Their currents turn awry, and lose the name of action. I: Fundamental limits to orbit reconstruction due to non-conservation of stellar actions

Arunima Arunima,[1★] Mark R. Krumholz,[1] Michael J. Ireland[1], Chuhan Zhang[1] and Zipeng Hu[2]
[1]Research School of Astronomy and Astrophysics, Australian National University, Canberra ACT 2601, Australia
[2]Kavli Institute for Astronomy and Astrophysics, Peking University, Beijing 100871, China

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

The conservation of stellar actions is a fundamental assumption in orbit reconstruction studies in the Milky Way. However, the disc is highly dynamic, with time-dependent, non-axisymmetric features like transient spiral arms and giant molecular clouds (GMCs) driving local fluctuations in the gravitational potential on top of the near-axisymmetric background. Using high-resolution magnetohydrodynamic simulations that incorporate gas dynamics and star formation, we quantify the rate at which these effects drive non-conservation of the actions of young stars from Myr to Gyr timescales. We find that action evolution is well described as a logarithmic random walk, with vertical action evolving more rapidly than radial action; the diffusion rate associated with this random walk is weakly dependent on the stellar birth environment and scales approximately linearly with the galactic orbital frequency at a star's position. The diffusion rates we measure imply a fundamental limit of ~ 100 Myr as the timescale over which stellar orbits can be reliably reconstructed using methods that assume action conservation. By comparing diffusion rates for younger stars to those measured for an older and more vertically-extended control population, we conclude that radial action evolution is driven primarily by transient spiral arms, while vertical action evolution is driven by gravitational scattering off gaseous structures. Our results have significant implications for galactic archaeology and disc dynamics studies, necessitating a closer look at the timescales over which actions are assumed to be conserved in the disc.

Key words: Galaxy: kinematics and dynamics – astrometry

1 INTRODUCTION

Understanding the dynamical evolution of stellar orbits is fundamental to reconstructing the past history of galaxies. The use of actions – adiabatic invariants in an axisymmetric potential – has been quite successful for this purpose in galactic studies. In the Milky Way, action-space analyses have been widely employed in studies identifying merger remnants and accreted substructures, as well as in globular cluster studies (Helmi et al. 2018; Myeong et al. 2019; Feuillet et al. 2020; Lane et al. 2022; Malhan et al. 2022; Limberg et al. 2022; Callingham et al. 2022; Chen & Gnedin 2022; Cabrera Garcia et al. 2024). While these approaches have been highly effective in the Galactic halo, where the environment is relatively dynamically quiet, applying action-based reconstruction methods to the disc presents additional challenges.

The orbits of stars within the galactic disc are shaped by both external perturbations (Antoja et al. 2018; Bland-Hawthorn et al. 2019; Li 2021; Antoja et al. 2023; Darragh-Ford et al. 2023; Frankel et al. 2023) and by secular internal processes such as interactions with giant molecular clouds (GMCs), spiral arms and the bar (Sellwood & Binney 2002; Roškar et al. 2012; Vera-Ciro et al. 2014; Mackereth et al. 2019; Tremaine et al. 2023). These perturbations lead to deviations of the gravitational potential away from the time-invariant, axisymmetric state required for strict action conservation. Despite these complexities, efforts have been made to use actions for disc studies (Trick et al. 2019; Coronado et al. 2020, 2022).

One prominent application of such methods is in cluster reconstruction studies that employ traceback techniques to infer the past positions and birth environments of stars. These methods typically assume a static, axisymmetric Galactic potential and rely on the conservation of stellar actions to extrapolate stellar trajectories backward in time. Numerous studies have used traceback methods to reconstruct the dispersal history of open clusters in the Milky Way disc, particularly in the Solar neighbourhood (Miret-Roig et al. 2018, 2020, 2022; Galli et al. 2018; Crundall et al. 2019; Squicciarini et al. 2021; Heyl et al. 2021, 2022; Schoettler et al. 2020, 2022; Ma et al. 2022; Zucker et al. 2022; Galli et al. 2023; Couture et al. 2023; Pelkonen et al. 2024). However, the validity of these reconstruction studies depends critically on the assumption of individual stellar actions remaining conserved over the timescales involved.

Most existing studies of action evolution over time focus on the global distribution of actions, particularly in the context of radial migration and dynamical heating. Simulations have explored how secular evolution redistributes angular momentum and heats the disc over gigayear timescales (Roškar et al. 2008; Vera-Ciro et al. 2014; Monari et al. 2016; Vera-Ciro & D'Onghia 2016; Halle et al. 2018; Mikkola et al. 2020; Okalidis et al. 2022). Observational studies, on the other hand, study action evolution by comparing the present-day properties of stars of different ages, often in the context of, again, disc heating and radial migration (Frankel et al. 2018; Ting & Rix

★ E-mail: arunima.arunima@anu.edu.au

2019; Frankel et al. 2020). However, these studies primarily describe population-wide trends, while what matters for reconstruction studies is the time evolution of the individual stars' actions. This is analogous to the distinction between the microscopic motion of gas molecules and the macroscopic evolution of temperature in a thermodynamic system – such a system may have a completely constant, predictable temperature, yet it may still be impossible to backtrace the trajectory of any individual molecule over any substantial period.

To quantify the time evolution of individual stars' actions, action diffusion provides a useful statistical framework. Prior studies have explored the long-term evolution of action due to scattering by spiral arms and bar-driven resonances (e.g., Solway et al. 2012; Daniel & Wyse 2015; Halle et al. 2018; Kawata et al. 2021). However, short-term diffusion, particularly in newly formed stars, remains largely unexplored. Fujimoto et al. (2023) investigated GMC-driven scattering of stars on short timescales but did not compute actions, leaving open questions about perturbations in action space. Moreover, existing studies of action conservation have largely been conducted in N-body simulations, which lack gas dynamics and star formation (e.g., Solway et al. 2012; Vera-Ciro & D'Onghia 2016; Mikkola et al. 2020). Stars form in dense, turbulent gas clouds (Federrath & Klessen 2012; Krumholz 2014), where gravitational potential evolves rapidly due to gas accretion, stellar feedback and local dynamical instabilities. Hence, the initial conditions of the stars are inherently linked to an evolving potential, making a self-consistent treatment of gas dynamics essential for studying early action evolution. Hydrodynamic plus N-body simulations that include self-gravity, radiative cooling and galactic scale gas flows, such as spiral arms and galactic shear, are the ideal tool for the purpose of exploring these effects.

These considerations motivate the present study, in which we use a high-resolution magnetohydrodynamics (MHD) simulation of a Milky Way-like disc galaxy to study the evolution of stellar actions. By computing these at high temporal resolution, we aim to measure the rate of action diffusion and characterise the timescales over which the actions deviate from conservation in the disc, providing insights into the reliability of traceback methods that rely on these assumptions. The remainder of this paper is organised as follows. In section 2, we describe the galactic simulations used in this study, section 3 details our method used to calculate stellar actions and introduces key notation for studying their time evolution, section 4 presents our main results, including the distribution of actions, their temporal evolution, and dependence of the evolution on stars' birth environment, and in section 5 we discuss the broader implications of our findings before summarising our conclusions. The second paper in this series will examine the use of actions to identify and reconstruct dissolved star clusters.

2 SIMULATION

We analyse simulations of an isolated Milky Way-like disc galaxy with flocculent spiral structure and no bar taken from Zhang et al. (2025, hereafter Z25). This simulation is an extension of the full galaxy zoom-in simulations described by Wibking & Krumholz (2023) and Hu et al. (2023) (hereafter WK23 and H23, respectively). Here, we summarise the details of this simulation and direct readers to WK23, H23 and Z25 for further information.

2.1 Numerical method

The simulations solve the equations of ideal magnetohydrodynamics using the GIZMO code (Hopkins 2015, 2016; Hopkins & Raives 2016), with gas and metal line cooling implemented via the GRACKLE library (Smith et al. 2017) (see Appendix A in WK23 for justification of this choice). The treatment for star formation in the simulation is as follows: for gas particles with density ρ_g exceeding a critical threshold ρ_{crit}, the local star formation density rate is calculated as

$$\dot{\rho}_{SFR} = \epsilon_{ff} \frac{\rho_g}{t_{ff}}, \tag{1}$$

where ϵ_{ff} is the star formation efficiency, ρ_g is the local gas density and

$$t_{ff} = \sqrt{\frac{3\pi}{32 G \rho}} \tag{2}$$

is the local gas free-fall time. We set $\epsilon_{ff} = 0.01$, consistent with the mean value observed across a wide density range (Krumholz et al. 2019). The critical threshold density ρ_{crit} depends on the simulation resolution; as we discuss below, the simulation takes place in two stages, a low-resolution one to allow the galaxy to settle to statistical steady-state, followed by a higher resolution, shorter stage to capture more detail. During the initial phase $\rho_{crit} = 100 \mu m_H$ cm^{-3}, where m_H is the mass of a hydrogen atom and $\mu = 1.4$ is the mean mass per H nucleus for gas of standard cosmic composition; this value is chosen such that, for gas of density ρ_{crit} and at the equilibrium temperature implied by our cooling curve, the Jeans mass is nearly equal to the simulation mass resolution $\Delta M = 859.3$ M$_\odot$ – see WK23 for details. During the second stage of the simulations, we increase ρ_{crit} to $1000 \mu m_H$ cm^{-3}, which roughly maintains this condition at the increased resolution. To avoid excessively small time steps and limit the computational expense in following very dense regions, we increase ϵ_{ff} to 10^6 for gas particles with $\rho_g \geq 100 \rho_{crit}$, forcing them to be converted instantly into collisionless star particles. Star formation is implemented stochastically, such that the probability of a gas particle being converted to a star particle in a time step of size Δt is

$$P = 1 - \exp(-\epsilon_{ff} \Delta t / t_{ff}). \tag{3}$$

See Z25 for further details.

Once star particles form, they interact with the galactic environment only through gravity and feedback. Since the resolution is high enough during the high-resolution simulation phase that each star particle does not represent an entire cluster sampling the full initial mass function (IMF), we cannot use the GIZMO treatment of IMF-integrated stellar feedback. Instead, stellar feedback is determined on a star-by-star basis by forming each star particle from a number of individual stars, stochastically drawing a synthetic stellar population from a Chabrier (2005) IMF using the SLUG stellar population synthesis code (da Silva et al. 2012; Krumholz et al. 2015). The evolution of each star follows the Padova stellar tracks (Bressan et al. 2012). The star's atmosphere is modelled using SLUG's "starburst99" spectral synthesis method (Leitherer et al. 1999). This provides us with each star's ionising luminosity as a function of its mass and age, which is injected back into the simulation using a Strömgren volume method. We also determine which stars end their lives as supernovae or asymptotic giant branch (AGB) stars, and inject feedback from these events following the recipe for handling partially-resolved SNe described by Hopkins et al. (2018). For full details on the treatment of the feedback see Armillotta et al. (2019). The mass and metal return in each time step for each star particle are based on its mass and evolutionary stage, following Sukhbold et al. (2016) for type II supernovae, Karakas & Lugaro (2016) for AGB stars, and Doherty et al. (2014) for super-AGB stars.

2.2 Initial conditions

As mentioned above, the simulation takes place in two stages. The first, described by WK23, begins from the isolated Milky Way-analog AGORA project initial conditions (Kim et al. 2016) and run using an IMF-integrated treatment of feedback; Z25 then use the 600 Myr snapshot from this simulation as the initial condition for their simulations. This snapshot exhibits a stable gas fraction similar to that of the present-day Milky Way. The WK23 simulations have gas particles with a mass of 859.3 $M_\odot$, dark matter particles with a mass of 1.254×10^5 $M_\odot$, and stellar disc and bulge particles with a mass of 3.4373×10^3 $M_\odot$. Star particles formed during the simulation inherit the mass of the gas particle from which they were created. To enhance resolution from the original $\Delta M = 859.3$ $M_\odot$ to $286.4 M_\odot$, Z25 first run without enhancing the resolution for 100 Myr, but using the star-by-star feedback prescription, to generate a realistic population of stellar particles for feedback and metal enrichment. They then increase the resolution using the particle splitting method described and used in H23. As shown in Fig. 3 of Z25, the star formation rate initially spikes due to the resolution increase, but it then stabilises within approximately 100 Myr.

We begin the analysis we describe below at this point, which for the purposes of this paper we define as $t = 0$. The simulations run for 464 Myr from that point, with output snapshots written at 1 Myr intervals, providing us with a data set of this duration and ≈ 300 $M_\odot$ mass resolution to study the star particle dynamics. At $t = 464$ Myr, we have approximately 1.32 million star particles. In Figure 1 we show 1% of the total star sample, selected at random. The greyscale background shows the gas surface density, with darker regions indicating higher densities. The overlaid stars are colour-coded according to their ages, as indicated by the accompanying colour bar, showcasing the spatial distribution and age variation of the stellar population. It is important to note that these "star particles" do not represent individual stars. Simulations resolving down to individual stars that sample the entire IMF are not feasible at present. However, they are small enough aggregates of stars that we can analyse them to examine the dynamical information that they retain about their birth properties as if they were individual stars.

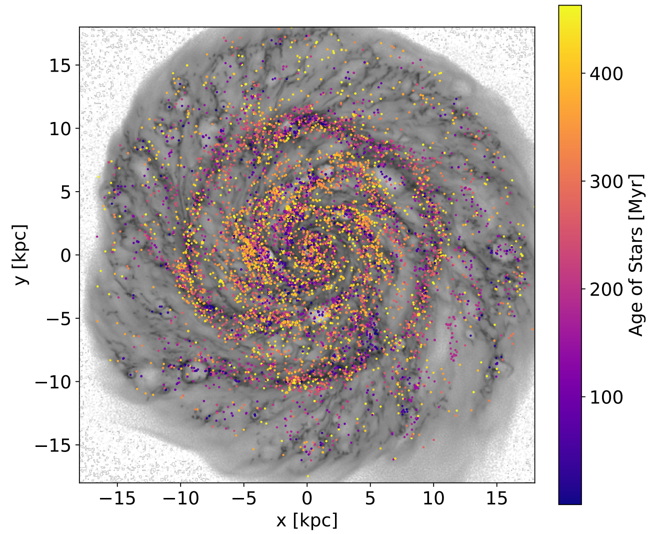

Figure 1. A snapshot of the simulation at $t = 464$ Myr (the final snapshot). We show the log of the gas surface density in greyscale in the background, with darker shades corresponding to higher densities. Overlaid on this is a random subsample representing 1% of the total star sample we use for our analysis, with the colour of each particle indicating its age as shown by the accompanying colour bar.

3 CALCULATING STELLAR ACTIONS

This section outlines how we calculate stellar actions from simulation outputs and study their evolution. A basic outline of our procedure is that we use the potential output for each particle to build an axisymmetric model for the galaxy's gravitational potential, and then use the time-dependent position and velocity of each star particle to compute its actions. The following subsections provide full details of the method.

3.1 Gravitational potential profile

To accurately model the time-dependent gravitational potential of the galaxy, we first define an appropriate reference frame. We cannot simply use the simulation frame because, although the simulation is initialised with the galactic plane at rest at $z = 0$ and cylindrically symmetric about the origin, over the duration of the simulation supernova explosions drive an asymmetric wind off the galaxy, imparting a non-negligible momentum to the remaining gas. To correct for this shift, and thereby ensure that we are calculating our actions with respect to the rest frame of the galaxy, we calculate the centre of mass (COM) of cold gas particles (those with temperature $\leq 10^4$ K) at each snapshot; we plot the resulting position in z-direction as a function of time in Figure 2. We observe a velocity ~ 2 km/s in the z direction and ~ 1 km/s in the x and y directions, and this movement of the COM of the cold gas is consistent with what we expect: the simulation produces galactic winds with a mass flux ~ 1 $M_\odot$/yr and a velocity of a few hundred km/s (WK23), so over the simulation timescale of ~ 500 Myr winds eject $\sim 5 \times 10^8 M_\odot$, roughly 10% of the total cold gas mass $\approx 4.7 \times 10^9 M_\odot$ at the start of the simulation. If the winds were completely one-sided, this would therefore be sufficient to accelerate the gas disc to several tens of km s^{-1}, but since the winds are only somewhat asymmetric some of the momentum cancels. In addition, gravitational forces transfer some of the momentum from the gas to the stellar disc, which is roughly ten times as massive. Consequently, the net velocity is reduced to only a few km/s. The z velocity is greater than the x and y velocities due to the greater escape of winds normal to the galactic plane.

The combination of lower velocity in the xy plane and the much larger extent of the disc in this direction mean that the displacements in the plane can be safely ignored. By contrast, this is not true in the z direction, with the displacement of ≈ 0.8 kpc over the course of the simulation is several times larger than the scale height of the thin disc. To remove this drift, we carry out a linear fit to the z position as a function of time, which we show as the straight orange line in Figure 2. This line defines a time-dependent COM position and constant velocity for the galaxy relative to the simulation frame, and our first step in computing the galactic potential is therefore to shift all particle positions and velocities for all time snapshots into this frame. However, we note that the linear fit is clearly not a perfect representation of the actual COM position as a function of time, which is not surprising, since acceleration of the galaxy by recoil from galactic winds applies a stochastically-varying acceleration.

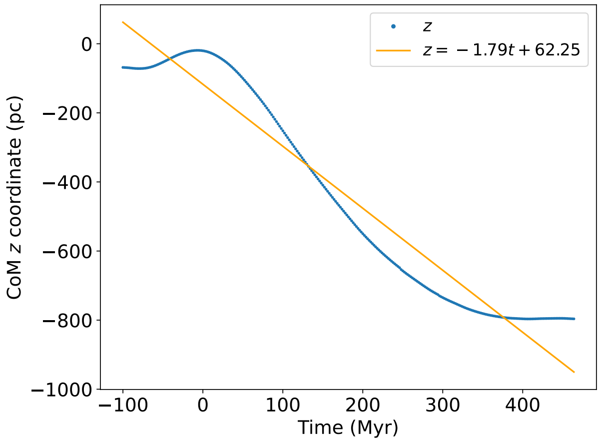

Figure 2. Coordinates of the centre of mass of cold gas in the simulation with time in the z direction. We also show a least-squares linear fit to the z position (solid line), the functional form for which is provided in the legend (with position in units of pc and times in units of Myr); the R^2 value for the fit is 0.94.

This potentially complicates the situation, as it suggests the absence of an inertial frame in which the galactic plane remains at rest. Such a scenario foreshadows the non-conservation of stellar actions, which would be a fundamental issue since most analyses rely on the assumption of a stable inertial frame for dynamical calculations (e.g., Bovy 2015; Sanders & Binney 2016). We return to this point below.

Having established our reference frame, we are now ready to compute our best estimate of an azimuthally symmetric galactic potential from large but sparse samples at our star particle positions. To do so, for each simulation snapshot we define a grid in cylindrical coordinates that spans the entire galaxy, with radial coordinates extending from 0.1 pc to 20 kpc, a vertical range of ± 1 kpc, and a resolution of 1 pc in both directions. The grid in the azimuthal direction has a resolution of $2\pi/25$ radians. For each point on the defined grid, we identify the nearest-neighbour particle and assign its potential value (which is computed by the GIZMO gravity solver) to that grid point. We then average over the ϕ direction to create an axisymmetric 2D potential grid in (R, z). Figure 3 shows some sample slices through the 2D potential we derive for one time snapshot.

3.2 Action calculation

Our next step is to calculate the actions of each star. Performing this calculation in full generality requires computationally-expensive numerical integration. However, we are solely concerned with stars that are both young – age $\lesssim 0.5$ Gyr – and near the galactic plane – root mean square height $\langle z^2 \rangle^{1/2} < 200$ pc, and $|z| < 1$ kpc for all stars. For stars with these properties, the epicyclic approximation, whereby we decompose stellar orbits into independent radial and vertical oscillations about a guiding center, is highly accurate (Hunt & Vasiliev 2025). Quantitatively, Solway et al. (2012) find that, even in the presence of spiral structures, for stars up to 1 kpc off the plane, and over a timescale of a few Gyr, vertical actions computed using the epicyclic approximation change by 20.7%, compared to 15.6% with a more exact calculation; this small difference confirms that the epicyclic approximation is appropriate for our purposes. This is advantageous, because in the epicyclic approximation actions are far cheaper to compute.

Given these considerations, we proceed using the expressions for

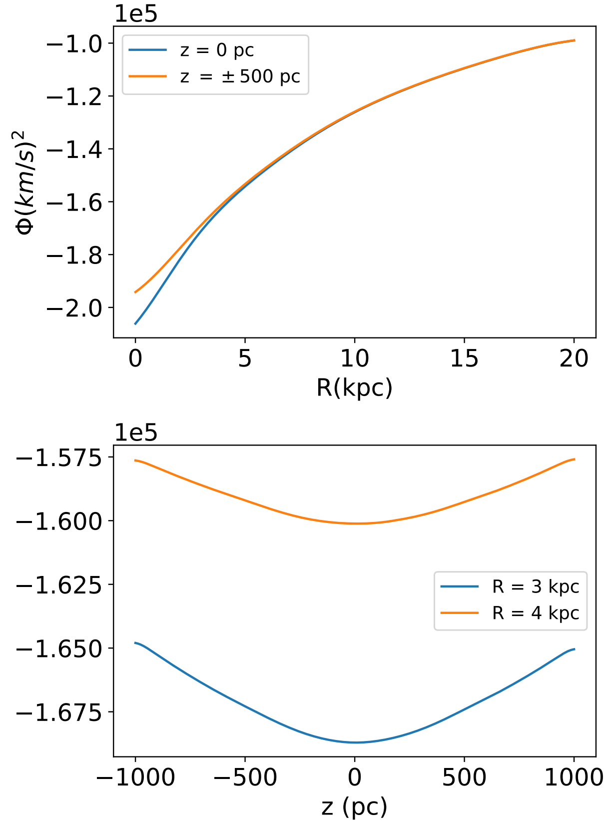

Figure 3. Variation in potential Φ with R at sample z-values (top), and with z at sample R values (bottom), for the $t = 100$ Myr simulation snapshot.

the epicyclic approximation provided by Binney & Tremaine (2008). In this approximation, stellar oscillations are characterised by the radial (κ) and vertical (ν) epicyclic frequencies, while the azimuthal rotation about the galactic centre is defined by the z-component of angular momentum L_z. The guiding centre radius R_g is related to the angular momentum L_z as

$$|L_z| = R_g^2 \Omega(R_g) \tag{4}$$

where

$$\Omega(R) = \sqrt{\frac{1}{R} \left. \frac{\partial \Phi}{\partial R} \right|_{R,z=0}} \tag{5}$$

is the circular frequency. The epicyclic frequencies are related to the potential as

$$\kappa^2 = \left(R \frac{\partial \Omega^2}{\partial R} + 4\Omega^2 \right)_{(R=R_g, z=0)} \tag{6}$$

$$\nu^2 = \left. \frac{\partial^2 \Phi}{\partial z^2} \right|_{(R=R_g, z=0)}. \tag{7}$$

We compute R_g, κ and ν for each star as follows. We first use the star's radial position R and azimuthal velocity v_ϕ to evaluate its angular momentum $L_z = R v_\phi$, and then to use Equation 4 together with the galactic rotation curve $\Omega(R)$ to evaluate the star's guiding radius R_g. Once R_g is known, we can then evaluate $\kappa(R_g)$ and $\nu(R_g)$ from Equation 6 and Equation 7 using the known rotation curve and potential. The primary challenge to executing this strategy is that the gravitational potential from the simulation is discrete and subject to numerical noise, and thus considerable care is required when evaluating the derivatives that appear in the expressions above. We provide a detailed description of our full methodology in section A.

Finally, we find expressions for radial, vertical, and azimuthal actions in the epicyclic approximation. In an axisymmetric potential, the azimuthal action J_ϕ is defined as

$$J_\phi = \frac{1}{2\pi} \oint v_\phi R \, d\phi. \tag{8}$$

Using $v_\phi = R\dot{\phi}$ and $L_z = R^2\dot{\phi}$, and solving the above integral, we get

$$J_\phi = L_z. \tag{9}$$

For the purpose of computing the radial and vertical actions, we note that in the epicyclic approximation we treat the radial and vertical motion as two independent harmonic oscillators, so that the Hamiltonian for the perturbation about the circular orbit is

$$H_{\rm pert} = \frac{v_R^2}{2} + \frac{\kappa^2}{2}(R - R_g)^2 + \frac{v_z^2}{2} + \frac{\nu^2}{2}z^2 \tag{10}$$

where v_R and v_z are the velocities in the radial and vertical directions, respectively, and z is the vertical coordinate of the star. The potential energy terms $(\kappa^2/2)(R - R_g)^2$ and $(\nu^2/2)z^2$ correspond to the radial and vertical oscillations around the equilibrium orbit. The usual simple harmonic oscillator has Hamiltonian $H = \frac{1}{2}(v^2 + \omega^2 x^2)$ for velocity v, displacement x, and oscillator frequency ω, and the corresponding action is E/ω, where E is the energy. By analogy, for our epicyclic approximation Hamiltonian the radial energy E_R is

$$E_R = \frac{v_R^2}{2} + \frac{\kappa^2}{2}(R - R_g)^2. \tag{11}$$

and the corresponding radial action is

$$J_R = \frac{E_R}{\kappa} = \frac{v_R^2 + \kappa^2(R - R_g)^2}{2\kappa}, \tag{12}$$

Similarly, the vertical energy E_z is

$$E_z = \frac{v_z^2}{2} + \frac{\nu^2}{2}z^2, \tag{13}$$

and the vertical action J_z is

$$J_z = \frac{v_z^2 + \nu^2 z^2}{2\nu}. \tag{14}$$

Thus given the radial and vertical epicyclic frequencies and guiding radii computed above, together with the stellar radial and vertical velocities, we can compute the radial and vertical actions for all stars.

3.3 Action evolution study

Once we have computed the actions for all stars at each snapshot, we are in a position to investigate their temporal evolution. Initial analysis reveals that a significant fraction of stars exhibit rapid variations in their actions over short (few Myr or less) timescales, likely the result of being part of gravitationally bound structures such as star clusters; we discuss these structures further in section 5, and they will be the principal focus of Paper II in this series. For the moment we wish to focus on the longer term secular evolution of actions, and to isolate these from the high-frequency variations we apply a Butterworth filter with a cut-off frequency of 1/30 Myr^{-1} to the time series of actions, effectively suppressing variations occurring

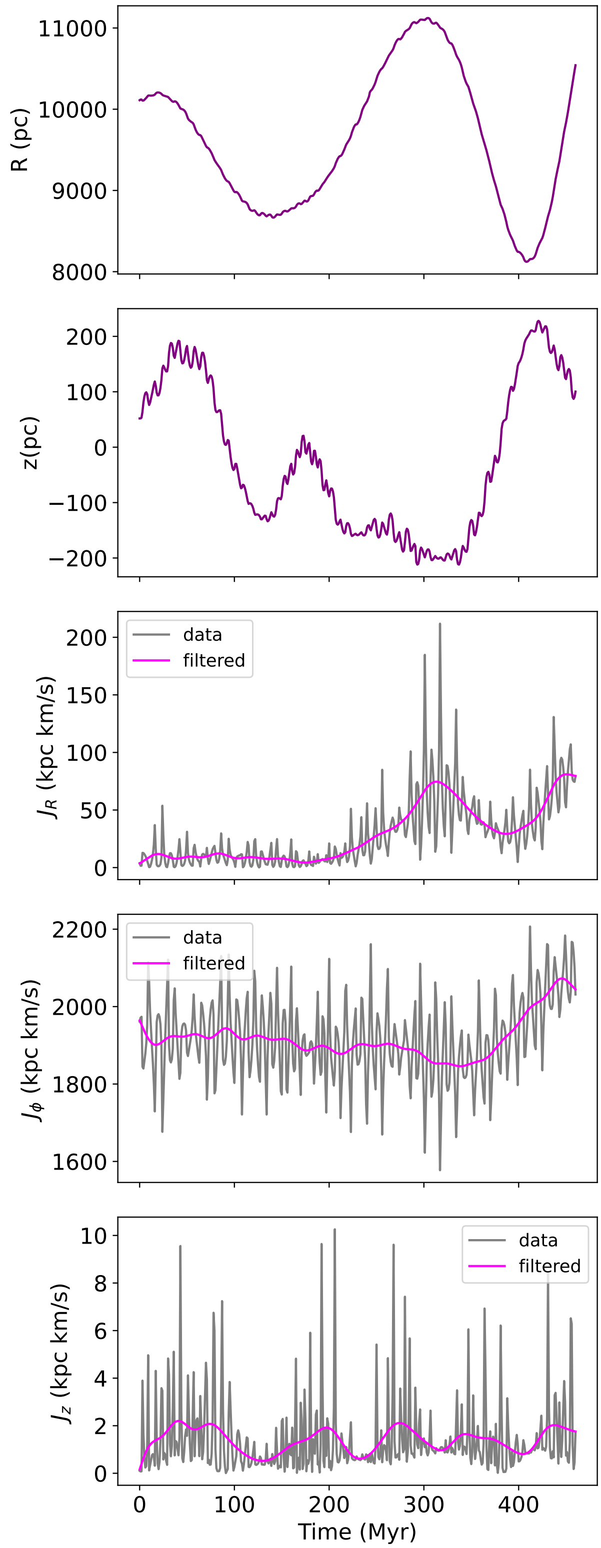

Figure 4. An example of a star that exhibits rapid variation in actions, likely because it is part of a star cluster. The top two panels show the R and z position as a function of time, while the bottom three show action versus time. The grey lines represent the actual actions we compute, while the pink lines show the actions after applying the low-pass filter as discussed in subsection 3.3.

on timescales shorter than 30 Myr. Figure 4 illustrates an example of a star displaying rapid variations in its actions, along with the impact of the low-pass filter in smoothing these fluctuations.

To characterise the change in action over time, we compute the absolute and relative differences in each component of the action between every possible pair of snapshots. Formally, for each star we define the absolute difference as

$$\Delta J_i(t, \Delta t) = |J_i(t + \Delta t) - J_i(t)| \tag{15}$$

where $i = (R, z, \phi)$ is the component, t is the time of the snapshot, Δt is the time interval between the two snapshots and $J_i(t)$ is the i component of the star's action action at time t in the simulation. Similarly, we define the relative difference as

$$\delta_i(\Delta t) = \frac{\Delta J_i(t, \Delta t)}{J_i(t)}. \tag{16}$$

We have ≈ 450 time snapshots, and therefore $\approx 10^5$ snapshot pairs, and $\sim 10^6$ stars, so our total sample for analysis consists of $\sim 10^{11}$ action differences.

4 RESULTS

We are now prepared to characterise the time-evolution of the actions. We first examine the evolution of the distribution of actions in subsection 4.1, and then study the properties of the time series describing individual stars rather than the full population in subsection 4.2. We discuss how the evolution varies with the environmental conditions under which stars are born in subsection 4.3 and subsection 4.4. In subsection 4.5, we compare the action evolution for two different population of stars.

4.1 Distribution of actions

Before examining the evolution of individual stellar actions over time, we first study the distribution of stellar actions and its time evolution. We do so by selecting our final snapshot ($t = 464$ Myr) and examining the action distribution as a function of stellar age, since this provides a snapshot of the data that can be compared directly to observations. We place all stars at the last snapshot in 1 kpc-wide bins of radial position R, and in Figure 5 we show the distributions of radial (J_R), azimuthal (J_ϕ), and vertical (J_z) actions for stars as a function of stellar age for the 4-5 kpc, 9-10 kpc, and 12-13 kpc radial bins. The solid central lines represent the medians of these distributions, while the shaded areas show the 16th to 84th percentile range.

Our results reveal that the overall distribution of actions remains relatively stable over the simulation period. We find that J_ϕ is essentially static, which is not surprising, given that our ≈ 500 Myr run time is too short for significant radial mixing. There is some weak evolution of J_R and J_z distributions, and the latter is particularly useful because there have been several observational studies of this evolution to which we can compare. Ting & Rix (2019) examine this quantity for stars of age up to 8 Gyr and observe significant broadening of the distribution, but this timescale is much larger than that for which our simulations run and is therefore not directly comparable. More recently, Garzon et al. (2024) performed a similar analysis for stars younger than 0.5 Gyr, well-matched to the timescale of our simulation. We overlay their results in the bottom row of our plot (in red). Since their analysis is limited to stars with the Galactocentric radii within 8 to 13 kpc, we only have comparisons for two of the radial bins plotted. For the 9-10 kpc radial bin, we find that our results agree very well with their mean J_z values across ages. However, in the 12-13 kpc radial bin, we find significantly lower actions and less evolution. Garzon et al. (2024) suggest that the very large change in action evolution from $9 - 10$ kpc to $12 - 13$ kpc that they find may be due to the warping of the outer Milky Way disc (e.g., Uppal et al. 2024, and references therein), a feature that is absent from our simulations, since we begin with a flat disc and do not include perturbations from dwarf galaxies or dark matter sub-halos that might induce a warp. This is likely the reason that we do not reproduce this aspect of the observations in our simulations. This also means that the results we derive below for the rates at which individual stellar actions evolve are likely to be underestimates in the far outer disc beyond ~ 10 kpc where the warp is significant. We should therefore think of the results we obtain as lower limits that describe action evolution in the absence of large-scale external perturbations to the galactic potential.

4.2 Change in actions over time

We now investigate the evolution of individual stellar actions over time as characterised by the absolute changes ΔJ_R, ΔJ_z and ΔJ_ϕ (Equation 15). These quantities are functions of both the current stellar age t_* and the time lag Δt, and we therefore place all stars in 2D bins of these quantities and compute the median in each bin; we use medians rather than means to ensure robustness against outliers. We show the result in Figure 6; in this plot, stellar age appears on the $x-$axis and the time lag Δt on the $y-$axis. The expected behaviour is immediately visible – small variations at short time intervals appear as the darkest colours in a horizontal band at the bottom of the plot. We see only a weak dependence on stellar age, as there are no very obvious vertical features in the heatmap.

To take a closer look, we extract vertical slices from the heatmap at ages 5, 10, 50, 100, 200, 300 and 400 Myr, as shown by the dashed lines in Figure 6. These slices, illustrated in Figure 7, show how the absolute change in actions depends on the time lag for stars of different ages. We see here that there is a weak dependence of ΔJ on stellar age, but in a direction that is opposite what one might naively expect – we expect younger stars to exhibit larger changes in their actions than older stars because they are more likely to reside near their natal molecular clouds (Lada & Lada 2003), where ongoing interactions with dense gas or other stars can lead to greater dynamical perturbations. Conversely, older stars have typically migrated away from their birth regions and are less influenced by such perturbations. Contrary to these expectations, we instead find that older stars undergo more rapid changes to their actions, as indicated by the larger slopes in Figure 7.

To understand why this occurs, we next examine relative rather than absolute changes in actions, focusing on J_R and J_z since J_ϕ is dominated by circular orbital motion and is less sensitive to dynamical perturbations. Figure 8 shows the median square relative changes of these two actions, δJ_R^2 and δJ_z^2, as a function of time lag for the same stellar ages as shown in Figure 7. In this plot, we see that the different stellar age bins collapse onto nearly a single line, suggesting a universal behaviour among stars of all ages.

The linear increase in δJ_R^2 and δJ_Z^2 with time before 100 Myr suggests that the process of stellar action change can be described as a random walk in the logarithm of the action, and that the process can therefore be approximated as diffusive. As discussed earlier, this diffusion is due to dynamical perturbations in the smooth, axisymmetric gravitational potential caused by structures such as giant molecular clouds and spiral arms, together with the large-scale acceleration of the disc due to galactic winds.

This clearly highlights a key result: actions are not conserved for

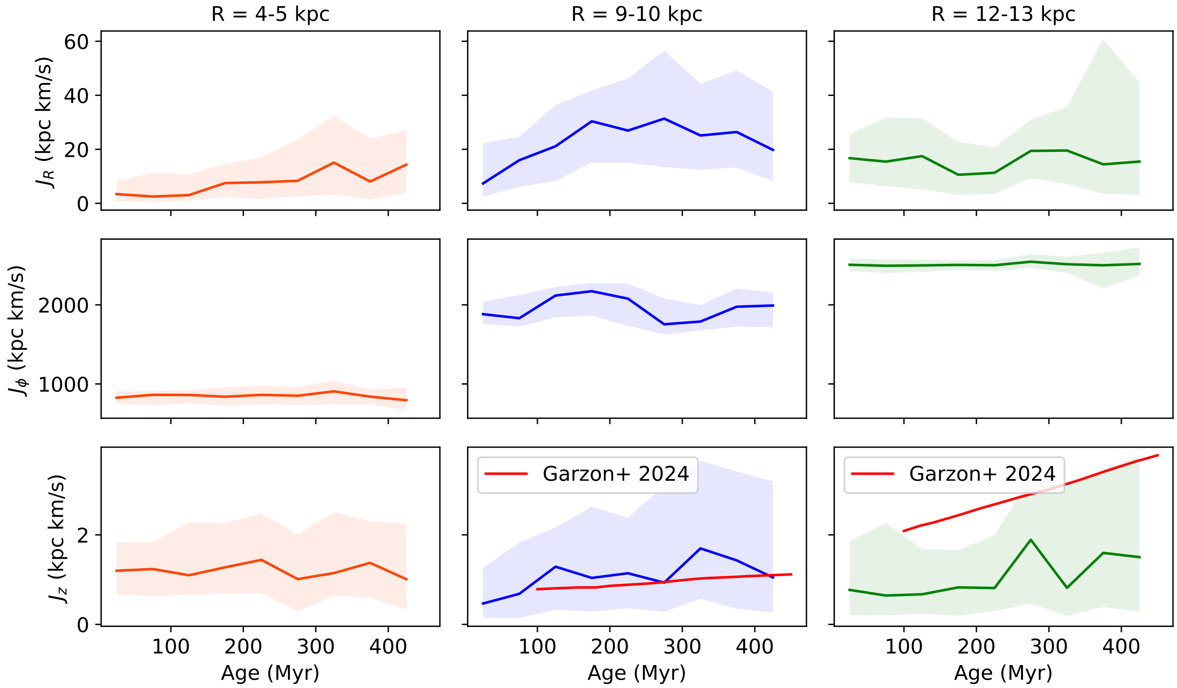

Figure 5. Distribution of actions (row-wise in order: J_R, J_ϕ and J_z) for stars of different ages at the last snapshot, $t = 464$ Myr, in the radial bins 4-5 kpc (left), 9-10 kpc (middle), and 12-13 kpc (right). The orange, blue and green lines represent the medians of the distributions, with shading indicating the 16th and 84th percentile range. The red lines in the two right-most columns of the bottom row show the mean J_z values as a function of age reported by Garzon et al. (2024, their Figure 6).

individual stars, placing constraints on reconstructing stellar orbits using present-day actions. To quantify the rate of diffusion, we carry out a least-squares find of the data shown in Figure 8 to a diffusion model of the form

$$\left\langle \delta J_R^2 \right\rangle = 2 D_R \, \Delta t, \tag{17}$$

and similarly for z, where here D is the diffusion coefficient; we carry out this fit from $\Delta t = 5 - 50$ Myr for D_R and for $5 - 40$ Myr for D_z, covering the age range when the behaviour is close to linear, and for the purpose of the fit we include all stars regardless of age t_*. We show our fits, together with the data, in Figure 9. Our best-fitting values for the diffusion coefficients are $D_R = 2.5$ Gyr^{-1} and $D_z = 4.8$ Gyr^{-1}. The inverses of these values correspond to characteristic diffusion timescales of approximately 400 Myr for J_R and 200 Myr for J_z, meaning that over these timescales individual stars forget their individual actions. The shorter timescale for J_z reflects stronger dynamical perturbations in the vertical direction. In both cases, however, we see that the curve in Figure 8 flattens beyond ≈ 100 Myr, indicating that actions have reached maximum decorrelation by this point, corresponding to around 50% median square relative change. By this age stellar actions are essentially drawn at random from the distribution of actions for stars of the appropriate age, and retain no memory of their values ≈ 100 Myr earlier.

While this is a good first-order description of the results, close examination of Figure 9 does reveal a weak age dependence of action diffusion: younger stars (represented by lighter colours) exhibit steeper slopes compared to older stars (darker colours) for both J_R and J_z, as indicated by the fact that the lighter-coloured lines corresponding to the younger stars lie predominantly above the red dashed line representing the entire dataset, while the darker lines for older stars lie below it. This is consistent with our expected age dependence for action diffusion: younger stars are more susceptible to perturbations near their birth sites, changing their actions rapidly, while older stars' actions change comparatively slowly. However, we can now see that this is true only for *relative actions*, while older stars experience faster changes in their *absolute* actions because they have larger radial and vertical actions on average. The difference is not large, however.

4.3 Dependence on local birth density

The slightly more rapid diffusion of relative actions that we measure for younger stars suggests an environmental effect – younger stars may diffuse more quickly because they are in closer proximity to denser regions with stronger local, non-axisymmetric gravitational pulls. However, tracking a star's density history over its entire lifetime or orbit on the timescales considered is complex, as any given star may transverse different regions of varying densities. Instead, to test whether density matters, we restrict our focus to the contribution of local density of the birth environment of the star to changes in its action.

To achieve this, we define 'newborn star cohorts' at each snapshot, consisting of all stars formed in the past 1 Myr. For each star in this cohort, we find the fifth nearest neighbour distance (d_5) to another star in the cohort as a proxy for the local density at which each star formed. Smaller d_5 values correspond to stars forming in denser environments where stars are packed more closely together. By combining these measurements across snapshots, we obtain a distribution of d_5 values, allowing us to assess the effect of environments in which stars are born. We find that the central parts of the distribution of d_5 are

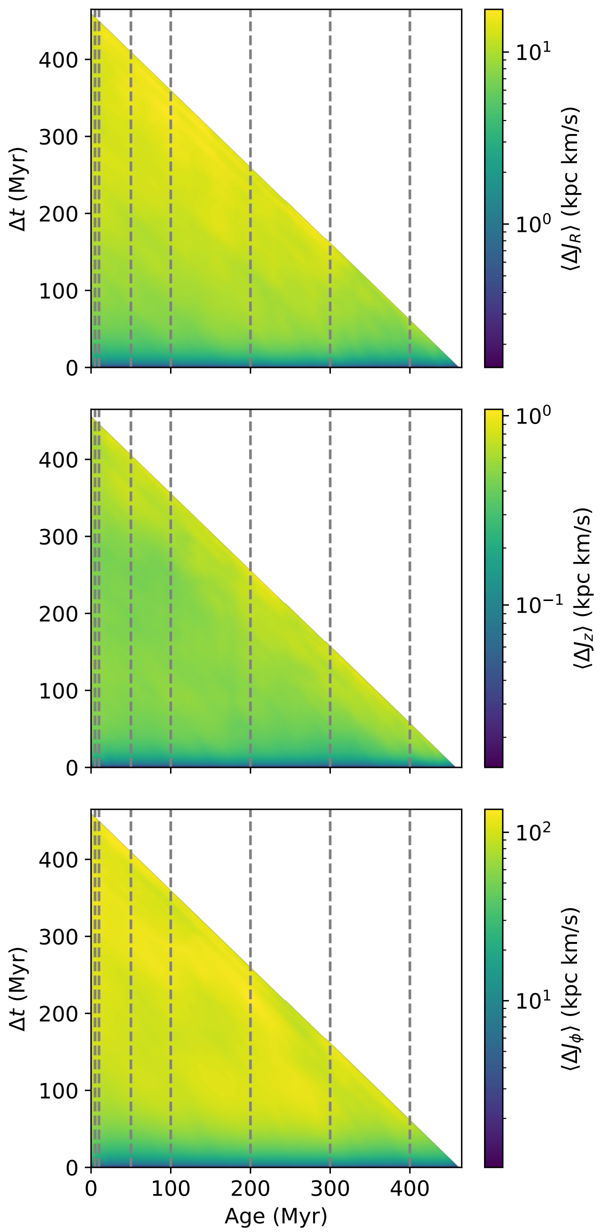

Figure 6. Heatmap of the median of the absolute change in stellar actions (row-wise in order: ΔJ_R, ΔJ_z, and ΔJ_ϕ) as a function of the time lag and stellar age. The dotted lines indicate stellar ages of 5, 10, 50, 100, 200, 300 and 400 Myr, corresponding to the time slices plotted in Figure 7.

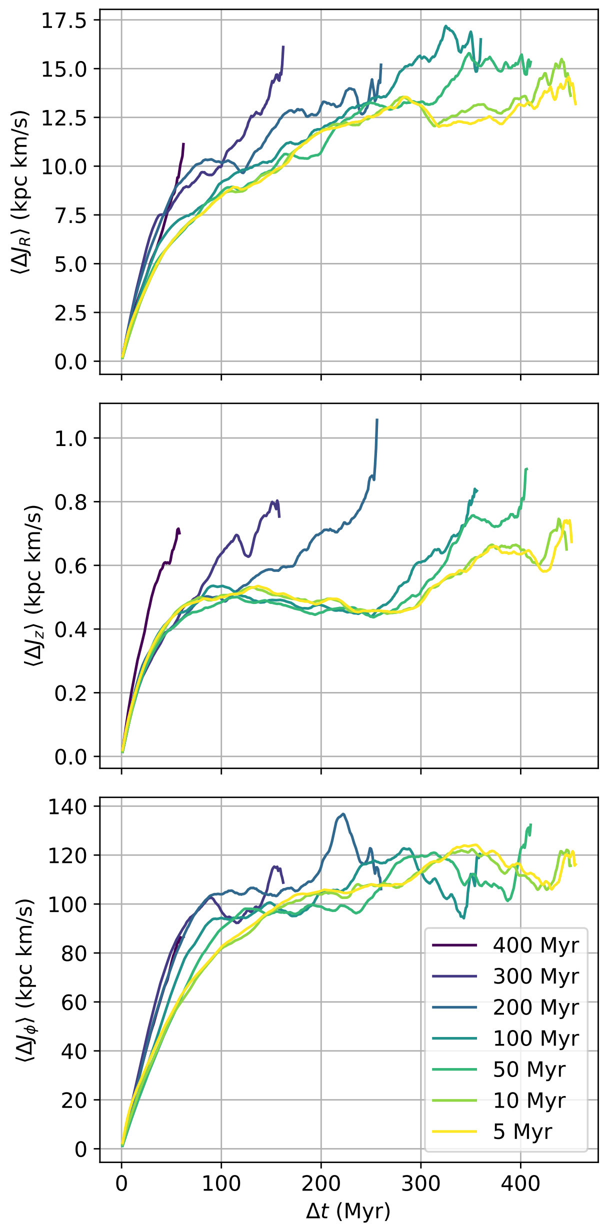

Figure 7. Median of the absolute change in stellar actions (row-wise in order: ΔJ_R, ΔJ_z, and ΔJ_ϕ) for stars of different ages across the simulation as a function of time lag. Different stellar age groups are shown in different colours as indicated in the legend.

relatively smooth, with a 10th percentile value of 0.34 pc and 90th percentile value of 2.64 pc, thus spanning roughly a factor of 500 in stellar density (which scales as d_5^{-3}). There is also a low-density tail to the distribution, so that the 95th and 99th percentiles increase sharply to 87.2 pc and 1323.53 pc, respectively. Given that the smoothing length in our simulation is ~ 0.5 pc, we define 'dense' birth regions as those where $d_5 < 0.5$ pc, moderately sparse regions as those for which d_5 is between 89th and 91st percentile value, and 'very sparse' regions defined as those with d_5 beyond the 99th percentile, i.e., $d_5 > 1323.5$ pc.

To see whether birth density affects the rate of action diffusion, we plot the square median relative change in radial and vertical actions, $\langle \delta J_R^2 \rangle$ and $\langle \delta J_z^2 \rangle$, for stars in our three sample density bins in Figure 10. The solid line represents stars that formed in dense environments, while the dashed lines correspond to those originating in sparse and very sparse regions – dark blue for the 90th percentile in d_5 (2.64 pc) and cyan for the 99th percentile in d_5 (1323.5 pc). For reference, we also overlay the best-fit diffusion relation from the previous section (shown as star markers), which we compute for all stars regardless of their birth density. The results show that the stars formed in dense and sparse regions both follow the same trend as the overall dataset, closely matching the linear fit, and indicating no significant dependence on density. Stars born in very sparse regions exhibit smaller changes in action, particularly at lower Δt

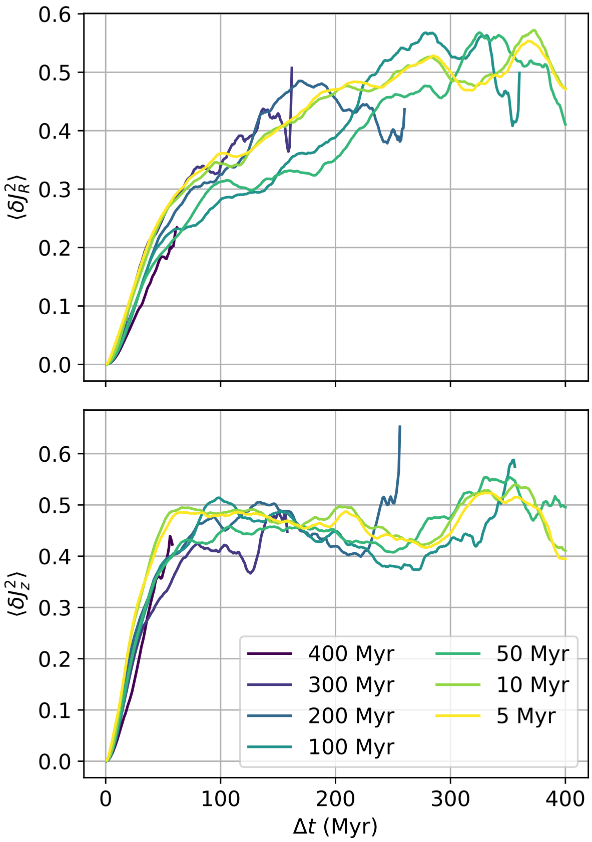

Figure 8. Same as the top two panels of Figure 7, but now showing the median square relative change in stellar actions (row-wise: δJ_R^2 and δJ_z^2).

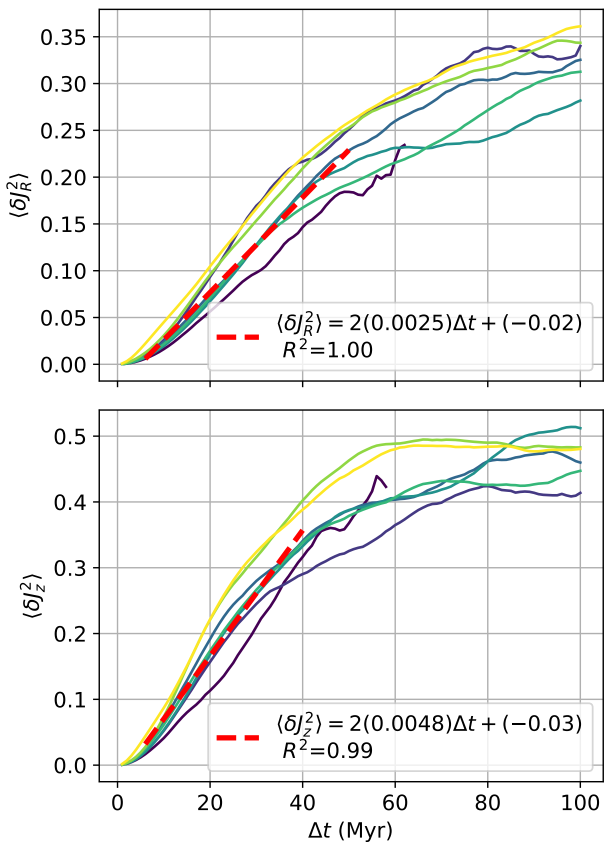

Figure 9. Same as Figure 8, but now zooming in on ages < 100 Myr. The red dashed lines indicate linear fits to the data, as described in the main text. The expressions and R^2 values for the fits are indicated in the legend, with times in units of Myr.

values, consistent with the hypothesis that stars born in denser environments suffer more environmental perturbations and thus diffuse more rapidly. However, the effect is extremely weak, and affects only a small fraction of stars: there is no visible difference for stars at the 90th, or even 95th (not shown), percentile of density. Only the $\sim 1\%$ of stars born in the quietest, most isolated environments show significantly reduced rates of action diffusion. It is worth noting that while birth density does not appear to strongly influence action changes, this does not rule out the possibility that the density of the environment a star encounters throughout its orbit could play a role. A more comprehensive analysis would involve tracking the density of regions that stars pass through over time and correlating this with changes in their actions. However, such an investigation is beyond the scope of this paper.

4.4 Dependence on birth radius

Since disc structure and density, and hence frequency of perturbations, should vary with galactocentric radius, we next examine how rates of action diffusion vary with this quantity. We bin our stars by the galactocentric radius at which they are born, using 1 kpc-wide bins, and show the square of the median relative change in radial and vertical actions, $\langle \delta_R^2 \rangle$ and $\langle \delta J_z^2 \rangle$, for different radial bins in Figure 11. The left panels present the result in absolute time, while the right panels show the same with the time normalised by the orbital period at the centre of the radial bin.

In absolute time, the rate of action change varies significantly with radius. The outer disc regions consistently show slower evolution compared to the inner disc for both R and z actions. This would suggest that the processes driving action diffusion are less efficient at larger radii. However, we must keep in mind that the orbital period is higher at higher radii. When time is normalised by the orbital period ($\tau_{\rm orb}$), this trend becomes much less pronounced as seen in the right panels of Figure 11, suggesting that to a first-order approximation, the rate of action diffusion is constant when time is measured in units of the orbital period. However, a residual difference still remains for cohorts of stars born in different radial bins even after dividing by the orbital period, indicating that the gravitational perturbations responsible for action evolution are less frequent in the outer disc than in the inner regions even accounting for the longer orbital period. This is not surprising given that the stellar density is higher and gravitational interactions with GMCs and spiral arms are more frequent in the inner disc regions. Additionally, it is worth noting that within 60% of the orbital period, squares of relative change in actions are already between 30 − 50%. This clearly highlights the limitation of using methods that rely on the long-term conservation of action for reconstructing past orbits or predicting future trajectories. Hence, for practical applications of actions, careful consideration of the timescales over which they are assumed to be invariant is required.

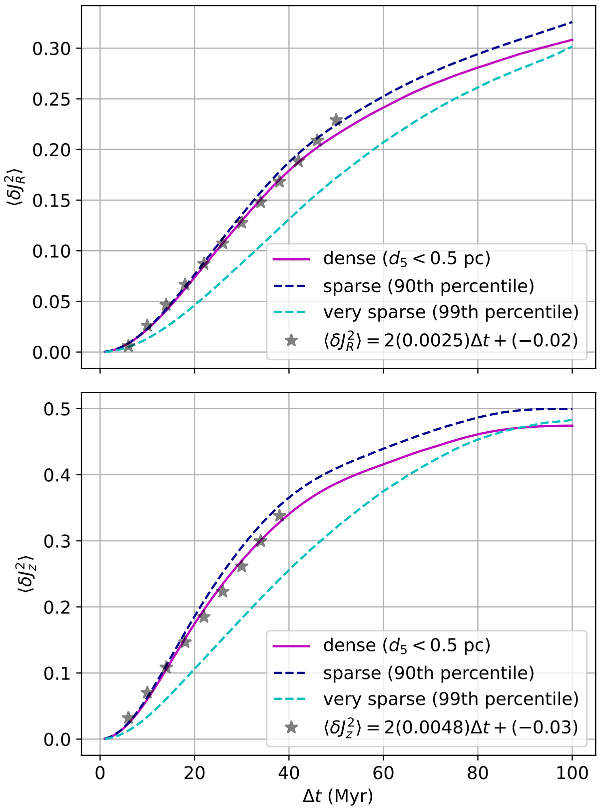

Figure 10. Square median relative change in stellar actions (row-wise: δJ_R^2 and δJ_z^2) for stars born in regions of different density. The solid line shows stars born in the 'dense' regions, while the dashed lines show those born in 'sparse' (dark blue) 'very sparse' (cyan) regions – see main text for full definitions. The star markers indicate our linear fits to the full sample without subdividing by density, as shown in Figure 9; the functional form of the fit is provided in the label, with Δt in units of Myr.

4.5 Comparison with initial stars

As described in subsection 2.2, the WK23 simulations include a population of stellar disc particles that are already present in the initial conditions of Z25. We refer to these pre-existing particles as 'initial stars'. In contrast, the star particles discussed so far are those born during the simulation and hence are all younger than 464 Myr. To understand the underlying mechanisms driving action evolution, it is illuminating to compare these two populations. Since the potential grid used to calculate action (see section 3 and section A) extends only to the disc, this comparison is restricted to stars within the disc region.

The key differences between these two populations are their vertical height distributions and their distribution relative to dense gas. Initial stars have positions that are uncorrelated with gas, while stars that form in the simulation are necessarily in close proximity to dense gas structures at least at birth. Similarly, the root mean square height $\langle z^2 \rangle^{1/2}$ of initial stars remains between 380 – 405 pc during the simulation, whereas, as already mentioned, for stars born during the simulation, it is much lower, between 130 – 200 pc. Thus the initial stars spend much more time away from the midplane while their younger counterparts remain confined to the midplane.

In Figure 12, we show the median square relative changes in radial and vertical actions, δJ_R^2 and δJ_Z^2, as a function of time lag for stars of different stellar ages born in the simulation (same as in Figure 8) along with the same for initial stars, represented by the blue dashed line. Despite the differences between the two populations, the change in radial action over time is quite similar. The fact that even the oldest stars experience a comparable level of radial action change to newly born stars suggests that this change is driven by a mechanism that is present throughout the disc and does not depend strongly on birth conditions and proximity to small-scale gas structures. The natural candidate that meets this requirements are transient spiral arms, which are known to be a key feature of disc dynamics and which extend over a range of vertical heights, since for a Milky Way-like disc they involve strongly-coupled perturbations in both gas and stars (e.g., Romeo & Mogotsi 2017).

On the other hand, the evolution of the vertical action differs significantly between the two populations. For initial stars, the vertical action remains nearly conserved, flattening out to about 20% change by the end of the simulation. In contrast, as seen already in previous sections, the stars born during the simulation show much larger changes in the vertical action. This suggests that the mechanism for out-of-plane action evolution is different from the mechanism for in-plane action evolution, and relies on proximity to gaseous structures in the midplane. This leaves scattering by GMCs, which are concentrated near the midplane, as the natural candidate: initial stars, which spend most of their time at greater heights, experience fewer encounters with GMCs and thus retain more of their initial vertical actions.

To further confirm our hypotheses, we examine phase-space orbits of sample initial stars and stars born in the simulation in Figure 13. We select for this purpose stars whose action evolution lies close to the median values of their respective populations, ensuring that these orbits are typical trajectories for both the categories. Since we applied a low-pass filter to the action time series (see subsection 3.3), we handpick stars that do not seem to be part of clusters to minimise the impact of filtering. The figure shows that orbits do confirm the expected trend: initial stars with larger actions have nearly-closed orbits, as expected for a conserved action, while the stars born during the simulation follow more irregular trajectories, which is consistent with stronger vertical action evolution for the younger population.

5 CONCLUSION

In this paper, we use high-resolution MHD simulations to study the evolution of actions in young stars from Myr to Gyr timescales while self-consistently incorporating gas dynamics, star formation, and a live stellar disc. Our key findings can be summarised as follows:

- The distributions of actions over our simulation time of ≈ 500 Myr remains almost constant, broadening only slightly. For parts of the disc interior to ≈ 12 kpc, and thus interior to the warp in the Milky Way disc, we find good agreement between the rate of broadening in the simulation and that observed in the Milky Way.
- Individual stellar actions, however, are not conserved and instead undergo relatively rapid evolution. The effect is strong enough that over timescales of a few hundred Myr stellar actions become fully decorrelated, and it is no longer possible to trace stars backward under the assumption of an axisymmetric time-steady potential. Contrary to expectation, absolute actions change more rapidly for older stars than for younger ones. However, we show that if we instead consider actions normalised by their initial values, younger stars indeed experience more rapid evolution as expected from the stronger perturbations near their birth environments.

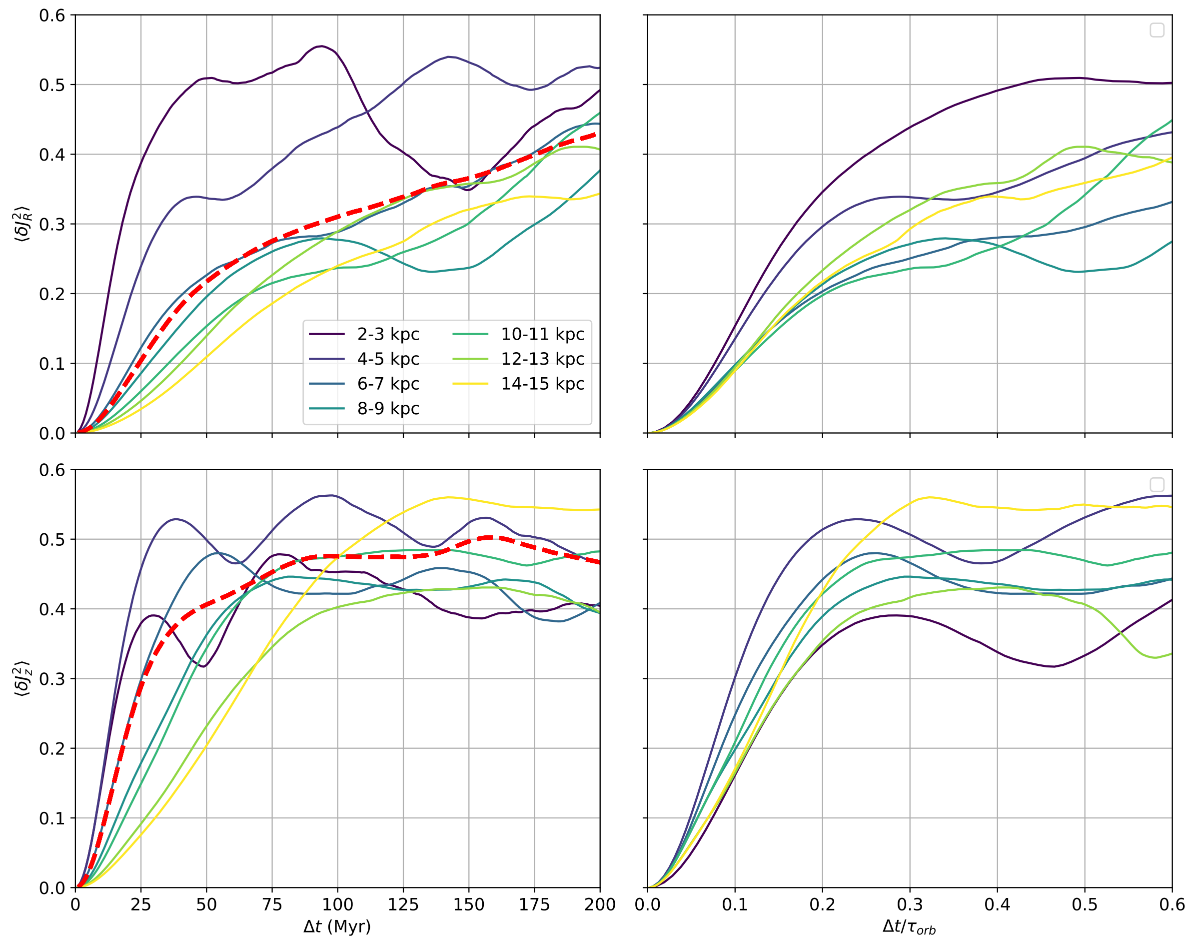

Figure 11. Square median relative change in stellar actions (row-wise: δJ_R^2 and δJ_z^2) for stars born in different radial bins. The birth radial bins are represented by different colours indicated in the legend. The red dashed line shows the values for the full sample without subdividing into radial bins. The left panels show the result in absolute time, while the right panels display the same data with time normalised by the orbital period ($\tau_{\rm orb}$) in each radial bin.

- At ages $\lesssim$ 100 Myr before stellar actions are fully decorrelated, the square median relative change in action increases roughly linearly with time lag, suggesting a random walk in logarithmic action space. We fit a diffusion model to these trends and find diffusion coefficients of $D_R = 2.5$ Gyr^{-1} and $D_z = 4.5$ Gyr^{-1} for radial and vertical actions respectively.

- We find that for the great majority of stars, the density of the stellar birth environment has little effect on the rate of action diffusion. However, this rate does vary with galactocentric radius. To first order, the diffusion coefficient simply scales as the inverse of the local galactic rotation period, and to second order, action diffusion is somewhat slower at larger radii even normalised to the orbital period. This is expected, given the lower stellar density and reduced frequency of encounters with GMCs and spiral arms in the outer disc.

- Comparing newborn stars to an older population with a factor of $\approx 2-3$ larger scale height, we find both populations have similar rates of evolution for radial action evolution, but that the vertical action evolves significantly more slowly for the older, more vertically-extended population. This suggests that different mechanisms drive changes in radial and vertical actions: predominantly transient spiral arms for the radial action, and predominantly scattering of the thin gas disc for the vertical action.

While the rapid action evolution we measure might at first seem surprising, it is important to remember that many previous studies of stellar action conservation have been conducted in the context of a purely collisionless stellar disc, and even in these cases actions are not perfectly conserved (e.g., Solway et al. 2012; Mikkola et al. 2020). Other studies that include gas find that its gravitational influence strongly scatters young stars (e.g., Fujimoto et al. 2023), and the general importance of gas becomes apparent if we recall that, while gas makes up only $\sim$ 15% of the mass of the Milky Way disc, its smaller scale height means that at the midplane gas and stars contribute about equally to the vertical gravitational acceleration (e.g., McKee et al. 2015; Krumholz et al. 2018). In this context it is worth considering as an example the Local Bubble, which is effectively a void in the local ISM, which otherwise has a mean density of ~ 1 atom cm^{-3}, several hundred pc across (Linsky et al. 2022; Zucker et al. 2022; O'Neill et al. 2024). The free-fall time of this structure is (Equation 2) ≈ 51.5 Myr, comparable to the epicyclic periods in the Solar neighbourhood ($T_\kappa \approx 170$ Myr and $T_\nu \approx 85$ Myr; Binney & Tremaine 2008), and its size means stars passing through this region experience a non-

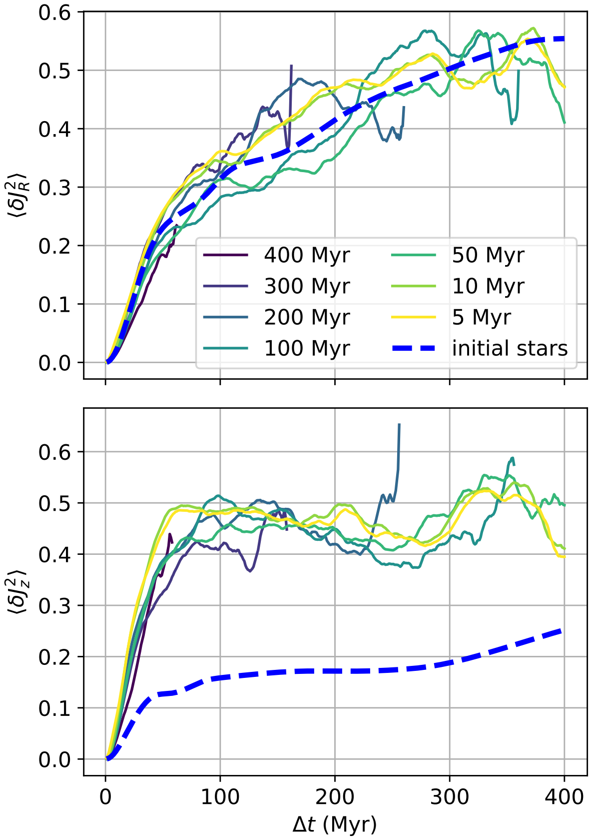

Figure 12. Same as Figure 8, overlaid with a blue dashed line that represents median square relative change in stellar actions for the initial stars (row-wise: δJ_R^2 and δJ_z^2)

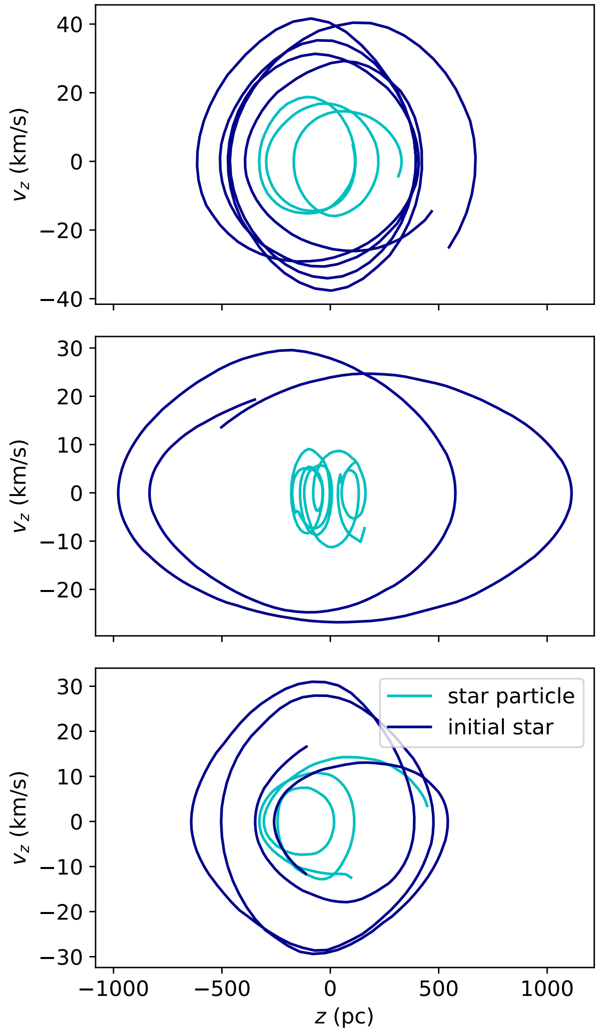

Figure 13. Examples of phase-space orbits for initial stars (cyan) and star particles born in the simulation (dark blue). The x-axis shows their vertical positions in parsecs, while the y-axis shows vertical velocities in km/s.

axisymmetric gravitational perturbations over a significant portion of their orbit. It is therefore not surprising that a structure of this size should be able to perturb stellar orbits significantly.

Our results imply that disc stars, unlike halo stars, do not follow regular, predictable orbits, and therefore that reliable stellar orbit reconstruction relative to the disc past $\sim$ 100 Myr is not feasible. However, this does not necessarily mean that star clusters are entirely irrecoverable. While individual stars undergo a random walk in logarithmic action space, our preliminary analysis indicates that stars born together tend to remain correlated with *each other* in action space over time. This correlation is not surprising, given that much of the change in action appears to be driven by relatively large-scale perturbations such as the Local Bubble – stars born closer to one another than the characteristic sizes of such structures are subjected to similar perturbations, and therefore should experience comparable changes in their actions and remain together in action space. In essence, while stars drift randomly in action space, those that form together “hold hands” as they move, thus preserving clustering signatures in action space that could still be used to reconstruct star clusters and associations, if not to trace them back to their formation locations. We will explore this idea further in Paper II of this series.

ACKNOWLEDGEMENTS

This research was undertaken with the assistance of resources from the National Computational Infrastructure (NCI Australia) and the Pawsey Supercomputing Centre, NCRIS-enabled capabilities supported by the Australian Government, through award jh2. AA and MRK acknowledge support from the Australian Research Council through Laureate Fellowship FL220100020. MI wishes to thank HW Rix for lunchtime discussions in 2019 that motivated some of this work, and the Australia-Germany Joint Research Co-operation Scheme of Universities Australia for funding this travel.

DATA AVAILABILITY

The codes and a subset of selected intermediate data used in this study are available at `https://github.com/aruuniima/stellar-actions-I` and `https://www.mso.anu.edu.au/~arunima/stellar-actions-I-data/`. The initial simulation outputs and potential fits are not included in the public repository due to their size, but will be provided by the authors upon reasonable request.

REFERENCES

Antoja, T., Helmi, A., Romero-Gómez, M., et al. 2018, Nature, 561, 360
Antoja, T., Ramos, P., García-Conde, B., et al. 2023, A&A, 673, A115
Armillotta, L., Krumholz, M. R., Di Teodoro, E. M., & McClure-Griffiths, N. M. 2019, MNRAS, 490, 4401
Binney, J. & Tremaine, S. 2008, Galactic Dynamics: Second Edition
Bland-Hawthorn, J., Sharma, S., Tepper-Garcia, T., et al. 2019, MNRAS, 486, 1167
Bovy, J. 2015, ApJS, 216, 29
Bressan, A., Marigo, P., Girardi, L., et al. 2012, MNRAS, 427, 127
Cabrera Garcia, J., Beers, T. C., Huang, Y., et al. 2024, MNRAS, 527, 8973
Callingham, T. M., Cautun, M., Deason, A. J., et al. 2022, MNRAS, 513, 4107
Chabrier, G. 2005, in Astrophysics and Space Science Library, Vol. 327, The Initial Mass Function 50 Years Later, ed. E. Corbelli, F. Palla, & H. Zinnecker, 41
Chen, Y. & Gnedin, O. Y. 2022, MNRAS, 514, 4736
Coronado, J., Fürnkranz, V., & Rix, H.-W. 2022, ApJ, 928, 70
Coronado, J., Rix, H.-W., Trick, W. H., et al. 2020, MNRAS, 495, 4098
Couture, D., Gagné, J., & Doyon, R. 2023, The Astrophysical Journal, 946, 6
Crundall, T. D., Ireland, M. J., Krumholz, M. R., et al. 2019, MNRAS, 489, 3625
da Silva, R. L., Fumagalli, M., & Krumholz, M. 2012, ApJ, 745, 145
Daniel, K. J. & Wyse, R. F. G. 2015, MNRAS, 447, 3576
Darragh-Ford, E., Hunt, J. A. S., Price-Whelan, A. M., & Johnston, K. V. 2023, ApJ, 955, 74
Doherty, C. L., Gil-Pons, P., Lau, H. H. B., Lattanzio, J. C., & Siess, L. 2014, MNRAS, 437, 195
Federrath, C. & Klessen, R. S. 2012, ApJ, 761, 156
Feuillet, D. K., Feltzing, S., Sahlholdt, C. L., & Casagrande, L. 2020, MNRAS, 497, 109
Frankel, N., Bovy, J., Tremaine, S., & Hogg, D. W. 2023, MNRAS, 521, 5917
Frankel, N., Rix, H.-W., Ting, Y.-S., Ness, M., & Hogg, D. W. 2018, ApJ, 865, 96
Frankel, N., Sanders, J., Ting, Y.-S., & Rix, H.-W. 2020, ApJ, 896, 15
Fujimoto, Y., Inutsuka, S.-i., & Baba, J. 2023, MNRAS, 523, 3049
Galli, P. A. B., Joncour, I., & Moraux, E. 2018, MNRAS, 477, L50
Galli, P. A. B., Miret-Roig, N., Bouy, H., Olivares, J., & Barrado, D. 2023, MNRAS, 520, 6245
Garzon, D. N., Frankel, N., Zari, E., Xiang, M., & Rix, H.-W. 2024, ApJ, 972, 155
Halle, A., Di Matteo, P., Haywood, M., & Combes, F. 2018, A&A, 616, A86
Helmi, A., Babusiaux, C., Koppelman, H. H., et al. 2018, Nature, 563, 85
Heyl, J., Caiazzo, I., Richer, H., & Miller, D. R. 2021, arXiv e-prints, arXiv:2110.04296
Heyl, J., Caiazzo, I., & Richer, H. B. 2022, ApJ, 926, 132
Hopkins, P. F. 2015, MNRAS, 450, 53
Hopkins, P. F. 2016, MNRAS, 462, 576
Hopkins, P. F. & Raives, M. J. 2016, MNRAS, 455, 51
Hopkins, P. F., Wetzel, A., Kereš, D., et al. 2018, MNRAS, 477, 1578
Hu, Z., Wibking, B. D., & Krumholz, M. R. 2023, MNRAS, 521, 5604
Hunt, J. A. S. & Vasiliev, E. 2025, arXiv e-prints, arXiv:2501.04075
Karakas, A. I. & Lugaro, M. 2016, ApJ, 825, 26
Kawata, D., Baba, J., Hunt, J. A. S., et al. 2021, MNRAS, 508, 728
Kim, J.-h., Agertz, O., Teyssier, R., et al. 2016, ApJ, 833, 202
Krumholz, M. R. 2014, Phys. Rep., 539, 49
Krumholz, M. R., Burkhart, B., Forbes, J. C., & Crocker, R. M. 2018, MNRAS, 477, 2716
Krumholz, M. R., Fumagalli, M., da Silva, R. L., Rendahl, T., & Parra, J. 2015, MNRAS, 452, 1447
Krumholz, M. R., McKee, C. F., & Bland-Hawthorn, J. 2019, ARA&A, 57, 227
Lada, C. J. & Lada, E. A. 2003, ARA&A, 41, 57
Lane, J. M. M., Bovy, J., & Mackereth, J. T. 2022, MNRAS, 510, 5119
Leitherer, C., Schaerer, D., Goldader, J. D., et al. 1999, ApJS, 123, 3
Li, Z.-Y. 2021, The Astrophysical Journal, 911, 107
Limberg, G., Souza, S. O., Pérez-Villegas, A., et al. 2022, ApJ, 935, 109
Linsky, J., Redfield, S., Ryder, D., & Moebius, E. 2022, Space Sci. Rev., 218, 16
Ma, Y., De Rosa, R. J., & Kalas, P. 2022, AJ, 163, 219
Mackereth, J. T., Bovy, J., Leung, H. W., et al. 2019, MNRAS, 489, 176
Malhan, K., Ibata, R. A., Sharma, S., et al. 2022, ApJ, 926, 107
McKee, C. F., Parravano, A., & Hollenbach, D. J. 2015, ApJ, 814, 13
Mikkola, D., McMillan, P. J., & Hobbs, D. 2020, MNRAS, 495, 3295
Miret-Roig, N., Antoja, T., Romero-Gómez, M., & Figueras, F. 2018, A&A, 615, A51
Miret-Roig, N., Galli, P. A. B., Brandner, W., et al. 2020, A&A, 642, A179
Miret-Roig, N., Galli, P. A. B., Olivares, J., et al. 2022, A&A, 667, A163
Monari, G., Famaey, B., Siebert, A., et al. 2016, MNRAS, 461, 3835
Myeong, G. C., Vasiliev, E., Iorio, G., Evans, N. W., & Belokurov, V. 2019, MNRAS, 488, 1235
Okalidis, P., Grand, R. J. J., Yates, R. M., & Springel, V. 2022, MNRAS, 514, 5085
O'Neill, T. J., Zucker, C., Goodman, A. A., & Edenhofer, G. 2024, ApJ, 973, 136
Pelkonen, V.-M., Miret-Roig, N., & Padoan, P. 2024, Astronomy & Astrophysics, 683, A165
Romeo, A. B. & Mogotsi, K. M. 2017, MNRAS, 469, 286
Roškar, R., Debattista, V. P., Quinn, T. R., Stinson, G. S., & Wadsley, J. 2008, ApJ, 684, L79
Roškar, R., Debattista, V. P., Quinn, T. R., & Wadsley, J. 2012, MNRAS, 426, 2089
Sanders, J. L. & Binney, J. 2016, MNRAS, 457, 2107
Schoettler, C., de Bruijne, J., Vaher, E., & Parker, R. J. 2020, MNRAS, 495, 3104
Schoettler, C., Parker, R. J., & de Bruijne, J. 2022, MNRAS, 510, 3178
Sellwood, J. A. & Binney, J. J. 2002, MNRAS, 336, 785
Smith, B. D., Bryan, G. L., Glover, S. C. O., et al. 2017, MNRAS, 466, 2217
Solway, M., Sellwood, J. A., & Schönrich, R. 2012, MNRAS, 422, 1363
Squicciarini, V., Gratton, R., Bonavita, M., & Mesa, D. 2021, MNRAS, 507, 1381
Sukhbold, T., Ertl, T., Woosley, S. E., Brown, J. M., & Janka, H. T. 2016, ApJ, 821, 38
Ting, Y.-S. & Rix, H.-W. 2019, ApJ, 878, 21
Tremaine, S., Frankel, N., & Bovy, J. 2023, MNRAS, 521, 114
Trick, W. H., Coronado, J., & Rix, H.-W. 2019, MNRAS, 484, 3291
Uppal, N., Ganesh, S., & Schultheis, M. 2024, MNRAS, 527, 4863
Vera-Ciro, C. & D'Onghia, E. 2016, ApJ, 824, 39
Vera-Ciro, C., D'Onghia, E., Navarro, J., & Abadi, M. 2014, ApJ, 794, 173
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261
Wibking, B. D. & Krumholz, M. R. 2023, MNRAS, 521, 5972
Zhang, C., Li, Z., Hu, Z., & Krumholz, M. R. 2025, MNRAS[[:arXiv]2411.01518]
Zucker, C., Goodman, A. A., Alves, J., et al. 2022, Nature, 601, 334

APPENDIX A: COMPUTING EPICYCLIC FREQUENCIES

In this appendix, we describe in detail the procedure that we use for calculating epicyclic frequencies κ and ν for each star particle. We begin with the 2D potential grid that spans from 0.1 pc to 20 kpc in the radial direction and from -1 to $+1$ kpc in the vertical direction, computed as described in subsection 3.1, and we must evaluate partial derivatives of this potential to compute κ and ν. Naive numerical evaluation of these derivatives yields extremely noisy estimates, and we must therefore adopt a smoothing procedure.

The guiding radius R_g and radial epicyclic frequency depend on $\partial\Phi/\partial R$ and $\partial^2\Phi/\partial R^2$ evaluated at $z = 0$. Our first step to estimate these quantities is to extract the potential grid values within ±2 pc of the midplane, corresponding to five grid points in the vertical direction, and average them to produce a one-dimensional function of R. This averaging reduces the noise. We next apply a Gaussian filter[1] with a standard deviation of 30 pc using `'nearest'` mode to prevent artificial edge effects; this further smooths the data. Finally, we fit the smoothed data with a B-spline with smoothing parameter $s = 200$ to obtain a continuous, differentiable function $\Phi(R, z = 0)$.

With this smooth function at hand, we are now prepared to evaluate the radial derivatives of Φ and the quantities that depend upon them. We first define a new radial grid with a resolution of 0.5 pc with its inner edge placed at the larger of 2 kpc and 5 kpc inward from the smallest radial position of any star at the time we are sampling, and its outer edge at the smaller of 17.5 kpc and 5 kpc outward from the largest stellar radial position; this ensures that we have good coverage over the full disc, but that we are not extrapolating to radii where we sample the potential poorly. We then evaluate Ω_i at each radial grid point R_i from Equation 5 using our B-spline representation of $\Phi(R)$. The corresponding specific angular momentum at each grid point is $L_{z,i} = R_i^2 \Omega_i$, and thus our grid represents a set of $(R_i, L_{z,i})$ pairs. Since the guiding radius R_g for each star is defined implicitly by its angular momentum $L_{z,*}$ and the condition $L_{z,*} = R_g^2 \Omega(R_g)$, we can use the $(R_i, L_{z,i})$ pairs as a lookup table to find the guiding radius corresponding to any R_g. In practice we do this by defining a new B-spline function $R_g(L_z)$ from our $(R_i, L_{z,i})$ pairs using a smoothing factor $s = 2$, and use the resulting function to evaluate R_g for every star. We then drop from our sample stars for which R_g lies outside the range $2 - 16$ kpc, on the grounds that they lie too close to our grid edge to be reliable.

Once the guiding radii are determined, we use them to compute the radial epicyclic frequencies via an analogous strategy. We calculate $\kappa(R_g)$ using Equation 6 at each grid point, again using our smooth B-spline approximation to $\Phi(R, z = 0)$ to evaluate derivatives. This yields a set of (R_i, κ_i) pairs, to which we apply another B-spline fit with smoothing factor $s = 3$ to obtain a smooth function $\kappa(R)$. We plug the guiding radii for each star in our sample into this function to obtain κ values for each of them.

The calculation of the vertical epicyclic frequency ν follows a similar but slightly more complex process, since we require the second derivative of $\Phi(z)$ at each star's guiding radius (see Equation 7), rather than for the radial case which only requires derivatives of $\Phi(R)$ evaluated at $z = 0$. To achieve this, we construct a sparse radial grid with a resolution of 200 pc, spanning the range $2 - 16$ kpc to match the guiding radius constraints that we imposed earlier. At each radial grid point, we extract the potential from the 2D grid by averaging over all radial values within ±1 pc of the sparse grid point. This averaging, as in the radial case, reduces local noise and provides a better estimate of the galactic potential as a function of z. We then smooth the resulting $\Phi(z)$ data by applying a Savitzky-Golay filter with a window length of 101 pc and a polynomial order of 2, which smooths the data while preserving the underlying curvature, followed by a Gaussian filter with standard deviation of 2 pc, which we apply three times. Finally, we use the smoothed data as input to a B-spline fit with $s = 50$ to generate a smooth, continuously-differentiable functional form of $\Phi(z)$ at each sparse radial grid point. We use this smooth function to evaluate $\partial^2\Phi/\partial z^2$ at $z = 0$, the quantity required to evaluate ν using Equation 7. This yields a set of (R_i, ν_i) pairs on our sparse radial grid. We fit these data with a B-spline fit with $s = 1000$ to obtain a continuous function $\nu(R)$. Finally, we determine ν for each star by inputting its guiding radius into this function.

[1] This and all subsequent operations use the implementations provided in SciPy version 1.11.4; see Virtanen et al. (2020).

This paper has been typeset from a TEX/LATEX file prepared by the author.