

Understanding the mechanisms behind the distribution of galactic metals

Chuhan Zhang,^{1★} Zefeng Li^{2★}, Zipeng Hu^{3★} and Mark R. Krumholz^{1,4}

¹Research School of Astronomy and Astrophysics, Australian National University, Cotter Road, Weston, ACT 2611, Australia

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

³Kavli Institute for Astronomy and Astrophysics, Peking University, Beijing 100871, China

⁴ARC Centre of Excellence for Astronomy in Three Dimensions (ASTRO 3D), Canberra, ACT 2601, Australia

Accepted 2025 January 29. Received 2025 January 24; in original form 2024 November 2

ABSTRACT

The evolution and distribution of metals within galaxies are critical for understanding galactic evolution and star formation processes, but the mechanisms responsible for shaping this distribution remain uncertain. In this study, we carry out high-resolution simulations of an isolated Milky Way-like galaxy, including a star-by-star treatment of both feedback and element injection. We include seven key isotopes of observational and physical interest, and which are distributed across different nucleosynthetic channels – primarily asymptotic giant branch stars (N, Ba, Ce), supernovae (O, Mg, S), and Wolf–Rayet stars (C) show measurably different correlation statistics in space and time and their fluctuations. This difference arises from the distinct ejection mechanisms associated with each nucleosynthetic process. The large-scale properties ensure that different elements, despite having different nucleosynthetic origins, are highly correlated with one another (>0.85 for all, >0.99 for same origins), and their spatial correlations vary together in time. However small-scale variations naturally break elements into distinct nucleosynthetic familiars, with elements originating from the same channels correlating better with each other than with elements from different origins. Our findings suggest both challenges and opportunities for ongoing efforts to use chemical measurements of gas and stars to unravel the history and physics of galaxy assembly.

Key words: galaxies: abundances – galaxies: evolution – galaxies: ISM.

1 INTRODUCTION

The distribution and evolution of metals (elements heavier than helium) in galaxies is a key focus for understanding how galaxies form and evolve. Metals are produced primarily through stellar nucleosynthesis and redistributed into the interstellar medium (ISM) through stellar winds, supernovae (SNe), and other feedback processes. Once in the ISM, these metals are mixed by transportation processes including radial migration of metal-enriched gas across the galaxy, as well as turbulence, affecting future star formation and drives the chemical evolution of galaxies (for reviews, see Tinsley 1980; Maiolino & Mannucci 2019; Sánchez et al. 2021).

Metallicity, the oxygen abundance in ionized regions, can be measured using emission-line diagnostics (for a review, see Kewley, Nicholls & Sutherland 2019). The measurements of metallicities have been extended from integrated single fibre technique (e.g. Tremonti et al. 2004) to spatially resolved technique (e.g. Marmol-Queraltó et al. 2011; Croom et al. 2012; Sánchez et al. 2012; Bundy et al. 2015; Erroz-Ferrer et al. 2019; López-Cobá et al. 2020; Emsellem et al. 2022). The deployment of integral field units (IFUs) enables measurements of the spatially resolved two-dimensional (2D) distributions of oxygen abundance across nearby galaxies (e.g. Rosales-Ortega et al. 2011; Sánchez-Menguiano et al. 2016). These

observations reveal the existence of metallicity gradients, typically showing that metal abundance decreases from the centre of galaxies outwards (e.g. Belfiore et al. 2017; Ho et al. 2018; Poetrodjojo et al. 2018; Sánchez-Menguiano et al. 2018; Kreckel et al. 2019). A range of theoretical studies have aimed to explain the origin of these gradients and situate them in the broader context of galaxy formation (e.g. Di Matteo et al. 2009; Ma et al. 2017; Sharda et al. 2021; Tissera et al. 2022).

However, gradients represent a significant simplification of the data, since they collapse complex 2D maps down to a single linear fit. To exploit the full power of IFU metallicity maps, higher order statistics are in need to decode the detailed information of the data, which in turn can advance our understanding of how metals are injected and mixed in galaxies. One of the simplest statistical descriptions for a 2D map is the two-point correlation function, which describes the characteristic size scales over which maps vary. Krumholz & Ting (2018, hereafter KT18), provide a minimal theoretical model to predict two-point correlations of galaxy metallicities based on the competition between metal injection and diffusion processes. This prediction has motivated various observational studies to examine two-point correlations (or similar statistics), in different samples of nearby galaxies (e.g. Kreckel et al. 2020; Li et al. 2021, 2023, 2025; Metha, Trenti & Chu 2021; Williams et al. 2022). All these studies find that metallicity maps of nearby galaxies contain statistically significant spatial structure on top of the overall gradient, and that the two-point correlation functions describing this structure generally follow the shape predicted by KT18. These studies indicate that

* E-mail: chuhan.zhang@anu.edu.au (CZ); zefeng.li@durham.ac.uk (ZL); zphu.charles@gmail.com (ZH)

nearby galaxy metallicity distributions are correlated on characteristic scales of $\sim\text{kpc}$, but with significant systematic variations with galaxy properties such as stellar mass and star formation rate.

While this analytic and observational work has begun the study of metallicity distribution statistics, there have been few simulation efforts to date. The only published study thus far to focus on two-point statistics instead of just gradients is from Li et al. (2024b), who post-process the Auriga cosmological simulations (Grand et al. 2017) to produce metallicity maps comparable to those accessible via observations. They find that the simulations successfully reproduce the correlation lengths observed in local galaxies, suggesting that they capture the most important processes associated with metal mixing in the ISM. The metal distributions produced in the simulations are also in reasonably good agreement with the predictions of KT18.

All studies to date, however, have significant limitations. The observations generally have at best resolutions of a few hundred pc – the main exception is studies of very nearby galaxies with PHANGS-MUSE, which resolution reaches ~ 50 pc, but only for samples of $\lesssim 10$ galaxies (e.g. Kreckel et al. 2020; Williams et al. 2022). Moreover, the studies published to date exclusively focus on oxygen, which is easier to measure in the gas phase than other elements given the sensitivity and spectral coverage of current IFUs. This will only begin to change for large samples once BlueMUSE comes online (Richard et al. 2024) and provides access to key diagnostic lines for nitrogen.

The main limitation for simulations to date is also resolution. While the Auriga simulations (Grand et al. 2017) seem to capture the rough outlines of metal mixing, their effective resolution of ~ 100 pc may be adequate to resolve mixing driven by large-scale mechanisms such as bars (Di Matteo et al. 2013), gravitational instabilities (Petit et al. 2015), cosmological accretion (Ceverino et al. 2016), and SN-driven turbulence (de Avillez & Mac Low 2002; Colbrook et al. 2017), but is clearly insufficient to resolve some other possibly important processes such as spiralarm (Grand et al. 2016; Orr et al. 2023) and thermal instability-driven mixing (Yang & Krumholz 2012). Indeed, as KT18 point out, in a modern spiral the mean ISM density is ~ 1 hydrogen atom per cm^3 and the typical gas scale height is ~ 100 pc, so simulations with mass resolutions of a few thousand M_\odot – typical of even zoom-in cosmological simulations – resolve the scale height of the ISM by only ~ 1 – 2 fluid particles. Even marginal resolution of the internal structure of the ISM requires mass resolutions $\sim 100 M_\odot$, or spatial resolutions ~ 10 pc for Eulerian codes. To date no cosmological simulations of large spiral galaxies satisfy this requirement.

A further limitation of both current theory and observations is that they tell us little about the spatial relationships between different elements. As noted above, due to the limited wavelength coverage of currently available high spatial resolution IFUs, observations of gas-phase elements to date focus almost exclusively on oxygen. By contrast, however, a much larger set of elements are available in stellar spectra. These measurements do not directly provide the spatial distribution of metals in the ISM, because as stars migrate after their formation their records of elemental abundances decorrelate with their current location in space. However, any spatial correlations between elements that were present at the time of star formation remain frozen into the distribution of stars in chemical abundance space, which persists for long times. There have been extensive efforts to understand the structure of this chemical space, since it matters for a wide variety of studies that rely on stellar abundances (e.g. Bland-Hawthorn, Krumholz & Freeman 2010; Bland-Hawthorn & Sharma 2016; Krumholz, McKee & Bland-Hawthorn 2019; Weinberg et al. 2019, 2022; Ting & Weinberg

2022). In recent years, however, these efforts, however, have for the most part focussed on empirical attempts to deduce the structure of chemical space from stellar spectra. There are no first-principles predictions, for example about how well different elements correlate, and almost none (beyond some general arguments in KT18) about how the spatial statistics of elements differ depending on their nucleosynthetic origin.

These situations motivate us to perform high-resolution simulations with both (1) the ability to capture the detailed structure at or better than the best-resolved current gas-phase metallicity maps, in anticipation of future higher resolution facilities capable of capturing even more detailed structure, and (2) the ability to track multiple elements and study the relationships between them. We focus on an isolated Milky Way-like galaxy at $z = 0$, following the production of multiple elements injected using a star-by-star treatment of feedback and nucleosynthesis, and tracing the subsequent transport of those elements through the ISM. We carry out these simulations until the statistics of their fluctuation distributions reach steady states, and then use the resulting steady-state 2D metal fluctuation maps to study the spatial statistics of multiple elements.

The outline of this paper is as follows. In Section 2, we describe our isolated galaxy simulation, including the numerical methods and initial conditions, along with the statistical tools we use to analyse the metallicity distributions in space and time. We describe the outcome of our simulations and the statistical properties of the metal field that we derive from them in Section 3. In Section 4, we discuss the results, drawing several conclusions about both the simulations and existing models for metallicity distributions. Finally, we draw conclusions in Section 5.

2 SIMULATION

In this paper, we simulate an isolated, magnetized Milky Way-like disc galaxy. To produce a realistic gaseous metal distribution, we follow the return of metals from every single star individually, following the injection of metals back into the surrounding gas as a part of stellar feedback. Once they are injected we treat the metals as passive scalars. In the remainder of this section we describe our simulation methods, initial conditions, and statistical analysis techniques.

2.1 Numerical method

Our simulation is an extension of the full galaxy zoom-in simulations described by Wibking & Krumholz (2023, hereafter WK23) and Hu, Wibking & Krumholz (2023, hereafter H23), and with the exception of some aspects of the treatment of star formation, feedback, and metals, our numerical methods are identical to theirs. For this reason we simply summarize the parts of our method that are the same here, referring readers to those papers for full details, and focus most of our attention on the modifications we have made to trace metals.

Our simulations use the GIZMO code (Hopkins 2015). We use GIZMO’s meshless finite mass (MFM) method for magnetohydrodynamic (MHD), and we implement gas cooling using the GRACKLE library (Smith et al. 2017); as discussed in WK23, GIZMO’s default implementation of cooling does not correctly produce a multiphase ISM, and this is not suitable for a simulation such as ours that resolves the phase structure of the ISM. We also enable GIZMO’s subgrid turbulent eddy mixing model (Colbrook et al. 2017; Hopkins et al. 2018).

Our simulation uses a customized treatment of star formation and stellar feedback. Our treatment of star formation is that if the density

of a gas particle exceeds a critical density ρ_c , we assign it a local volumetric star formation rate $\dot{\rho}_{\text{SFR}} = \epsilon_{\text{ff}} \rho_g / t_{\text{ff}}$, where ρ_g is the gas particle density, ϵ_{ff} is the star formation efficiency parameter, and $t_{\text{ff}} = \sqrt{3\pi/32G\rho_g}$ is the local gas free-fall time. We adopt $\epsilon_{\text{ff}} \approx 0.01$ as shown by a wide range of observations (Krumholz et al. 2019), and set the critical density to $\rho_c = 10^3 \text{ H cm}^{-3}$, which given our mass resolution and cooling curve is roughly equivalent to setting ρ_c such that Jeans mass is equal to the mass resolution for particles with density ρ_c and temperatures at the equilibrium value for that density – see WK23 for details. We similarly adopt a minimum softening length of gas particles of 0.1 pc, which is roughly the Jeans length at ρ_c and the equilibrium temperature. To avoid spending too much computational time following very dense structures, for particles with $\rho_g > 10^2 \rho_c$ we set $\epsilon_{\text{ff}} = 10^6$ so that they are converted to stars nearly instantaneously. Thus our composite expression for the volumetric star formation rate is

$$\dot{\rho}_{\text{SFR}} = \begin{cases} 0 & \rho_g < \rho_c, \\ \epsilon_{\text{ff}} \rho_g / t_{\text{ff}} & \rho_c \leq \rho_g < 10^2 \rho_c, \\ 10^6 \rho_g / t_{\text{ff}} & \rho_g \geq 10^2 \rho_c. \end{cases} \quad (1)$$

As usual, we implement this probabilistically, so that during a time-step of size Δt , a particle of density ρ_g has a probability $P = 1 - \exp(-\dot{\rho}_{\text{SFR}} \Delta t / \rho_g)$ of being converted to a star particle.

Once star particles form, we carry out star-by-star tracking of feedback and metal injection; this is in contrast to the default GIZMO treatment of an IMF-integrated stellar population, which is not suitable for the resolutions we reach where the expected number of SNe per star particle is ~ 1 . In our simulation, when a star particle forms we draw a synthetic stellar population for that star particle using the SLUG stochastic stellar population synthesis code (da Silva, Fumagalli & Krumholz 2012; Krumholz et al. 2015). The stars are drawn from a Chabrier (2005) IMF (Chabrier 2005), using a fully stochastic treatment for all stars above $1 M_\odot$ in mass. Each star follows an individual evolutionary track, the Padova stellar tracks (Bressan et al. 2012), and at each time-step we model its atmosphere using SLUG’s ‘starburst99’ option for stellar atmosphere models (Leitherer et al. 1999). This allows us to calculate the instantaneous ionizing luminosity of each individual star particle, taking into account the unique properties of the stars that contribute to the population, which we inject back into the simulation domain using a Strömgren volume method to calculate the effects of ionization feedback. Similarly, we track which stars end their lives as SNe, injecting mass, energy, and metals and which end their lives as asymptotic giant branch (AGB) stars, injecting mass and metals but no energy.

Our treatment of SN energy injection follows the approach described in Armillotta et al. (2019), which is a variant of the common approach of injecting 10^{51} erg of energy in regions where the density is low enough that the Sedov–Taylor phase of SN blast wave expansion can be resolved, and gradually changing over to injecting radially outward momentum as the resolution degrades; see Armillotta et al. (2019) for details.

With regard to mass and metal return, we track the distribution of several key isotopes, treating each one as a passive scalar. The isotopes we include are ^{12}C , ^{14}N , ^{16}O , ^{32}S , ^{24}Mg , ^{138}Ba , and ^{140}Ce . Among these, ^{12}C , ^{14}N , ^{16}O , and ^{32}S are dominant in the ISM, while ^{24}Mg , ^{138}Ba , and ^{140}Ce are crucial in tracing processes within stars. Initially, all isotope abundances are set to zero across the simulation. As the simulation progresses, isotopes are injected into the gas as part of the stellar feedback process. Specifically, in each time-step for each star particle, we calculate the real-time yields of each isotope released by stars based on their mass and evolutionary stage.

To determine these yields, we rely on three different yield tables: Sukhbold et al. (2016) for type II SNe, Karakas & Lugaro (2016) for AGBs, and Doherty et al. (2014) for super-AGBs. After obtaining the isotope yields for each stellar particle, we enrich the surrounding gas particles. The newly released total mass and isotopes are distributed to the gas particles around the star following the same algorithm that is used for distributing SN energy distribution (Hopkins et al. 2018). In essence, the mass is distributed in proportion to the overlap between the stellar particle’s position and the neighbouring gas particles’s weighting kernels.

2.2 Initial conditions

The initial condition for our simulation is a snapshot taken from the simulation of an isolated Milky Way-analogue galaxy by WK23, which was also performed using GIZMO and including a very similar treatment of star formation and feedback to the one we adopt here, but with lower resolution, an IMF-integrated rather than a star-by-star treatment of feedback to accord with this lower resolution, and no tracking of metals. To generate an initial condition suitable for resimulation at higher resolution, we follow the procedure outlined in H23.

We start from the 600 Myr snapshot of WK23, which has a gas fraction similar to the present-day Milky Way, and we switch from the IMF-integrated treatment of feedback in that simulation to the star-by-star treatment described above. We then advance the simulation for 100 Myr, during which time we continue to disable metals and we leave the resolution unchanged; our goal during this phase is to build-up a realistic population of stellar particles that can provide feedback and metals. Because the integrated amount of feedback is the same for our IMF-averaged and star-by-star approaches, the star formation rate and other properties of the simulation remain stable during this time. After this 100 Myr interval, we subdivide particles in the simulation to increase the resolution from the original 859.3 M_\odot used in WK23 to 286.4 M_\odot , following the particle splitting method described by H23; we also turn on metal injection and diffusion at this point. We will refer to the state that exists at this point as time $t = 0$ from this point on. Due to the sudden increase in mass resolution that occurs at this time, the star formation rate (SFR) in the simulation undergoes an initial fluctuation, but then settles back into a steady state at nearly the same SFR as before we increased the resolution within $\lesssim 100$ Myr. We use the properties of the simulation after this point as the basis for our statistical analysis below.

Due to the splitting, the smoothing length of the gas particles needs to be adjusted by GIZMO in the initial few Myr. To illustrate the initial conditions of our simulation, we demonstrate the snapshots at 5 Myr in Fig. 1 when the smoothing length has been fixed.

2.3 Analysis methods

The central object of our study is the statistics of metallicity fluctuations in galaxies. Here, we explain the steps in the analysis pipelines we use to derive these statistics from the raw simulation outputs; unless otherwise stated, we carry out these steps for each time snapshot.

2.3.1 Characterising the metal distribution

The first step in our analysis chain is to convert the three-dimensional simulated metal distribution into a face-on projection of the metallicity Z for each element as a function of position in the plane of the




Figure 1. Snapshot of the simulation at $t = 5$ Myr. The panels from left to right shows the total surface density of the gas (expressed as number of H nuclei per unit area) and stars, respectively. The dimensions of each map are $30 \text{ kpc} \times 30 \text{ kpc}$, while the pixel size of each map is 20 pc .

galactic disc. To this end, we use YT (Turk et al. 2011) plus MESH0ID¹ to calculate the surface density of total gas mass and the mass of each element projected along the z -direction, which is defined as the direction orthogonal to the disc plane; formally, we write these surface densities as $\Sigma = \int \rho \, dz$ and $\Sigma_Z = \int \rho_Z \, dz$, where Σ is the total gas surface density, Σ_Z is the surface density of some element, and both of these quantities are functions of the position (x, y) in the galactic plane. We carry out this calculation on a 2D grid of size $30 \text{ kpc} \times 30 \text{ kpc}$ with spacing $\Delta x = \Delta y = 37.5 \text{ pc}$ (corresponding to a resolution of 800×800), centred on the centre of the galaxy. We then define the abundance field for each element by $Z = \Sigma_Z / \Sigma$.

Because the primary variation of the metal field is simply a radial gradient caused by the higher star formation rate in the galactic centre, when characterizing the statistics of metal fields it is common to subtract off this gradient (Kreckel et al. 2019, 2020; Li et al. 2021, 2023, 2024b). We therefore also compute an azimuthally averaged metallicity (again for each element) \bar{Z}_r in annular bins of width Δx centred on the galactic centre, and define a metallicity fluctuation map $Z' = Z - \bar{Z}_r$ with the gradient removed.

2.3.2 Auto-, cross-, and time-correlation

The next step in our algorithm is to calculate the two-point auto and cross-correlations of the metal fields. Formally, we define these as an expanded form of Li et al. (2021)

$$\xi_{a,b}(r) = \left\langle \frac{\langle Z'_a(\mathbf{r} + \mathbf{r}') Z'_b(\mathbf{r}') \rangle_{\mathbf{r}'}}{\sqrt{\langle Z'^2_a(\mathbf{r}') \rangle_{\mathbf{r}'} \langle Z'^2_b(\mathbf{r}') \rangle_{\mathbf{r}'}}} \right\rangle_{\theta} \quad (2)$$

where Z'_a and Z'_b are the metallicity fields for two isotopes a and b ; here the angle brackets indicate averages, with $\langle \cdot \rangle_{\mathbf{r}'}$ indicating an average over the dummy position variable \mathbf{r}' and $\langle \cdot \rangle_{\theta}$ an average over the angle θ of the lag vector \mathbf{r} . If a and b are the same isotopes this represents the autocorrelation, while if they are different it represents the cross-correlation. Note that the autocorrelation is exactly equal to unity at lag $r = 0$.

In practice, we compute the auto and cross-correlation by placing all pixel pairs in our maps into bins of lag, and averaging over bins (Li et al. 2021). Let r_n be the central lag of the n th bin, and $r_{n-1/2}$ and $r_{n+1/2}$ be the minimum and maximum lag for that bin. We then compute the correlation function at lag r_n for metal fields a and b as

$$\xi_{a,b}(r_n) = \left(\frac{1}{\sigma_a \sigma_b} \right) \frac{1}{N_n} \sum_{r_{n-1/2} < r_{ij} < r_{n+1/2}} Z'_a(\mathbf{r}_i) Z'_b(\mathbf{r}_j), \quad (3)$$

where \mathbf{r}_i and \mathbf{r}_j are the positions of pixels i and j , $r_{ij} = |\mathbf{r}_i - \mathbf{r}_j|$ is the distance between them, the sum runs over all N_n pixel pairs (i, j) for which $r_{n-1/2} < r_{ij} < r_{n+1/2}$, and $\sigma_{a,b}^2 = \langle Z'^2_{a,b} \rangle$ are the variances of fields a and b across the whole map.

In addition to the static galaxy metal fields that are accessible through observations, we can also study metal fluctuations over time in our simulations. We would therefore like to compute the two-point correlation function between two snapshots at different times. The main difficulty with this is that, just as the spatial correlation is dominated by the metallicity gradient, which we must remove if we are to study higher order statistics, the temporal correlation will be dominated simply by the overall rotation of the galaxy. To remove this effect, we trace all gas particles back through time following the overall galactic rotation curve. We divide the distance r from 0 to 15 kpc into 400 bins, so the bin width matches the sampling size in our metallicity maps. We compute the rotation curve v_ϕ at each galactocentric radius as the mass-weighted mean of the circular velocities of the gas particles in that radial bin, $v_\phi = \sum m_i v_{\phi,i} / \sum m_i$, where the sum runs over all gas particles i in a radial bin, m_i is the particle mass, and $v_{\phi,i}$ is the ϕ component of velocity in a polar (r, ϕ) coordinate system centred on the galactic centre. We show the rotation curve we derive from one of our snapshots (at $t = 564 \text{ Myr}$ – see Section 3) in Fig. 2, and we find that there is very little variation over time, so we adopt the rotation curve shown in this figure for all times.

Having established the rotation curve, we are now in a position to define a correlation in time. While in principle we could compute the correlation as a function of both time and space lag (see KT18), and between different elements, in this work we will limit ourselves to lags in time and for a single element only. We therefore define the

¹<https://github.com/mikegrudic/meshoid>

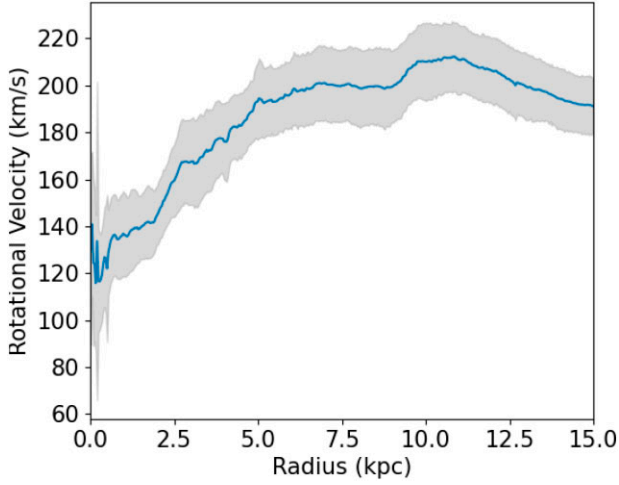


Figure 2. Galactic rotation curve computed as described in Section 2.3.2. The blue curve indicates the mean rotational velocity as a function of galactocentric radius from 0 to 15 kpc at 564 Myr, while the grey area shows the 1σ scatter from particle to particle.

time correlation as

$$\xi(t) = \frac{\langle Z'(\mathbf{r}, t') Z'(\mathbf{r} - v_\phi t \hat{\phi}, t' + t) \rangle_{\mathbf{r}, t'}}{\sqrt{\langle Z'^2(\mathbf{r}, t') \rangle_{\mathbf{r}, t'} \langle Z'^2(\mathbf{r}, t' + t) \rangle_{\mathbf{r}, t'}}}. \quad (4)$$

Here, $Z'(\mathbf{r}, t)$ is the projected metal fluctuation field (for some element) at time t and position \mathbf{r} , $\hat{\phi}$ is a unit vector in the ϕ direction, and the notation $\langle \cdot \rangle_{\mathbf{r}, t'}$ indicates an average over all positions \mathbf{r} and over the dummy time variable t' . The term in the numerator $v_\phi t \hat{\phi}$, where v_ϕ is the rotation velocity evaluated at the particle's radius, removes the effects of overall rotation of the galaxy.

In practice we implement evaluation of equation (4) on a collection of simulation snapshots by selecting pairs of outputs separated by a fixed time interval t .² We then rotate the particle positions in the later snapshot in each pair by an angle $\Delta\phi = -t(v_\phi/r)$ about the axis defined by the galactic plane and galactic centre, generate metal fields from the rotated positions as described in Section 2.3.1, and then evaluate the two-point correlation of the two fields using equation (3) for a spatial lag of zero. This yields a value of the correlation for each snapshot pair with time lag t , and to obtain the final value of $\xi(t)$ we simply take the average over all pairs. We limit our analysis to lags of no more than 50 Myr, since we expect the derotation procedure to become increasingly inaccurate over longer times. Moreover, because rotating the particle positions and generating the metallicity fluctuation maps from those rotated positions is somewhat computationally intense, and we output snapshots at relatively high cadence and thus the number of possible pairs is very large, we do not use all possible snapshot pairs when evaluating equation (4). Instead, we only perform the rotation operation on snapshots at intervals of 10 Myr, though we then compare these rotated snapshots to earlier snapshots at finer time cadences (i.e. so for example we only rotate snapshots at 100 Myr, 110 Myr, and so forth, but we still compute a correlation at a lag of 1 Myr by comparing the rotated 100 and 99 Myr snapshots, the 110 and 109 Myr snapshots, and so on). We have verified that all the results we present do not change qualitatively if we alter the cadence of 10 Myr.

²Since the outputs are uniformly spaced in time, there is no need for binning.

2.3.3 Parametric parameter estimates from the correlation functions

In addition to the raw correlation functions, it is helpful to extract some values by fitting them against a parametric model, for which purpose we use the model of KT18, which has been shown to provide very good fits to both observations (Li et al. 2021, 2023) and lower resolution cosmological simulations (Li et al. 2024b). The model predicts that the two-point correlation function as a function of spatial lag r and time lag t is

$$\xi(r, t) = \frac{2}{\sqrt{\ln\left(1 + \frac{2\kappa t_*}{x_0^2}\right) \ln\left[1 + \frac{2\kappa(t_* - t)}{x_0^2}\right]}} \int_0^\infty e^{-(\kappa t + x_0^2)a^2} \left[1 - e^{-2\kappa(t_* - t)a^2}\right] \frac{J_0(ar)}{a} da, \quad (5)$$

where J_0 is the Bessel function of order zero. The parameters appearing in this expression are the diffusion coefficient κ , the injection time-scale t_* , and the injection width x_0 .

Since the correlations we have computed are evaluated either at the same time (i.e. the time lag $t = 0$) or at the same position (spatial lag $r = 0$), we can specialize equation (5) to these two cases. This gives a spatial correlation at zero time lag

$$\xi(r) = \frac{2}{\ln(1 + 2\phi^2)} \int_0^\infty e^{-l_{\text{corr}}^2 a^2 / \phi^2} \left(1 - e^{-2l_{\text{corr}}^2 a^2}\right) \frac{J_0(ar)}{a} da, \quad (6)$$

and a temporal correlation at zero spatial lag

$$\xi(t) = \frac{\ln(1 + 2\phi^2 - \phi^2 t / t_{\text{corr}}) - \ln(1 + \phi^2 t / t_{\text{corr}})}{\sqrt{\ln(1 + 2\phi^2) \ln(1 + 2\phi^2 - 2\phi^2 t / t_{\text{corr}})}}, \quad (7)$$

where we have defined $l_{\text{corr}} = \sqrt{\kappa t_*}$ as the correlation length, $t_{\text{corr}} = l_{\text{corr}}^2 / \kappa$ as the correlation time, and $\phi = l_{\text{corr}} / x_0$ as the ratio of the correlation and injection lengths. We fit the measured autocorrelation function for element for each snapshot in our simulation to the functional form given by equation (6), and we fit the time correlations for each element to equation (7). We carry out these fits using the PYTHON package EMCEE (Foreman-Mackey et al. 2013), an implementation of an affine-invariant ensemble sampler for Markov chain Monte Carlo (MCMC). For the spatial autocorrelation we use x_0 and l_{corr} as our fit parameters, while for the temporal correlation we use ϕ and t_{corr} ; we adopt flat priors for all values > 0 on all these quantities. We take the (log) likelihood function for the fit to be a χ^2 form given by

$$\ln L = -\frac{1}{2} \sum_{i=1}^N \left[\frac{(\xi_{\text{mod}} - \xi_{\text{sim}})^2}{\sigma_{\xi, \text{sim}}^2} + \ln(\sigma_{\xi, \text{sim}}^2) \right], \quad (8)$$

where the sum is over all the bins of space or time lag at which we measure the simulation correlation function, ξ_{mod} is the KT18 model-predicted correlation function evaluated from equation (6) or equation (7), ξ_{sim} is the auto or time-correlation measured directly from the simulations, and $\sigma_{\xi, \text{sim}} = \sqrt{\text{var}(\xi_{\text{sim}}) / N_{\text{pair}}}$ is a standard error that we set equal to the variance $\text{var}(\xi_{\text{sim}})$ of all the N_{pair} pixel or snapshot pairs that contribute to a given bin. In the MCMC fits, we use 100 walkers and run them for 1000 steps, discarding the first 500 steps for burn-in; we choose this interval by visually examining the flatness of the chains for all parameters. We estimate posterior PDFs from the remaining 500 time-steps.




Figure 3. Star formation rate as a function of time in the simulation; time $t = 0$ corresponds to the time at which we increase the resolution and turn on metal injection and diffusion, and the transient increase in star formation rate that occurs at this time is a result of the change in resolution.

3 RESULTS

We first describe the overall evolution of our simulation in Section 3.1. We examine spatial statistics of the individual element metal fields in Section 3.2 and temporal statistics in Section 3.3. We finally consider the relationships between different elements in Section 3.4.

3.1 Overview of simulation results

We evolve the simulation from the initial condition described in Section 2.2 for a total of 664–100 Myr to $t = 0$ at the original WK23 resolution with no metal injection, and then to $t = 564$ Myr at increased resolution and with metals enabled. We show the star formation rate in the simulation as a function of time in Fig. 3. We see that there is a brief transient when we increase the resolution, but that after ≈ 50 Myr the simulation settles back to near steady-state, giving us ≈ 500 Myr of steady-state evolution to analyse. There is a small secular decrease in star formation rate over this time driven by gas consumption, but this is only a tens of per cent-level effect.

We show the state of the simulation at the final time in Fig. 4. We caution that the absolute metallicities are substantially lower, and the amplitude of metallicity residuals slightly larger, than would typically be expected of spiral galaxy in the present-day Universe. This is simply the result of our having started our simulations from zero metallicity and having run them for only 0.66 Gyr: the simulations have simply not formed enough stars over this time to match the total metallicities produced over 13 Gyr of cosmological star formation. However, this does not matter for our purpose in studying the statistics of metallicity fluctuations, because adding a constant or even a radially varying (but azimuthally symmetric) metallicity would not affect the statistics of interest to us (cf. Section 2.3).

We can see in that the total metallicity and metallicity fluctuations partly mirror the spiral structure visible in the gas (and to a lesser extent stellar) surface density maps, but that the correspondence is not perfect. For example there are several low-density bubbles blown driven by supernovae in the gas surface density map that show high abundances in the total metallicity and fluctuation maps. We also see that the metallicity maps for different elements are quite similar in their overall appearance, a point to which we shall return in Section 3.4.

3.2 Spatial correlations of elemental abundances

We next examine the two-point autocorrelations of individual elements (Section 2.3.2). We show an example measurement of this for

^{14}N at the final time in our simulations in Fig. 5. As the plot shows, we measure a clear autocorrelation, which is qualitatively similar to the ones seen in observed galaxies (e.g. fig. 4 of Li et al. 2023) and in simulations (e.g. fig. 3 of Li et al. 2024b). On the other hand, we also see clear oscillatory structure in the two-point correlation at separations of ≈ 3 –8 kpc that was not visible in earlier observational or theoretical work on metallicity autocorrelations. We return to the question of what physical mechanism is responsible for this structure in Section 4.2.

Given that the KT18 model does not include structures such as the one shown, and in fact the model never predicts negative values for the two-point correlation function, there is some ambiguity in how best of fit the model to the data to extract the parametric quantities discussed in Section 2.3.3. To handle this, we choose to fit to the KT18 model only using the two-point correlation values at separations smaller than the first crossing of 0.1 of the two-point correlation function. We find that doing so yields a significantly better fit to the measured data at small lags, since the fit is no longer attempting to reproduce the negative correlations at large lags. We illustrate this in Fig. 5, where we compare the outcomes of the two fitting procedures.

We show the median injection width and correlation length computed by our MCMC fits as a function of time in Fig. 6; confidence intervals from the MCMC results are not shown to minimize clutter, but are generally very small, such that for most isotopes at most times they would not be visible even if we did plot them. The plot shows that correlation length of all isotopes quickly settles at a few kpc, with factor of ~ 2 oscillations on time-scales of a few 100 Myr. We see that ^{12}C consistently has the smallest correlation length and injection width, and that it stands out as different from all the other isotopes, which are generally clustered close to one another in both x_0 and l_{corr} . Among the other isotopes, ^{16}O , ^{24}Mg , and ^{32}S have on average slightly smaller correlation lengths and slightly larger injection widths compared to ^{14}N , ^{138}Ba , and ^{140}Ce . Each of these groups of three fall nearly on top of the other members of that group. However, even between groups, oscillations of correlation length and injection width are highly correlated, with all generally increasing or decreasing in near lockstep.

Given that the KT18 model does not reproduce some of the major qualitative features of the data, one might worry about parameters extracted based on it. In order to evaluate whether this is a problem, we can also carry out a *non-parametric* fit. We note that for the functional form predicted by the KT18 model at zero time lag (equation 5), the value of the correlation at 1 correlation length is always around 0.15 for physically reasonable values of the injection width, and the value of the correlation function evaluated at 1 injection width is always around 0.97 independent of the value of the correlation length. We can therefore determine model-independent estimates of the injection width and correlation length simply by measuring for what lags the measured autocorrelations pass through 0.97 and 0.15, respectively; the results we obtain by doing so should closely match the parametric values obtained by fitting to the KT18 in cases where that model is a good description of the data. In practice, we implement this measurement by constructing a cubic spline interpolation of the correlation curve and measuring the lags for which this interpolated correlation curve drops to 0.97 and 0.15. We plot the results of this exercise in Fig. 7.

We can see that the injection widths we obtain are significantly smaller than those that emerge from the parametric fit, which suggests a systematic issue with the KT18 model, a topic we will discuss further in Section 4.3. By contrast the correlation lengths we obtain from this model-independent fitting procedure is quite similar to those that result from the parametric fits; typical differences are tens of per cent, we see the same ordering of isotopes from smallest to




Figure 4. Snapshot of the simulation at $t = 564$ Myr. The left column shows the total surface density of gas (expressed as number of H nuclei per unit area) and stars. The upper middle panel shows vertically averaged ^{12}C abundance Z , expressed in $12 + \log(^{12}\text{C}/\text{H})$ units, while the lower middle panel shows the abundance fluctuation map left after subtracting off the mean metallicity in annular bins (see Section 2.3.1 for details). The right column is the same as the middle column, but for ^{16}O . The dimensions of each map are $30 \text{ kpc} \times 30 \text{ kpc}$, while the pixel size of each map is 20 pc .

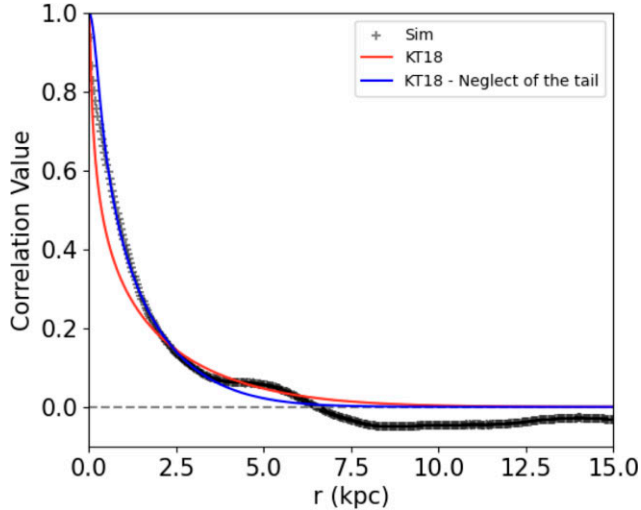


Figure 5. Autocorrelation curve of ^{14}N at $t = 450$ Myr. Black crosses indicate the values measured directly from the simulation, while the red and blue lines are fits to the KT18 model (equation 5 in this paper) over the full 15 kpc shown and over the range of lags smaller than the first crossing of 0.1 of the measured values at $\approx 3 \text{ kpc}$, respectively; the fits shown are evaluated using the median values of each free parameter returned by the MCMC. We see that the fit over the restricted region more closely matches the data at small lags.

largest correlation length, and we see the same major qualitative shape of increase and decrease over time. We therefore conclude that our parametric correlation lengths are robust, and reflect real features present in the underlying spatial autocorrelations.




Figure 6. Best-fitting injection width x_0 and correlation length $l_{\text{corr}} = \sqrt{\kappa t_*}$ from parametric fits to the autocorrelations as a function of time; the values shown are the medians of the posterior PDFs returned by the MCMC fit. Colours correspond to different isotopes, as indicated in the legend.

3.3 Temporal correlations of elemental abundances

We next examine correlations in time, limiting our analysis to lags of $0\text{--}50 \text{ Myr}$ as described in Section 2.3.2. We evaluate the




Figure 7. Same as Fig. 6, but now showing injection widths and correlation lengths derived using a model-independent, non-parametric fitting method; see main text for details.

temporal correlations over this range of lags using data from 304 to 564 Myr, the time period over which our correlation lengths seem to have settled to statistical steady-state (cf. Fig. 6). However, we caution that we still find a slow secular increase in temporal correlation over this period, so for example if we compute the time correlation at a fixed lag using snapshots taken from 304 to 434 Myr versus from 434 to 564 Myr, the latter are typically ~ 10 per cent larger for all isotopes. This suggests that our simulation may not evolve long enough to fully converge on the time correlation.

With this caveat in mind, we show the time-correlation we compute in Fig. 8. We see that elemental abundance patterns decorrelate over time-scales of tens of Myr. As we saw for the correlation length, ^{12}C stands out as the least-correlated of the isotopes we follow, while the remaining isotopes are bunched more closely together. In addition, we found that the **KT18** model slightly overestimates the correlation time, which may be due to the differential rotation of the gas, leading to additional decorrelation beyond the effects of metal diffusion itself.

We report the results of the MCMC fit to equation (7) in Table 1. The time-scale t_{corr} for isotopes to be fully decorrelated [$\xi(t_{\text{corr}}) = 0$] is smallest for ^{12}C at ~ 111 Myr, while all other isotopes take about the same amount of time, ~ 170 Myr. In the Table 1, we also compare the ratio of correlation length and injection width ϕ we derive from fitting the time correlation, which we denote ϕ_t , with the value derived from the fits to the spatial correlation presented in Section 3.2, which we denote ϕ_r ; for the latter, the quantity we report in the table is the median and 16th to 84th percentile range of the medians at each snapshot from 304 to 564 Myr, the same as the time period over which we fit the temporal correlation. We find that the confidence intervals for ϕ_t and ϕ_r are not entirely consistent, with the values derived from the spatial correlation larger. The discrepancy is largest for ^{12}C .

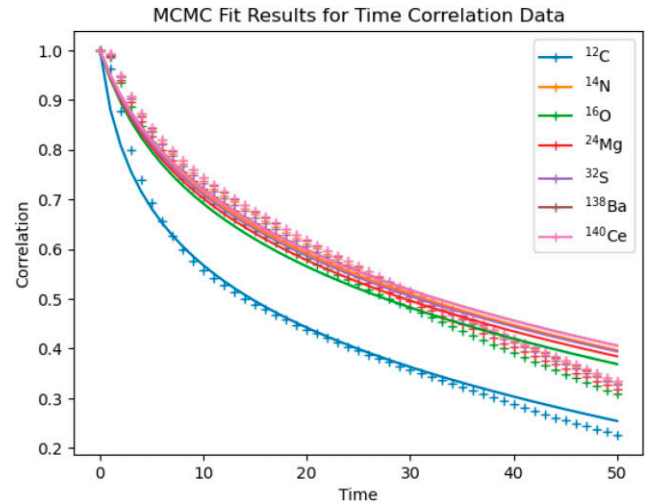


Figure 8. Crosses show the measured time correlation over 0–50 Myr lags while the centre are the mean values using snapshots from 304 to 564 Myr, following the procedure described in Section 2.3.2. Lines show **KT18** model fits to these data, computed using the median values of the posterior PDFs derived by fitting the average measured time correlations to the functional form given by equation (5). Colours correspond to different isotopes, as indicated in the legend.

Table 1. Result from fitting a **KT18** model to the measured time correlations shown in Fig. 8 and the results of spatial correlations shown in Fig. 6 over the corresponding time range for each isotope; see Section 3.3 for details on the fitting procedure. The values for both parameters of the temporal correlation (ϕ_t and t_{corr}) are the median and 68 per cent confidence intervals of the posterior PDF returned by the fit, while the value we report for the spatial correlation (ϕ_r) is the median and 68 per cent range over the time 304–564 Myr used to fit the time correlation (cf. Fig. 6).

Isotope	$\log \phi_t$	$\log \phi_r$	Correlation time t_{corr} (Myr)
^{12}C	$1.00^{+0.33}_{-0.25}$	$1.86^{+0.18}_{-0.26}$	111^{+92}_{-36}
^{14}N	$0.83^{+0.42}_{-0.24}$	$1.27^{+0.30}_{-0.18}$	174^{+295}_{-68}
^{16}O	$0.84^{+0.38}_{-0.25}$	$1.40^{+0.20}_{-0.16}$	152^{+202}_{-57}
^{24}Mg	$0.83^{+0.40}_{-0.24}$	$1.30^{+0.26}_{-0.16}$	170^{+250}_{-67}
^{32}S	$0.85^{+0.41}_{-0.25}$	$1.34^{+0.24}_{-0.16}$	166^{+265}_{-65}
^{138}Ba	$0.83^{+0.40}_{-0.25}$	$1.26^{+0.31}_{-0.16}$	178^{+271}_{-70}
^{140}Ce	$0.84^{+0.43}_{-0.24}$	$1.25^{+0.31}_{-0.17}$	182^{+323}_{-74}

3.4 Cross-correlation between elements

Finally, we compute the cross-correlations between all the isotopes we follow. Although we are free to compute this cross-correlation at any lag, the scientifically interesting value is the cross-correlation at zero lag, since this represents the pattern of gas-phase elemental abundances that will be frozen into newly formed stars, and thus will shape the eventual stellar abundance distribution in chemical space. We show this quantity in Fig. 9. Recall that a value of unity indicates perfect correlation, negative one indicates perfect anticorrelation, and zero indicates no correlation.

The first remark to make based on the figure is that all of the elements are quite well-correlated, with even the least-correlated elements we follow showing correlations of 0.85. We remind readers that this is *not* simply a result of the overall metallicity gradient in the galaxy, since we have subtracted that off and are comparing only




Figure 9. This map shows the cross-correlation among ¹²C, ¹⁴N, ¹⁶O, ²⁴Mg, ³²S, ¹³⁸Ba, and ¹⁴⁰Ce.

the fluctuation maps that remain after doing so. None the less, we find that to first-order all of the isotopes we follow vary up and down together.

A second observation is that, although all elements are highly correlated, we can still roughly divide them into three groups, which are even better-correlated internally, and somewhat less correlated with elements from other groups. These groupings are the same ones that we have seen when examining spatial and temporal correlation statistics. One consists of ¹⁶O, ²⁴Mg, and ³²S, which form the central block in Fig. 9 and are nearly perfectly correlated with each other (correlation coefficient 0.99). The second group consists of ¹⁴N, ¹³⁸Ba, and ¹⁴⁰Ce, which are almost perfectly correlated with each other, but slightly less correlated with the first group. Finally, we have ¹²C in a group of its own, showing the least correlation with the other two groups. This division of the isotopes into three groups mirrors the divisions we saw in the spatial correlation and, at least partially, in the temporal correlation. We attempt to understand the origin of this behaviour in Section 4.1.

4 DISCUSSION

In Section 3, we saw two major themes in the results. First, there are systematic differences between isotopes in their spatial statistics, temporal statistics, and correlations with each other. In these categories the isotopes we have included appear to separate into three groups. Second, while a simple diffusion model like KT18 appears to describe the zeroth-order spatial statistics, to first-order they show significant structure that appears to correlate with the large-scale structures in the galaxy, and they deviate significantly from the KT18 prediction on small scales. We now seek to understand the physical origins of these results.

4.1 On the statistics of element families

We have seen that the spatial statistics of the isotopes we include in our simulation break into three rough groups: ¹²C is alone in one group and is the most different from the others, and then the remaining two consist of ¹⁶O, ²⁴Mg, and ³²S in one group and ¹⁴N, ¹³⁸Ba, and ¹⁴⁰Ce in the other. In order to understand the origins of

this grouping, it is helpful to examine which stars are responsible for producing which isotopes at which times.

For an IMF dN/dM and a set of stellar evolution models that predict the total cumulative mass $M_X(M, t)$ of some isotope X that is returned to the ISM by a star of initial mass M and age t , we can write the IMF-integrated mass return as a function of time as

$$M_X(t) = \frac{\int_{M_0}^{M_1} M_X(M, t) \frac{dN}{dM} dM}{\int_0^\infty \frac{dN}{dM} dM}, \quad (9)$$

where $M_0 = 0.08 M_\odot$ to $M_1 = 120 M_\odot$ are the minimum and maximum stellar masses for our chosen Chabrier (2005) IMF. Similarly, we can write the cumulative contribution to metal return at age t from stars of mass $< M$ as

$$M_X(< M, t) = \frac{\int_{M_0}^M M_X(M, t) \frac{dN}{dM} dM}{\int_{M_0}^{M_1} \frac{dN}{dM} dM}. \quad (10)$$

To evaluate equation (9) for our case, we run SLUG in its non-stochastic mode (i.e. numerically integrating all quantities over the IMF rather than randomly drawing individual stars) using the same Chabrier (2005) IMF and same set of yield tables as in our simulations, integrating to a maximum time $t = 664$ Myr, the duration of our simulation. Similarly, we evaluate equation (10) by running SLUG using an IMF that is a δ function at some initial M , again running to $t = 664$ Myr and recording the cumulative yield at this time. We carry out this procedure for every mass from $2.7 M_\odot$, the minimum stellar mass for which the return is non-zero at 664 Myr given our choice of stellar tracks and yield tables, to $M_1 = 120 M_\odot$ in steps of $0.1 M_\odot$. We then numerically evaluate the integral in equation (10). We plot the cumulative yield $M_X(< M, t)$ at $t = 664$ Myr and the IMF-integrated yield $M_X(t)$ as a function of time in the left and right panels of Fig. 10. For convenience we can also divide the total yield into three distinct channels: AGB injection from stars with initial masses $2.7\text{--}8 M_\odot$, SN injection by stars from 8 to $30 M_\odot$ and Wolf–Rayet (WR) injection from stars $> 30 M_\odot$, though we caution that these divisions are not entirely precise, neither the boundary between stars that do and do not pass through a WR phase nor the boundaries between stars that do and do not explode successfully as SNe lie at a single mass for our chosen tracks and SN models. With this caveat in mind, Table 2 gives the contribution ratio of three nucleosynthetic channels to each of the isotopes we follow and for a stellar population age $t = 664$ Myr – the latter is important, because it means that we are not including a large contribution to carbon from lower mass stars with lifetimes longer than this; we return to this point below.

The most important point to take from Fig. 10 and Table 2 is that the three natural groupings we found in the spatial statistics of the different isotopes we follow are also visible in the yields as a function of mass and time, and the dominant nucleosynthetic channel. At the age $t = 664$ Myr corresponding to the run time of our simulation, ¹²C is mostly produced by stars with initial masses $> 40 M_\odot$ during their WR phase at $t \approx 3\text{--}4$ Myr; ¹⁴N, ¹³⁸Ba, and ¹⁴⁰Ce are dominated by AGB stars and emerges at stellar ages $\gtrsim 50$ Myr (and even long for Ba and Ce); and ¹⁶O, ²⁴Mg, and ³²S are mainly contributed by SNe at stellar ages $\approx 10\text{--}20$ Myr.

The clear link between nucleosynthetic site (or equivalently age of stars that produce a particular isotope) and spatial statistics suggests a simple interpretation: in our simulations ¹²C is injected mostly by very rare events associated with the formation of the most massive stars, and injection happens almost immediately after these stars form. Because few clusters inject it and because the injection precedes the dispersal of much of the circumstellar gas by SNe,




Figure 10. Cumulative yield $M_X(< M, t)$ as a function of initial mass at $t = 664$ Myr (equation 10; left panel) and IMF-integrated yield normalized to the yield at $t = 664$ Myr, $M_X(t)/M_X(664 \text{ Myr})$ (equation 9; right panel), computed using SLUG for the same IMF and yield tables used in our simulations. Colours correspond to different isotopes, as indicated in the legend.

Table 2. Fraction of each isotope that we follow in our simulations produced via the AGB, SN, and WR nucleosynthetic channels for a stellar population of age $t = 664$ Myr aged cluster. Bold numbers highlight the dominant yield channel for this isotope.

isotopes	AGB	SN	WR
^{12}C	0.10	0.11	0.79
^{14}N	0.57	0.23	0.20
^{16}O	0.09	0.53	0.38
^{24}Mg	0.23	0.62	0.15
^{32}S	0.14	0.74	0.12
^{138}Ba	0.89	0.06	0.05
^{140}Ce	0.87	0.07	0.06

^{12}C winds up with both the smallest correlation length and the least correlation with other isotopes. Injection of ^{12}C is followed by injection of ^{16}O , ^{24}Mg , and ^{32}S , which are all coincident with SNe, which dramatically re-arrange the gas compared to the configuration that existed at the time of ^{12}C injection. This explains the lower cross-correlation between these elements and ^{12}C , as well as these elements longer correlation lengths. Finally, ^{14}N , ^{138}Ba , and ^{140}Ce emerge from AGB stars and are not associated with energetic events that re-arrange the gas. These wind up with very slightly larger correlation lengths and somewhat reduced cross-correlations with the SN-associated events because the stars that inject these elements have finite time to drift from their birth sites before giving up their metals, but because their injection is not associated with a dramatic rearrangement of the gas by energetic feedback, the differences are smaller than those between ^{12}C and the others.

We note that this argument raises the question of whether ^{12}C would be as different from other elements as we find if we were to continue the simulation for multiple Gyr, such that we could capture the return of C by $\approx 1\text{--}2.5 M_\odot$ stars on these very long time-scales. This would presumably shift the behaviour of ^{12}C somewhat

closer to that of the other AGB elements. On the other hand, over such long times the stars responsible for element injection may have so thoroughly phase-mixed that their contribution is effectively azimuthally uniform, in which case this contribution would be important for the total yield but unimportant for spatial fluctuation statistics. Determining where between these two possibilities reality lies will require longer time-scale simulations, but maintaining something approaching the high resolution we have available here rather, rather than moving to the much lower resolution used in cosmological simulations.

4.2 Galactic spiral structure imprinted on metallicity distributions

As illustrated in Fig. 4, the spiral structure that is visible in the galactic gas and stellar distribution is reflected in both the metallicity and the metallicity fluctuation maps. This structure, is, in turn, likely responsible for the both the non-monotonic behaviour and the negative values that we see in the autocorrelation at some lags (Fig. 5). The physical origin of this correlation is likely just that spiral arms are where the majority of the stars form, and that supernovae and metal injection occur shortly thereafter. As a result there is a large-scale pattern imprinted on metal injection, which in turn leads to a similar large-scale pattern on the fluctuation distribution.

To demonstrate that these effects are indeed a result of spiral structure, we can instead of studying the full autocorrelation, which involves translations of the map in arbitrary directions, instead study the autocorrelation under rotations about the galactic centre. To be precise, we define the azimuthal autocorrelation of a metallicity fluctuation field Z' as

$$\xi_{\text{az}}(\varphi) = \frac{\langle Z'(\mathbf{r})Z'(\mathbf{r} + \varphi \hat{\phi}) \rangle_{\mathbf{r}}}{\langle Z'^2(\mathbf{r}) \rangle_{\mathbf{r}}}, \quad (11)$$

where $\langle \cdot \rangle_{\mathbf{r}}$ indicates an average over our usual $15 \text{ kpc} \times 15 \text{ kpc}$ region centred on the galactic centre at $\mathbf{r} = 0$. Intuitively, ξ_{az} simply




Figure 11. Azimuthal correlations $\xi_{az}(\phi)$ (equation 11) as a function of rotation angle ϕ at 564 Myr. Colours correspond to different isotopes, as indicated in the legend.

measures the strength of the correlation between the metal field and a version of the metal field that has been rotated through an angle ϕ . This quantity is interesting because of its ability to pick out the presence of large-scale rotationally symmetric structures such as spiral arms. In the absence of such structures, we would expect $\xi_{az}(\phi)$ to approach a constant, non-negative value for large rotation angles. However, if the metal field contains large-scale rotationally symmetric structures, there will be some rotation angles where regions of high Z' in the rotated and non-rotated maps overlap, leading to a high correlation, and other rotation angles where high Z' regions in the rotated map preferentially align with low Z' regions in the non-rotated map, yielding a negative correlation. We plot $\xi_{az}(\phi)$ at $t = 564$ Myr in Fig. 11. We see a positive but monotonically decreasing correlation from 0° to 70° , a small bump at roughly 70° to 100° , and then a negative correlation at 100° to 180° . The presence of these features is consistent with the idea that the non-monotonic behaviour we observe in the full spatial autocorrelation (Fig. 6) is a signature of spiral features with large-scale rotational symmetry.

Spiral-induced features in gas-phase metallicity distributions have not been explored quantitatively before in theoretical models and simulations. The KT18 model explicitly assumes that metal injection is a Poissonian process, with no spatial structure. Previous N -body simulations of isolated galaxies have found spiral arm-associated variations in *stellar* metallicity that they have attributed to stellar radial motions induced by spiral arms (Grand et al. 2015, 2016; Sánchez-Menguiano et al. 2016), but this mechanism depends on the ability of collisionless stars to move radially inward and outward in response to gravitational perturbations from a passing arm. It is clearly inapplicable to the gas. Cosmological simulations have also found evidence for spiral features in gas-phase metallicity maps (Bellardini et al. 2021; Orr et al. 2023; Li et al. 2024b), but at much lower resolution than in our work, and the authors of these works either have not searched for spiral effects on metallicity correlation functions, or have not found convincing evidence for them (e.g. compare fig. 3 of Li et al. 2024b to our Fig. 5). Thus our simulations appear to be the first to identify spiral structures in higher order statistics.

The fact that we see these patterns while previous simulations have not may be a resolution effect: the previous simulations that have found spiral features, from the Feedback In Realistic Environments-

2 (FIRE-2) and Auriga projects, have mass resolutions of $\gtrsim 5 \times 10^3 M_\odot$, more than an order of magnitude lower than our $\approx 300 M_\odot$. At the mean density of the Milky Way's ISM, $\approx 1 \text{ cm}^{-3}$, the former corresponds to ≈ 60 pc resolution, while the latter corresponds to ≈ 20 pc. This difference means that the spiral arms in our simulations are significantly sharper than in the Auriga or FIRE-2 simulations, which may be why we see a stronger spiral arm effect.

Interestingly, there does appear to be some observational evidence for spiral-aligned features in high-resolution metallicity maps. Sánchez-Menguiano et al. (2016) find evidence for spiral structure in high-resolution metallicity fluctuation maps derived from MUSE observations of NGC 6754. Ho et al. (2017, 2018) report similar structures in two galaxies observed as part of the TYPHOON survey, and Wenger et al. (2019) find evidence for spiral features in the metallicity distribution of H II regions in the Milky Way. The features we identify here appear to be at least qualitatively consistent with these observations.

4.3 The injection width and small-scale structures

In addition to the large-scale features in the elemental autocorrelation imprinted by galactic-scale structures like spiral arms, we also see behaviour at small scales that is not fully captured by models. In the context of the original KT18 model, the injection width is the effective radius over which a single event injects metals, before those metals begin to be transported and mixed by general ISM turbulence. If this view were correct, we would expect much smaller injection widths for the isotopes that come predominantly from AGB stars, where injection is not accompanied by an explosive release of energy, than for those injected by primarily by SNe, where it is. By contrast, the correlation length is assumed to be reflective of the properties of ISM turbulence, and thus should presumably be similar for all elements. Neither of these assumptions are fully borne out by our data: ^{12}C , injected by WR stars, has both the smallest correlation length and injection width, while the AGB-dominated isotopes that we follow (^{14}N , ^{138}Ba , and ^{140}Ce) generally have slightly larger injection widths than the SN-dominated ones (^{16}O , ^{24}Mg , ^{32}S).

It is suggestive that, while these isotopes' injection widths do not appear to be correlated with the present of absence of energetic events at the time of injection, they do appear to form a sequence in time: the injection width is smallest for the isotope that is injected first (and prior to the onset of SNe) at stellar population ages of $\approx 3\text{--}4$ Myr – ^{12}C – and largest for the AGB-produced isotopes that are injected on the longest time-scales of ~ 100 Myr. The SN-injected ones injected at $\sim 10\text{--}20$ Myr sit between these two. One possible explanation for this finding is that the primary determinant of injection width is not so much the size of the bubble blown by the energy accompany element return, but whether injection happens prior to dispersal of star-forming clouds, and – closely related – whether stars and gas have time to move relative to one another prior to element returns. The return of significant ^{12}C on short time-scales makes it unique in that return happens before either SNe or stellar drift have had time to induce significant gas-star separation – although our simulations include photoionization, Jeffreson, Semenov & Krumholz (2024) point out that this is insufficient to disperse massive molecular clouds, and these these massive clouds disproportionately drive galactic-scale correlations of SNe. It seems reasonable to hypothesise that the same is true for element distributions. By contrast, SNe-injected elements necessarily accompany gas dispersal, and AGB elements are injected after both gas dispersal and significant drift of stars away from their birth sites have taken place, and thus are injected over an

even wider area than SN-produced ones. Such a picture is consistent with the ordering of injection widths that we have uncovered in our simulations.

5 CONCLUSION

We carry out high-resolution simulations of an isolated, Milky Way-like galaxy in order to provide comprehensive insights into the mechanisms that govern the spatial and temporal distribution of metals in galaxies. By simulating the injection of key isotopes – such as ^{12}C , ^{14}N , ^{16}O , ^{24}Mg , ^{32}S , ^{138}Ba , and ^{140}Ce – using a detailed star-by-star feedback approach, we have been able to observe how different nucleosynthetic processes contribute to galactic metal distributions, and how these distributions reflect elements’ diverse nucleosynthetic origins. Our simulations are the first to study this question at a resolution sufficient to capture the vertical structure and thus the turbulence in the interstellar medium.

We find that the metallicity fluctuation distributions – the residuals left once we remove the overall radial metallicity gradient – of all the isotopes settle to statistical steady state over time-scales of ~ 200 Myr, roughly one galactic rotation. In this steady state the metallicity distributions are correlated on $\sim \text{kpc}$ scales, comparable to the correlation lengths observed in nearby Milky Way-mass galaxies (e.g. Li et al. 2021, 2023). In this steady state the spatial autocorrelation functions of element spatial distributions are reasonably well-described to zeroth-order by the predictions of simple injection–diffusion models such as that proposed by Krumholz & Ting (2018), and the spatial patterns of the metallicity field remain correlated on time-scales of ~ 100 Myr. The abundances of different isotopes are also highly correlated with one another, such that the effective dimensionality of the chemical element space spanned by young stars is likely to be much smaller than one might have guessed under a naive assumption that correlations in different elements – or even different element groups – are independent.

On top of this zeroth-order picture, however, our very high resolution allows us to detect significant additional structure in the spatial statistics of isotope distributions. The different isotopes we follow naturally fall into three groups depending on their dominant nucleosynthetic origins. Specifically, ^{12}C , which over the ≈ 0.7 Gyr duration of our simulations is primarily produced by WR stars, has the smallest spatial correlation length, shortest temporal correlation, and is most weakly correlated with the other elements we follow. These features are likely a result of its very rapid injection, with both preceding the onset of supernovae and occurs before stars have time to undergo significant drift relative to the gas from which they formed. The isotopes we follow that are primarily produced by AGB stars – ^{14}N , ^{138}Ba , and ^{140}Ce – have the longest correlation lengths and are extremely well correlated with one another, likely because the stars that produce them have plenty of time to disperse before they die. Finally, the isotopes we follow whose nucleosynthetic origin is primarily in type II SNe – ^{16}O , ^{24}Mg , and ^{32}S – have an intermediate correlation length and correlate best with each other, slightly less with AGB-produced elements, and still less with ^{12}C , although we emphasize that all of these correlations are still strong in an absolute sense. The different nucleosynthetic origin sites also affect the small-scale structure of the spatial correlation, the quantity that is described in the Krumholz & Ting (2018) model as the ‘injection width’ that characterizes the size of the region into which each stellar source deposits its metals before they begin to be mixed into the ISM. In the Krumholz & Ting this size scale is expected to be determined by the amount of energy that accompanies element release, so SNe, which inflate large bubbles in the ISM, should have large injection

widths than AGB stars, which do not. We find that this expectation is not satisfied in our simulations, and that AGB-produced elements are more correlated on small scales (i.e. have larger injection width) than other isotopes. This suggests that stellar drift and the timing of element injection relative to SNe is a dominant factor in establishing the small-scale behaviour of metallicity correlations.

In addition to these isotope-by-isotope analyses, thanks to our high resolution that captures spiral features in galaxies very well, we are able for the first time to see clear imprints of these galactic-scale structure in statistics of galactic gas-phase metallicity distributions. Spiral structures lead to large-scale non-monotonic and oscillatory patterns in the autocorrelations of metallicity distributions that differ from those predicted by simpler diffusion models. The origin of these features is characteristic fluctuations of increased metallicity that align with regions of enhanced star formation, and the existence of such features underscores the need to incorporate the effects of galactic dynamics when interpreting observed metallicity maps.

Future research should continue to explore the impact of different feedback processes, improve the modelling of stellar drift, and integrate these findings with observational data to enhance our understanding of chemical evolution in galaxies. In addition to increasing the number of distinct isotopes that we follow overall, incorporating feedback and metal injection from type Ia SNe and neutron stars mergers will also be crucial to extending our statistical analysis to include the iron peak and r -process elements that are primarily produced by events of these types. Another substantial upgrade that we intend to include in future work is for newborn stars to inherit the chemical information of their gas particles, allowing us to study the spatial and element-to-element correlations of the stars themselves, and the coupling between stars and gas, following the stars past the moment of their formation. Doing so will be crucial to understanding how stars of that are born near to each other and thus have similar chemical compositions disperse in space, broadening the distribution of abundance patterns at fixed position. Understanding how this process happens is a crucial step toward efforts at chemical tagging – using stellar abundances to reconstruct the star formation history of galaxies (e.g. Bland-Hawthorn & Sharma 2016; Krumholz et al. 2019). By incorporating both observational insights and theoretical models, we can refine our understanding of how metals evolve in galaxies, ultimately providing a more complete picture of galactic evolution and star formation processes.

ACKNOWLEDGEMENTS

The authors acknowledge helpful conversations with Yuan-Sen Ting, Melissa Ness, and Trey Wenger. We acknowledge high-performance computing resources provided by the Australian National Computational Infrastructure (award jh2) through the National and ANU Computational Merit Allocation Schemes. ZL acknowledges the Science and Technology Facilities Council (STFC) consolidated grant ST/X001075/1. MRK acknowledges support from the Australian Research Council through Laureate Fellowship FL220100020.

DATA AVAILABILITY

The data underlying this article will be shared upon reasonable request to the corresponding author.

REFERENCES

- Armillotta L., Krumholz M. R., Di Teodoro E. M., McClure-Griffiths N. M., 2019, *MNRAS*, 490, 4401

- Belfiore F. et al., 2017, *MNRAS*, 469, 151
- Bellardini M. A., Wetzel A., Loebman S. R., Faucher-Giguère C.-A., Ma X., Feldmann R., 2021, *MNRAS*, 505, 4586
- Bland-Hawthorn J., Sharma S., 2016, *Astron. Nachr.*, 337, 894
- Bland-Hawthorn J., Krumholz M. R., Freeman K., 2010, *ApJ*, 713, 166
- Bressan A., Marigo P., Girardi L., Salasnich B., Dal Cero C., Rubele S., Nanni A., 2012, *MNRAS*, 427, 127
- Bundy K. et al., 2015, *ApJ*, 798, 7
- Ceverino D., Sánchez Almeida J., Muñoz Tuñón C., Dekel A., Elmegreen B. G., Elmegreen D. M., Primack J., 2016, *MNRAS*, 457, 2605
- Chabrier G., 2005, *ASSL*, 327, 41
- Colbrook M. J., Ma X., Hopkins P. F., Squire J., 2017, *MNRAS*, 467, 2421
- Croom S. M. et al., 2012, *MNRAS*, 421, 872
- da Silva R. L., Fumagalli M., Krumholz M., 2012, *ApJ*, 745, 145
- de Avillez M. A., Mac Low M.-M., 2002, *ApJ*, 581, 1047
- Di Matteo P., Pipino A., Lehnert M. D., Combes F., Semelin B., 2009, *A&A*, 499, 427
- Di Matteo P., Haywood M., Combes F., Semelin B., Snaith O. N., 2013, *A&A*, 553, A102
- Doherty C. L., Gil-Pons P., Lau H. H. B., Lattanzio J. C., Siess L., 2014, *MNRAS*, 437, 195
- Emsellem E. et al., 2022, *A&A*, 659, A191
- Erroz-Ferrer S. et al., 2019, *MNRAS*, 484, 5009
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Grand R. J. J., Bovy J., Kawata D., Hunt J. A. S., Famaey B., Siebert A., Monari G., Cropper M., 2015, *MNRAS*, 453, 1867
- Grand R. J. J. et al., 2016, *MNRAS*, 460, L94
- Grand R. J. J. et al., 2017, *MNRAS*, 467, 179
- Ho I. T. et al., 2017, *ApJ*, 846, 39
- Ho I. T. et al., 2018, *A&A*, 618, A64
- Hopkins P. F., 2015, *MNRAS*, 450, 53
- Hopkins P. F. et al., 2018, *MNRAS*, 480, 800
- Hu Z., Winking B. D., Krumholz M. R., 2023, *MNRAS*, 521, 5604
- Jeffreson S. M. R., Semenov V. A., Krumholz M. R., 2024, *MNRAS*, 527, 7093
- Karakas A. I., Lugaro M., 2016, *ApJ*, 825, 26
- Kewley L. J., Nicholls D. C., Sutherland R. S., 2019, *ARA&A*, 57, 511
- Kreckel K. et al., 2019, *ApJ*, 887, 80
- Kreckel K. et al., 2020, *MNRAS*, 499, 193
- Krumholz M. R., Ting Y.-S., 2018, *MNRAS*, 475, 2236
- 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
- Leitherer C. et al., 1999, *ApJS*, 123, 3
- Li Z., Krumholz M. R., Wisnioski E., Mendel J. T., Kewley L. J., Sánchez S. F., Galbany L., 2021, *MNRAS*, 504, 5496
- Li Z. et al., 2023, *MNRAS*, 518, 286
- Li S.-l., Li Z., Wisnioski E., Krumholz M. R., Sánchez S. F., 2025, *MNRAS*, 536, 430
- Li Z. et al., 2024b, *MNRAS*, 528, 7103
- López-Cobá C. et al., 2020, *AJ*, 159, 167
- Ma X., Hopkins P. F., Feldmann R., Torrey P., Faucher-Giguère C.-A., Kereš D., 2017, *MNRAS*, 466, 4780
- Maiolino R., Mannucci F., 2019, *A&AR*, 27, 3
- Mármol-Queraltó E. et al., 2011, *A&A*, 534, A8
- Metha B., Trenti M., Chu T., 2021, *MNRAS*, 508, 489
- Orr M. E. et al., 2023, *MNRAS*, 521, 3708
- Petit A. C., Krumholz M. R., Goldbaum N. J., Forbes J. C., 2015, *MNRAS*, 449, 2588
- Poetrodjojo H. et al., 2018, *MNRAS*, 479, 5235
- Richard J. et al., 2024, preprint ([arXiv:2406.13914](https://arxiv.org/abs/2406.13914))
- Rosales-Ortega F. F., Díaz A. I., Kennicutt R. C., Sánchez S. F., 2011, *MNRAS*, 415, 2439
- Sánchez-Menguiano L. et al., 2016, *ApJ*, 830, L40
- Sánchez-Menguiano L. et al., 2018, *A&A*, 609, A119
- Sánchez S. F. et al., 2012, *A&A*, 538, A8
- Sánchez S. F., Walcher C. J., Lopez-Cobá C., Barrera-Ballesteros J. K., Mejía-Narváez A., Espinosa-Ponce C., Camps-Fariña A., 2021, *Rev. Mex. Astron. Astrofis.*, 57, 3
- Sharda P., Krumholz M. R., Wisnioski E., Forbes J. C., Federrath C., Acharyya A., 2021, *MNRAS*, 502, 5935
- Smith B. D. et al., 2017, *MNRAS*, 466, 2217
- Sukhbold T., Ertl T., Woosley S. E., Brown J. M., Janka H. T., 2016, *ApJ*, 821, 38
- Ting Y.-S., Weinberg D. H., 2022, *ApJ*, 927, 209
- Tinsley B. M., 1980, *Fund. Cosmic Phys.*, 5, 287
- Tissera P. B., Rosas-Guevara Y., Sillero E., Pedrosa S. E., Theuns T., Bignone L., 2022, *MNRAS*, 511, 1667
- Tremonti C. A. et al., 2004, *ApJ*, 613, 898
- Turk M. J., Smith B. D., Oishi J. S., Skory S., Skillman S. W., Abel T., Norman M. L., 2011, *ApJS*, 192, 9
- Weinberg D. H. et al., 2019, *ApJ*, 874, 102
- Weinberg D. H. et al., 2022, *ApJS*, 260, 32
- Wenger T. V., Balser D. S., Anderson L. D., Bania T. M., 2019, *ApJ*, 887, 114
- Winking B. D., Krumholz M. R., 2023, *MNRAS*, 521, 5972
- Williams T. G. et al., 2022, *MNRAS*, 509, 1303
- Yang C.-C., Krumholz M., 2012, *ApJ*, 758, 48

This paper has been typeset from a \LaTeX file prepared by the author.