*Shared stellar sources leave related chemical patterns in gas and newborn stars.*

MNRAS **000**, 1–20 (2026) Preprint 9 September 2026 Compiled using MNRAS LATEX style file v3.3

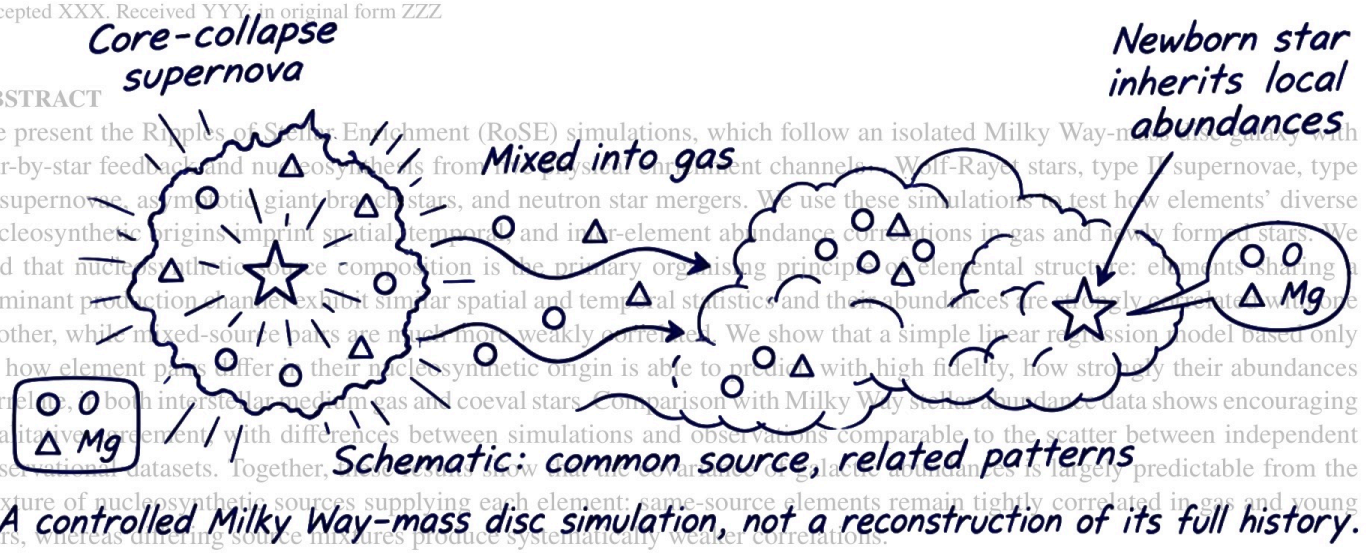
# Ripples of Stellar Enrichment (RoSE) – simulating element production and mixing in a Milky Way-mass galactic disc star-by-star

Chuhan Zhang[1★], Mark R. Krumholz[1], Melissa K. Ness[1], Yuan-Sen Ting[2,3], Zefeng Li[4], and Zipeng Hu[5]

[1]*Research School of Astronomy and Astrophysics, Australian National University, Cotter Road, Weston ACT 2611, Australia*
[2]*Department of Astronomy, The Ohio State University, Columbus, OH 43210, USA*
[3]*Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany*
[4]*Centre for Extragalactic Astronomy, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK*
[5]*Kavli Institute for Astronomy and Astrophysics, Peking University, 5 Yiheyuan Road, Haidian District, Beijing 100871, People's Republic of China*

arXiv:2603.29242v2 [astro-ph.GA] 7 Sep 2026

Accepted XXX. Received YYY; in original form ZZZ




## ABSTRACT

We present the Ripples of Stellar Enrichment (RoSE) simulations, which follow an isolated Milky Way-mass galaxy with star-by-star feedback and nucleosynthesis from five physical enrichment channels: Wolf-Rayet stars, type II supernovae, type Ia supernovae, asymptotic giant branch stars, and neutron star mergers. We use these simulations to test how elements' diverse nucleosynthetic origins imprint spatial, temporal, and inter-element abundance correlations in gas and newly formed stars. We find that nucleosynthetic source composition is the primary organising principle of elemental structure: elements sharing a dominant production channel exhibit similar spatial and temporal statistics and their abundances tracked closely on top of one another, while mixed-source pairs are much more weakly correlated. We show that a simple linear regression model based only on how element pairs differ in their nucleosynthetic origin is able to predict, with high fidelity, how strongly their abundances correlate, in both interstellar medium gas and coeval stars. Comparison with Milky Way stellar abundance data shows encouraging qualitative agreement, with differences between simulations and observations comparable to the scatter between independent observational datasets. Together, these results show that the covariance of galactic abundances is largely predictable from the mixture of nucleosynthetic sources supplying each element: same-source elements remain tightly correlated in gas and young stars, whereas differing source mixtures produce systematically weaker correlations.

**Key words:** galaxies: abundances — galaxies: ISM — ISM: abundances — nucleosynthesis — stars: abundances

## 1 INTRODUCTION

The distribution of metals (elements heavier than helium) within galaxies provides a critical window into the physics of the interstellar medium (ISM) and galaxy assembly. Metals are produced primarily through stellar nucleosynthesis, with stars of different masses producing different elements over different timescales, and then injected into the ISM through stellar winds, supernovae (SNe), and compact object mergers. Once in the ISM, these metals are mixed by turbulence, transported by large-scale flows, and incorporated into new generations of stars, driving the chemical evolution of galaxies and shaping the conditions for future star formation (for reviews, see Tinsley 1980; Maiolino & Mannucci 2019; Sánchez et al. 2021).

The past two decades have witnessed a revolution in our ability to map metal distributions spatially. The metallicity of the ISM, most commonly traced via oxygen abundance in ionised regions, can be measured using emission line diagnostics (for a review, see Kewley et al. 2019). The deployment of integral field units (IFUs) enables measurements of spatially resolved two-dimensional distributions of oxygen abundance across nearby galaxies (e.g., Rosales-Ortega et al. 2011; Sánchez-Menguiano et al. 2016; Huang et al. 2026). These observations reveal metallicity gradients, typically showing that metal abundance decreases from centres of galaxies outward (e.g., Belfiore et al. 2017; Ho et al. 2018; Poetrodjojo et al. 2018). A range of theoretical studies have aimed to explain 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).

However, gradients represent a significant simplification of the data, since they collapse complex two-dimensional (2D) maps down to a single inside-out linear fit. To exploit the full information content of IFU metallicity maps, higher-order statistics are needed to decode the detailed processes of metal injection and mixing. One of the simplest yet most informative statistics for 2D maps is the two-point correlation function, which quantifies characteristic length scales over which metallicity fluctuations persist. Krumholz & Ting (2018, hereafter KT18) provide a minimal theoretical framework to predict two-point correlations of galaxy metallicities based on the competition between metal injection and diffusion. This prediction has motivated various observational studies examining two-point correlations (or similar statistics) in nearby galaxies (e.g., Kreckel et al. 2020; Li et al. 2021, 2023; Metha et al. 2021; Williams et al. 2022; Li et al. 2025a). These studies find that metallicity maps of nearby galaxies

★ E-mail: chuhan.zhang@anu.edu.au

© 2026 The Authors

Existing theory predicts metal mixing statistics, but previous simulations lacked delayed sources and stellar chemical inheritance.

contain statistically significant spatial structure beyond the overall gradient, and that the two-point correlation functions describing this structure generally follow the shape predicted by KT18, with correlation lengths of $\sim 1$ kpc, though with significant systematic variations with galaxy properties such as stellar mass and star formation rate.

Very recently, two-dimensional metallicity maps have also started to become available for other elements, most notably nitrogen and sulfur (Bresolin et al. 2025; Li et al. 2025b), enabling both comparisons of the spatial statistics of different elements and calculations of the extent of element-element cross-correlation. Krumholz et al. (2025) extended the original KT18 theory to incorporate such statistics. The models are qualitatively consistent with the data thus far, but to date there are only two galaxies with measured two-dimensional abundance maps for multiple elements. This situation is likely to improve as facilities such BlueMUSE come online (Richard et al. 2024), providing access to key diagnostic lines for nitrogen and potentially other elements.

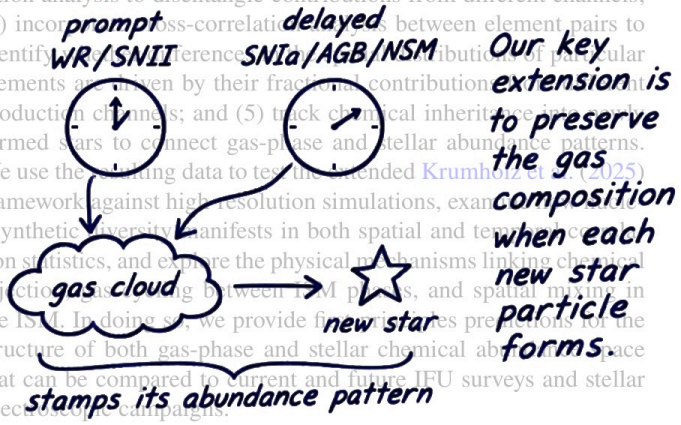
Stars within the Milky Way provide another, complementary, observational window on elemental abundance distributions. Only a small number of elements, all produced primarily by either core collapse SNe or AGB stars, can be measured in the gas. By contrast, stellar spectra provide access to a much larger suite of elements. While such measurements do not directly trace the spatial distribution of elements in galaxies, since stars migrate after formation, any spatial correlations between elements that were present at formation remain frozen into the distribution of stars in chemical abundance space (excluding fully convective stars whose surface abundances can be altered by internal processes). There have been extensive efforts to understand this chemical space structure, since it matters for a wide variety of studies relying on stellar abundances (Bland-Hawthorn et al. 2010; Bland-Hawthorn & Sharma 2016; Krumholz et al. 2019; Weinberg et al. 2022; Ting & Weinberg 2022; Manea et al. 2025; Mead et al. 2025), and the data are again qualitatively consistent with the analytic Krumholz et al. (2025) model.

While analytic and observational work has begun the study of metallicity distribution statistics, simulation efforts have been limited. Li et al. (2024) post-processed the Auriga cosmological simulations (Grand et al. 2017) to produce metallicity maps comparable to those accessible via observations, finding that simulations successfully reproduce observed correlation lengths and are in reasonable agreement with KT18 predictions. In Zhang et al. (2025, hereafter Paper I) we extended this work using isolated galaxy simulations with mass resolution more than an order of magnitude better than Auriga, and star-by-star tracking of feedback and nucleosynthesis, enabling us to follow metal return by individual stellar sources. This additional resolution allows us to model potentially important mixing associated with small-scale phenomena such as spiral arms (Grand et al. 2016; Orr et al. 2023) and thermal instability (Yang & Krumholz 2012) more explicitly, complementing cosmological simulations that capture galaxy assembly and environmental effects. Paper I demonstrated that small-scale metallicity fluctuations naturally organize into nucleosynthetic families, with elements from the same production channel (e.g., O, Mg, S from core-collapse supernovae) correlating more strongly with each other than with elements from different channels, consistent with the results of observations of both gas and stars (Ting & Weinberg 2022; Li et al. 2025b; Mead et al. 2025).

A half-gigayear run misses much of the iron-peak and neutron-capture story.

However, the work in Paper I left significant gaps. First, due to its limited duration the simulation used in that paper could only be run for $\sim 0.5$ Gyr, and as a result the paper omitted elements produced via nucleosynthetic channels that require longer timescales to return their products, most prominently iron peak and *s*- and *r*-process nuclides. Second, the simulation code used in that paper only recorded elemental abundances in gas instead of newly-produced stars limiting its utility for making comparisons with the growing body of stellar chemical abundance space measurements.

These gaps motivate the present study, which extends the high-resolution isolated Milky Way-mass disc simulation framework established in Paper I to address multiple outstanding questions simultaneously. Building on our star-by-star treatment of feedback and nucleosynthesis, we now (1) expand coverage to six modelled nucleosynthetic channels including Type Ia supernovae and neutron star mergers, enabling complete tracking of nine isotopes spanning short and long delay times; (2) extend the simulation duration to 600 Myr, which is long enough for the correlation statistics we analyse to settle into a quasi-steady state; (3) implement component-resolved correlation analysis to disentangle contributions from different channels; (4) incorporate cross-correlations between element pairs to identify differences in the distribution of particular elements arising from their fractional contribution from different production channels; and (5) track chemical inheritance into newly-formed stars to connect gas-phase and stellar abundance patterns. We use the resulting data to test the extended Krumholz et al. (2025) framework against high-resolution simulations, examine how nucleosynthetic diversity manifests in both spatial and temporal correlation statistics, and explore the physical mechanisms linking chemical injection, mixing between ISM phases, and spatial mixing in the ISM. In doing so, we provide first principles predictions for the structure of both gas-phase and stellar chemical abundance space that can be compared to current and future IFU surveys and stellar spectroscopic campaigns.




## 2 SIMULATION

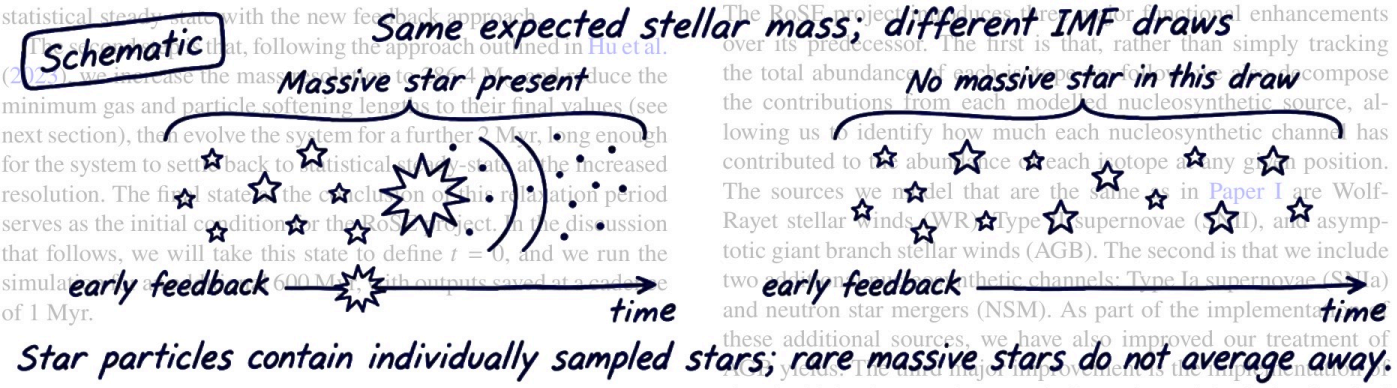
We present the second generation of our numerical simulations, Ripples of Stellar Enrichment (RoSE), developed from the isolated, magnetised, hydrodynamic $L_*$ disc galactic chemical simulation presented in Paper I. This new simulation retains the star-by-star tracking method for chemical feedback introduced in our previous work but includes a broader range of elements and nucleosynthetic sources. In this section, we only summarise most aspects of the simulations, and go into detail only when discussing specific improvements and key differences between the two generations of simulations. For full details regarding the numerical setup and methodology of the first-generation simulations, we refer the reader to Paper I.

### 2.1 Initial conditions

As in the first-generation simulations, we begin from the 600 Myr snapshot from Wibking & Krumholz (2023); these authors used the GIZMO code (Hopkins 2015) to simulate an isolated galaxy with initial conditions taken from the AGORA project (Kim et al. 2016), which was designed to be a close match to the bulk properties of the Milky Way. The simulation includes a gaseous disc, a stellar disc, a stellar bulge, and a dark matter halo, all of which are simulated live rather than using fixed potentials. However, we caution that, while the simulation matches gross Milky Way properties such as stellar mass, gas mass, and bulge-to-disc ratio, we do not attempt to reproduce more detailed aspects of the Milky Way's structure such as the properties of the bar, interactions with satellite galaxies, or the effects of external accretion; we defer a full discussion of these caveats and their implications to Section 5.4.

Wibking & Krumholz (2023) used the native implementation of stellar feedback in GIZMO (Hopkins 2015), which was appropriate for their $\approx 10^3$ $M_\odot$ resolution, but is insufficient for the much higher resolution we seek to reach. Our first step therefore is to restart the simulation replacing the GIZMO feedback module with one that uses star-by-star-tracking based on the SLUG (Stochastically Lighting Up Galaxies; da Silva et al. 2012; Krumholz et al. 2015) stellar population synthesis code. In this approach for each star particle formed we draw a stellar population for that particle star-by-star from the initial mass function, and then inject ionising and supernova feedback based on star-by-star calculations of ionising luminosity and timing of supernova explosions. We use this method to continue the simulation for an additional 100 Myr, allowing it to re-settle into statistical steady state with the new feedback approach.

The second change is that, following the approach outlined in Hu et al. (2023), we increase the mass resolution to 286 $M_\odot$ and reduce the minimum gas and particle softening lengths to their final values (see next section), then evolve the system for a further 3 Myr, long enough for the system to settle back to statistical steady state at the increased resolution. The final state of the conclusion of this relaxation period serves as the initial condition for the RoSE project. In the discussion that follows, we will take this state to define $t = 0$, and we run the simulation for a further 600 Myr, with outputs saved at a cadence of 1 Myr.




## 2.2 Numerical methods

### 2.2.1 Simulation physics: MHD, cooling, star formation, and feedback

We run our simulation using GIZMO's meshless finite-mass (MFM) Godunov solver with ideal MHD enabled and constrained-gradient divergence control (Hopkins 2015, 2016; Hopkins & Raives 2016). The simulation includes multiple metal fields (see below for details), each of which is treated as a passive scalar. In addition to advective transport with the gas, these scalars are diffused using GIZMO's default Smagorinsky turbulent model (Colbrook et al. 2017; Hopkins et al. 2018). The simulation includes a simple non-equilibrium chemical network to follow H and He ionisation, and uses the GRACKLE module (Smith et al. 2017) to handle radiative cooling and heating; we use this rather than the default GIZMO cooling module because the latter does not properly produce a two-phase neutral medium (see Wibking & Krumholz 2023 for details). Gravity and hydrodynamics use adaptive softening with a minimum 0.0001 code units for gas ($\approx 0.1$ pc physical), 0.005 for "active" star particles formed during the simulation, 0.05 for dark matter, and 0.02 for "passive" star particles present from the start of the simulation.

Star formation in the simulation begins once gas reaches a density threshold $\rho_c = 10^3$ H atoms per cm$^3$, which at our final resolution corresponds to a density for which the particle mass is approximately equal to the Jeans mass for gas in thermal equilibrium. For gas above this threshold, the per-freefall star formation efficiency is $\epsilon_{\rm ff} = 0.01$, consistent with observations (Krumholz et al. 2019, and references therein). To conserve computational resources, we increase this to 100 (i.e., immediate conversion to stars guaranteed) for gas exceeding 100 times the density threshold. As mentioned above, when a star particle forms, we use SLUG (da Silva et al. 2012; Krumholz et al. 2015) to draw a population of individual stars that will populate that star particle, and whose evolution will determine its feedback thereafter. We use the same IMF, tracks, and stellar atmospheres in SLUG as in Paper I. Stars feed back on the gas via photoionisation and supernovae. The photoionising luminosities of each star particle are computed on-the-fly by SLUG, and we implement photoionisation using a Strömgren volume approximation. Similarly, SLUG identifies which individual stars explode as supernovae using the models of Sukhbold et al. (2016), and provides the time after formation at which these explosions occur; each explosion adds $10^{51}$ erg to the gas, which we implement using a mixed momentum-energy feedback model since even at our high resolution the Sedov-Taylor phase is often unresolved. See Armillotta et al. (2019) and Wibking & Krumholz (2023) for details on how feedback is implemented.

### 2.2.2 Metal injection and tracing

The RoSE project introduces three major additional enhancements over its predecessor. The first is that, rather than simply tracking the total abundance of each isotope, we also decompose the contributions from each modelled nucleosynthetic source, allowing us to identify how much each nucleosynthetic channel has contributed to the abundance of each isotope at any given position. The sources we model that are the same as in Paper I are Wolf-Rayet stellar winds (WR), Type II supernovae (SNII), and asymptotic giant branch stellar winds (AGB). The second is that we include two additional nucleosynthetic channels: Type Ia supernovae (SNIa) and neutron star mergers (NSM). As part of the implementation of these additional sources, we have also improved our treatment of stellar yields. The third major improvement is the implementation of chemical inheritance, where newly formed stars inherit the chemical abundances of the gas from which they form. This provides a self-consistent framework within which we can simultaneously study the chemical evolution of the stellar and gaseous components within the galaxy. In the remainder of this section we detail the implementation of each of these improvements.

Within the RoSE framework, we independently track nine isotopes – which we list in Table 1 – produced by six distinct nucleosynthetic channels. We track each combination of element and channel as a distinct scalar field, so we in the analysis that follows we can reconstruct not just the abundances of all elements, but the relative contributions to those abundances made by different channels. Given that our simulation runs for $\lesssim 1$ Gyr, we can broadly divide these channels into two categories: those that return elements on timescales comparable to or smaller than the duration of our simulation, and where therefore all of the element return is due to stars that form self-consistently during the simulation, and those that operate on longer timescales and where element return must therefore be attributed to stars that were present in the simulation initial condition. We use distinct methods to treat these two channels. In the discussion that follows we will refer to star particles that form self-consistently over the course of the simulation, which have SLUG stellar populations associated with them, and "old star particles" to refer to star particles that were present prior to $t = 0$, and which do not have associated SLUG stellar populations.

Source contributions remain separable even after their metals mix.

**2.2.2.1 Short-timescale nucleosynthetic channels.** Our short-timescale nucleosynthetic channels include enrichment from WR stars, SNII explosions, and massive AGB stars. For these channels element return comes from SLUG particles, and our numerical implementation is as follows. Each SLUG particle has an associated stellar population drawn from the IMF, and each star drawn in turn has an individual nucleosynthetic yield, which is taken from tabulations of the yields for particular channels. We take our yields for WR winds and SNII explosions from Sukhbold et al. (2016), while we take our

Prompt sources follow newly formed stars, while delayed sources are attached statistically to the pre-existing stellar population.

yields for AGB stars from Doherty et al. (2014) and Karakas & Lugaro (2016); see Paper I for details. Note that only some AGB stars return yields on timescales shorter than our simulation duration; we discuss how we divide the AGB population based on its lifetime below.

For WR winds and SNII, element return occurs over a short timescale relative to the dynamical time of the simulation, and we therefore retain the scheme used in Paper I: when a star reaches the end of its life, we instantaneously add the elements in returns to the surrounding gas particles using the same weighting kernel as that used to distribute supernovae energy and momentum. In contrast, AGB stars contribute feedback through a continuous mass loss process over extended timescales. Implementing this in a truly continuous manner would be computationally prohibitive, so we use an approximate method wherein the AGB yield is calculated and injected into the ISM only at the moment its parent SLUG particle is dynamically awakened by the simulation code. During this awakening phase, the code checks the particle's evolutionary history and computes the total integrated AGB yield accumulated since the last awakening event. These awakenings happen on a locally-determined dynamical timescale, ensuring that gas deposition happens on the same timescale as gas or stellar motion.

**2.2.2.2 Long-timescale nucleosynthetic channels.** The long-timescale return channels are SNe Ia, NSMs, and low-mass AGB stars, and we assign mass return via these channels to old star particles. For SNe Ia and NSMs, we implement a stochastic model whereby we assign a specific occurrence rate (i.e., rate per unit mass) for each event, which we apply probabilistically. For SNIa, the observed rate in Milky Way-mass galaxies is approximately 0.5 event per century (Li et al. 2011; Wiseman et al. 2021), and our simulation initial condition has a stellar mass of $3.7 \times 10^{10}$ $M_\odot$, so we set the rate per unit time per unit stellar mass to $\Gamma_{\rm Ia} = 1.4 \times 10^{-13}$ $M_\odot^{-1}$ $yr^{-1}$. Thus during each update through time $\Delta t$ for an old star particle of mass $M_*$, we assign it a probability $P = 1 - \exp(-\Gamma_{\rm Ia} M_* \Delta t)$ of hosting a SNIa.

We boost the neutron-star-merger rate only for sampling, then reduce each event's yield to preserve total production.

We handle NSMs in exactly the same way, except that the rate of NSMs is considerably lower. Although the most commonly cited estimate for the rate is one event per 100 events per million years (Kalogera et al. 2004), only a fraction of stars with $8 - 20$ $M_\odot$ eventually producing NSMs (Hirai et al. 2022), we enhance this rate to 500 events per million years in our model, corresponding to $\Gamma_{\rm NSM} = 1.4 \times 10^{-14}$ $M_\odot^{-1}$ $yr^{-1}$ to ensure that a statistically significant sample of NSM events occur during the simulation timeframe, allowing for a robust investigation of their distinct contribution to the galactic chemical inventory, particularly for elements like r-process nuclei; we reduce the yield per event correspondingly, so that the total element production rate remains the same.

For both NSM and SNIa, once an explosion event has been triggered in an old stellar particle, we then inject $10^{51}$ erg of energy, along with 1.4 $M_\odot$ of mass for SNIa and 0.18$M_\odot$ for NSM, to the surrounding ISM. For the SNIa channel, we adopt element-by-element yields per event $\Delta M_i$ from Seitenzahl et al. (2013), specifically their Table 2 for the case of a progenitor metallicity $Z = 0.02$. For the NSM channel, we base the yield calculations on Holmbeck & Andrews (2024) again diluting down by the same factor by which we have enhanced the rates so that the total yield is the same. We summarise the mass return per event for each element for each channel in Table 1. We distribute the injected mass from these channels to fluid elements surrounding the triggering particle exactly as we do for SNII.

WR / SNII / young AGB

SNIa / NSM / old AGB

birth → time

For old AGB stars, which inject mass near-continuously, we inject mass from each old star particle each time it is awakened during the simulation evolution, using the same strategy as for AGB yields from SLUG particles. To determine the rate of element return, we use SLUG to calculate the cumulative yield $\Delta M_i$ of each element $i$ from a simple stellar population of mass $M_{\rm SSP}$ over a time $t_{\rm H} = 13.6$ Gyr, using the same IMF, tracks, and yield tables as for SLUG particles, and disabling all nucleosynthetic mechanisms except AGB stars. We then set the element return rate per unit mass for old star particles to $\dot{y}_i = \Delta M_i / M_{\rm SSP} t_{\rm H}$, which is equivalent to assigning each old star particle a mass return rate equal to the time-averaged return rate over a Hubble time. This is a crude approximation, since it ignores the detailed star-formation history of the galaxy, so in Appendix A we quantify its impact by repeating the elemental cross-correlation analysis after multiplying the old-star AGB contribution by factors $f = 0$, 0.5, 1, and 2. This post-processing test brackets both removal and a factor-of-two enhancement of the adopted rate.

The old-star AGB rate is deliberately crude, so Appendix A stress-tests it rather than hiding the approximation.

We assign old AGB yields in Table 1 to the long-timescale channel associated with old star particles, and we track a separate scalar field for each long-timescale channel, and reconstruct the isotope-by-isotope contribution by multiplying that scalar by the fixed yield ratios in Table 1. By contrast, the short-timescale channels associated with newly formed SLUG particles are sampled stochastically from the IMF, so two star particles with the same total mass need not return elements in identical ratios. For those channels we therefore track separate scalar fields for each combination of isotope and nucleosynthetic channel.

**Table 1.** Old star yield table used in RoSE. For SNIa and NSM we report the mass returned per event; SNIa values are taken from Seitenzahl et al. (2013), and NSM values from Holmbeck & Andrews (2024). For old AGB stars, the quantity listed is the mass return per unit stellar mass unit time.

| Isotope | SNIa [$M_\odot$] | NSM [$M_\odot$] | old AGB [$Myr^{-1}$] |
|---|---|---|---|
| $^{12}$C | $3.04 \times 10^{-3}$ | 0 | $1.00 \times 10^{-7}$ |
| $^{14}$N | $3.21 \times 10^{-6}$ | 0 | $4.71 \times 10^{-8}$ |
| $^{16}$O | $1.01 \times 10^{-1}$ | 0 | $1.15 \times 10^{-7}$ |
| $^{24}$Mg | $1.52 \times 10^{-2}$ | 0 | $1.14 \times 10^{-8}$ |
| $^{32}$S | $1.11 \times 10^{-1}$ | 0 | $6.05 \times 10^{-9}$ |
| $^{56}$Fe | $6.22 \times 10^{-1}$ | 0 | $2.45 \times 10^{-8}$ |
| $^{138}$Ba | 0 | $9.2 \times 10^{-7}$ | $3.30 \times 10^{-12}$ |
| $^{140}$Ce | 0 | $2.5 \times 10^{-7}$ | $9.21 \times 10^{-13}$ |
| $^{153}$Eu | 0 | $1.1 \times 10^{-7}$ | $6.27 \times 10^{-15}$ |

iron channel

neutron-capture products

## 3 ANALYSIS METHODS

This paper extends the two-point correlation toolkit developed in Li et al. (2021) and Paper I to (i) quantify spatial and temporal scales of mixing and (ii) diagnose how nucleosynthetic channels imprint distinct correlation signatures across elements. In this section we detail the steps in our analysis pipeline.

### 3.1 Map-making and pre-processing

Our raw outputs consist of 600 snapshots files, saved at a cadence of 1 Myr; each file contains fields storing the abundance of each element in each gas and SLUG particle. Our first processing step is that, for each snapshot and element, we construct a face-on 2D map

We remove the radial gradient, then measure the size and lifetime of the remaining chemical patches.



over a 20 kpc × 20 kpc square centred on the galaxy centre[1], sampled on an 800 × 800 grid ($\Delta x = \Delta y = 25$ pc). For each pixel in the map we compute surface densities of total gas mass and mass of each element, $\Sigma$ and $\Sigma_Z$. We compute the elemental surface densities both for the total mass of each element listed in Table 1 summed over all nucleosynthetic sources, and for the individual fields broken down by source. For each element we define the log-metallicity field as

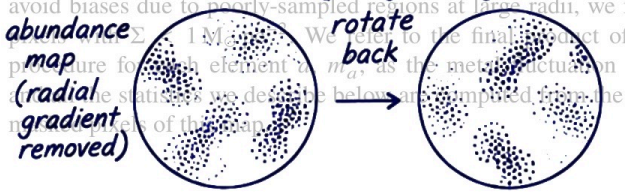
$$m = \log \frac{\Sigma_Z}{\Sigma}. \tag{1}$$

To remove the dominant radial trend, we subtract the azimuthal median profile in annuli of width $\Delta r = 25$ pc to obtain the metal fluctuation field,

$$m' = m - \overline{m}(r), \tag{2}$$

Without subtracting the radial profile, the galaxy's inside-out gradient would masquerade as long-range mixing.

where $\overline{m}(r)$ is the median metallicity in each radial bin. Finally, to avoid biases due to poorly-sampled regions at large radii, we mask pixels with $\Sigma < 1\,M_\odot$ pc$^{-2}$. We refer to the final product of this procedure for each element as $m'_a$ as the metal fluctuation map, and the statistics we describe below are sampled from the non-masked pixels of this map.




### 3.2 Spatial autocorrelation

To quantify spatial structure of each element, we next compute the autocorrelation for element $a$ as

$$\xi_{a,a}(r) = \left\langle \frac{\langle m'_a(\mathbf{r}' + \mathbf{r})\, m'_a(\mathbf{r}')\rangle_{\mathbf{r}'}}{\langle m'^2_a(\mathbf{r}')\rangle_{\mathbf{r}'}} \right\rangle_\theta, \tag{3}$$

where the outer average over $\theta$ is the mean over all orientations of the position vector $\mathbf{r}$, and the inner averages over $\mathbf{r}'$ are averages over all positions. In practice we evaluate this quantity by considering every pixel pair in the map and binning pairs by separation. Let $r_n$ be the central separation for bin $n$ with edges $r_{n\pm1/2}$. Then we compute the autocorrelation at separation $r_n$ as

$$\xi_{a,a}(r_n) = \frac{1}{\sigma_a^2}\frac{1}{N_n}\sum_{r_{n-1/2}<r_{ij}<r_{n+1/2}} m'_a(\mathbf{r}_i)\, m'_a(\mathbf{r}_j), \tag{4}$$

where $\sigma_a^2 = \langle m'^2_a\rangle$ is the variance over the total metal fluctuation field and $N_n$ is the number of pairs in the bin. We evaluate the autocorrelation over bins of width $\Delta r = 25$ pc from separations of zero (where $\xi_{a,a} = 1$ by construction) up to the maximum scale allowed by the 20 kpc map. We estimate uncertainties in the autocorrelation using a jackknife procedure of dividing the map into 12 equal-area azimuthal sectors; the jackknife variance across sectors provides conservative errors on $\xi_{a,a}(r)$.

In addition to the raw auto-correlation, we also compute parameters for a parametric fit to the autocorrelation data. In Paper I, we showed that the autocorrelation data are reasonably well-fit by the analytic model proposed by KT18, which predicts a functional form for the autocorrelation function

$$\xi(r) = \frac{2}{\ln(1+2\phi^2)}\int_0^\infty e^{-l_{\rm corr}^2 a^2/\phi^2}\left(1 - e^{-2l_{\rm corr}^2 a^2}\right)\frac{J_0(ar)}{a}\,da, \tag{5}$$

where $J_0$ is the zeroth-order Bessel function, $l_{\rm corr}$ is the correlation length describing the large-scale structure of the autocorrelation function, and $\phi = l_{\rm corr}/x_0$ is the ratio of the correlation length to $x_0$, the effective size of the region of which elements are initially deposited by stellar feedback.

For each element, source, and snapshot we fit $x_0$ and $l_{\rm corr}$ by comparing our measured autocorrelations to the predictions of equation 5 using the EMCEE MCMC sampler (Foreman-Mackey et al. 2013). We adopt uniform priors $0 < x_0/{\rm pc} < 10^4$, $0 < l_{\rm corr}/{\rm pc} < 10^5$, and take our log likelihood function to be $\log\mathcal{L} = -\sum_n [\xi_{a,a}(r_n) - \xi_{\rm mod}(r_n)]^2/2\sigma_n^2$, where $\xi_{a,a}(r_n)$ is the autocorrelation we measure at separation $r_n$ from the map, $\xi_{\rm mod}(r)$ is evaluated from equation 5, $\sigma_n$ is the variance for bin $n$ that we determine via our jackknife procedure, and the sum runs over all bins up to the first zero-crossing of $\xi_{a,a}(r_n)$. These parts of our fitting procedure are identical to that in Paper I, so we refer readers to that paper for further details. To obtain our final posterior distributions, we run EMCEE using 100 walkers with 1000 steps, discarding the first 100 steps for burn-in and thinning the remaining parts of the chains by a factor of 5.

### 3.3 Time-correlation

To measure how rapidly metal structures decorrelate in time, we compute a time correlation, compensated for the effects of differential rotation. To compute this quantity, we first tabulate the rotation curve $v_c(r)$ from the simulation by computing the mean azimuthal circular velocity in 25 pc radial bins; we take the mean over all snapshots for this purpose, but the variance from one snapshot to another is very small. Then in order to compute the correlation between the metal fluctuation fields $m'_a(\mathbf{r}, t)$ and $m'_a(\mathbf{r}, t+\tau)$ separated by a time lag $\tau$, we shift the positions of all particles in the later snapshot to remove the effects of galactic rotation; in practice we accomplish this by converting the particles' projected positions $(x, y)$ in the galactic plane to a polar coordinate system $(r, \theta)$ with its origin at the galactic centre, changing the $\theta$ coordinate to $\theta' = \theta - v_c(r)\tau/r$, and then converting the resulting $(r, \theta')$ coordinates back to Cartesian positions. We then compute a de-rotated metal fluctuation field $\widetilde{m}'_a(\mathbf{r}, t+\tau)$ from the updated positions, using the same procedure as described in Section 3.1.

Derotation lets us ask whether a pattern changed, rather than merely rotated past us.

We now define the (zero-lag-in-space) time correlation for element $a$ as

$$\xi_a(\tau) = \frac{\langle m'_a(\mathbf{r}, t)\, \widetilde{m}'_a(\mathbf{r}, t+\tau)\rangle_{\mathbf{r}}}{\sigma_a(t)\,\sigma_a(t+\tau)}, \tag{6}$$

where $\sigma_a(t)$ is the total variance of the metal fluctuation map $m'_a$ at time $t$. In practice we compute this quantity from our discrete, pixelised maps as

$$\xi_a(\tau) = \frac{\sum_i m'_{a,i}(t)\widetilde{m}'_{a,i}(t+\tau)}{\sqrt{\sum_i [m'_{a,i}(t)\widetilde{m}'_{a,i}(t+\tau)]^2}}, \tag{7}$$

where $m_{a,i}$ is the value the metal fluctuation map in pixel $i$, and the sums run over all non-masked pixels. While in principle we can compute this quantity for any pair of snapshots, the number of possible snapshot pairs is very large ($\approx 2\times10^5$), and we wish to omit both early snapshots (because, as we show below, the simulations have not yet settled to steady-state at early times) and snapshots where the lag is so large that our de-rotation procedure becomes questionable. We therefore limit ourselves to considering comparisons between a later snapshot from $t = 500 - 600$ Myr at intervals of 10 Myr, and an earlier snapshot that is $1 - 50$ Myr prior to that. This gives us a total of 11 samples at each possible lag from $\tau = 1 - 50$ Myr. We take the

[1] Note that in Paper I we used a 30 kpc × 30 kpc region. However, for the present paper we found that this introduces some noise coming from regions further from the galactic centre where the particle distribution is more poorly sampled. We have therefore restricted our analysis region to be somewhat closer to the galactic centre.

Cross-correlations ask whether two elements occupy the same places in gas or vary together among coeval stars.



time correlation coefficient $\xi_a(\tau)$ to be the mean value of these 11 samples, and the uncertainty to be the variance of these samples.

As with the spatial autocorrelation, we fit the measured time-correlation data to the functional form predicted by KT18,

$$\xi(t) = \frac{\ln\left(1+2\phi^2-\phi^2 t/t_{\rm corr}\right) - \ln\left(1+\phi^2 t/t_{\rm corr}\right)}{\sqrt{\ln\left(1+2\phi^2\right)\ln\left(1+2\phi^2-2\phi^2 t/t_{\rm corr}\right)}}, \tag{8}$$

where $t_{\rm corr}$ is the correlation time to be fit. As with the spatial correlation, we fit this model to the measured time correlations using MCMC sampling. We adopt flat priors in $\log t_{\rm corr}$ and $\log\phi$ for $\phi \in [1, 100]$ and $t_{\rm corr} \in [50, 1.5\times10^4]$ Myr, and a Gaussian log-likelihood function $\log\mathcal{L} = -\sum_n[\xi_a(\tau_n) - \xi_{\rm mod}(\tau_n)]^2/2\sigma_n r$, where the sum is over our $n = 11$ time lag measurements, $\xi_a(\tau_n)$ and $\sigma_n$ are the mean correlation and variance measured at lag $\tau_n$, and $\xi_{\rm mod}(\tau_n)$ is the model prediction given by equation 8 for a given pair $(\phi, t_{\rm corr})$. We gain using EMCEE to conduct MCMC sampling (Foreman-Mackey et al. 2013), using 200 walkers and 1000 steps, with first 100 steps discarded for burn-in and the chains thinned by a factor of 5. The MCMC posteriors show a strong covariance between $\phi^2$ and $t_{\rm corr}$, motivating the use of the derived timescale $T_{\rm corr} \equiv t_{\rm corr}/\phi^2$, which is much less covariant with $\phi$; we evaluate this quantity directly from the MCMC samples to preserve covariance. In what follows we will report both $t_{\rm corr}$ and $T_{\rm corr}$.

### 3.4 Element-element cross-correlation

To measure the degree of similarity between the distributions of different elements, we make use of the cross-correlation. We can define elemental cross-correlations for both the gas and for stars, consistent with the two primary ways of measuring the cross correlation in observations. For gas, we employ the two-point cross-correlation formula from Paper I (their equation 3), at zero lag, which we reproduce here for convenience:

$$\xi_{a,b} = \left(\frac{1}{\sigma_a\sigma_b}\right)\frac{1}{N_n}\sum_i m'_{a,i}m'_{b,i}. \tag{9}$$

Here $m'_{a,i}$ and $m'_{b,i}$ are the metal fluctuation fields for elements $a$ and $b$ in pixel $i$, the sum runs over all $N_n$ pixels, and $\sigma^2_{a,b} = \langle(m'_{a,b})^2\rangle$ are the variances of fields $a$ and $b$ across the whole map.

For stars, we focus on coeval stellar populations, and compute mass-weighted Pearson correlations between elemental abundances in narrow age-selected groups. For each snapshot at time $t$, we bin stars by their formation time (birth epoch) using 1 Myr bins: stars born during $[t_{\rm form}, t_{\rm form} + 1\,{\rm Myr})$ constitute one age cohort. This 1 Myr bin width ensures that all stars within a group formed from ISM gas with nearly identical metallicity distributions, enabling a direct measurement of how element-to-element abundance scatter reflects the local ISM composition at their shared birth epoch. For element pair $(a, b)$ and each age bin, we define the correlation as

$$\rho^\star_{a,b} = \frac{\sum_s w_s\,(X_{a,s}-\mu_a)(X_{b,s}-\mu_b)}{\sqrt{\sum_s w_s\,(X_{a,s}-\mu_a)^2}\,\sqrt{\sum_s w_s\,(X_{b,s}-\mu_b)^2}}, \tag{10}$$

where the sum runs over all stars $s$ in a given age bin, $w_s$ is the mass of star $s$, $X_{a,s}$ and $X_{b,s}$ are the abundances of elements $a$ and $b$ in star $s$, $\mu_a = \left(\sum_s w_s X_{a,s}\right)/\left(\sum_s w_s\right)$ is the weighted mean abundance of element $a$ over all stars in the bin, and similarly for element $b$. We compute $\rho^\star_{a,b}$ for each 1 Myr age bin across the full simulation epoch ($t = 600$ Myr), yielding time series of cell correlations that reflects how the ISM's nucleosynthetic imprint evolves and is preserved in successive stellar generations.

We compute stellar cross-correlations in 1 Myr age bins across the full 1–600 Myr simulation.

## 4 RESULTS

We first present a qualitative overview of the simulation outcome in Section 4.1. We then present results for the spatial and temporal correlations of elements in Section 4.2 and Section 4.3, and finally examine elemental cross-correlations in Section 4.4.

### 4.1 Qualitative overview

Figure 1 shows the state of the simulation at $t = 600$ Myr. The top row displays hydrogen column density $N_{\rm H}$ (left), total $^{56}$Fe abundance (i.e., $^{56}$Fe contributed by all nucleosynthetic channels; middle), and total $^{12}$C abundance (right). The bottom row shows stellar surface density (left) and gas metallicity fluctuation maps, $m'$, for the same four species: total Fe (middle) and C (right). The gas distribution exhibits vigorous star-forming activity with peak column densities exceeding $\sim 10^{23}$ cm$^{-2}$ in star-forming regions. On large scales, metallicity exhibits a strong correlation with gas distribution, such that regions of low gas density often display higher metallicity. This arises because the cavities blown out by supernova explosions are invariably more metal-rich. Metal abundance maps also reveal small-scale chemical heterogeneities superimposed on these large-scale structures, reflecting the discrete injection events and subsequent turbulent mixing that characterize the ISM enrichment process. The stellar disc traces the gas spiral arms with a time delay corresponding to stellar drift. Metallicity fluctuations (panels e and f) quantify departures from radial profiles, revealing chemical inhomogeneities on scales of a few hundred parsecs, the size scale of individual star formation episodes and feedback events that are the subject of our chemical analysis.

Figure 1 illustrates how spatial distributions for the contributions from different nucleosynthetic sources vary, using $^{56}$Fe as an example. In this figure, we use the red, green, and blue channels to represent $^{56}$Fe produced by SNIa, AGB stars (both young and old), and SNII and WR stars, respectively. At early times ($t = 50$ Myr), discrete supernova-driven bubbles are clearly visible as blue (SNII+WR) and red (SNIa) structures, their sizes and morphologies reflecting the local gas density – denser nuclear regions host more compact and numerous bubbles, while lower-density outer regions show larger, more extended cavities. Individual events imprint distinct signatures: the early Type II and Type Ia explosions carve out structures that persist as recognizable features throughout hundreds Myr evolution.

With advancing time, turbulent mixing progressively blends the three colour channels, and by $t \approx 400$ Myr the colour distribution seems to have settled into a rough statistical steady-state in which the large-scale metal distribution has become relatively homogeneous. However, even at the final epoch ($t = 600$ Myr), the RGB decomposition retains significant structure at small scales, revealing that chemical differentiation persists in patches corresponding to recent star-forming regions and their feedback-driven bubbles. This persistence of such structures is the natural consequence of continuous nucleosynthetic injection: as older metal-enriched regions mix and homogenize, newly-injected metals from ongoing star formation create fresh spatial gradients and correlations.

A critical observation we can already make from Figure 1 and Figure 2 is the radial dependence of the efficiency with which gas is mixed in our simulation. Within the inner region ($r \lesssim 5$ kpc), the source-decomposed colour distinctions are less pronounced than in the outer disc, indicating more efficient local mixing in the simulated gas. In contrast, the outer disc retains more pronounced source-to-source variation, reflecting both lower turbulent activity and the geometrically larger separation of injection sites relative to the dif-

high cross-correlation

low

birth-time slice

One-megayear stellar cohorts isolate chemical inheritance before Galactic history mixes different generations together.

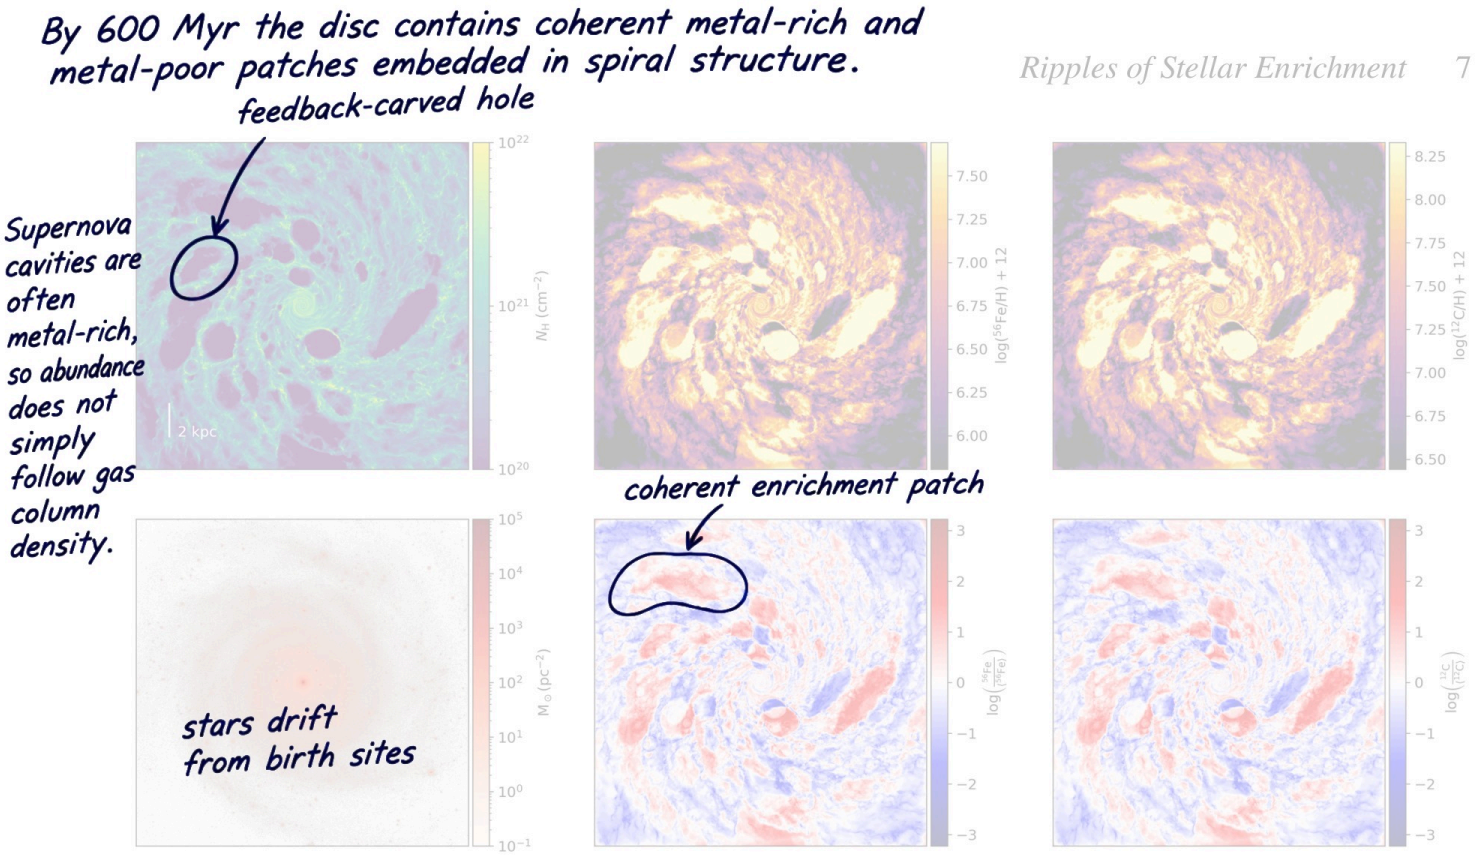



**Figure 1.** Overview of the RoSE simulation at $t = 600$ Myr. Top row: (a) hydrogen column density $N_{\rm H}$ (cm$^{-2}$); (b) $^{56}$Fe abundance $\log(^{56}{\rm Fe/H}) + 12$; (c) $^{12}$C abundance $\log(^{12}{\rm C/H}) + 12$. Bottom row: (d) stellar surface density $\Sigma_\star$ (M$_\odot$ pc$^{-2}$); (e) $^{56}$Fe fluctuation field $m'(^{56}{\rm Fe})$; (f) $^{12}$C fluctuation field $m'(^{12}{\rm C})$. Each panel covers a 20 kpc × 20 kpc field centred on the galaxy; a 2 kpc scale bar is shown in panel (a). Note that all abundances shown are measured in the gas rather than in the stars.

The observable signal is the residual patchwork, not the much larger radial trend we removed.

fusion length scale. This radial gradient in chemical homogeneity has important consequences for interpreting abundance patterns in different galactic environments.

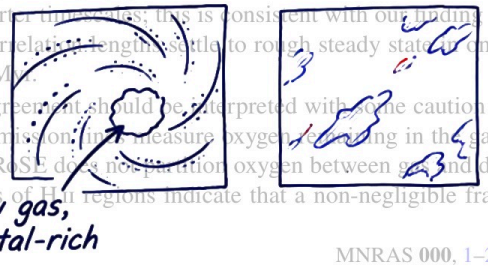
Finally, Figure 3 shows the fractional contribution of each modelled nucleosynthetic channel to each of the nine isotopes we track as a function of time in the simulation. As expected, the composition of each element reflects the nucleosynthetic pathways encoded in the stellar yields: WR winds produce most C, SNII dominates the $\alpha$ elements (O, Mg, S), SNIa contributes the majority of iron-peak elements (Fe), AGB stars contribute significantly to neutron-capture elements (Ba, Ce) and moderately to carbon and nitrogen, and the rare neutron star merger channel leaves an imprint primarily on the $r$-process element Eu, but is negligible for all other elements. However, we can also see an important artefact over the first ≈ 50 Myr of the simulation: a transient enhancement of WR and SNII source fractions. This feature reflects our progressively-increasing resolution initialisation procedure as described in Section 2.1: when we increase the resolution, this leads to a temporary burst of star formation, which in turn leads to a transient injection of SNII and WR elements a short time later. We will therefore generally ignore the first ≈ 100 Myr of the simulation in the analyses the follow.

### 4.2 Spatial scales and elemental clustering

We next quantify the spatial structure and its time evolution via the autocorrelation fits described in Section 3.2. In Figure 4, we plot our parametric fits to the injection scale $x_0$ and correlation length $l_{\rm corr}$ for all elements versus time, summing over all nucleosynthetic sources. We see that, after a transient period in the first ≈ 200 Myr of the simulation, the injection scale fluctuates around tens of parsecs and the correlation length oscillates around 1–3 kpc for all elements. The transient spikes visible in these plots correspond to periods of upward fluctuations in the star formation rate (compare Figure 4 to the upper panel in Figure 3), during which newly-injected supernova-enriched material temporarily compresses the spatial scale of the metal distribution.

The correlation lengths $l_{\rm corr}$ and injection scales $x_0$ that we find for $^{16}$O, the element whose spatial structure is most easily measurable in the gas phase, are in good agreement with the results of extragalactic studies. For example, studies of large samples of galaxies ($\gtrsim 100$) with star formation rates $\dot{M}_* \sim 1$ M$_\odot$ yr$^{-1}$, comparable to our simulated galaxy, are observed to have $l_{\rm corr} \sim 1$ kpc (e.g., Kreckel et al. 2019, 2020; Metha et al. 2021; Li et al. 2021; Li et al. 2022; Williams et al. 2022), and higher-resolution follow-ups that have measured $x_0$ typically find values ∼ 50 pc (e.g., Li et al. 2022), consistent with the results shown in Figure 4. More subtly, by comparing metallicity correlations in galaxies with and without active nuclei, Li et al. (2025a) conclude that correlation lengths are not primarily determined by galaxy assembly histories or other processes acting over cosmological timescales, but instead reach an equilibrium determined primarily by galaxies' instantaneous star formation rates on much shorter timescales; this is consistent with our finding that that galaxy correlation lengths settle to rough steady states on only a few hundred M

This agreement should be interpreted with some caution because nebular emission lines measure oxygen abundance in the gas phase, whereas RoSE does not partition oxygen between gas and dust. Observations of H II regions indicate that a non-negligible fraction of



Different iron sources begin as distinct colored bubbles, then mix toward white while fresh local differences keep appearing.

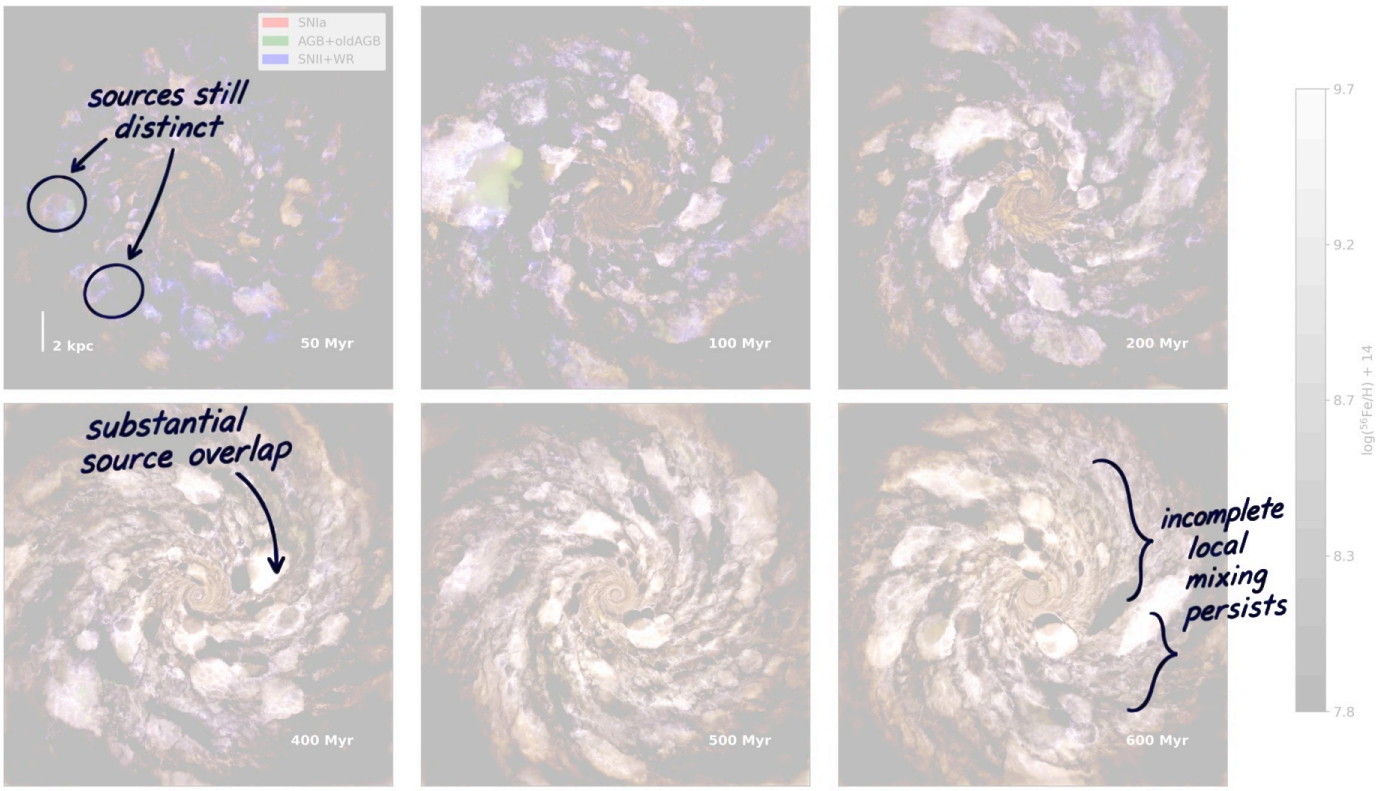
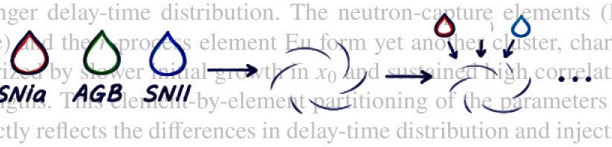



**Figure 2.** Maps of $^{56}$Fe abundance in interstellar gas decomposed into three colour channels corresponding to different nucleosynthetic origin sites: SNIa (red), AGB stars (green, including both short- and long-timescale AGB contributions), and SNII+WR winds (blue). Each channel's intensity encodes its fractional abundance relative to hydrogen, as indicated in the colourbar. Each panel is 20 kpc × 20 kpc in size, and shows the simulation state at times from 50 to 600 Myr as indicated by the labels; the colour scale is the same in all six panels, and is set so the range from zero to maximum intensity corresponds to the 10–90th percentiles of the total $^{56}$Fe/H field at $t = 600$ Myr.

oxygen is depleted onto grains and that the depletion increases with metallicity (Peimbert & Peimbert 2010). A spatially uniform depletion would shift the abundance zero-point and leave our inferred $l_{\rm corr}$ and $x_0$ unchanged, but variations with metallicity or local ISM conditions could introduce or suppress spatial fluctuations in gas-phase oxygen and thereby bias the inferred $l_{\rm corr}$ and $x_0$ relative to the total-oxygen values predicted by RoSE. We return to this limitation in Section 5.4.

White means source contributions overlap; it does not mean the iron abundance is spatially uniform.

If we now examine the relationship between elements rather than focusing on $^{16}$O alone, a striking feature of Figure 4 is the clear *elemental clustering* in parameter space. Carbon ($^{12}$C), primarily sourced from WR winds, consistently exhibits the lowest injection scale and highest correlation length throughout the simulation. In contrast, the core-collapse products ($^{16}$O, $^{24}$Mg, $^{32}$S from SNII) show intermediate values that vary nearly identically in time. Iron ($^{56}$Fe), dominated by the delayed SNIa channel, traces a distinct path offset to larger injection scale and shorter correlation length, reflecting its longer delay-time distribution. The neutron-capture elements (Ba, Ce) and the $\alpha$-process element Eu form yet another cluster, characterised by slower initial growth in $x_0$ and sustained high correlation lengths. This element-by-element partitioning of the parameters directly reflects the differences in delay-time distribution and injection environment between channels.




To confirm that the differences between elements are indeed driven mainly by their different origin sites, in Figure 5 we disentangle the contributions of individual nucleosynthetic channels to the spatial structure of $^{56}$Fe, presenting parametric fits to the part of the $^{56}$Fe distribution contributed by each of the five nucleosynthetic channels that make a non-zero contribution to it (all except NSM), alongside the total (multi) field. It is clear from this figure that, once the galaxy settles into near steady-state after a few hundred Myr, the parameters for the *total* field of a given element (grey lines) closely track those of its dominant nucleosynthetic source (SNIa; beige line). While the figure shows results of $^{56}$Fe only, we find the same qualitative trend for other isotopes. Thus we conclude that the spatial statistics of the chemical enrichment in galaxies is predominantly determined by the spatial distribution of the dominant channel.

The outer disc retains stronger color separation because injection sites are farther apart relative to the mixing scale.

### 4.3 Temporal correlations

We next examine the temporal correlations of the metal field, which we compute following the procedure outlined in Section 3.3. We show the measured correlation as a function of time lag for each element in Figure 6, together with the corresponding best-fitting KT18 model (equation 8).[2] It is clear that the functional forms predicted by the

[2] Careful readers may notice some difference from the results in Paper I, even for elements that were included in that paper. This was due to a coding error in Paper I, which was discovered in the course of preparing this work, whereby we computed the time correlation on the non-logged rather than the logged metallicity fluctuation field. That error has been corrected in the present paper.

Element tracks cluster by dominant source, with injection widths of tens to hundreds of parsecs and kiloparsec-scale correlations.

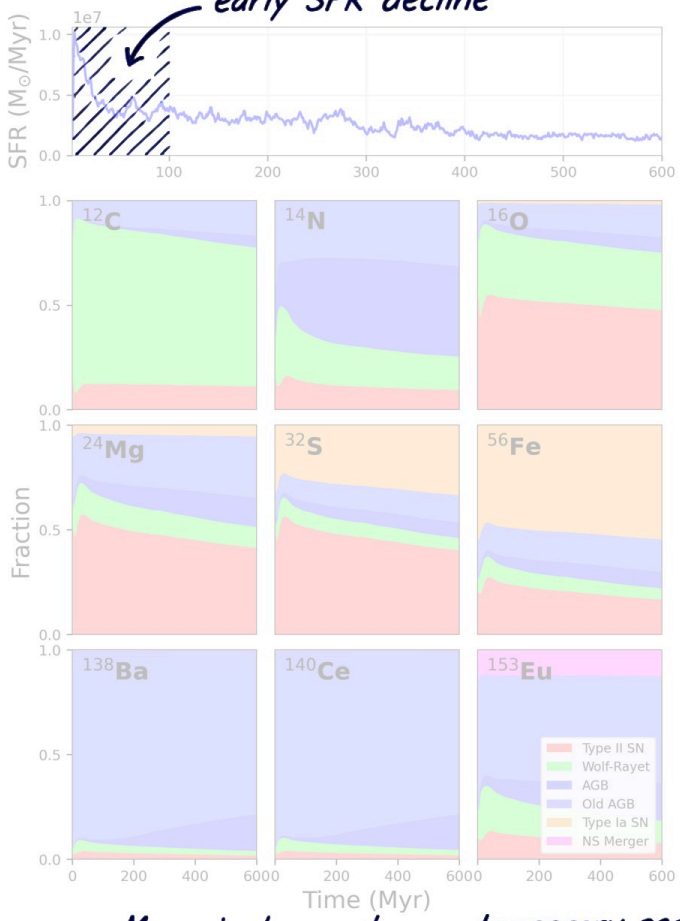
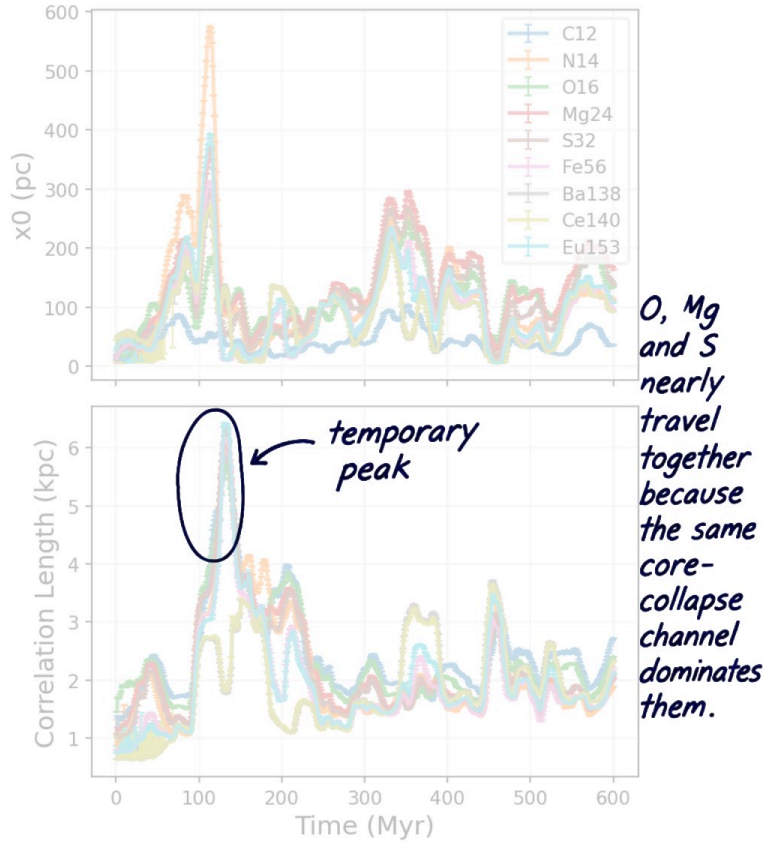



**Figure 3.** The nine lower panels show the fractional cumulative contribution of each nucleosynthetic channel to each of nine tracked isotopes as a function of time; for comparison the top panel shows the star formation rate (SFR) over the same time period. The stacked horizontal bars in each of the nine element panels show the fraction of the mass of the indicated element that was produced up to that time by each of the six channels indicated in the legend. Note that as the SFR approaches equilibrium at later times, the elemental fractions also stabilize, reflecting the steady-state balance between different nucleosynthetic channels.




O, Mg and S nearly travel together because the same core-collapse channel dominates them.

Many isotopes show a temporary peak near 120 Myr in both Figure 4 panels.

**Figure 4.** Autocorrelation parameters versus time evolution for all isotopes, computed using the total metal field fluctuation field summed over all nucleosynthetic channels. The top panel shows the injection width $x_0$, the bottom shows the correlation length $l_{\rm corr}$. Each coloured curve corresponds to one isotope, as indicated in the legend. Points show the 50th percentile value and error bars indicate the 5–95 percentile ranges from the MCMC fits (see Section 3.2).

models provide a good description of the data. We report our fit parameters and their uncertainties in Table 2.

Examining the table and figure, we again see a trend between temporal correlation and nucleosynthetic origin site. The elements whose production is dominated by SNII – O, Mg, and S – cluster tightly together and have the longest correlation time. The elements produced by SNIa, AGB stars, and NS mergers – N, Fe, Ba, Ce, and Eu – form a clearly distinct group with a shorter correlation time than the SNII elements. Finally, the one element whose production is largely via WR stars – C – stands out as quite distinct from all other elements in that its correlation drops rapidly in time, but then flattens considerably.

Figure 7 confirms that these differences can indeed be attributed to nucleosynthetic origin site. In analogy to Figure 5, in this figure we show the temporal correlation calculated separately for each of the nucleosynthetic sources that produce Fe; the results for other elements are again similar. We report KT18 model fit parameters in Table 3. In this figure and table the differences in correlation time for elements produced at different astrophysical sites becomes extremely clear. Production by SNII, which tends to produce large, long-lived bubbles of enrichment, yields the longest time correlation, while NSMs, which are rare events that do not produce large bubbles because they inject limited energy, decorrelate extremely rapidly. Other sources are intermediate.

The quantitative agreement between spatial and temporal KT parameters reinforces the self-consistency of the diffusion-injection framework: elements with larger spatial correlation lengths (e.g., Ba, Ce, Eu) also exhibit longer temporal decorrelation times, while elements with shorter spatial scales (O, Mg, S) decorrelate more rapidly in time. This consistency validates the physical picture wherein nucleosynthetic delay times set both the characteristic injection scale (through the spatial extent of parent stellar populations) and the mixing timescale (through the interval between successive injections at a given location).

### 4.4 Elemental cross-correlations

Our final analysis examines the cross-correlations between different elements in both the gas and stellar phases, computed as described in Section 3.4. Figure 8 shows the former; in this diagram, the numerical value and colour for each element-element pair show the mean cross-correlation averaged over the last 100 Myr of the simulation,

An element's dominant production channel controls both its spatial structure and how quickly that structure forgets itself.

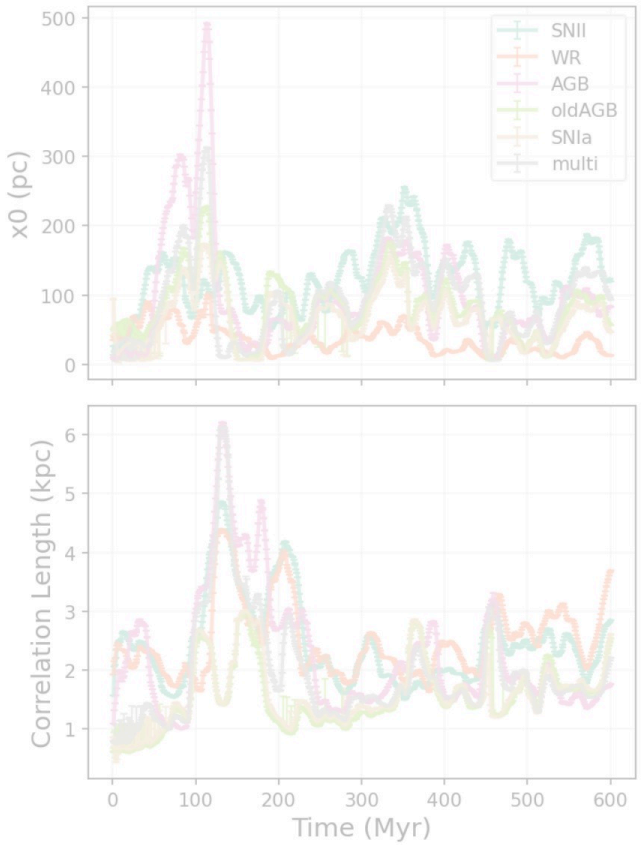


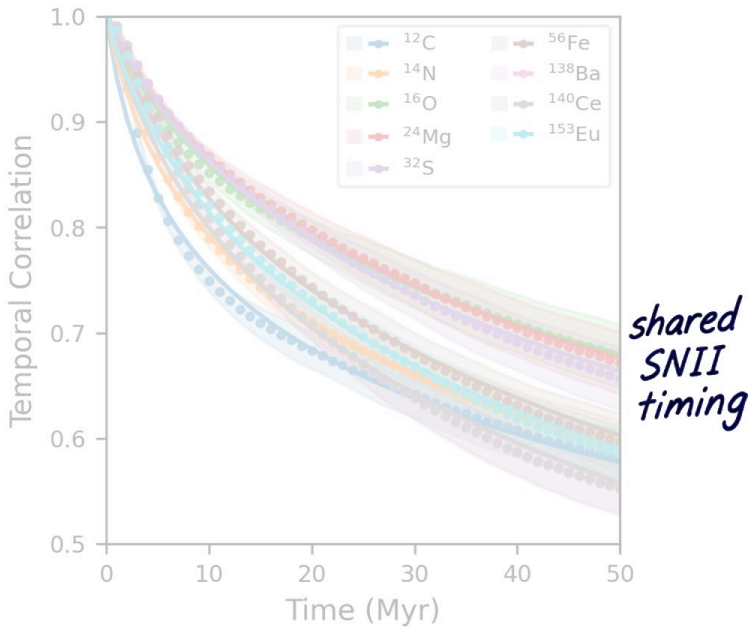

**Figure 5.** Same as Figure 4, but now showing parameters for fits to just the $^{56}$Fe fluctuation field decomposed by nucleosynthetic source – SNII, WR, AGB, oldAGB, SNIa – as indicated in the legend. We also show the results for the total $^{56}$Fe field (labelled "multi") for comparison; this curve is the same as the $^{56}$Fe shown in Figure 4.

**Table 2.** KT18 model fit parameters from time-correlation analysis for all tracked isotopes using the total (multi-source) metal field. $\phi$ is the dimensionless shape parameter; $t_{\rm corr}$ (Myr) is the correlation time; $T_{\rm corr} \equiv t_{\rm corr}/\phi^2$ (Myr) is the characteristic temporal decorrelation timescale. Uncertainties are 16–84 percentile ranges from the MCMC posterior.

| Isotope | $\phi$ | $t_{\rm corr}$ (Gyr) |
|---|---|---|
| $^{12}$C | $41^{+33}_{-22}$ | $2.8^{+3.2}_{-1.7}$ |
| $^{14}$N | $25^{+32}_{-14}$ | $1.8^{+3.6}_{-1.2}$ |
| $^{16}$O | $24^{+24}_{-14}$ | $3.5^{+5.8}_{-2.4}$ |
| $^{24}$Mg | $15^{+22}_{-9}$ | $1.8^{+4.7}_{-1.3}$ |
| $^{32}$S | $13^{+22}_{-8}$ | $1.5^{+4.1}_{-1.0}$ |
| $^{56}$Fe | $12^{+21}_{-7}$ | $8.4^{+2.4}_{-0.5}$ |
| $^{138}$Ba | $11^{+17}_{-6}$ | $0.6^{+1.3}_{-0.3}$ |
| $^{140}$Ce | $11^{+18}_{-6}$ | $0.6^{+1.4}_{-0.4}$ |
| $^{153}$Eu | $13^{+23}_{-8}$ | $0.9^{+2.3}_{-0.6}$ |

while white lines within each panel show the cross correlations as a function of time, with the horizontal axis running over the full 600 Myr simulation duration. The vertical axis goes from 0 to 1 in correlation strength. Some elements (e.g., C, O, Mg, S – all from SNII) show very strong correlations, while pairs mixing sources (e.g., WR-dominated C with AGB-dominated Ba, Ce) exhibit smaller correlations. This is a direct signature of the dif-

Decomposing one isotope by source is the cleanest test: chemistry is held fixed while the injection physics changes.




shared SNII timing

**Figure 6.** Data points and shaded bands show the measured mean temporal correlation and its variance as a function of time lag, with values computed via the procedure outlined in Section 3.3. Solid lines show best-fitting KT18 models, with fit parameters shown in Table 2.

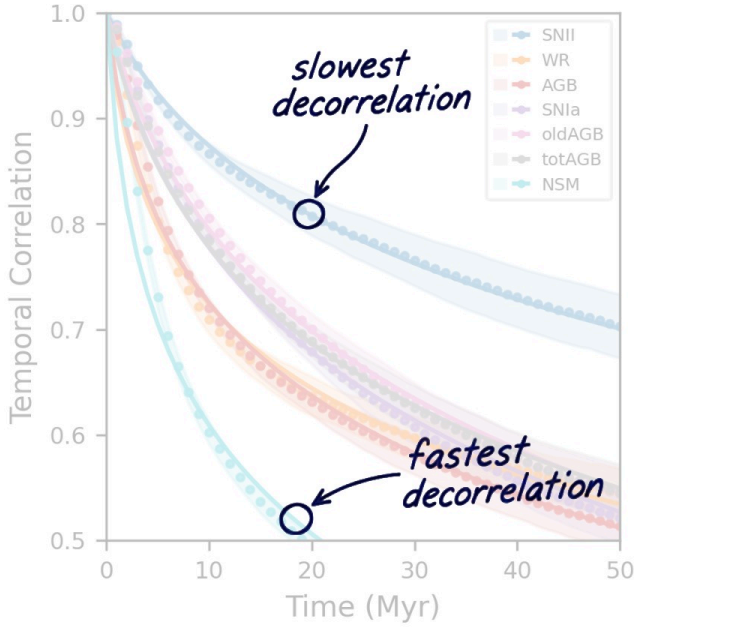



**Figure 7.** Same as Figure 6, but now showing results for Fe only, decomposed by the contribution to Fe from different nucleosynthetic origin sites. Parameters for the best-fitting KT18 models (solid lines) are given in Table 3.

Rare neutron-star mergers decorrelate fastest; large core-collapse-supernova bubbles leave the longest-lived patterns.

Same-source elements remain tightly correlated, while differing delay times tend to weaken gas and stellar correlations.

**Table 3.** Same as Table 2, but now showing the results of KT18 model fits to the temporal correlation of the $^{56}$Fe metal field decomposed into contributions from individual nucleosynthetic sites.

| Component | $\phi$ | $t_{\rm corr}$ (Gyr) |
|---|---|---|
| SNII | $23^{+21}_{-13}$ | $3.8^{+5.9}_{-2.7}$ |
| WR | $41^{+33}_{-21}$ | $2.0^{+2.1}_{-1.1}$ |
| AGB | $24^{+29}_{-13}$ | $1.0^{+1.5}_{-0.6}$ |
| SNIa | $11^{+17}_{-6}$ | $0.5^{+1.0}_{-0.3}$ |
| oldAGB | $10^{+16}_{-5}$ | $0.5^{+1.1}_{-0.3}$ |
| totAGB | $15^{+22}_{-8}$ | $0.8^{+1.6}_{-0.4}$ |
| NSM | $18^{+21}_{-8}$ | $0.2^{+0.3}_{-0.1}$ |

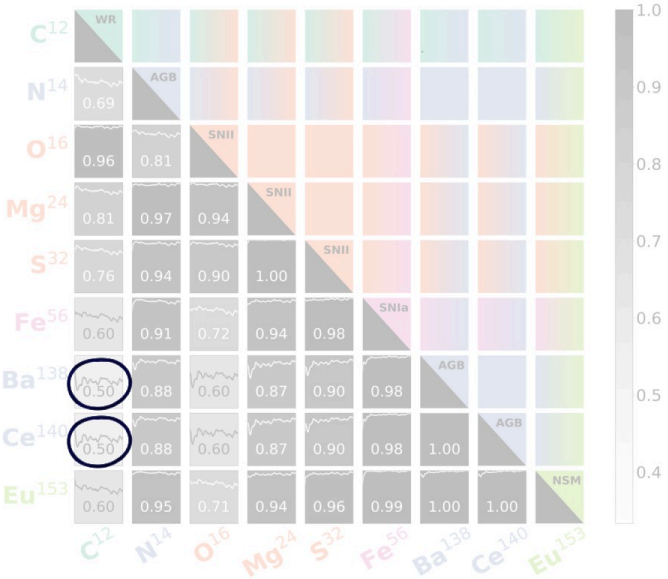



**Figure 8.** Gas-phase cross-correlation matrix for all nine tracked isotopes spanning the full 1–600 Myr simulation. Each cell $(i, j)$ shows the time evolution of the two-point cross-correlation coefficient for element pair $i$ versus $j$ (plotted as a white curve), with the background color (red colormap) and numerical value indicating the time-averaged correlation strength over the final 100 Myr quasi-equilibrium epoch. The lower-triangle layout shows all unique pairs; the diagonal is unity (self-correlation). To guide the eye, the shaded upper-triangle highlights the dominant production channel of each element (except Eu as it is the only element has significant NSM contribution) with the colours indicating the different channels (see legend along the diagonal). Note that elements sharing the same dominant source (e.g., O, Mg, S from SNII) show consistently high correlation; pairs mixing sources exhibit reduced correlation reflecting their distinct delay-time distributions.

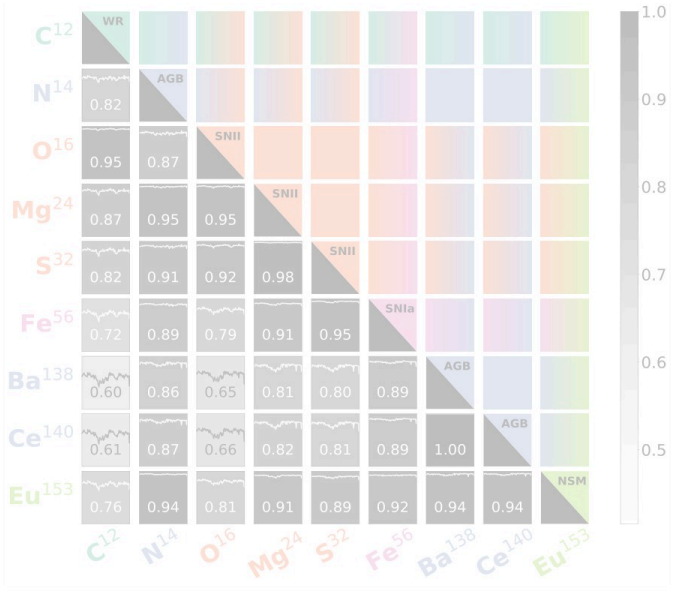



**Figure 9.** Same as Figure 8, but now showing the cross-correlations in stars. The white lines showing time evolution measure the cross-correlation in stars born at a given time. The correlations are weighted by stellar particle mass following equation 10, and the curves and reported means cover birth epochs $t$ = 501–600 Myr.

Mixed-source pairs are still correlated; their ordering retains information about relative delay times.

SNII family: near-unity

widely separated source histories

ference in injection sites: the two distinct nucleosynthetic sources can still [illegible] correlated, but are progressively erased over timescales of the correlation time. The final level of correlation is set by the competition between ongoing mixing and erasure by [illegible] Krumholz et al. (2025), we find that for element pairs from mixed nucleosynthetic sources, the degree of correlation appears to be related to how far apart the two nucleosynthetic channels are in their delay times. The example of C, which is returned on the shortest timescales due to the dominance of WR stars in its production, is illustrative: C correlates best with the SNII elements (O, Mg, S), which have the next shortest delay time, followed by N from short-lived AGB stars, then Fe from SNIa, and finally the elements produced primarily by old AGB stars and NSM elements (Ba, Ce, Eu), which have the longest delay times. Thus the degree of correlation between two elements appears to be closely linked to how different their delay times are between star formation and element return.

Stars preserve the gas composition at birth, even after leaving their natal cloud.

same family structure survives in stars

Figure 9 presents the corresponding stellar cross-correlations. Each time point as indicated by the white lines in this [illegible] cross-correlations for stars born at the same 1 Myr cohort; for example, the correlation at $t$ = 550 Myr is computed from stars born from 550–551 Myr into the simulation. We see that the overall correlation structure is preserved in the gas phase, demonstrating that the nucleosynthetic source imprints present in the ISM at any given moment are faithfully recorded in the stellar populations formed during that epoch and persist as permanent chemical signatures. However, quantitative differences emerge when comparing gas and stellar correlation strength element-pair-by-element-pair: prompt-source pairs (e.g., C–O from WR+SNII) show stronger correlation in stars than in gas, because stars formed early in the simulation inherited the ISM's high correlation established during the initial rapid enrichment phase, and this signature is preserved indefinitely; conversely, pairs involving delayed sources (e.g., Fe correlations with alpha elements) exhibit higher gas-phase values, reflecting the ongoing SNIa enrichment that continues to modify the ISM but has not yet fully imprinted the cumulative stellar disc. This phase-dependent variation encodes information about the nucleosynthetic progenitor timescales and the time-integrated enrichment history. The stabilisation timescale for stellar cross-correlations is also considerably shorter than decorrelation timescales measured in Section 4.3. This rapid convergence arises because each stellar cohort samples only a narrow time slice of the ISM, avoiding the temporal averaging that broadens the gas-phase correlation timescales; however, the finite delay between metal ejection, gas cooling, and star formation introduces a lag, so stellar cross-correlations reach equilibrium slightly later than their gas counterparts.

As shown in Figure 10, these quantitative comparisons between

A six-number description of how two elements differ in source mixture predicts most of their correlation.

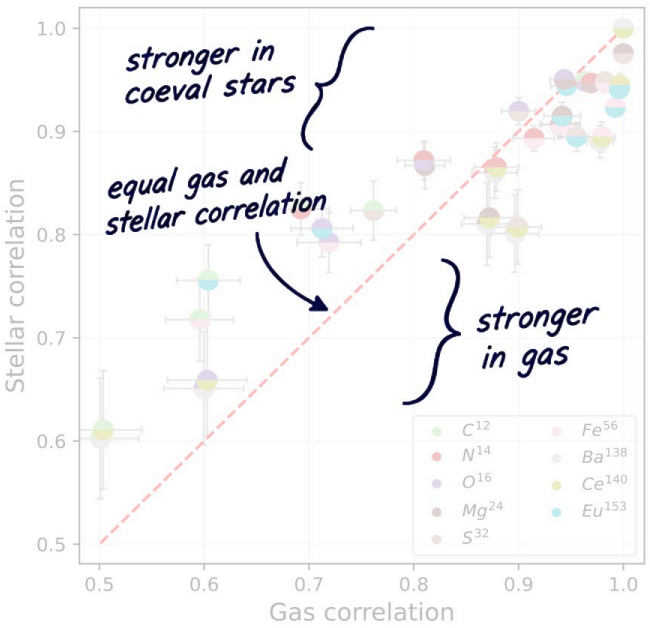



**Figure 10.** Direct comparison of gas and stellar cross-correlations across element pairs. Each point represents an element pair (colored half-circles indicate the two constituent elements), computed from gas metallicity fluctuations over the time interval $t = 501 - 600$ Myr ($x$-axis) and stellar age bins $0 - 100$ Myr at the final snapshot ($y$-axis). Error bars denote standard deviations across individual time steps (gas) and age bins (stellar). The red dashed line marks $y = x$ as a reference.

gas and stellar phase correlations reveal a fundamental principle: element pairs with prompt nucleosynthetic sources show stronger correlations in stars than in gas, because stellar populations preserve the high correlations established during rapid early enrichment; conversely, pairs involving delayed sources show higher gas correlations, reflecting the ongoing enrichment that continues to modify the ISM but has not yet fully imprinted the stellar populations. This phase-dependent behavior directly encodes information about nucleosynthetic timescales and enrichment histories, motivating a deeper mechanistic analysis in the Discussion section.




## 5 DISCUSSION

Our results demonstrate that numerous aspects of elemental distributions within galaxies – spatial structure, correlation in time, and element-element cross-correlation in both gas and stars – appear to be directly traceable to those elements diverse nucleosynthetic origin sites. In this discussion we seek to understand this behaviour in more detail.

### 5.1 A simple linear model for cross-correlation

As a first step, we investigate to what extent a very simple linear model can predict how correlated two elements will be solely based on the amounts by which different nucleosynthetic channels contribute to their production. To this end for each element pair $(e_1, e_2)$ and snapshot, we define

$$\Delta_k(e_1, e_2) = |f_k(e_1) - f_k(e_2)|\,, \tag{11}$$

where $f_k(e)$ is the fractional contribution of nucleosynthetic channel $k \in \{\mathrm{SNII}, \mathrm{WR}, \mathrm{AGB}, \mathrm{SNIa}, \mathrm{NSM}, \mathrm{oldAGB}\}$ to element $e$'s total abundance at that time, i.e., the quantity shown in Figure 3. Thus at each time and for each element pair we obtain a 6-dimensional vector $\mathbf{\Delta} = [\Delta_{\mathrm{SNII}}, \Delta_{\mathrm{WR}}, \Delta_{\mathrm{AGB}}, \Delta_{\mathrm{SNIa}}, \Delta_{\mathrm{NSM}}, \Delta_{\mathrm{oldAGB}}]$ that describes how different two elements are in their nucleosynthetic origins; element pairs with very similar origins will have all elements of $\mathbf{\Delta}$ close to zero, while for two elements with completely disjoint nucleosynthetic origins the entries in $\mathbf{\Delta}$ will sum to two. We calculate $\mathbf{\Delta}$ for times $t = 501$–$600$ Myr at 1 Myr intervals.

To test how well $\mathbf{\Delta}$ can predict element-element cross correlations, we now consider a simple model of the form

$$\xi^{\mathrm{pred}}_{e_1,e_2} = \mathbf{c}^{\mathrm{T}}\mathbf{\Delta} + b, \tag{12}$$

The model knows no maps, orbits or temperatures — only the absolute source-fraction differences.

where $\xi^{\mathrm{pred}}_{e_1,e_2}$ is the predicted cross-correlation between elements $e_1$ and $e_2$ (in either the stellar or gas phase), and $\mathbf{c}$ and $b$ are the source-wise regression coefficients and intercept, which we leave as free parameters to be fit. Before proceeding to fit the model, however, we improve the underlying data by introducing three "pure" basis elements, $\mathrm{Pure}_{\mathrm{SNIa}}$, $\mathrm{Pure}_{\mathrm{NSM}}$, and $\mathrm{Pure}_{\mathrm{oldAGB}}$, defined by $f_k(\mathrm{Pure}_j) = \delta_{kj}$ corresponding to the single fluids described in Section 2.2. These synthetic elements are helpful to include because they provide a longer baseline in $\mathbf{\Delta}$, thereby improving fit quality. This is especially important for channels that do not provide a majority contribution to more than one or two of the elements we follow (e.g., SNIa and NSM). For these pure elements, we compute $\mathbf{\Delta}$ and measure the stellar and gas correlations in the simulations exactly as we do for the other nine elements.

We now fit for $\mathbf{c}$ and $b$ over all element pairs (including our newly-introduced "pure" elements) and all times, using as input data to the fit the measured cross-correlations between each element pair at the same times for which we have measured $\mathbf{\Delta}$; we carry out separate fits for gas and stars. We estimate central values for $\mathbf{c}$ and $b$ using ordinary least squares fitting, while to estimate uncertainties we use a bootstrapping method: for each iteration, we independently sample $N$ elements pairs with replacement from the full dataset (where $N$ is the total number of pairs), then fit the constrained regression model equation 12 using a minimum-loss procedure with bounds $c_k \in (-\infty, 0]$. From the resulting 200 coefficient vectors, we extract lower and upper confidence bounds using the 16th and 84th percentiles, corresponding to $\pm 1\sigma$ in a normal distribution. We report the fit coefficients and uncertainties obtained from this procedure in Table 4, and compare the predictions made by our best-fitting models for gas and stars in Figure 11; the model predictions have been clipped to the range $[0, 1]$.

Synthetic pure-source elements extend the training baseline where nature gives us too few clean examples.

This exercise allows us to make a number of interesting observations. First, despite its simplicity, the model is remarkably good at predicting how well different element pairs correlate: the $R^2$ values that characterise the quality of fit are 0.84 and 0.87 for gas-phase and stellar element-element correlations, respectively. Second, the fact that the fit coefficients are all negative encodes the fundamental principle that larger differences in nucleosynthetic origin suppress element correlation: element pairs with small $\mathbf{\Delta}$ yield predictions closer to unity (high correlation), while pairs with large compositional divergence yield smaller correlations. Third, the relative sizes of the coefficients of $\mathbf{c}$ reflects the relative importance of each channel: WR shows the strongest discriminative power (largest $|c_k|$), followed by SNII and NSM, while SNIa contributes weakly. This hierarchy arises because WR and SNII dominate carbon and $\alpha$ elements respectively, which have very distinct spatial and temporal signatures, followed by rare NSM merger events. By contrast both SNIa and AGB (both old and young) channels have smaller effects on element-element correlation,

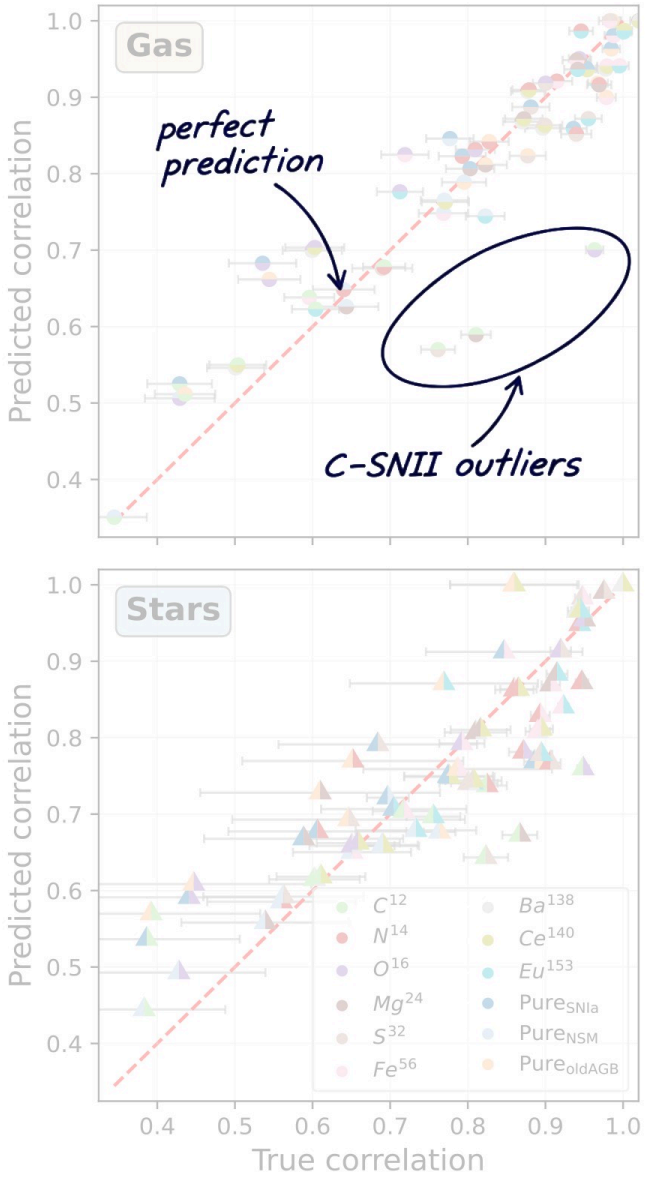



**Figure 11.** These panels show the true element-element correlations measured from the simulations on their horizontal axes in the gas phase (top) and stars (bottom), while the vertical axes of these panels show the correlation predicted by our simple linear regression model (equation 12). Each point corresponds to a particular element pair, with the colours in the half-circle or half-triangle markers indicating the elements being compared (see legend). The horizontal coordinates of the points represent the time average over the period from 501–600 Myr in the simulations, and the horizontal error bars show the standard deviation over this time. The red dashed line is the one-to-one line for comparison.

likely due to the smaller spatial and temporal variations they induce in the distributions of the elements they produce.

Notably, cross-correlations between C and elements produced by SNII are the most significant outliers in both the gas and stellar fits; for these, the model systematically over-predicts the cross-correlation, with a larger bias in the gas phase. No comparable bias is seen for cross-correlations between C and elements not dominated by SNII production. We therefore attribute the failures of the model for C-SNII correlations to the extremely short injection intervals

**Table 4.** Best-fit linear regression coefficients $c_k$ for predicting gas-phase and stellar element-pair cross-correlations from total mass-weighted source-fraction differences. Quoted uncertainties are derived from bootstrapping as discussed in the main text. More negative $c_k$ values imply stronger decorrelation for greater compositional divergence.

*largest absolute coefficient*

| Nucleosynthetic Channel | Decorrelation Coefficient | |
|---|---|---|
| | $c_{k,\mathrm{gas}}$ | $c_{k,\mathrm{star}}$ |
| WR | $-0.77^{+0.07}_{-0.07}$ | $-0.53^{+0.10}_{-0.08}$ |
| SNII | $-0.36^{+0.07}_{-0.06}$ | $-0.40^{+0.09}_{-0.06}$ |
| AGB | $-0.16^{+0.07}_{-0.06}$ | $-0.25^{+0.11}_{-0.10}$ |
| oldAGB | $-0.11^{+0.03}_{-0.03}$ | $-0.19^{+0.05}_{-0.04}$ |
| SNIa | $-0.05^{+0.03}_{-0.02}$ | $-0.15^{+0.04}_{-0.03}$ |
| NSM | $-0.23^{+0.02}_{-0.02}$ | $-0.24^{+0.03}_{-0.02}$ |

between WR and SNII. Carbon injected by WR winds has little time to establish an independent spatial pattern before nearby SNII ejecta overwrite it. This dynamical overwriting can make the carbon yield be governed by stochastic SNII-driven transport rather than by the smoother co-production picture assumed by the model. In a gas-phase two-point statistic, this dynamical overwriting can make the measured C-SNII cross-correlation exceed the model prediction, yielding the observed negative residual; the same tendency is present but weaker for stellar abundances, a phenomenon that we show in the next section is part of a general trend for differences between stellar and gas-phase cross-correlations among elements injected in SNII explosions.

*WR source differences have the largest fitted penalty; SNIa differences the smallest here.*

### 5.2 Differences between elemental cross-correlations in gas and stars

The comparison between gas-phase and stellar-phase model coefficients reveals that, while the linear regression model achieves similarly good fits for both phases ($R^2 = 0.84$ for gas versus 0.73 for stars), the fit coefficients themselves differ substantially, as do the values of the cross-correlations themselves. This suggests that although source-fraction differences drive correlations in both gas and stars, additional physical processes modulate how these source compositional differences manifest. The sensitivity test in Appendix A shows that this overwriting is not specific to the adopted old-star return rate: across $j = 0$–2 the model retains $R^2 = 0.70$–0.86 in gas and 0.71–0.77 in stars.

*WR carbon drizzle* → *SNII blast* → *dynamical overwrite*

To investigate the origin of differences between gas and stellar correlations, we plot the differences for each element pair in the lower left part of Figure 12. This figure reveals a striking pattern: C and O are systematically more highly correlated with other elements in the stellar phase than in the gas phase, while Fe and Eu show the opposite trend, with little difference between stellar and gas correlation for most elements; intermediate elements exhibit characteristics grouped by nucleosynthetic origins. This element-dependent divergence between phases is not a random effect but reflects fundamental differences in how rapidly different enrichment channels inject materials and how quickly the ISM mixes them.

*A nearby SNII blast can rearrange fresh WR carbon before its pattern develops independently.*

The decorrelation of C in the gas phase arises from the rapid, stochastic injection of C is dominated by WR winds and O by SNII, both of which operate on short timescales (few Myr to tens of Myr following star formation). These rapid injection events establish strong spatial correlations in the ISM at early times. However, when these enriched materials are injected, they often go into large-scale supernova-driven bubbles that are initially hot

Stars inherit cold-gas correlations, driving systematic differences from correlations measured across all gas phases.

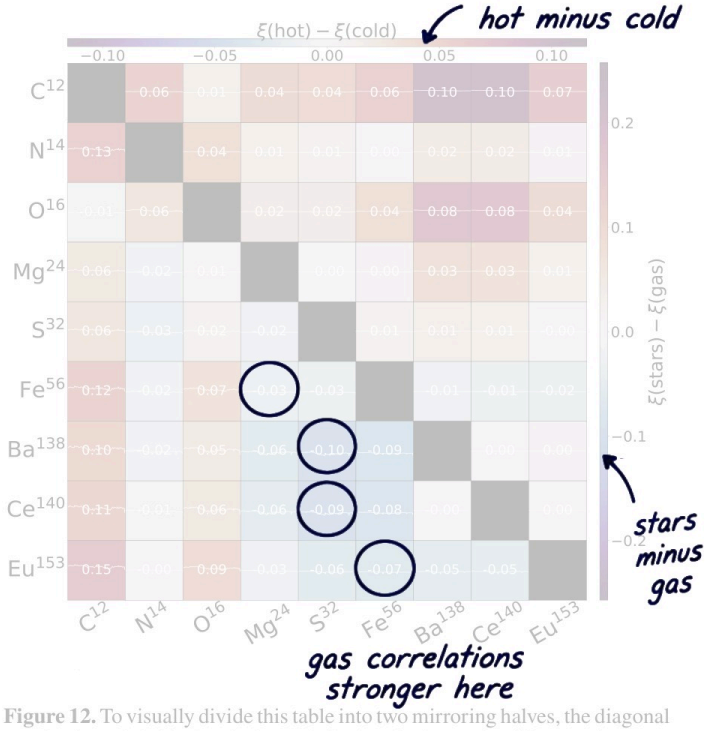
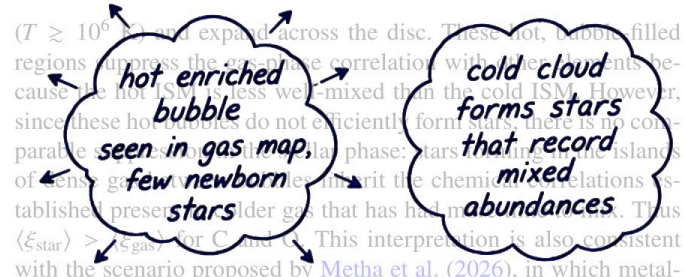



**Figure 12.** To visually divide this table into two mirroring halves, the diagonal row has been left blank. The lower-left triangle shows the difference between the stellar and gas-phase cross-correlations, $\xi_{\rm star} - \xi_{\rm gas}$ for each element pair; the white line in each cell shows the time evolution of the difference, with the horizontal axis running from 501–600 Myr, and the background colour and numerical value give the mean value over the final 100 Myr. The upper-right triangle of cells are identical, except that rather than the difference between stellar and gas-phase cross-correlations show the difference between the gas-phase correlations for hot ($T > 1000$ K) and cold ($T < 1000$ K) gas, $\xi_{\rm hot} - \xi_{\rm cold}$ – see main text for details. Note that the star-gas and hot-cold differences use different colour scales – the horizontal colour bar is for hot-cold, and the vertical one for star-gas – but that both colourmaps are symmetric about zero.




($T \gtrsim 10^6$ K) and expand across the disc. These hot, bubble-filled regions suppress the gas-phase correlation with other elements because the hot ISM is less well-mixed than the cold ISM. However, since these hot bubbles do not efficiently form stars, there is no comparable suppression in the stellar phase: stars continue to form in islands of dense gas that never interact and thus inherit the chemical correlations established present in older gas that has had more time to mix. Thus $\langle \xi_{\rm star} \rangle > \langle \xi_{\rm gas} \rangle$ for C and O. This interpretation is also consistent with the scenario proposed by Metha et al. (2026), in which metal-rich [α/Fe] bubbles produced by core-collapse supernovae dominate short-timescale metallicity fluctuations.

Conversely, Fe and Eu are delayed-source elements dominated by SNIa and NS mergers. These channels are distributed in both space and time, injecting material continuously over the full simulation and throughout the volume of the gas, regardless of whether it is cold and star-forming or hot and non-star-forming. Because they are not associated with the large-scale gaseous structures created by star formation feedback, they show relatively little difference between their correlations in gas and stars. Elements injected by younger AGB stars fall between these two limiting cases, with some mild association with structures established by star formation feedback.

To check the validity of this interpretation, we partition the gas into hot and cold components using a temperature threshold of 1000 K and compute Pearson correlations separately for each phase. Unlike the gas-phase correlations we have discussed thus far, which are computed on projections and thus directly comparable to observations, for this purpose we instead compute the elemental correlations exactly as we do for stars in Section 3.4, i.e., we collect the set of gas particles present at every time snapshot, divide them into hot and cold phases, and then for each phase we compute the cross-correlation between each element pair for the gas particles using equation 10. This quantity is not observable, but it is informative of how the correlation structure of elements differs between hot and cold gas. We show the result in the upper-right triangle of Figure 12). The figure shows that the pattern of differences between hot and cold gas correlations (upper-right triangle) mirrors the pattern of differences between stellar and total gas correlations (lower-left triangle) with striking fidelity. This correspondence confirms that the star-gas divergence is fundamentally a thermal effect: stars preferentially form in the cold ISM phase and thus inherit the chemical correlations present in cold gas, while the hot ISM contributes to the gas-phase spatial correlation but does not form stars. As a result, the elements that are more strongly correlated in stars than in gas are precisely those that are more strongly correlated in hot gas than in cold.

The cross-correlation evolution further supports this picture. Gas-particle internal correlations and stellar correlations converge quickly and remain comparatively stable, whereas gas two-point spatial correlations (see in Figure 8) show much larger temporal fluctuations. This difference indicates that two-point statistics are especially sensitive to large-scale transport and structural reconfiguration driven by feedback (such as the hot bubbles). Consistently, element pairs with large [illegible] in gas two-point correlation, linking thermal selection directly to the amplitude of temporal decorrelation.

The matching triangular patterns are the clue that thermal phase selection drives the star-gas difference.

### 5.3 Comparison with Observations

Our final discussion topic is a comparison of the stellar cross-correlations found in our simulations with those observed in Milky Way stars. For our observational comparison data, we use two data sets. The first is a collection of 560 Solar-type stars from Casali et al. (2020), who derive abundances and isochronal ages from high-resolution HARPS spectra (Mayor et al. 2003) retrieved from the European Southern Observatory (ESO) archive; the elements that overlap between their data and our simulations are C, Mg, S, Fe, Ba, Ce, and Eu. The second is a selection 199,289 red giant stars with abundances for 16 elements from Ness et al. (2026), which are selected using quality cuts on signal to noise, stellar parameters, flags and abundance uncertainty, from the SDSS V Galactic Genesis survey (SDSS Collaboration et al. 2025; Mészáros et al. 2025); the elements in common between this sample and our simulations are O, Mg, S, Fe, and Ce. The ages for each of the stars are from Stone-Martinez et al. (2025). Since carbon and nitrogen are the primary age-dating features in that work, we exclude C and N from this sample to avoid circularity. For both sets of data, we group the stars into age bins 1 Gyr wide and calculate the cross-correlations in each age bin using equation 10. We then take the average cross-correlation over all age bins as our final observational estimates. Because the datasets lack stellar mass measurements, we assign equal weight to every simulated star particle and every observed star for this comparison, rather than using the mass-weighted simulated statistics adopted elsewhere, ensuring that both axes correspond to the same object-weighted sampling.

We compare the simulated and observed element-pair correlations in Figure 13. Before even considering the comparison to simulations, it is important to note that the two observational datasets differ substantially from each other, even for the same element pair. For those

Milky Way data broadly reproduce the same-source versus mixed-source ordering, but survey systematics rival many simulation offsets.



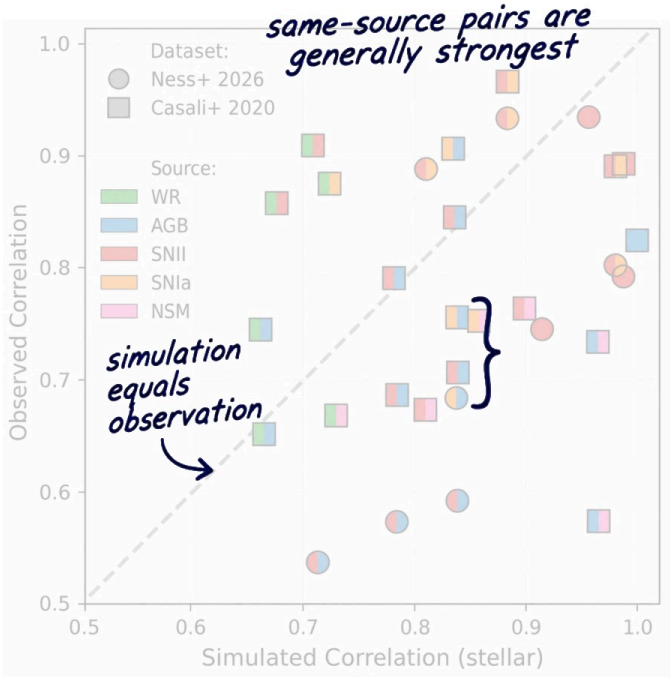



**Figure 13.** Comparison of simulated versus observed element-pair Pearson correlations. Each point plots the simulated stellar correlation (x-axis) against the observed correlation (y-axis) for the same element pair. The dashed diagonal is the 1:1 reference line, and departures from this line indicate mismatches in pairwise chemical correlation structure between the simulations and the observations. Marker shape identifies the observational dataset: circles for Ness et al. (2026), squares for Casali et al. (2020). The individual elements plotted are O, Mg, S, Fe, Ce for the Ness et al. data set and C, Mg, S, Fe, Ba, Ce, Eu for the Casali et al. one. Each marker is split into two colours to indicate the dominant nucleosynthetic channels producing each of the two elements being compared (see legend): WR stars for C, SNII for O, Mg, and S, AGB stars for N, Ba, and Ce, SNIa for Fe, and NSM for Eu. (Note that, although Figure 3 shows that NSM does not dominate over AGB for Eu, we nonetheless colour it as NSM because it is the only element in our sample for which NSM makes a significant contribution.)

Typical inter-survey scatter in correlation coefficient is 0.05–0.15, comparable to many simulation offsets.

benefit pairs in common between the two samples, the typical inter-dataset scatter in correlation coefficient is $\sigma \approx 0.05$–$0.15$, with estimations from Casali et al. generally higher than from Ness et al.. These discrepancies are comparable to, and often exceed, the simulation-observation offsets visible in the figure. This inter-dataset scatter reflects the combined effects of differing spectral resolution and wavelength coverage, distinct stellar populations and evolutionary stages being probed, heterogeneous stellar parameter pipelines, and different selection functions across the Galactic disc. Consequently, it is not immediately clear whether apparent tensions between simulation and observation reflect genuine deficiencies in our model or instead arise from systematic uncertainties inherent in the observational abundance measurements themselves. We explore these systematics and their possible origins in more detail in Appendix B.

With this caveat in mind, the overall level of agreement between simulations and observations is encouraging, and indicates that RoSE has successfully captured the primary processes influencing the chemical distribution structure of galaxies. In particular, in both the simulations and observations we generally find that correlations are highest between elements from the same nucleosynthetic group (i.e., solid-coloured points lie preferentially at the top-right of the plot), correlations between SNII and SNIa elements are on average the next-highest, and correlations between SNII and AGB elements are lowest.

The largest systematic tensions between the simulations and observations are for SNII-AGB pairs from the Ness et al. sample (blue and red circles), and for pairs involving WR or NSM elements from the Casali et al. data set (squares including pink or green). For the SNII-AGB pairs, the simulation tends to predict somewhat stronger correlations than the Ness et al. (2026) observations, with points falling below the one-to-one line. There is significantly better agreement with the Casali et al. results for SNII-AGB pairs, highlighting the importance of systematic disagreements within the observations. For the pairs including WR or NSM elements, when compared with the Casali et al. dataset RoSE appears to underestimate the correlation of WR element pairs while overestimating that of NSM pairs. Systematic physical and observational mechanisms could contribute to these discrepancies. On the observational side, AGB-sourced elements (Ba, Ce, Eu) are typically measured from weaker, more blended spectral features, and their formal abundance uncertainties are larger than for the SNII (O, Mg, S); measurement noise artificially suppresses observed correlations, preferentially affecting these mixed-source pairs. Additionally, the two observational samples probe different radial and age distributions within the Milky Way disc: the SDSS-V red giants span a broader range of Galactic radii and ages than the Solar-neighbourhood, Solar-type stars of Casali et al., introducing radial abundance gradients and age-metallicity scatter that further dilute pairwise correlations. On the simulation side, the RoSE model tracks only nine isotopes across six nucleosynthetic channels, potentially underestimating the decorrelation produced by the full chemical diversity of real stellar populations, and the isolated, relatively quiescent galaxy environment of RoSE does not capture bar-driven radial migration or satellite-driven kinematic heating, both of which would admix chemically distinct populations and weaken correlations in observed samples.

common pairs show inter-dataset scatter

That makes a precise failure diagnosis premature, especially for weak or blended AGB-element lines.

Taken together, however, the comparison between simulations and observations shown in Figure 13 demonstrates that the nucleosynthetic source framework established in this work, wherein same-channel pairs are strongly correlated and mixed-channel pairs are weaker, is qualitatively validated by current observational data. The residual quantitative offset in mixed-source pairs, while genuine, is no larger than the inter-dataset scatter between the two independent observational samples, and thus cannot presently be attributed unambiguously to shortcomings in RoSE. Future surveys delivering high-fidelity multi-element abundance measurements, together with precise stellar ages and ample sized age-matched subsamples, will support statistical analyses such as those presented here, will be essential for testing the chemical evolution models' subgrid physical processes in greater detail. Our RoSE results presented here also highlight challenges of extracting robust correlation measurements from observed stellar spectra.

Our isolated disc omits bar-driven migration, satellite heating and mergers that mix populations born at different radii and epochs.

### 5.4 Scope and caveats

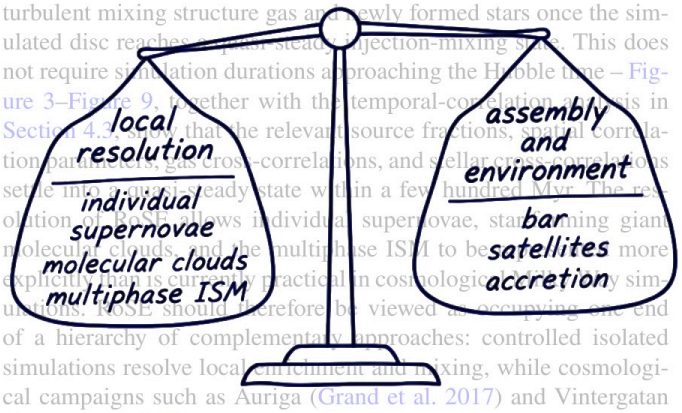
We close the discussion by making explicit the intended scope and limitations of RoSE. Our simulation is intended to be a controlled, high-resolution experiment on an isolated Milky Way-mass disc, and while the initial condition is designed to reproduce the bulk characteristics of the Milky Way, it is not tuned to reproduce its kinematic structure or assembly history in detail. For example, our simulation does not include a bar, perturbations from satellite galaxies, or ongoing cosmological gas accretion or mergers. These missing processes are important for determining the present-day distribution of old stars in both chemical and physical space: bar-driven radial migra-

RoSE is a controlled high-resolution mixing experiment, not a complete reconstruction of the Milky Way's history.

tion, satellite heating, and mergers all mix stellar populations born at different radii and epochs, thereby changing quantitative correlations (for example between abundances, age, and velocity dispersion) measured in a Galactic stellar sample that is selected over a fixed volume rather than at fixed stellar age. Moreover, even if we did include these processes, it is unclear that the 600 Myr duration of our simulation would be sufficient to capture their effects. It is for this reason that we have not attempted to examine the statistical properties of mixed-age stellar populations in this study, and have focused instead only on the gas and the abundances of mono-age stellar populations. Cosmological simulations specifically tuned to match the Milky Way (e.g., Grand et al. 2017; Agertz et al. 2021) include these assembly and environmental effects and are generally better suited to such questions, making them complementary to the controlled experiment studied here.

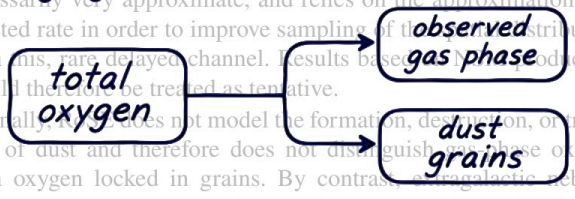
We trade cosmological context for more explicit individual supernovae, molecular clouds, and a multiphase ISM.

Conversely, the spatial and mass resolution required to resolve individual supernovae, molecular clouds, and a multiphase ISM in current cosmological campaigns generally does not permit the same star-by-star treatment of feedback and star formation. Instead, they rely on star-cluster formation recipes that average over many individual stars. The goal of this paper is to measure how source-dependent enrichment and turbulent mixing structure gas and newly formed stars once the simulated disc reaches a regular steady injection-mixing state. This does not require simulation durations approaching the Hubble time – Figure 3–Figure 9, together with the temporal-correlation analysis in Section 4.3, show that the relevant source fractions, spatial correlation patterns, gas cross-correlations, and stellar cross-correlations set in on a quasi-steady state within a few hundred Myr. The resolution of RoSE allows individual supernovae, star-forming giant molecular clouds, and the multiphase ISM to be resolved explicitly, which is currently impractical in cosmological campaigns or galaxy simulations. RoSE should therefore be viewed as occupying one end of a hierarchy of complementary approaches: controlled isolated simulations resolve local enrichment and mixing, while cosmological campaigns such as Auriga (Grand et al. 2017) and Vintergatan (Agertz et al. 2021) capture assembly history and environment.




Another important caveat regards our treatment of nucleosynthesis and element injection. Our set of nine isotopes and five nucleosynthetic channels is chosen to provide good coverage of the main enrichment channels for most elements, but it does omit a number of other nucleosynthetic pathways, such as classical novae, cosmic-ray spallation, and hypernovae, which may be dominant for certain isotopes. Moreover, we have explored only a single relatively simple set of yield tables, and have not explored either how the results would change if we were to assume different (or more complex) yield models that include, for example, multiple SNa subclasses. These omissions are most important for elements not tracked here, for detailed isotopic ratios for the elements tracked, or for elements whose real production is dominated by multiple enrichment channels. Our treatment of NSMs is also necessarily very approximate, and relies on the approximation of a boosted rate in order to improve sampling of the tail of the distribution from this rare, delayed channel. Results based on NSM production should therefore be treated as tentative.

Spatially varying dust depletion could alter observed gas-phase oxygen fluctuations without changing total oxygen.




Finally, RoSE does not model the formation, destruction, or transport of dust and therefore does not distinguish gas-phase oxygen from oxygen locked in grains. By contrast, extragalactic nebular abundance maps measure only gas-phase oxygen. Estimates for H II regions imply oxygen depletions of approximately 0.08–0.12 dex, with the depletion tending to increase with metallicity (Peimbert & Peimbert 2010); grain growth and destruction can also make depletion depend on local ISM conditions. Uniform depletion would have little effect on the fluctuation statistics used here, but spatially varying depletion could alter the measured oxygen autocorrelation, correlation length, and injection scale. Comparisons between RoSE's total-oxygen structure and observed gas-phase oxygen maps should therefore not be interpreted as exact one-to-one tests.

## 6 CONCLUSIONS

On the canvas of a galactic disc, stellar enrichment leaves measurable, rhythmically patterned imprints: stellar winds fall like fine drizzle, slowly and locally moistening their surroundings, while supernovae burst like sudden storms, instantaneously driving strong dynamical perturbations and wide-ranging mass transport; short-lived injection events behave like transient showers, whereas delayed channels resemble prolonged rainy seasons that accumulate and reshape the chemical landscape over long intervals. Different elements, governed by their nucleosynthetic channels, paint distinguishable spatial "ripples" across the disc – ripples that encode both source information and the timescales of mixing and transport. RoSE (Ripples of Stellar Enrichment) follows this process, revealing how nucleosynthetic channels, phase segregation, and dynamics jointly determine inter-element correlations and their persistence.

This study demonstrates that the spatial, temporal, and cross-correlation structure of elemental fluctuation distributions in galaxies is fundamentally organized by the diversity of nucleosynthetic sources. Our core findings are threefold:

(i) Nucleosynthetic source clustering is the primary organizing principle of elemental spatial structure. Elements produced by the same nucleosynthetic channel (e.g., SNII-dominated O, Mg, S or AGB-dominated Ba, Ce) cluster tightly in spatial statistics, with auto-correlation parameters closely tracking those of the dominant source. Individual nucleosynthetic channels imprint characteristic length scales and timescales into the ISM that persist long enough for the elemental spatial distribution to act as a fossil record of source-specific injection environments.

(ii) The differences or similarities between elements' nucleosynthetic origin similarly determine the cross-correlation between their abundances in both the interstellar medium (ISM) and young stars. A simple linear regression model relating the cross-correlation between element pairs to the level of similarity or difference in their nucleosynthetic origins predicts those cross-correlations remarkably well ($R^2 = 0.84$ in gas, 0.73 in coeval stars). The predictive power of this model is fundamentally unchanged between gas and stellar phases. Varying the assumed old-star AGB return rate between zero and twice its fiducial value changes individual correlations but preserves the element-pair ordering and the explanatory power of the source-fraction model. This result aligns qualitatively with the analytic Krumholz et al. (2025) model wherein the differences in delay time between star formation and element return for different astrophysical origin sites governs the spatial and correlation structure of elemental abundances.

(iii) There are, however, noticeably differences between the abundance correlations found in stars and the interstellar medium, with most element pairs showing higher correlation in stars than in ISM gas. We show that this difference arises from the temperature-dependent environment of star formation. Supernova feedback produces large-scale hot bubbles of enriched gas, which lower the correlation in the gas phase; however, these hot phases do not form stars, so young stellar populations form in cooler, more-mixed gas and inherit stronger correlations. The extent to which elements are better correlated in stars than in gas is therefore primarily a function of [illegible] and the [illegible] also produce large bubbles of very hot gas.

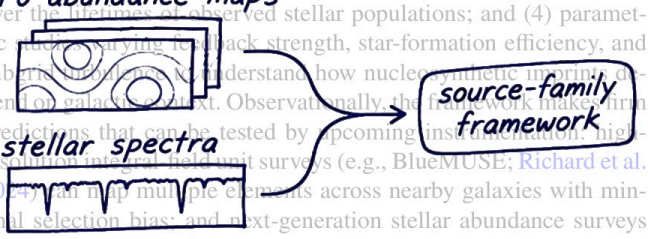
Core result: nucleosynthetic source diversity organizes spatial, temporal, and inter-element correlation structure.

We also compare our simulations results with observations of co-eval Milky Way stars. This comparison qualitatively validates the nucleosynthetic source framework, while the remaining quantitative discrepancies are modest, comparable to the level of disagreement between independent observational data sets, and plausibly explained by uncertainties in stellar age estimation in the observational samples. Consequently, current observations do not yet allow us to unambiguously distinguish whether these residual differences reflect physical processes missing from RoSE (such as bar-driven radial migration) or instead arise from abundance and age systematics in the observational data, underscoring the need for high-fidelity multi-element abundances and precise ages.

Looking forward, the success of the nucleosynthetic source framework motivates several promising research directions. In simulations, natural extensions include: (1) cosmological zoom-in simulations at comparable resolution to test how ripple patterns depend critically on galactic environment and feedback physics; (2) inclusion of additional physical processes (bar-driven radial migration and AGN feedback to see how they modify and whether processes account for any simulation-observation residuals; (3) extension to longer timescales (multi-Gyr) to examine the asymptotic mixing and correlation behaviour achieved over the lifetimes of observed stellar populations; and (4) parametric studies varying feedback strength, star-formation efficiency, and subgrid physics to understand how nucleosynthetic imprints depend on galaxy context. Observationally, the framework makes clear predictions that can be tested by upcoming instrumentation: high-resolution integral-field unit surveys (e.g., BlueMUSE; Richard et al. 2019) can map multiple elements across nearby galaxies with minimal selection bias; and next-generation stellar abundance surveys (e.g., 4MOST; de Jong et al. 2019, WEAVE-StePS; Iovino et al. 2023, SDSS-V; Kollmeier et al. 2026) can provide homogeneous multi-element abundances and precise ages for tens of thousands of Galactic stars with well-defined selection functions. These observational advances will enable detailed comparisons with nucleosynthetic source framework predictions, and ultimately verify the precise roles of thermal physics, dynamics, radial migration, and small-scale turbulent transport in shaping galactic chemical structure. The RoSE project demonstrates that high-resolution simulations tracking star-by-star nucleosynthetic feedback provide an essential bridge between the predictive power of analytic source-delay theory and the complexity of real observations, enabling us to decode the enrichment histories written in the spatial and chemical structure of galaxies.

*Key next tests are multi-element gas maps and precise, age-resolved stellar abundances.*

*BlueMUSE can test predicted source-dependent spatial ripples across multiple gas-phase elements.*




*Large stellar surveys can test the source-family pattern only if abundance and age systematics are controlled.*

## ACKNOWLEDGEMENTS

We acknowledge high-performance computing resources provided by the Australian National Computational Infrastructure (award jh2) through the National and ANU Computational Merit Allocation Schemes and Astronomy Supercomputer Time Allocation Committee (ASTAC) in Australia. Some of the text and code refinements were developed in micro-managed highly supervised engagement with AI language models (Claude, Anthropic and Gemini, Google), but all final text was generated directly or edited heavily by the human authors.

CZ acknowledges financial support from the Research School of Astronomy and Astrophysics supplementary scholarship, Australian National University postgraduate research scholarship and Chinese Scholarship Council. 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. ZH acknowledges support from National Natural Science Foundation of China (NSFC) through grant No. 12503026 and support from Boya Fellowship at Peking University. This research was supported by NSF Grant AST-2406729 and a Humboldt Research Award from the Alexander von Humboldt Foundation.

Thanks to the observations collected with the FLAMES instrument at VLT/UT2 (Paranal Observatory, ESO, Chile), for the Gaia-ESO Large Public Spectroscopic Survey (188.B-3002, 193.B-0936).

Funding for the Sloan Digital Sky Survey V has been provided by the Alfred P. Sloan Foundation, the Heising-Simons Foundation, the National Science Foundation, and the Participating Institutions. SDSS acknowledges support and resources from the Center for High-Performance Computing at the University of Utah. SDSS telescopes are located at Apache Point Observatory, funded by the Astrophysical Research Consortium and operated by New Mexico State University, and at Las Campanas Observatory, operated by the Carnegie Institution for Science. The SDSS web site is www.sdss.org.

SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration, including the Carnegie Institution for Science, Chilean National Time Allocation Committee (CNTAC) ratified researchers, Caltech, the Gotham Participation Group, Harvard University, Heidelberg University, The Flatiron Institute, The Johns Hopkins University, L'Ecole polytechnique fédérale de Lausanne (EPFL), Leibniz-Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Extraterrestrische Physik (MPE), Nanjing University, National Astronomical Observatories of China (NAOC), New Mexico State University, The Ohio State University, Pennsylvania State University, Smithsonian Astrophysical Observatory, Space Telescope Science Institute (STScI), the Stellar Astrophysics Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Illinois at Urbana-Champaign, University of Toronto, University of Utah, University of Virginia, Yale University, and Yunnan University.

## DATA AVAILABILITY

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

## REFERENCES

Agertz O., et al., 2021, MNRAS, 503, 5826
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
Bland-Hawthorn J., Sharma S., 2016, Astronomische Nachrichten, 337, 894
Bland-Hawthorn J., Krumholz M. R., Freeman K., 2010, ApJ, 713, 166
Bresolin F., et al., 2025, MNRAS, 539, 755
Casali G., et al., 2020, A&A, 639, A127
Colbrook M. J., Ma X., Hopkins P. F., Squire J., 2017, MNRAS, 467, 2421
Di Matteo P., Pipino A., Lehnert M. D., Combes F., Semelin B., 2009, A&A, 499, 427
Doherty C. L., Gil-Pons P., Lau H. H. B., Lattanzio J. C., Siess L., 2014, MNRAS, 437, 195
Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306
Grand R. J. J., et al., 2016, MNRAS, 460, L94
Grand R. J. J., et al., 2017, MNRAS, 467, 179

We vary the uncertain old-star AGB return from zero to twice the fiducial value without rerunning the dynamics.

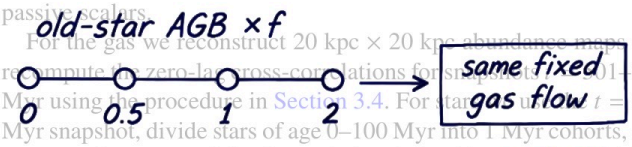


Hirai Y., Beers T. C., Chiba M., Aoki W., Shank D., Saitoh T. R., Okamoto T., Makino J., 2022, MNRAS, 517, 4856
Ho I. T., et al., 2018, A&A, 618, A64
Holmbeck E. M., Andrews J. J., 2024, ApJ, 963, 110
Hopkins P. F., 2015, MNRAS, 450, 53
Hopkins P. F., 2016, MNRAS, 462, 576
Hopkins P. F., Raives M. J., 2016, MNRAS, 455, 51
Hopkins P. F., et al., 2018, MNRAS, 480, 800
Hu Z., Wibking B. D., Krumholz M. R., 2023, MNRAS, 521, 5604
Huang R., et al., 2026, MNRAS
Iovino A., et al., 2023, A&A, 672, A87
Jofré P., Heiter U., Soubiran C., 2019, ARA&A, 57, 571
Karakas A. I., Lugaro M., 2016, ApJ, 825, 26
Kewley L. J., Nicholls D. C., Sutherland R. S., 2019, ARA&A, 57, 511
Kim J.-h., et al., 2016, ApJ, 833, 202
Kollmeier J. A., et al., 2026, AJ, 171, 52
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
Krumholz M. R., Ting Y.-S., Li Z., Zhang C., Mead J., Ness M. K., 2025, The Open Journal of Astrophysics, 8, 50356
Li W., Chornock R., Leaman J., Filippenko A. V., Poznanski D., Wang X., Ganeshalingam M., Mannucci F., 2011, MNRAS, 412, 1473
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., 2022, MNRAS, 518, 286–304
Li Z., et al., 2023, MNRAS, 518, 286
Li Z., et al., 2024, MNRAS, 528, 7103
Li S.-L., Li Z., Wisnioski E., Krumholz M. R., Sánchez S. F., 2025a, MNRAS, 536, 430
Li Z., et al., 2025b, ApJ, 987, L28
Ma X., Hopkins P. F., Feldmann R., Torrey P., Faucher-Giguère C.-A., Kereš D., 2017, MNRAS, 466, 4780
Maiolino R., Mannucci F., 2019, The Astronomy and Astrophysics Review, 27
Manea C., et al., 2025, ApJ, 993, 45
Mayor M., et al., 2003, The Messenger, 114, 20
Mead J., De La Garza Evia R., Ness M. K., 2025, ApJ, 994, 52
Mészáros S., et al., 2025, AJ, 170, 96
Metha B., Trenti M., Chu T., 2021, MNRAS, 508, 489
Metha B., Trenti M., Norman C., 2026, MNRAS,
Ness M. K., et al., 2026, Monthly Notices of the Royal Astronomical Society, p. stag637
Orr M. E., et al., 2023, MNRAS, 521, 3708
Peimbert A., Peimbert M., 2010, ApJ, 724, 791
Poetrodjojo H., et al., 2018, MNRAS, 479, 5235
Richard J., et al., 2024, in Bryant J. J., Motohara K., Vernet J. R. D., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 13096, Ground-based and Airborne Instrumentation for Astronomy X. p. 1309622 (arXiv:2406.13914), doi:10.1117/12.3019831
Rosales-Ortega F. F., Díaz A. I., Kennicutt R. C., Sánchez S. F., 2011, MNRAS, 415, 2439
SDSS Collaboration et al., 2025, arXiv e-prints, p. arXiv:2507.07093
Sánchez-Menguiano L., et al., 2016, ApJ, 830, L40
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
Seitenzahl I. R., et al., 2013, MNRAS, 429, 1156
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
Stone-Martinez A., Holtzman J. A., Lu Y., Hasselquist S., Imig J., Griffith E. J., Bellinger E. P., Saydjari A. K., 2025, AJ, 170, 66
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, Fundamentals Cosmic Phys., 5, 287
Weinberg D. H., et al., 2022, ApJS, 260, 32
Wibking B. D., Krumholz M. R., 2023, MNRAS, 521, 5972
Williams T. G., et al., 2022, MNRAS, 509, 1303
Wiseman P., et al., 2021, MNRAS, 506, 3330
Yang C.-C., Krumholz M., 2012, The Astrophysical Journal, 758, 48
Zhang C., Li Z., Hu Z., Krumholz M. R., 2025, MNRAS, 540, 3906
da Silva R. L., Fumagalli M., Krumholz M., 2012, ApJ, 745, 145
de Jong R. S., et al., 2019, The Messenger, 175, 3

## APPENDIX A: SENSITIVITY TO THE OLD-STAR AGB RETURN RATE

The prescription in Section 2.2 assigns old star particles a time-averaged AGB element return rate per unit stellar mass. To quantify the impact of this assumption, we introduce a multiplier $f$ applied to the old-star AGB contribution in the reconstructed correlation analysis for $f = 0$, 0.5, 1, and 2, where $f = 1$ is the fiducial model used throughout the main text. This is a post-processing sensitivity test: we vary only the old-star AGB scalar in the source-resolved simulation outputs and multiply the old-star AGB scalar contribution by $f$ when reconstructing total element abundances and source fractions. The test therefore isolates the dependence of the reported chemical statistics on the adopted return rate; it does not rerun the dynamics, which are unchanged because the tracked metal fields are passive scalars.

The metals are passive scalars, so this test changes enrichment weights while keeping the dynamics fixed.




For the gas we reconstruct 20 kpc × 20 kpc abundance maps and recompute the zero-lag cross-correlations for snapshots $t = 201$–600 Myr using the procedure in Section 3.4. For stars at $t = 600$ Myr snapshot, divide stars of age 0–100 Myr into 1 Myr cohorts, and compute the mass-weighted correlation defined by equation 10. Thus both phases sample the same birth epochs. We also repeat the regression analysis of equation 12, including the three pure basis elements $\mathrm{Pure}_{\mathrm{SNIa}}$, $\mathrm{Pure}_{\mathrm{NSM}}$, and $\mathrm{Pure}_{\mathrm{oldAGB}}$ whenever their corresponding channels are present. For $f = 0$, the old-star AGB contribution vanishes identically, so we remove both $\mathrm{Pure}_{\mathrm{oldAGB}}$ and the oldAGB predictor from the regression; this fit therefore contains 11 variables, five source channels, and 55 unique pairs, compared to 12 variables, six channels, and 66 pairs for $f > 0$.

Figure A1 shows that the response is smooth but element-pair dependent. Across the full range $f = 0$–2, the largest gas-phase excursion is 0.305 for C–Fe, while the median excursion over all 36 element pairs is 0.140. The stellar correlations are less sensitive: their largest excursion is 0.149 for Fe–Ba and their median excursion is 0.039. The largest changes generally involve pairs whose source fractions receive appreciable old-star AGB contributions, whereas pairs dominated by a shared prompt source remain close to unity. Importantly, pairs assigned to the same dominant nucleosynthetic channel remain more strongly correlated on average than mixed-source pairs for all values of $f$.

Test whether the source-family hierarchy survives, not whether every coefficient stays fixed.

The comparison in Table A1 separates quantitative sensitivity from robustness of the physical interpretation. Removing the old-star AGB contribution (Figure A1) can shift the fiducial correlation values, particularly in the gas, but the element-pair rank ordering remains strongly correlated with the fiducial result ($\rho = 0.926$ for gas and 0.908 for stars). For the less extreme $f = 0.5$ and 2 cases, $\rho \geq 0.977$ in both phases. The regression in Figure A2 retains substantial explanatory power throughout the tested range, with $R^2 = 0.70$–0.88 for gas and 0.71–0.77 for stars. We therefore conclude that the precise numerical values of some gas-phase correlations are sensitive to the assumed old-star AGB return rate, but the

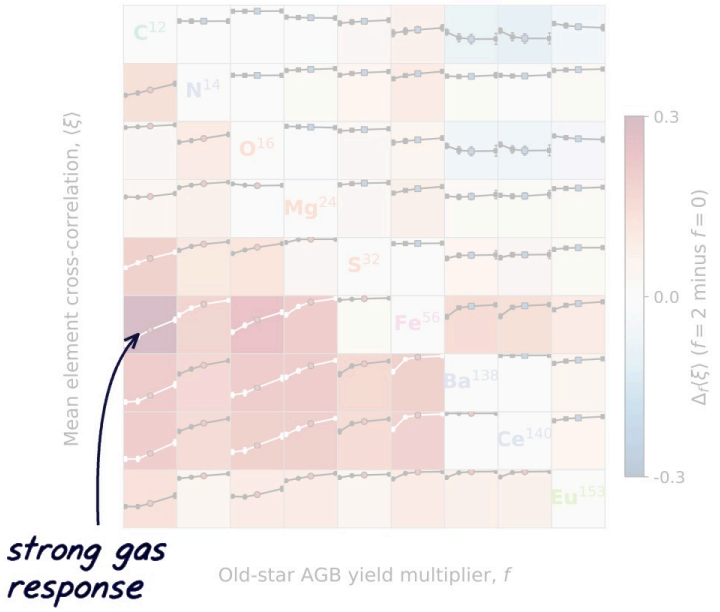
**Table A1.** Summary of the old-star AGB sensitivity test. Correlation changes and Spearman rank coefficients compare the 36 real-element pairs at each $f$ to the fiducial $f = 1$ result. The final columns report the coefficient of determination for the source-fraction regression, including the pure basis elements.

stars change less

| | Gas | | | | Mass-weighted stars | | | |
|---|---|---|---|---|---|---|---|---|
| $f$ | median $\lvert\Delta\xi\rvert$ | max $\lvert\Delta\xi\rvert$ | Spearman $\rho$ | $R^2$ | median $\lvert\Delta\xi\rvert$ | max $\lvert\Delta\xi\rvert$ | Spearman $\rho$ | $R^2$ |
| 0 | 0.078 | 0.181 | 0.926 | 0.70 | 0.028 | 0.132 | 0.908 | 0.77 |
| 0.5 | 0.028 | 0.082 | 0.993 | 0.81 | 0.007 | 0.023 | 0.993 | 0.71 |
| 1 | – | – | 1.000 | 0.84 | – | – | 1.000 | 0.73 |
| 2 | 0.041 | 0.135 | 0.979 | 0.88 | 0.010 | 0.028 | 0.977 | 0.77 |

ordering remains strong

Even removing old-star AGB entirely leaves a strong rank correlation with the fiducial result.




strong gas response

**Figure A1.** Sensitivity of the element-element cross-correlations to the old-star AGB yield multiplier $f$. Each off-diagonal panel shows the mean correlation over the final 100 Myr as a function of $f$, with error bars giving the standard deviation among the 100 individual snapshots (gas; circles in the lower triangle) or 1 Myr stellar age cohorts (mass-weighted stars; squares in the upper triangle). Black-edged points mark the fiducial $f = 1$ model. The background colour encodes $\Delta_f\langle\xi\rangle = \langle\xi\rangle_{f=2} - \langle\xi\rangle_{f=0}$ on a zero-centred symmetric scale: red denotes increasing correlation with $f$, while blue denotes decreasing correlation.

Gas correlations change more than mass-weighted stellar correlations in this sensitivity test.

main conclusions—that nucleosynthetic source similarity organises the correlation matrix, that same-source pairs correlate more strongly than mixed-source pairs, and that source-fraction differences predict much of the pair-to-pair variation—do not depend on the fiducial choice of this parameter. The mass-weighted stellar results are quantitatively more robust than the gas-phase results.

## APPENDIX B: SOURCES OF OBSERVATIONAL DISCREPANCIES

In this appendix we provide a more detailed analysis of the inter-element correlation structure in the Casali et al. (2020) and Ness et al. (2026) samples, and the reasons why the element-element correlations we extract from these two data sets might differ. Broadly speaking, these two data sets differ in that the former is a much smaller sample of stars selected to be both very nearby and very similar to the Sun, while the latter is a much broader sample, which offers far better statistics but at the price of larger statistical and systematic errors.

Given the distinction between the two data sets, the first potential source of discrepancy to consider is abundance precision. Uncertainties are important because errors reduce the level of correlation. To test how much this effect contributes to the discrepancy between the two data sets, we degrade the Casali et al. abundances (which have smaller quoted uncertainties) by adding Gaussian noise at the level of the median APOGEE uncertainties and recomputing the element-element correlations. We find that doing so moves the Casali et al. correlations towards those of Ness et al.. Quantitatively, the mean absolute deviation (MAD) between the two correlation matrices is reduced from 0.139 to 0.089 when we add artificial noise to the Casali et al. measurements. However, the fact that the residual remains 0.089 implies that precision alone cannot explain the offset.

A second potential source of systematic difference is the size of the window in [Fe/H] used to select the sample. The Casali et al. sample is highly uniform and near-Solar, while the APOGEE sample used in Ness et al. is considerably broader. We can either compute correlations on the full sample, or down-select to pick out narrow [Fe/H] windows around Solar. This down-selection reduces the dynamical range in abundance space and thus can suppress correlations; for example, the MAD with the Casali et al. sample is 0.14 for the full Ness et al. sample, but this increases to 0.72 and 0.56 if we select stars within [Fe/H] windows of ±0.05 and ±0.10 dex, respectively. Thus differences in [Fe/H] coverage and intrinsic chemical diversity are likely an important contributor to differences between the two data sets.

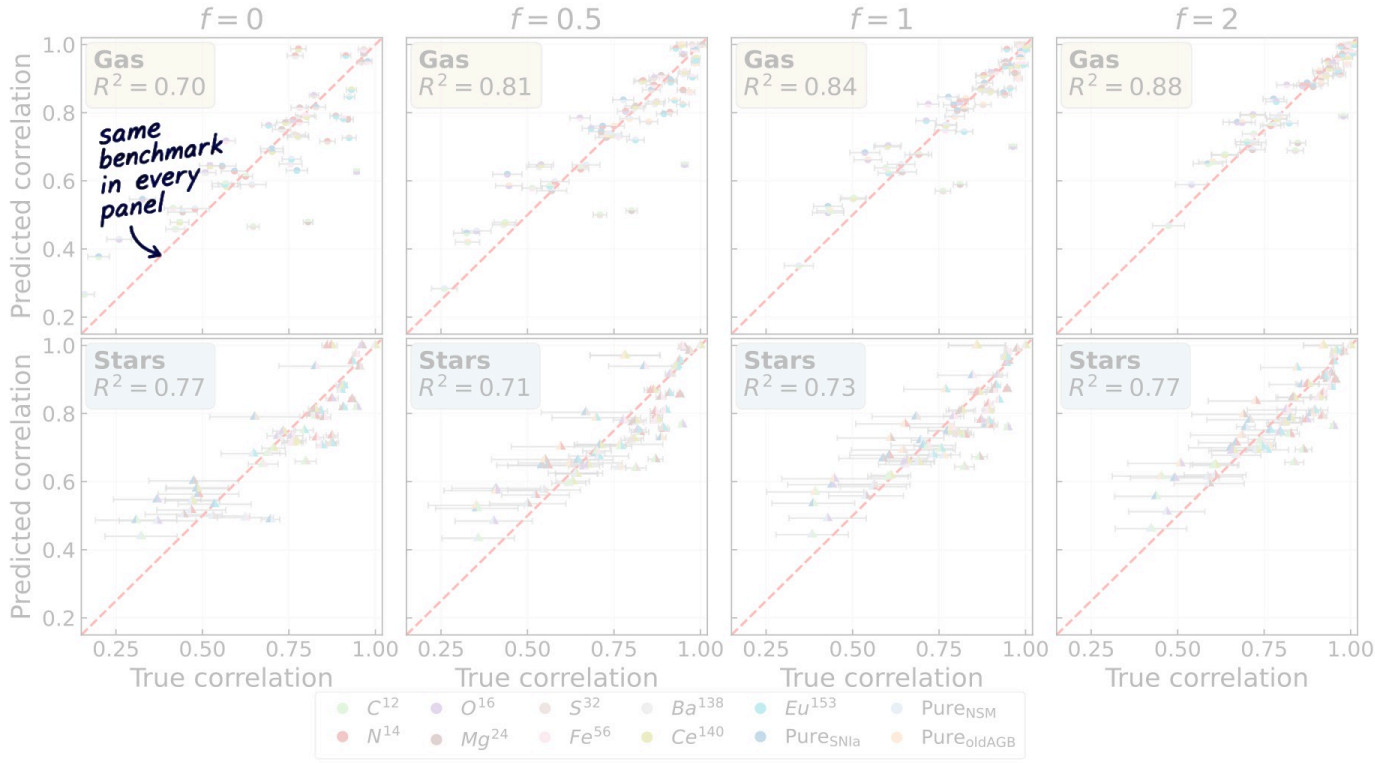
A third potential origin for the disagreement between the two data sets is known temperature-dependent systematics, particularly for Ce in the infrared. These can be mitigated but not eliminated by restricting both samples to a narrow $T_{\rm eff}$ range. Doing so reduces the difference in cross correlations between Ce and other elements: $\Delta\xi_{\text{Ce-Fe}}$ is reduced from −0.12 to −0.09, $\Delta\xi_{\text{Ce-Mg}}$ from −0.17 to −0.15, and $\Delta\xi_{\text{Ce-S}}$ from −0.15 to −0.14, where we define $\Delta\xi = \xi_{\text{Ness}} - \xi_{\text{Casali}}$). This exercise confirms that such systematics are real, yet they are not the dominant driver of the overall discrepancy.

The fourth and most significant difference lies in the age estimates, which matter because abundance measurement errors and age estimates are often tightly coupled. Figure B1 shows how the cross-correlations for six element pairs in common to the Casali et al. and Ness et al. samples vary as a function of inferred stellar age. We see that the two data sets broadly agree at intermediate ages, but that they diverge substantially for stars older than $\sim$ 10 Gyr. If we restrict the comparison to stars younger than 10 Gyr, the offsets between the two samples shrink considerably: $\Delta\xi_{\text{Ce-Fe}}$ changes from −0.12 to −0.01, $\Delta\xi_{\text{Ce-Mg}}$ from −0.17 to −0.06, $\Delta\xi_{\text{Ce-S}}$ from −0.15 to −0.06, $\Delta\xi_{\text{Fe-S}}$ from −0.08 to −0.05, and $\Delta\xi_{\text{Mg-S}}$ from −0.10 to −0.04; only $\Delta\xi_{\text{Fe-Mg}}$ remains unchanged at −0.03. This pattern indicates that much of the discrepancy between the two surveys is driven by the oldest stars and

The source-fraction regression retains substantial predictive power across every tested AGB prescription.



Gas $R^2$: 0.70–0.88




Stars $R^2$: 0.71–0.77

**Figure A2.** Regression predictions under the four old-star AGB yield multipliers. Columns show $f = 0$, 0.5, 1, and 2; rows show gas (half-circle markers) and mass-weighted stellar correlations (triangular markers). Marker colours, error bars, and the one-to-one line follow Figure 11. The three pure elements are included for $f > 0$; at $f = 0$, $\mathrm{Pure_{oldAGB}}$ and the inactive oldAGB regression term are omitted.

by the treatment of age. Moreover, the element pairs involving Ce in the Casali et al. sample show particularly strong, non-monotonic fluctuations with age. Perhaps most intriguingly, both samples independently exhibit a bump in cross-correlation for stars with ages $\sim$ 5–10 Gyr, suggesting that the correlation structure is responding to features in the Galactic star formation history. While our simulations do not follow a full multi-Gyr evolution, the observed age trends in stellar abundances provide a powerful window onto the sequence of quasi-equilibrium states experienced by the Milky Way (and by analogy other nearby disc galaxies) over cosmic time.

In this sense, the higher correlations found in Casali et al. compared to Ness et al. during the age range from $\sim$ 5–10 Gyr may best be interpreted not as a failure of the methodology in either study, but as evidence that different surveys sample stellar populations with different effective chemical dimensionalities and measurement response functions, consistent with broader discussions of abundance precision and systematics (e.g., Jofré et al. 2019).

The exact fit changes; the explanatory framework does not.

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

The two stellar surveys agree at intermediate ages but diverge sharply for the oldest stars, implicating age-dependent systematics.

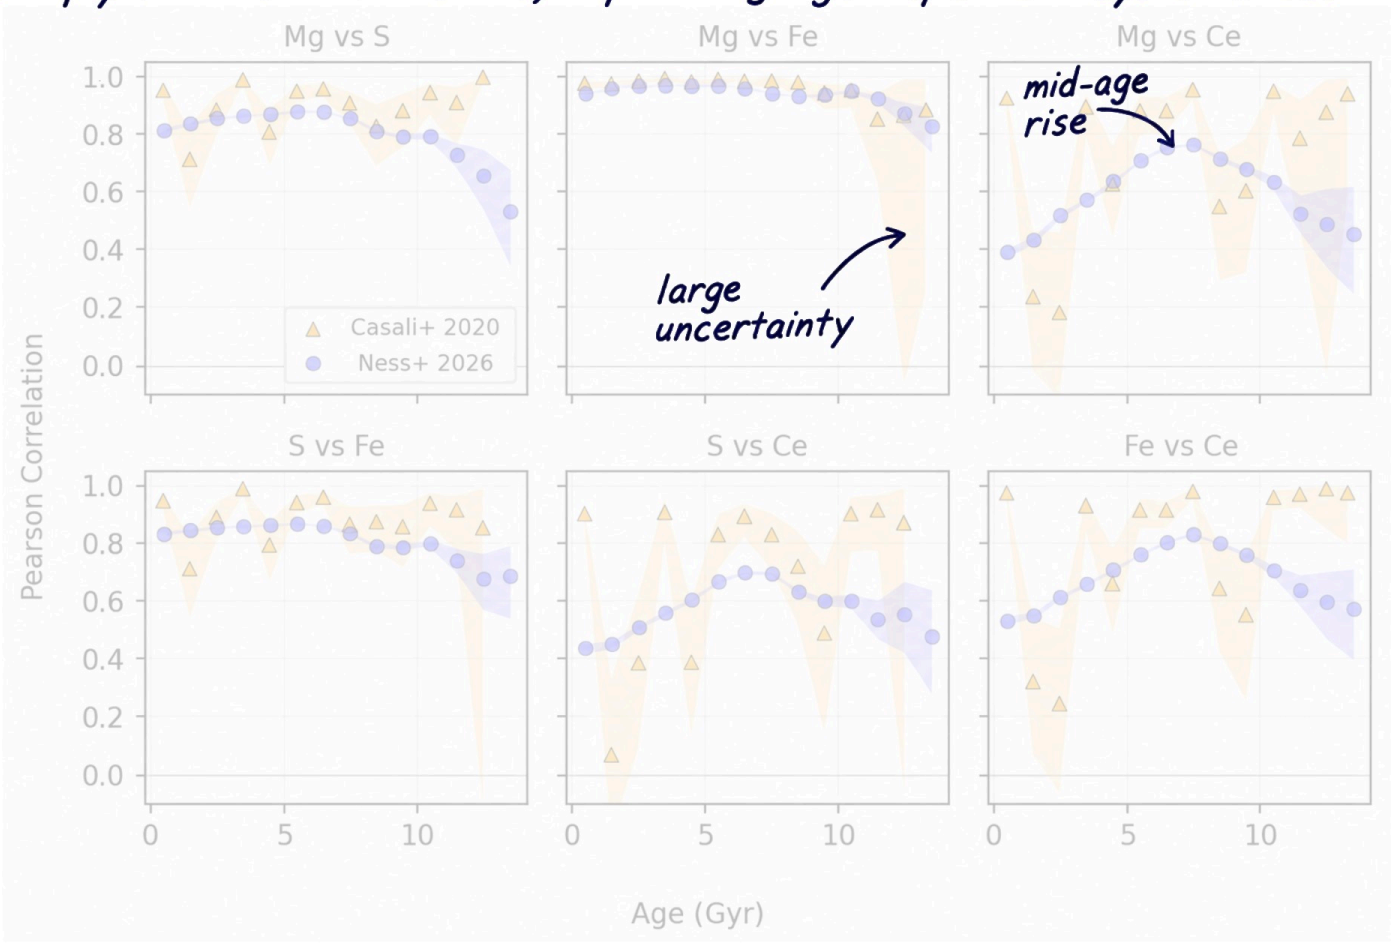



**Figure B1.** Cross-correlation as a function of stellar age for the six element pairs in common between the samples of Casali et al. (2020) and Ness et al. (2026). Points are computed in fixed 1 Gyr age bins (excluding bins with fewer than 5 samples). The shaded band around each sample trace shows the 95% confidence interval of the Pearson correlation, estimated with the Fisher $z$ transform in each age bin and mapped back to correlation space. All panels share the same axes.

The shared mid-age bump may be a real imprint of the Milky Way's changing star-formation history.

Sparse Casali age bins have broad intervals; avoid over-reading individual old-age excursions.