., Varela, A. M., & Muñoz-Tuñón, C. 2001, A&A, 367, 405
Quillen, A. C. 2002, AJ, 124, 722
Quillen, A. C., Minchev, I., Sharma, S., Qin, Y.-J., & Di Matteo, P. 2014, MNRAS, 437, 1284
Raha, N., Sellwood, J. A., James, R. A., & Kahn, F. D. 1991, Natur, 352, 411
Regan, M. W., & Elmegreen, D. M. 1997, AJ, 114, 965
Roy, A. E. 1991, Predictability, stability, and chaos in N-body dynamical systems. Proceedings (NATO Advanced Study Institute (ASI) Series B), Vol. 272 (New York, NY: Springer)
Sellwood, J. A., & Gerhard, O. 2020, MNRAS, 495, 3175
Sellwood, J. A., & Wilkinson, A. 1993, RPPh, 56, 173
Sheth, K., Elmegreen, D. M., Elmegreen, B. G., et al. 2008, ApJ, 675, 1141
Sheth, K., Regan, M., Hinz, J. L., et al. 2010, PASP, 122, 1397
Skokos, C., Patsis, P. A., & Athanassoula, E. 2002, MNRAS, 333, 847
Smirnov, A. A., Tikhonenko, I. S., & Sotnikova, N. Y. 2021, MNRAS, 502, 4689
Sparke, L. S., & Sellwood, J. A. 1987, MNRAS, 225, 653
Stadel, J. G. 2001, Ph.D Thesis, Univ. Washington
Stinson, G., Seth, A., Katz, N., et al. 2006, MNRAS, 373, 1074
Valluri, M., Debattista, V. P., Quinn, T., & Moore, B. 2010, MNRAS, 403, 525
Valluri, M., & Merritt, D. 1998, ApJ, 506, 686
Valluri, M., Shen, J., Abbott, C., & Debattista, V. P. 2016, ApJ, 818, 141
van Albada, T. S., & Sanders, R. H. 1982, MNRAS, 201, 303
Vasiliev, E. 2019, MNRAS, 482, 1525
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261
Wadsley, J. W., Stadel, J., & Quinn, T. 2004, NewA, 9, 137
Weinberg, M. D., & Katz, N. 2007, MNRAS, 375, 425
Wheeler, V., Valluri, M., Beraldo e Silva, L., Dattathri, S., & Debattista, V. P. 2023, arXiv:2306.07424
Widrow, L. M., & Dubinski, J. 2005, ApJ, 631, 838