The metal-poor tail adds a constraint on how the young Milky Way acquired its gas.

Draft version March 24, 2026
Typeset using LaTeX **twocolumn** style in AASTeX7

arXiv:2511.18901v2 [astro-ph.GA] 23 Mar 2026

The dawn is quiet II: Gaia XP constraints on the Milky Way's proto-Galaxy from very metal-poor MDF tails

Boquan Chen [1,2] Matthew Orkney [3,4] Yuan-Sen Ting [1,2] and Michael Hayden [5]

[1] Department of Astronomy, The Ohio State University, Columbus, OH 43210, USA
[2] Center for Cosmology and AstroParticle Physics (CCAPP), 191 West Woodruff Avenue, Columbus, OH 43210, USA
[3] Institut de Ciencies del Cosmos (ICCUB), Universitat de Barcelona (IEEC-UB), Martí i Franquès 1, E08028 Barcelona, Spain
[4] Institut d'Estudis Espacials de Catalunya (IEEC), E-08034 Barcelona, Spain
[5] Department of Physics and Astronomy, The University of Oklahoma, 440 W. Brooks St., Norman, OK 73019, USA

ABSTRACT

The earliest phase of the Milky Way's evolution involved a transition from a dispersion-supported proto-galaxy to a rotationally supported disk. A key chemical signature of this transition is the moderate rise in [α/Fe] near [Fe/H] ≈ -1.3, which we previously interpreted as evidence for α-enhanced gas accretion fueling early disk formation. However, this trend alone does not uniquely constrain the trade-off between initial gas mass, inflow rate, and star formation efficiency (SFE), leaving the physical condition of the proto-Milky Way uncertain. To break this degeneracy, we analyze the metal-poor tail ($-3 <$ [Fe/H] < -2) of the Milky Way's metallicity distribution function (MDF) using Gaia DR3 BP/RP (XP) metallicities from ten catalogs. After applying recommended quality cuts, all catalogs exhibit a single-slope exponential tail with slopes $k \simeq 0.6$–2.0. Comparison with one-zone galactic chemical-evolution (GCE) models that replicated the [α/Fe]-rise from Paper I shows that shallow tails ($k \simeq 0.6$) require a massive initial cold gas reservoir ($\gtrsim 10^9\,M_\odot$), while steeper tails ($k \gtrsim 1$) arise from small reservoirs that built up over time with weak inflow. MDFs with $k \simeq 1.0$ are best reproduced under our GCE framework, which favor a proto-Galaxy with a moderate gas reservoir (10^8–$10^9\,M_\odot$) sustained through weak continuous inflow ($\sim 2\ M_\odot\ yr^{-1}$) and SFE comparable to today's value (a few $\times 10^{-10}\,yr^{-1}$) during the first Gyr. This scenario is reinforced by MDFs of 30 Milky Way analogs in the Auriga simulations, which exhibit similar slopes ($k \approx 1.25$). The metal-poor MDF tail thus provides a quantitative constraint on the Milky Way's early gas accretion and star formation history.

The preferred quiet buildup depends on catalog selection and the chemical-evolution model.

Keywords: Galaxies (573)

1. INTRODUCTION

The earliest phase of the Milky Way's formation was dominated by a rapid and turbulent phase in the form of a compact proto-galaxy, preceding the emergence of the rotationally supported disk. Elemental abundances from APOGEE (Majewski et al. 2017) and kinematics from Gaia (Gaia Collaboration et al. 2016) indicate an in-situ, kinematically hot "Aurora" phase at low metallicity, followed by a fast "spin-up" of the Galaxy as its disk started to form between $-1.3 <$ [Fe/H] < -0.9 over ∼1–2 Gyr, with a simultaneous drop in abundance scatter (Belokurov & Kravtsov 2022). A central concentration of 18,000 giant stars with [Fe/H] < -1.5 have been found within ~ 5 kpc of the Galactic center and proposed to be the proto-Galactic nucleus. These stars suggest that the metal-poor "heart" contains $> 10^8\ M_\odot$ and formed > 12.5 Gyr ago (Rix et al. 2022). Multiple components belonging to the proto-Galaxy have also been identified and characterized with a mix of in situ vs accreted origins (Malhan & Rix 2024; Horta & Schiavon 2025).

The disk of the Milky Way commenced its assembly around 13 Gyr ago, with its star formation rate (SFR) peaking around 11 Gyr ago, coinciding with the Gaia-Sausage-Enceladus (GSE) (Xiang & Rix 2022; Vincenzo et al. 2021; Belokurov & Kravtsov 2022; Xiang et al. 2025). This used to be considered unusually early, but a cold rotating disk has been confirmed by an extragalactic observation at $z \sim 4.3$ when the universe was only 1.5 Gyr old (Neeleman et al. 2020). The quiescent evolution of the disk took over around 8 Gyr ago and continued to the present as evidenced by a clear division in the stellar age-metallicity distribution (Haywood et al. 2013; Snaith et al. 2015; Feuillet et al. 2019;



XP metallicities → metal-poor MDF → early gas history

−3 < [Fe/H] < −2



Sahlholdt et al. 2022). These three phases of evolution are confirmed by a sample of ten million giants from Gaia (Chandra et al. 2024). However, there is evidence that the kinematically defined thin and thick disk could have also co-evolved from 13 Gyr ago (Nepal et al. 2024; Borbolato et al. 2025).

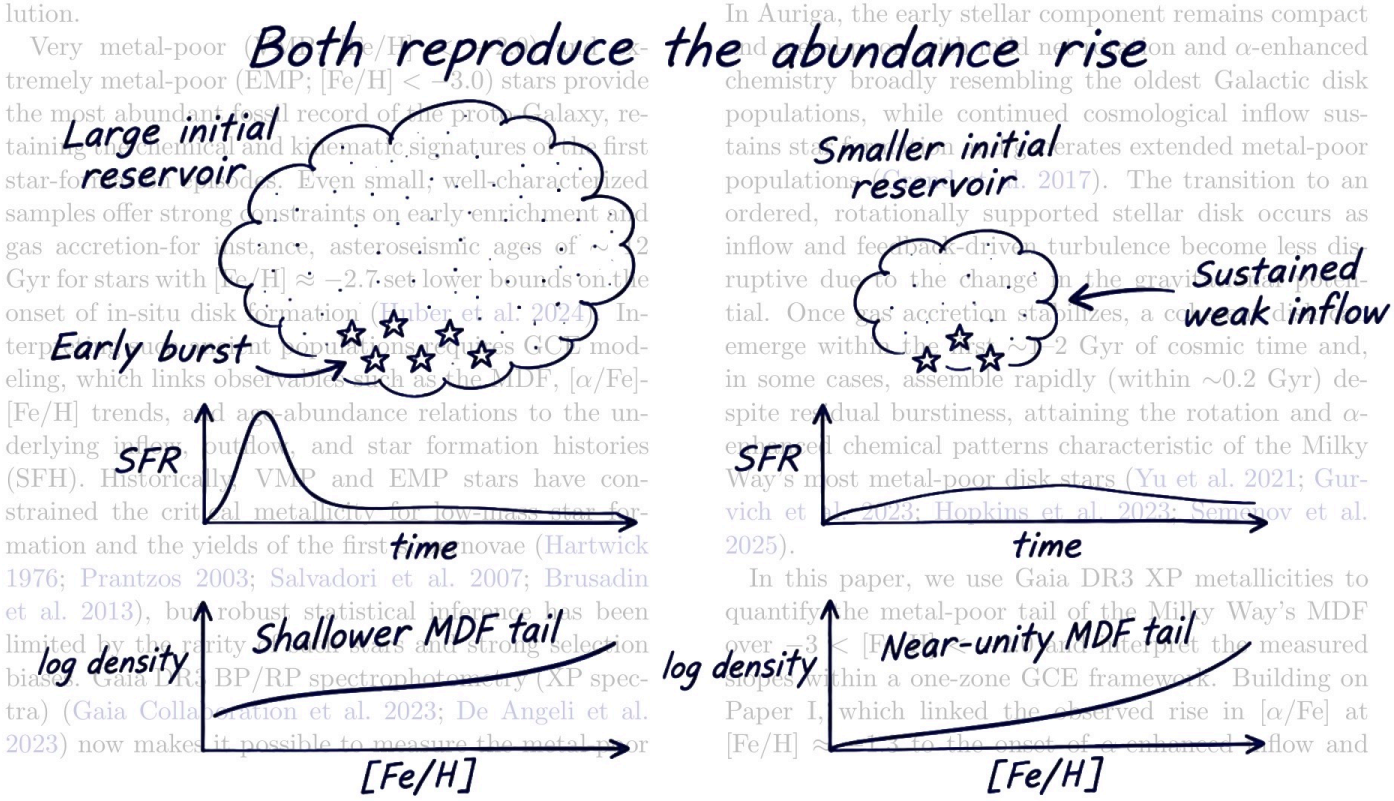
As the Galactic disk began to form, a moderate rise in [α/Fe] appeared near [Fe/H] ≈ -1.3, even among stars on low-eccentricity orbits (Conroy et al. 2022; Belokurov & Kravtsov 2022; Rix et al. 2022). This feature marks the transition from the turbulent, kinematically hot Aurora phase to the onset of ordered, rotationally supported star formation and has been interpreted as the chemical signature of the disk's birth. In Paper I (Chen et al. 2024), we showed that the observed [α/Fe] upturn is best explained by the sudden onset of α-enhanced gas inflow-pre-enriched by core-collapse supernovae in the halo-which supplied both angular momentum and fuel for renewed star formation after a preceding inflow-poor phase. The resulting enhancement thus reflects the coupled chemical and dynamical effects of the accreted material. However, the [α/Fe]-rise alone cannot uniquely determine the proto-Galaxy's initial gas mass, star-formation efficiency, or inflow timescale, as multiple parameter combinations yield similar patterns. To break this degeneracy, we now turn to the slope of the metal-poor tail of the metallicity distribution function (MDF), which traces the balance between gas supply and metal production during the Galaxy's earliest evolution.

Very metal-poor (VMP; [Fe/H] < -2.0) and extremely metal-poor (EMP; [Fe/H] < -3.0) stars provide the most abundant fossil record of the proto-Galaxy, retaining the chemical and kinematic signatures of the first star-forming episodes. Even small, well-characterized samples offer strong constraints on early enrichment and gas accretion-for instance, asteroseismic ages of $\sim$12 Gyr for stars with [Fe/H] ≈ -2.7 set lower bounds on the onset of in-situ disk formation (Huber et al. 2024). Interpretation of these faint populations requires GCE modeling, which links observables such as the MDF, [α/Fe]-[Fe/H] trends, and age-abundance relations to the underlying inflow, outflow, and star formation histories (SFH). Historically, VMP and EMP stars have constrained the critical metallicity for low-mass star formation and the yields of the first supernovae (Hartwick 1976; Prantzos 2003; Salvadori et al. 2007; Brusadin et al. 2013), but robust statistical inference has been limited by the rarity of these stars and strong selection biases. Gaia DR3 BP/RP spectrophotometry (XP spectra) (Gaia Collaboration et al. 2023; De Angeli et al. 2023) now makes it possible to measure the metal-poor tail of the Galactic MDF with homogeneous, all-sky stellar-parameter catalogs far larger than any previous spectroscopic sample.

Modern cosmological simulations provide a physically grounded framework for interpreting observational constraints on early disk assembly. The FIRE Project (Hopkins et al. 2014) is a suite of high-resolution zoom simulations that resolves stellar feedback processes and their impact on gas regulation. The Auriga simulations (Grand et al. 2017) follow Milky Way–mass galaxies in a ΛCDM context with detailed chemo-dynamical evolution and realistic disk formation. IllustrisTNG (Pillepich et al. 2019) provides large-volume statistics linking gas accretion, enrichment, and stellar mass growth across galaxy populations. Across these frameworks, Milky Way–mass progenitors at $z \gtrsim 2$ are predicted to experience bursty star formation fueled by cold, dense inflows (Dekel & Birnboim 2006; Muratov et al. 2015; Stern et al. 2021), with feedback-driven outflows regulating early enrichment and driving strong time variability in the star-formation history (Hopkins et al. 2014).

Although early star formation is centrally concentrated, repeated burst–feedback cycles can dynamically heat stellar orbits and disperse the oldest populations into spheroidal, halo-like configurations (El-Badry et al. 2016, 2018). The resulting proto-galaxies are typically built by one or two dominant progenitors plus a small number of lower-mass companions (Horta et al. 2024). In Auriga, the early stellar component remains compact and [illegible] metal-poor tail of in-situ and α-enhanced chemistry broadly resembling the oldest Galactic disk populations, while continued cosmological inflow sustains star formation and generates extended metal-poor populations (Grand et al. 2017). The transition to an ordered, rotationally supported stellar disk occurs as inflow and feedback-driven turbulence become less disruptive due to the change in the gravitational potential. Once gas accretion stabilizes, a coherent disk can emerge within the first $\sim$1–2 Gyr of cosmic time and, in some cases, assemble rapidly (within $\sim$0.2 Gyr) despite residual burstiness, attaining the rotation and α-enhanced chemical patterns characteristic of the Milky Way's most metal-poor disk stars (Yu et al. 2021; Gurvich et al. 2023; Hopkins et al. 2023; Semenov et al. 2025).

In this paper, we use Gaia DR3 XP metallicities to quantify the metal-poor tail of the Milky Way's MDF over $-3 <$ [Fe/H] < -1 and interpret the measured slopes within a one-zone GCE framework. Building on Paper I, which linked the observed rise in [α/Fe] at [Fe/H] ≈ -1.3 to the onset of α-enhanced inflow and



Low-resolution XP spectra turn rare-star searches into a population measurement.

disk spin-up, we now constrain the proto-Galaxy's parameters using a complementary diagnostic: the logarithmic slope of the metal-poor MDF tail. This approach leverages Gaia XP's homogeneous, Galaxy-wide sampling of VMP stars to test whether the Milky Way's earliest phase resembled the burst-regulated, gas-rich progenitors predicted by cosmological simulations. Together, these analyses aim to provide a quantitative, physically motivated description of the conditions under which the Galactic disk first emerged.

Broad spectral information carries metallicity clues for millions of stars.

This paper is organized as follows. Section 2 summarizes the Gaia DR3 XP metallicity catalogs used in this work and the adopted quality cuts, and presents basic CMD and spatial/kinematic diagnostics to characterize their selection footprints and potential contamination. Section 3 describes the GCE model grid (introduced in Paper I) and the procedure for predicting and measuring metal-poor MDF tail slopes. In Section 4 we measure the logarithmic MDF slopes over $-3 < [\mathrm{Fe/H}] < -2.0$ across catalogs, test their robustness, and interpret the inferred k values in terms of the proto-Galaxy's parameters. We then compare the observed slopes to Milky Way analogs in the Auriga simulations. Section 5 discusses systematic uncertainties and the connection to prior halo MDF and GCE studies. Section 6 summarizes our conclusions.

2. GAIA XP METALLICITY CATALOGS

2.1. *Before Gaia*

Before Gaia DR3, assembling large and homogeneous samples of very metal-poor stars relied on a sequence of survey strategies that progressively expanded both sky coverage and statistical accuracy. Early objective-prism work demonstrated that the Ca II H&K resonance lines provide a robust metallicity diagnostic at low spectral resolution; in particular, Bessell & Norris (1984) identified a star with $[\mathrm{Fe/H}] = -4.5$, establishing the feasibility of systematically probing the metal-poor tail. Subsequent discoveries pushed the metallicity frontier to the ultra-metal-poor regime ($[\mathrm{Fe/H}] < -4$) and revealed surprising findings on their elemental abundances (Christlieb et al. 2002; Frebel et al. 2005; Caffau et al. 2011).

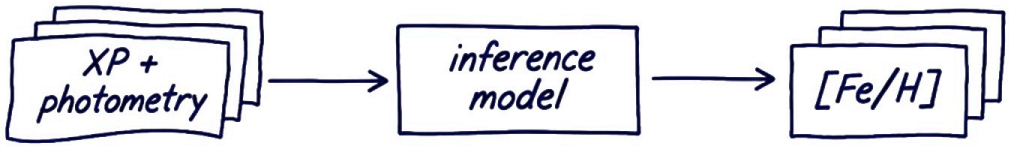
This approach was scaled to large-area surveys such as the HK survey (Beers et al. 1985, 1992) and the Hamburg/ESO (HES) survey (Wisotzki et al. 2000; Christlieb et al. 2008), transforming isolated discoveries into statistically meaningful halo samples. Medium-resolution spectroscopic campaigns, including SEGUE (Yanny et al. 2009) and RAVE (Smith et al. 2007; Matijevič et al. 2017), further expanded the census of metal-poor stars across wide sky areas. Complementary photometric programs such as SkyMapper (Keller et al. 2007) and Pristine (Starkenburg et al. 2017) employed metallicity-sensitive filters for efficient candidate pre-selection, while LAMOST (Zhao et al. 2012) provided extensive low-resolution spectroscopy from which metal-poor targets could be drawn. Compiled databases such as SAGA (Suda et al. 2008) and JINAbase (Abohalima & Frebel 2018) aggregated high-resolution abundance measurements, though with heterogeneous selection functions. Together, these efforts established the principal methodologies for identifying metal-poor stars—low-resolution spectral indicators of line blanketing, photometric metallicity proxies, and broad spectroscopic classification—and they form the historical foundation against which the homogeneous Gaia XP metallicity measurements can be interpreted.

2.2. *Training Methods for Gaia Catalogs*

The Gaia XP spectra, released in DR3, provide low-resolution ($R \sim 20$–100) spectrophotometry from the BP and RP instruments covering 330–1050 nm for sources down to $G \approx 17.65$ (De Angeli et al. 2023). Although the resolving power is significantly lower than that of most spectroscopic surveys, the all-sky coverage, homogeneous calibration, and continuous spectral sampling encode metallicity information for tens of millions of stars, including very metal-poor populations. XP shares methodological features with earlier objective-prism and photometric surveys—both rely on broad-band spectral energy distributions and weak metal-line signatures—but with the advantage of contiguous spectrophotometric data and a uniform instrument response.

We provide a brief summary of the various techniques used to derive metallicities from Gaia XP spectra here. The Gaia DR3 pipeline serves as the baseline of analyzing the XP spectra (see GaiaXPy[6] for more details). GSP-Phot fits XP spectra plus parallax and G to synthetic libraries in a Bayesian forward model, producing $T_{\rm eff}$, $\log g$, and a calibrated metallicity for ~ 471 M sources (useful mostly qualitatively for [M/H] per their own validation) (Andrae et al. 2023b). The signal-to-noise (S/N) ratio of the XP coefficients is a key factor in the reliability of these estimates. Building on that, Andrae et al. (2023a) trained gradient-boosted trees on APOGEE DR17 labels augmented with a curated very metal-poor set, and fed the model XP, Gaia, broad-band, narrowband (including synthetic Pristine CaHK) and CatWISE photometry (Marocco et al. 2021) as features; they predict $T_{\rm eff}$, $\log g$, [M/H] for ~ 175 M stars

[6] https://www.cosmos.esa.int/web/gaia/gaiaxpy



and also provide a carefully vetted bright RGB subset ($G \lesssim 16$) where the metallicity is most reliable.

SHBoost (Khalatyan et al. 2024) targets essentially "all stars with XP and parallaxes": it learns $T_{\rm eff}$, $\log g$, [M/H], A_V, and mass with XGBoost trained on StarHorse labels compiled from many spectroscopic surveys (plus curated hot/VMP/WD sub-samples) (Queiroz et al. 2018) using XP coefficients together with Gaia+2MASS(Skrutskie et al. 2006)+WISE(Wright et al. 2010); the team also applies a posteriori calibration for metallicity systematics. Zhang et al. (2023) go the other way—forward-modelling XP and broadband photometry with a data-driven model trained on LAMOST; they provide $T_{\rm eff}$, $\log g$, [Fe/H] (and revised distances/extinctions) for ~ 220 M stars, i.e., an all-sky, all-types sample anchored to LAMOST labels.

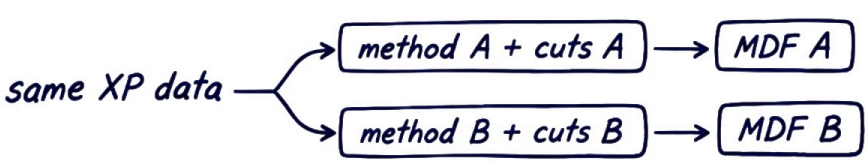
Several recent XP-based analyses apply more targeted strategies, either by improving the XP flux calibration before modeling or by restricting the training domain to regimes where the abundance information is strongest. Ye et al. (2025) first infer per-star flux corrections with an APOGEE-trained neural network to suppress small-scale XP wiggles and then fit the corrected spectra with `FERRE` against synthetic grids, producing a 68 M-star catalog of $T_{\rm eff}$, $\log g$, and [M/H] with enhanced control of systematics across the sky. AspGap (Li et al. 2024) focuses on red-giant branch stars and employs a neural network trained on APOGEE that includes an auxiliary component designed to reconstruct the corresponding high-resolution APOGEE spectrum. This additional reconstruction objective strengthens the link between XP features and abundance-sensitive spectral structure, yielding $T_{\rm eff}$, $\log g$, [M/H], and precise [α/M] for ~ 37 M RGB stars.

Hattori (2025) adopt a complementary strategy optimized for low-extinction sightlines. They apply explicit cuts in Galactic latitude and $E(B-V)$, impose a $G < 18.85$ limit, and train boosted trees on APOGEE DR17 augmented with a curated very metal-poor sample, publishing [M/H] and [α/M] for ~ 48 M stars in these dust-minimal windows. Yang et al. (2025) emphasize giants as well, dereddening and correcting XP spectra before using an uncertainty-aware, cost-sensitive neural network (UA-CSNet) trained on high-quality spectroscopic labels (PASTEL) to predict [Fe/H] for ~ 20 M giants with particular attention to the metal-poor tail.

Two efforts push specific elemental abundances beyond metallicity. Ardern-Arentsen et al. (2025) train a neural network on spectroscopic reference labels to infer both [Fe/H] and [C/Fe] from XP (plus auxiliary Gaia information) for > 10 M stars, then identify $\sim$ 2,000 bright carbon-enhanced metal-poor (CEMP) candidates; their tables link back to vetted giant selections to ensure reliability at the metal-poor end. Kane et al. (2025) focus on abundance patterns associated with "second-generation" globular cluster (GC) populations. They employ a multilayer perceptron (MLP) neural network trained on APOGEE cross-matches to predict $T_{\rm eff}$, $\log g$, [Fe/H], [N/O], and [Al/Fe] directly from normalized XP coefficients. Applied to the Andrae-vetted RGB sample, these predictions allow them to identify high-[N/O] field-star candidates that may trace a GC-origin channel. Finally, one catalog uses XP to synthesize a metallicity-sensitive narrow band: the Pristine DR1-XP release computes synthetic CaHK magnitudes for ~ 219 M stars and, combined with broadband Gaia (and deeper Pristine where available), publishes photometric metallicities for > 30 M high-S/N FGK stars, with strong performance below [Fe/H] ≈ -1 and extending to $\lesssim -3.5$ (Martin et al. 2024).

Most XP metallicity catalogs are supervised and anchored to big spectroscopic surveys (APOGEE, LAMOST, StarHorse), but they differ in who they target (all-sky vs. giants vs. low-extinction), how they treat the XP fluxes (raw vs. corrected vs. forward-modelled), and which labels they predict (some give [M/H], some [Fe/H], many now add [α/M] or [C/Fe]). These choices set the reliability envelope depending on the specific user case. We compile in Table 1 a summary of the major public catalogs derived from Gaia XP spectra mentioned above. The table lists the size of each release, the spectroscopic or synthetic training sets, the modeling framework (e.g., forward models, tree-based regressors, or neural networks), and the authors' recommended quality cuts that constrain each catalog to its validated regime which we adopt for our analysis. This overview is intended to clarify the heterogeneity among XP analyses and to guide the selection of appropriate subsamples for metal-poor studies.

Because the Gaia XP metallicity catalogs do not share a common selection function, we use the published metallicities primarily for differential analyses. We do not attempt to construct a volume-complete sample or combine all catalogs into a single homogenized dataset, but instead treat each catalog within its published quality-controlled regime. Several potential biases are relevant. First, metal-poor giants are intrinsically brighter than dwarfs, so magnitude-limited samples can over-represent giants and probe different Galactic volumes and component mixtures depending on the CMD coverage of each catalog. Second, because the Gaia XP samples span both disk and halo sightlines, contamination by metal-rich disk stars with underestimated metallicities may affect the metal-richer boundary of the in-



Compare each catalog within its validated regime; these samples are not a single volume-complete survey.

Gas supply, cooling, and star formation jointly shape the chemical history.

ferred metal-poor tail. Since these effects depend on the training set, inference domain, and recommended quality cuts of each public release, they must be evaluated empirically when interpreting the MDFs (see Section 4.2).

3. GALACTIC CHEMICAL EVOLUTION MODEL

We use the same 735 GCE models that successfully replicated the moderate rise in [Mg/Fe] as seen in the H3 survey (Conroy et al. 2019) from our Paper I. Our model framework is designed to investigate the early [α/Fe] behavior in the proto-Milky Way, simulating the first two Gyr of its evolution within a bin with 1 kpc radius. The model divides this period into two one-Gyr phases, with time steps of 30 Myr, reflecting the longest core-collapse supernovae (CCSNe) progenitor lifespans. Stellar populations, ranging from 0.5 to 40 $M_\odot$ and following an initial mass function (IMF) of $\xi(m) \propto m^{-2.3}$, are tracked alongside cold and warm interstellar medium (ISM) components. Nucleosynthetic yields from Nomoto et al. (2013) are interpolated across metallicities to determine the chemical enrichment from stars. Stars between 0.5 and 0.9 $M_\odot$ are formed and store cold ISM, but they don't contribute to nucleosynthesis due to the lack of nucleosynthesis yields.

The model incorporates key astrophysical processes. Star formation is governed by the Kennicutt-Schmidt law, converting cold ISM into new stellar populations. Stellar evolution dictates the return of enriched gas: CCSNe from stars more massive than 9 $M_\odot$ explode within one time step (30 Myr), while intermediate-mass stars (3–8 $M_\odot$) evolve into white dwarfs (WDs) that can later produce Type Ia supernovae (SNe Ia). The SNe Ia rate depends on the WD mass reservoir, a constant fraction ($f_{\rm SNIa}$), and an exponential delay timescale. Enriched material ejected from stars is distributed primarily to the warm ISM (79%) and a smaller fraction to the cold ISM (1%), with rest lost to the circum-galactic medium (CGM). The warm ISM cools to replenish the cold ISM, while feedback from star formation and supernovae heats a portion of the cold ISM into the warm ISM, regulated by mass-loading factors, $\eta_{\rm SF} = 2.0$ and $\eta_{\rm SN} = 5.0$. The outflow mass-loading factors are fixed and directly proportional to the mass involved in star formation and stellar evolution activities. The model also allows for primordial gas entry into the warm ISM.

To understand the conditions leading to an early [α/Fe] rise, we explore a range of five fundamental free parameters: the initial cold ISM mass ($m_{0,\rm cold}$), the SNe Ia fraction ($f_{\rm SNIa}$), the warm ISM cooling timescale ($\tau_{\rm cool}$), the star formation efficiency ($\tau_{\rm SF}$), and the gas inflow rate ($\eta_{\rm inflow}$). These parameters are varied across astrophysically motivated ranges in numerous model runs. Notably, to achieve an [α/Fe] rise after an initial decline (due to SNe Ia), the star formation efficiency and inflow rate, along with inflow gas composition, are allowed to change in the second Gyr of the simulation, facilitating increased α-element production through enhanced CCSNe activity or accretion of α-enriched gas.

4. RESULTS

4.1. *Zero Rotation Expected for the Proto-Galaxy*

These models were already selected to reproduce the early α/Fe rise.

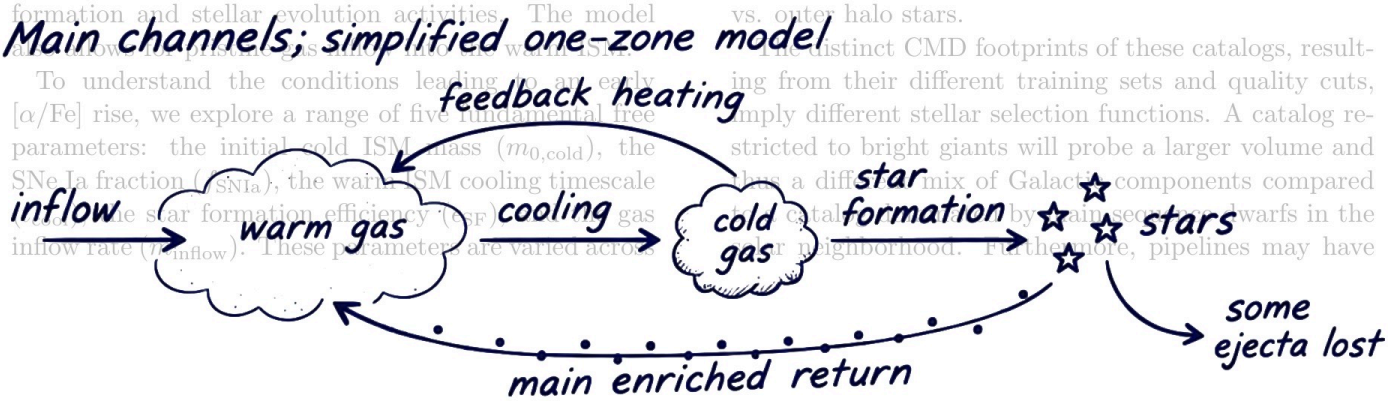
Figure 1 shows the stellar density distribution in the Gaia color-magnitude diagram (CMD) for various Gaia XP spectroscopic metallicity catalogs. Panels (a) use apparent magnitudes G and BP-RP colors, while the panels (b) use observed absolute magnitudes $G_{\rm obs}$ derived from parallax without extinction or reddening corrections.

$$G_{\rm obs} \equiv G + 5\log_{10}(\varpi_{\rm mas}) - 10, \quad (1)$$

Panels a) provide a visual representation of the quality cuts per the authors' recommendations listed in the last column of Table 1, while panels (b) reduce distance-driven smearing and makes the CMDs more comparable across catalogs, at the cost of relying on parallax quality. We do not apply an additional dereddening correction because the retained samples have already been filtered into relatively low- or moderate-extinction regimes by the adopted catalog-specific quality cuts ($E(B-V) < 0.1$–0.5), and several pipelines explicitly account for extinction or operate on corrected XP spectra (see Table 1).

XP spectra exist only for sources with $G < 17.65$. In addition, many catalogs apply explicit or implicit cuts that confine their inference domains to their training samples (e.g., RGB/RC giants or specific $T_{\rm eff}$-$\log g$ ranges). These choices produce visible differences in CMD occupancy across catalogs: giant-only catalogs emphasize the upper RGB and red clump stars, whereas catalogs trained on mixed luminosity classes populate the main-sequence turnoff (MSTO) and subgiant branch as well. The footprints of the catalogs on the CMD could potentially bias the metallicity distribution in two ways: 1) At fixed age and evolutionary phase, a more metal-poor star would be brighter; 2) the magnitude limits changes the spatial coverage and the mixture of inner vs. outer halo stars.

The distinct CMD footprints of these catalogs, resulting from their different training sets and quality cuts, imply different stellar selection functions. A catalog restricted to bright giants will probe a larger volume and have a different mix of Galactic components compared to a catalog dominated by main-sequence and dwarfs in the solar neighborhood. Furthermore, pipelines may have



The table counts selected metal-poor stars, not each full catalog release.

Table 1. Summary of Stellar Parameter Catalogs from Gaia XP Spectra

Reference	Number of Stars[a]	Training Data	Method	Adopted Recommended Quality Cuts
Andrae et al. (2023)	222,572	APOGEE DR17 + VMP stars (Li et al. 2022)	XGBoost: Gradient-boosted trees with weighted training.	Use the provided bright RGB sample ($G < 16$) vetted for high precision and purity.
Zhang et al. (2023)	3,326,913	LAMOST DR8 + "Hot Payne"	Empirical Forward Model: NN predicts spectra from labels, which are then inferred.	Use the `quality_flags` column (< 8), incorporating cuts on χ^2, prior probability, and XP–parallax consistency. We further restrict to stars within the training-set convex hull: $3500 < T_{\rm eff} < 8500$, $-0.5 < \log g < 5$.
Li et al. (2023)	155,145	APOGEE DR17	Neural Network (AspGap): Includes a "hallucinator" component.	Follow the recommended red-giant-only sample ($\sim$ 12 million stars). We also adopt $\sigma_{\rm [M/H]} < 0.15$ dex.
Martin et al. (2024)	1,785,328	Spectroscopic sample (SDSS, APOGEE) + Synthetic spectra (MARCS)	Photometric Grid Model: Calibrated color–metallicity grid using synthetic CaHK magnitudes.	Use flag `mcfrac > 0.8` to ensure Ca H&K magnitude is computed from valid Gaia XP coverage, `Pvar < 0.3`, `RUWE < 1.4`.
Fallows & Sanders (2024)	233,605	APOGEE DR17	Uncertain Neural Network: NN with explicit modeling of multiple uncertainty components.	No separate quality flag; only sources with $4000 < T_{\rm eff} < 7000$ K and $0 < \log g < 5$ appear in the released catalog.
Hattori (2024)	1,447,280	APOGEE DR17 + VMP stars	Quantile Regression Forests (QRF): Tree-based ensemble predicting full probability distributions.	Restrict to low-extinction sources with $E(B-V) < 0.1$. Use the `bool_flag_cmd_good` flag to select stars within the trained CMD domain.
Khalatyan et al. (2024)	279,126	StarHorse (Q23; compilation of many spectroscopic surveys)	XGBoost (SHBoost): Uses `xgboost-distribution` for uncertainty estimation.	Use the output quality flag (`xgb_met_outputflag = 0`) and magnitude cut ($G \leq 16.5$).
Kane et al. (2024)	80,696	APOGEE DR17	Neural Network (MLP): Heteroscedastic regression to classify GC-origin stars.	No specific quality cuts for the released catalog. Published sources satisfy S/N > 30, $4000 < T_{\rm eff} < 6500$ K, and $> 99\%$ GC membership.
Ardern-Arentsen et al. (2025)	92,352	LAMOST spectra re-analysed with FERRE + high-resolution VMP/CEMP literature; XP coefficients as inputs.	Neural Network (heteroscedastic): ANN on normalized XP coefficients; outputs $T_{\rm eff}$, $\log g$, [Fe/H], [C/Fe] with per-star variances; ensemble-averaged predictions.	Recommended "clean bright" sample: $G < 16$, $E(B-V) < 0.3$, Pvar < 0.5, $(\mathrm{BP-RP})_0 < 1.35$.
Ye et al. (2025)	712,821	APOGEE DR17 (for flux correction model)	Hybrid: NN for spectral correction, then χ^2 fitting to synthetic spectra.	Use the `dflux_per` flag (< 0.08). Recommended temperature range is 4000 K $\leq T_{\rm eff} \leq$ 7000 K. We restrict to $\log \chi^2 < 4$ quality cuts for precise metallicities.
Yang et al. (2025)	1,888,361	PASTEL catalog	Uncertainty-aware Cost-Sensitive Neural Network (UA-CSNet).	The reliable sample uses $E(B-V) < 0.5$. The model is for giants with $0.7 < (\mathrm{BP-RP})_0 < 1.9$. Additional CMD cut with $M_G < 3.5 + 2.5(\mathrm{BP-RP})$.

[a] With [Fe/H] < -1.3 after the adopted quality cuts.

Counts apply after quality cuts and [Fe/H] < −1.3.

Catalog choices leave visible footprints in the color-magnitude diagram.

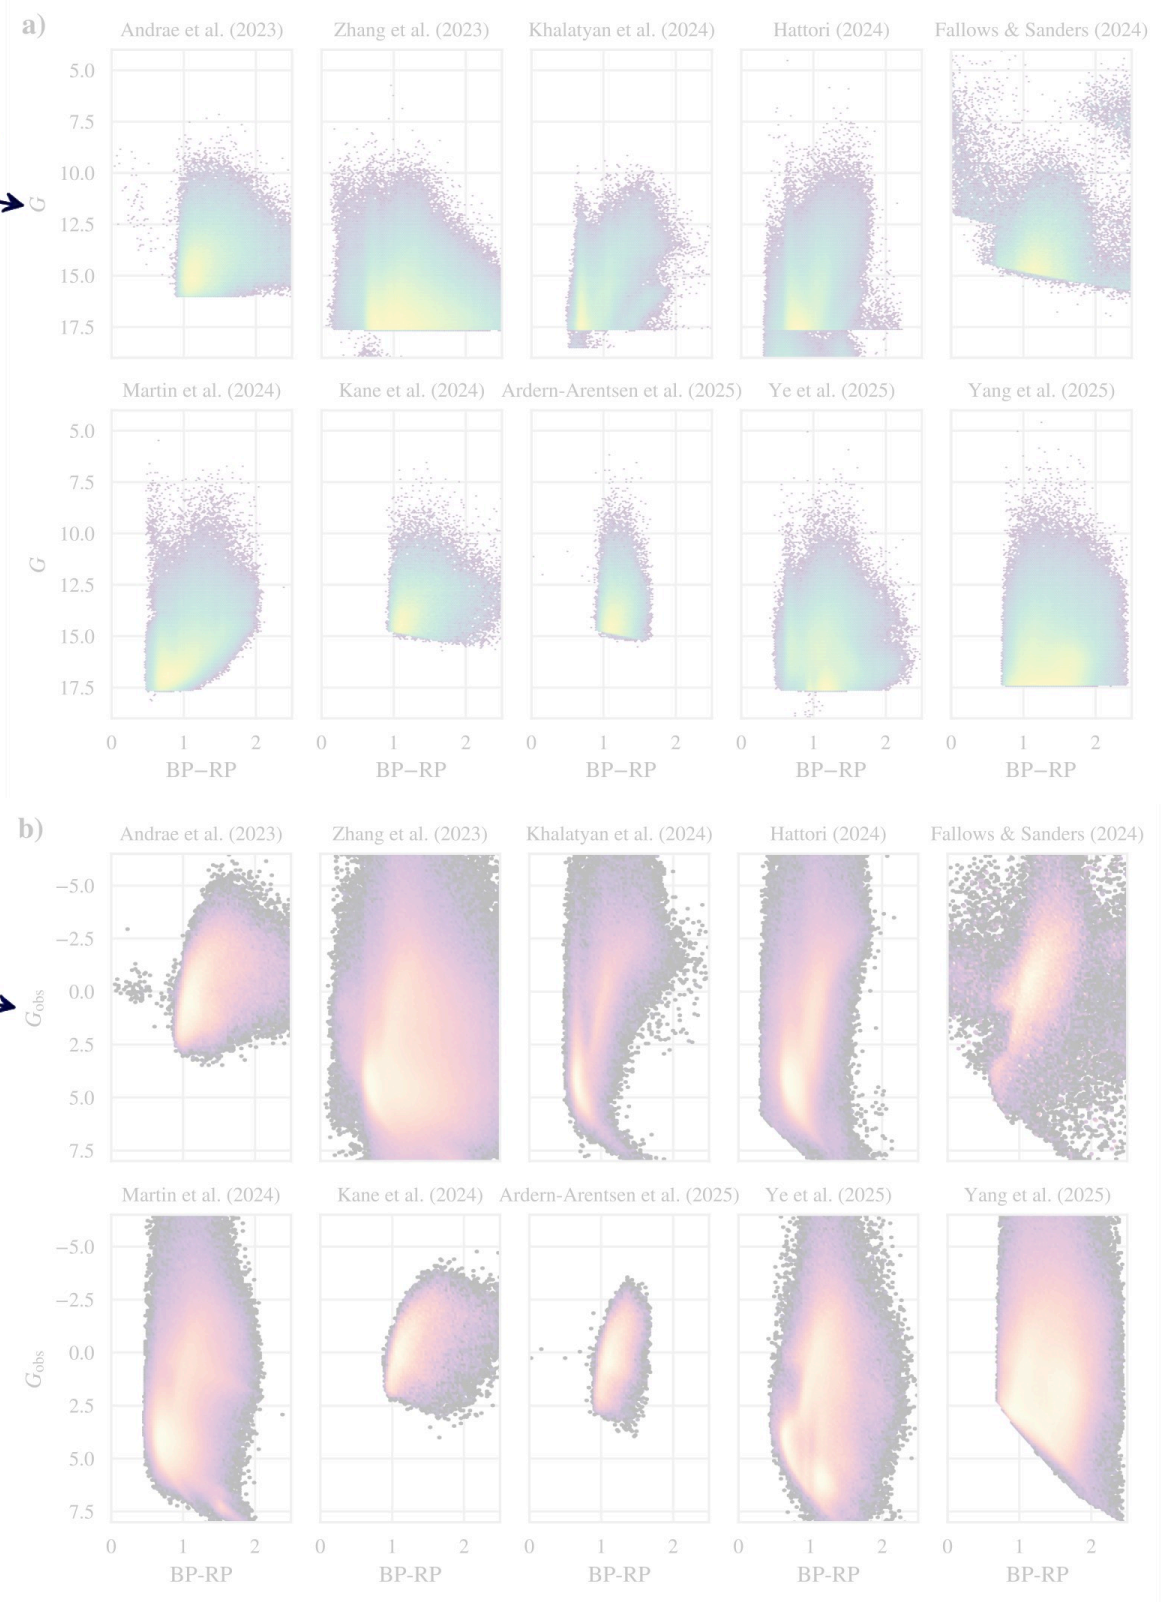


Figure 1. Stellar density distributions in the Gaia CMD. Panels in (a) display densities based on apparent magnitude, while panels in (b) show densities using observed absolute magnitudes derived from Gaia parallaxes. Data are sourced from various Gaia XP metallicity catalogs following their quality cuts listed in Table 1. The quality cuts restrict several catalogs to mostly giants due to training pipelines and output quality control.

Observed absolute magnitudes use parallaxes here, without an added extinction correction.

Radial velocities check a subset; they do not select the full MDF sample.

Only about 5–10% have the radial velocities used in this diagnostic.

varying completeness and contamination rates; for instance, a model trained primarily on red giants might misclassify metal-rich dwarfs as metal-poor if their colors are similar, or fail to recover genuine metal-poor stars in sparsely sampled regions of the CMD. These effects can alter the apparent normalization and shape of the MDF by changing the relative contributions of different Galactic components and by introducing metallicity-dependent systematic errors.

To assess possible contamination, we derived Galactocentric positions and velocities only for the subset of stars with published Gaia DR3 RVS radial velocities, which constitutes approximately 5–10% of each Gaia XP metallicity catalog after the adopted quality cuts. These kinematic data are used solely as an external diagnostic of the metal-poor population and were not used to define the MDF samples or to impose kinematic cuts on the full catalogs. Stars with [Fe/H]< -1.5 are expected to have near-zero mean rotation with few objects reaching disk-like azimuthal velocities $v_\phi \gtrsim 100$ km s^{-1} (Chiba & Beers 2000; Carollo et al. 2007; Deason et al. 2011; Myeong et al. 2018; Yuan et al. 2024). To test this expectation and diagnose catalog-dependent artifacts, we transformed the Gaia astrometry together with the available Gaia DR3 RVS radial velocities for this subset into Galactocentric coordinates (R, ϕ, z) and velocities (v_r, v_ϕ, v_z) using `astropy` (Astropy Collaboration et al. 2022). Because Gaia RVS radial velocities are available for only a minority of XP sources, these kinematic distributions should be interpreted as a validation subset rather than as a complete description of the parent metallicity catalogs.

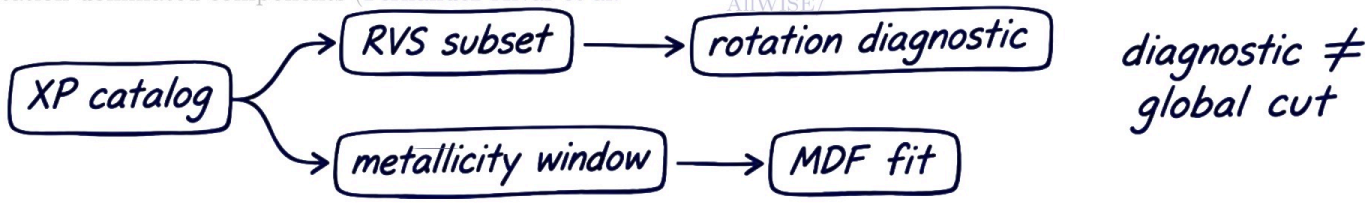
Figure 2 presents the two-dimensional stellar density distribution of v_ϕ versus [Fe/H] for $-3 \leq$ [Fe/H] < -1.5 in the metallicity catalogs. The running median v_ϕ is marked in each panel with a red solid line and the 16th and 84th percentile with dotted lines. The catalog by Kane et al. (2024) applied a cut of $v_\phi < 160$ km s^{-1}. In several catalogs, we observe a concentration of stars with $v_\phi \approx 220$ km s^{-1} as [Fe/H] approaches our boundary (typically $-1.8 \lesssim$ [Fe/H] < -1.5), most prominently in Zhang et al. (2023); Fallows & Sanders (2024); Martin et al. (2024); Yang et al. (2025). At the very metal-poor end with [Fe/H] < -2.5, the median v_ϕ differs across catalogs: in Fallows & Sanders (2024); Ye et al. (2025); Yang et al. (2025) it exceeds 50 km s^{-1}, while in the remaining catalogs it stays closer to 0. The high-v_ϕ features near the metallicity boundaries can arise from: (1) misclassification of metal-rich disk stars as metal-poor; (2) genuine metal-poor stars on fast, disk-like orbits associated with the metal-weak thick disk or ancient, rotation-dominated components (Fernández-Alvar et al. 2021; Bellazzini et al. 2024; Yuan et al. 2024; Xiang et al. 2025).

A physically motivated channel for genuinely fast, metal-poor stars could be an early, gas-rich merger that seeded the proto-disk's angular-momentum axis, producing a prograde, low-[Fe/H] tail (Santistevan et al. 2021). An alternative is that later mergers, such as the GSE, could have torqued pre-existing stellar populations, imparting net rotation (McCluskey et al. 2024). Whether the observed prograde motions are primordial or merger-induced, their kinematic coherence suggests a link to the early building blocks of the proto-Galaxy (Sestito et al. 2019; Carter et al. 2021). However, as shown in Fig. 2, the bulk of stars with [Fe/H] < -2.0 exhibit low v_ϕ, consistent with a dispersion-dominated proto-Galaxy. Because only $\sim$ 5–10% of stars in each catalog have Gaia RVS radial velocities, kinematic cuts alone cannot be applied to the entire catalogs. An alternative approach is to apply isochrone-based selections, which can be done with the available photometry and helps mitigate contamination. We restrict our MDF analysis to [Fe/H] < -2.0 to minimize the inclusion of disk stars. Since most stars already follow the observed isochrone for an old, metal-poor stellar population and most RVS stars have low rotation, we leave a detailed kinematic decomposition for future work.

Figure 3 shows the two-dimensional number density in Galactocentric radius R and height |z| for the XP catalogs we used within our metallicity range. Unlike many previous studies that targeted high-|z| regions, our working samples are overwhelmingly dominated by stars in the disk plane: most stars lie at |z| < 2 kpc because the underlying Gaia XP selections are magnitude-limited and largely all-sky. The vetted RGB catalog by Andrae et al. (2023a) and the catalogs by Kane et al. (2025) and Yang et al. (2025) target giants, which results in more even R coverage but naturally misses the dwarfs at low-|z|. The pipeline by Andrae et al. (2023a) require WISE $W1/W2$ and even use M_{W1} and $G - W$ cuts to purify giants. CatWISE was adopted precisely to improve completeness over AllWISE[7] in crowded fields. Consequently, some residual suppression in the dustiest mid-plane lines of sight is expected.

Ardern-Arentsen et al. (2024) applied their CEMP-focused network to the same vetted RGB sample so their R–|z| distribution follows the same footprint. Zhang et al. (2023) assumed a single extinction law and applied it across the full XP sample, which yields the strongest mid-plane concentration at a given R because

[7] https://irsa.ipac.caltech.edu/data/WISE/docs/release/AllWISE/



Rotation reveals suspicious populations near the metal-rich boundary.

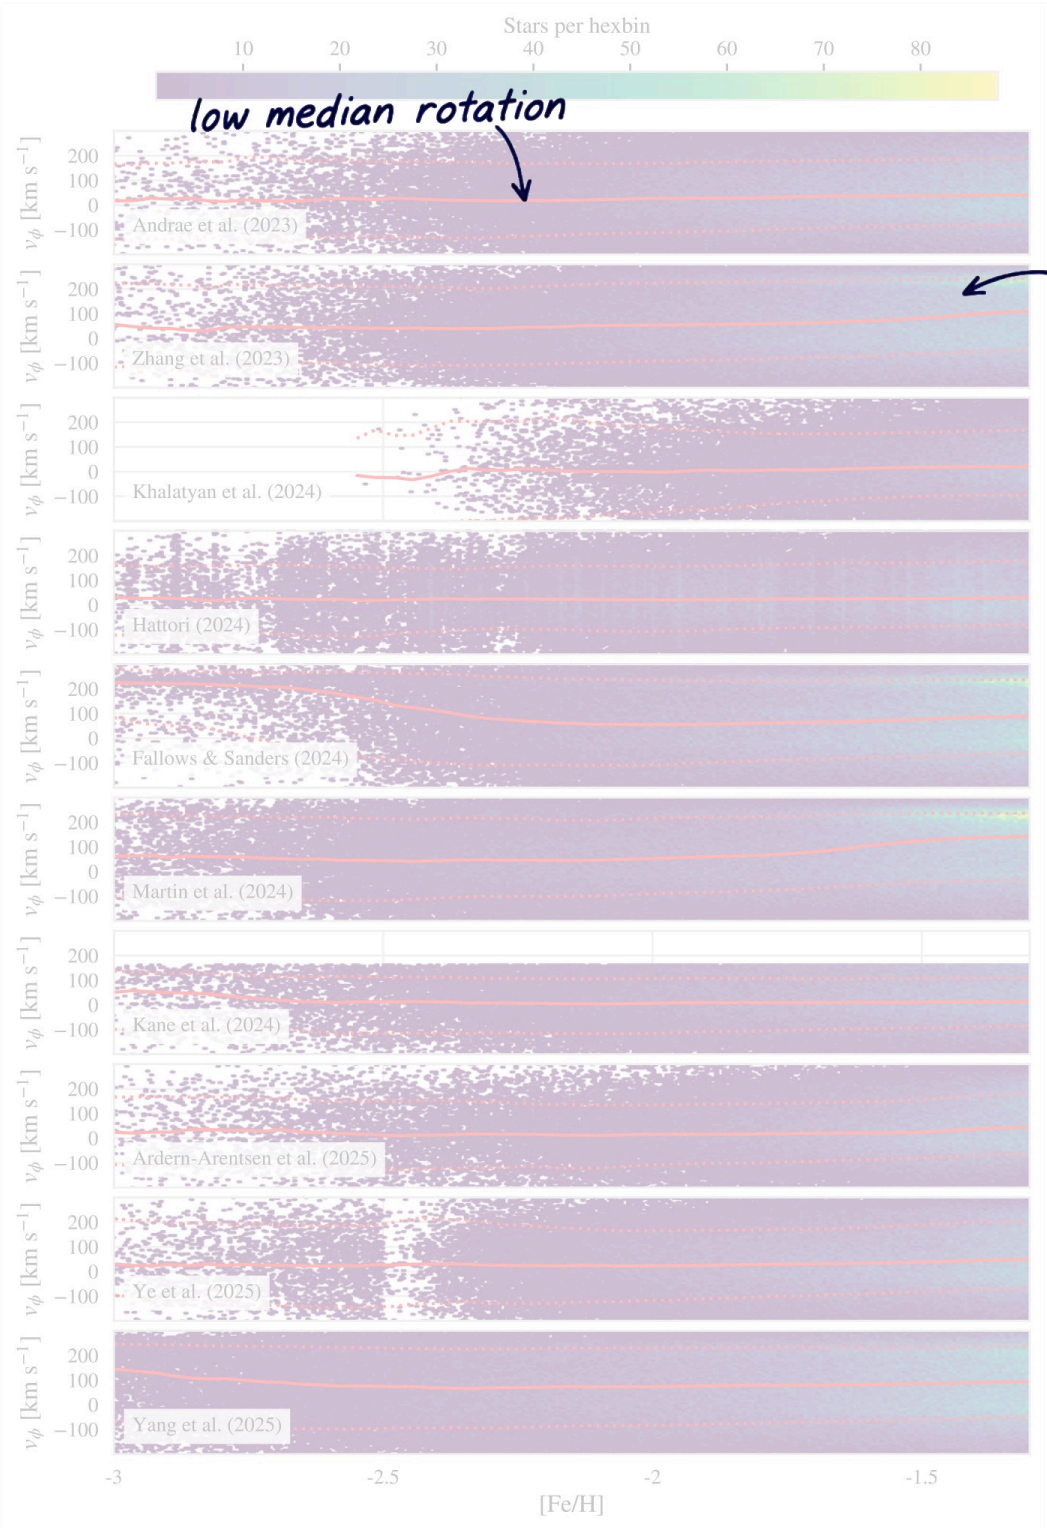


Figure 2. Stellar distribution in Galactocentric rotational velocity versus metallicity for various Gaia XP metallicity catalogs. The solid red curve shows the running median v_ϕ as a function of metallicity, computed in a sliding window of width 0.2 dex in [Fe/H]. The dotted lines show the 16th and 84th percentile range of v_ϕ respectively within the same window. Most metal-poor stars have low rotational velocities as expected for the proto-Galaxy, but false metal-poor stars that likely belong to the disk become common as [Fe/H] exceeds -2.0. Similar contamination also appears in Fallows & Sanders (2024) and Yang et al. (2025) as [Fe/H] drops below -2.6.

These panels show the RVS subset, not every star used for the MDF.

Different catalog footprints sample different Galactic volumes.

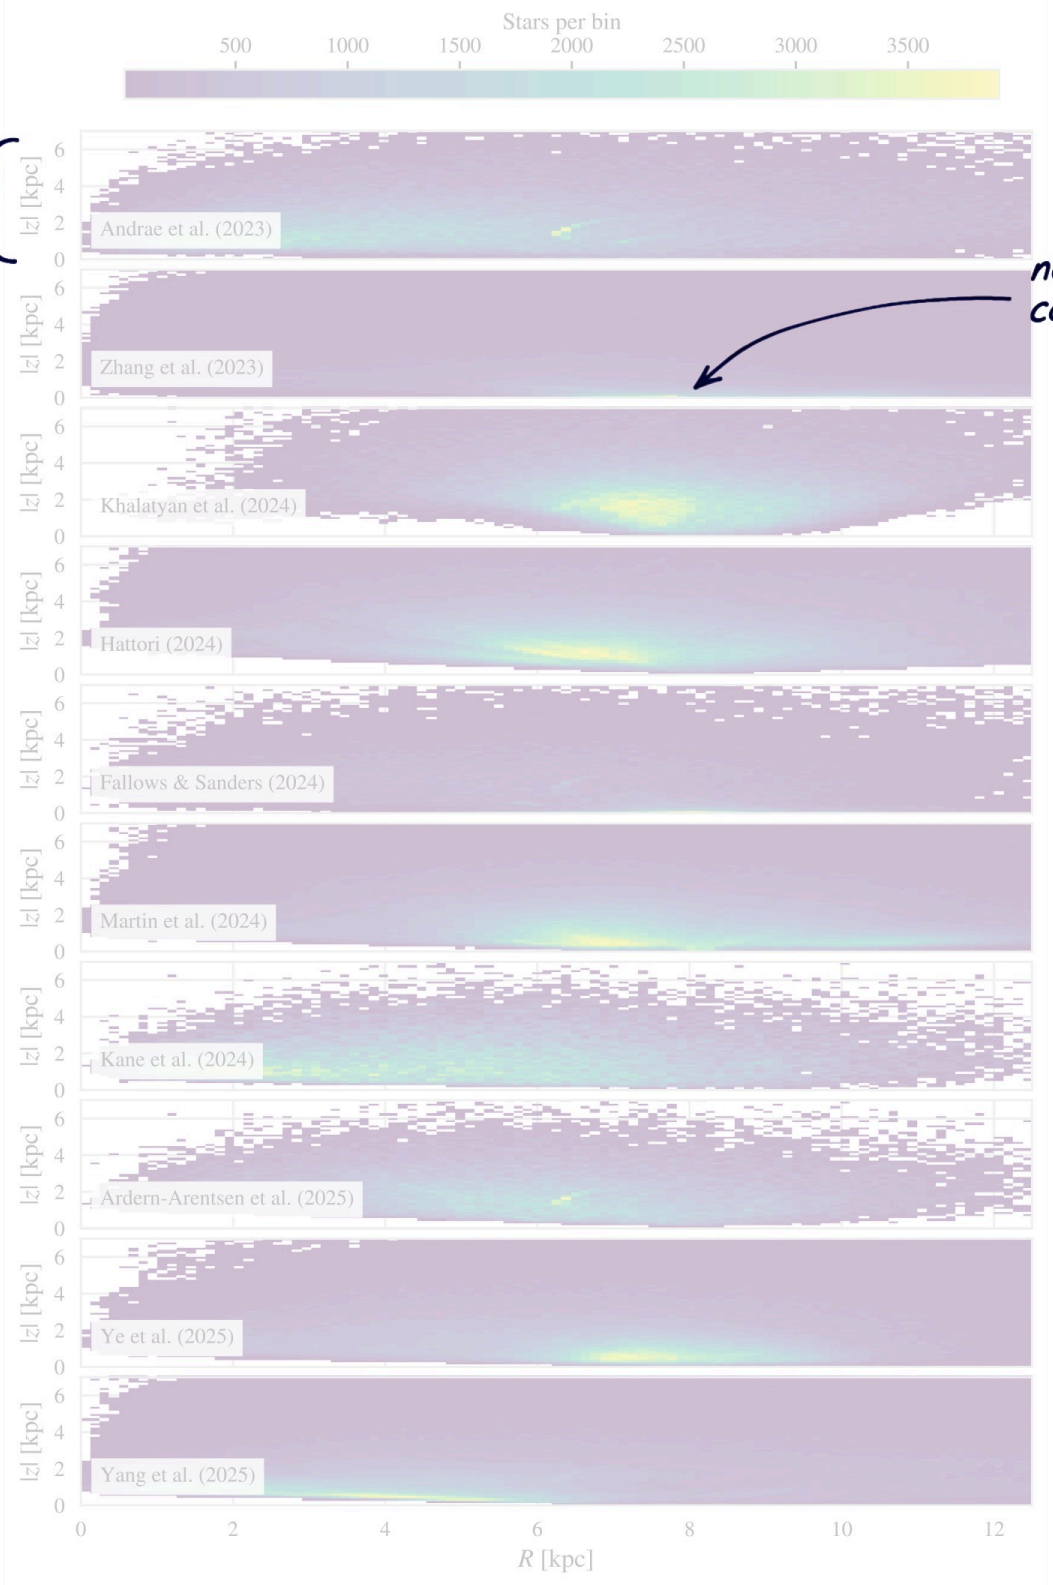


Figure 3. Stellar distribution in Galactocentric radius (R) versus height above the Galactic plane (|z|) for the Gaia XP metallicity catalogs. Differences in training sets and quality controls produce distinct spatial footprints. Catalogs restricted to giants extend to larger R and |z|, whereas others more closely follow the native XP selection function as being concentrated in the solar neighborhood. A notable outcome of Gaia XP is the recovery of large numbers of metal-poor stars at low |z|, an area that was sparsely sampled before Gaia XP.

Low height above the plane does not by itself make a star a thin-disk member.

The CMD helps locate where inferred metallicities may be unreliable.

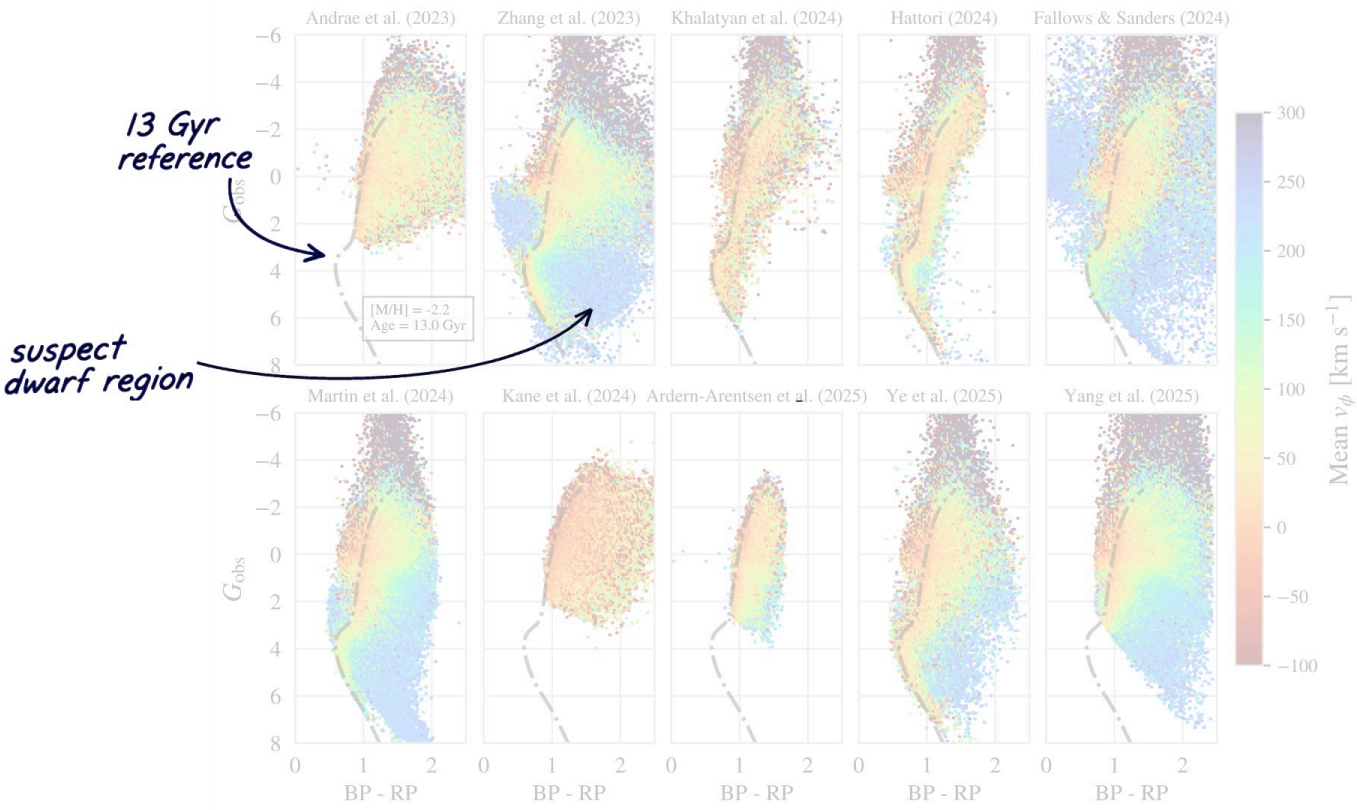


Figure 4. Gaia CMD colored by median Galactocentric rotational velocity (v_ϕ) in various Gaia XP metallicity catalogs. The magnitudes are observed absolute magnitudes as in b) in Fig. 1. The slow-rotating stars highlight a PARSEC 1.2S isochrone (Bressan et al. 2012) for a stellar population 13 Gyr old and with [M/H] = -2.2. Most stars with radial velocities in the various catalogs also tend to reside along such an isochrone on the CMD as seen in Figure 1. This suggests that at least most stars with low v_ϕ in Gaia RVS are genuinely old metal-poor stars.

it largely follows the raw Gaia density. Similarly, Fallows & Sanders (2024) included broad-band (2MASS/WISE) colors with XP and publish an all-sky catalog the mirrors XP density. SHBoost by Khalatyan et al. (2024) was applied to all Gaia XP sources but their quality flags assign the highest metallicity reliability to stars with low extinction. Hattori (2025) restricts to low-extinction sightlines, which removes the opaque mid-plane and accentuates higher-|z| at fixed R. The Pristine-Gaia synthetic catalog by Martin et al. (2024) iteratively corrects for extinction and was explicitly built to cover all Galactic components on the sky, which yields good plane coverage with only modest suppression in the dustiest lines of sight. Ye et al. (2025) cautioned that reliability for their metallicity measurements is the strongest for roughly $4000 \lesssim T_{\rm eff} \lesssim 7000$ K, effectively emphasizing the FGK-rich mid-plane region in practice.

Figure 4 maps median v_ϕ across the CMD (using $G_{\rm obs}$), localizing, via the colour scale spanning from -100 to 300 km s^{-1}, where prograde stars concentrate in evolutionary phase. Stars with median $|v_\phi| < 50$ km s^{-1} trace a PARSEC 1.2S isochrone for an old, metal-poor population ([M/H] $= -2.2$, Age $= 13$ Gyr), as expected, while stars that deviate from the isochrone—typically toward the blue turnoff and red subgiant/dwarf edge—show elevated v_ϕ over 200 km s^{-1}. These concentrations with high v_ϕ on the CMD coincide with regions where training set coverage is sparse because more than half the catalogs are trained on or anchored to APOGEE. For example, Fallows & Sanders (2024) trained their NNs on APOGEE data and Ye et al. (2025) used APOGEE to train their NNs for flux correction. Even broader "transfer" models like SHBoost by Khalatyan et al. (2024) still draw much of their label scale from spectroscopic surveys that under-sample parts of the CMD and thus inherit the coverage gaps. Therefore, the high-v_ϕ patches are most naturally interpreted as thin-disk contaminants whose metallicities are severely underestimated by the pipelines.

A similar trend is seen in independent photometric-metallicity work by Martin et al. (2024): CaHK-syn metallicity estimates degrade toward hotter stars and lose sensitivity for dwarfs. Substituting their default (giant-calibrated) metallicities with the compan-

The old isochrone is a population check, not an individual stellar age measurement.

Fit individual metallicities; the histogram is only a display.

ion dwarf-model values did not alter the predominance of high-v_ϕ dwarfs. The difference in [Fe/H] between the two models is generally 0.2–0.4 dex, which is small compared to our metallicity range. Catalogs that confine their labels to well-calibrated regions of the CMD generally avoid spurious high-v_ϕ patches. For example, Andrae et al. (2023a) released a vetted RGB-only catalog. Ardern-Arentsen et al. (2025) also selected their objects partially from this RGB sample. More broadly, a training-set-dependent CMD mask-such as the flag provided by Hattori (2025)-offers a practical way to refine samples and minimize contamination.

4.2. *Robust MDF Slope Across Catalogs and Selections*

To measure the slope of the metal-poor tail, we fit each catalog over the fixed interval $-3 \leq [\mathrm{Fe/H}] \leq -2$. Rather than fitting a straight line to binned histograms, we use the individual stellar metallicities directly. This avoids sensitivity to bin size and reduces the impact of noisy low-count bins.

We model the MDF tail as an exponential function,

$$p(x) \propto \exp(\beta x), \qquad x \equiv [\mathrm{Fe/H}], \tag{2}$$

which is equivalent to a straight line in logarithmic space,

$$\log_{10} p(x) = kx + c, \qquad k = \beta / \ln 10. \tag{3}$$

slope in log density

We determine k with an unbinned maximum-likelihood fit using only stars in the chosen metallicity range. To estimate the uncertainty, we use bootstrap resampling. For each catalog, we repeatedly resample the stars with replacement, refit the slope, and take the median and 16th–84th percentile range of the resulting slope distribution. We report these as the best-fit value and uncertainty for each catalog in Table 2.

Figure 5 shows the MDFs over $-3 < [\mathrm{Fe/H}] < -1.8$ on a logarithmic scale for the Gaia XP catalogs listed in Sec. 2. Using log density highlights the near-exponential falloff of the metal-poor tail, as in Fig. 4 of Rix et al. (2022) for $-2 < [\mathrm{Fe/H}] < -1$. For visual guidance, dashed lines indicate three reference slopes ($k = 0.6, 1.0, 1.4$) that bracket the range we observe across catalogs in this interval. We excluded Li et al. (2024), whose metallicities are truncated at $[\mathrm{Fe/H}] = -2.0$, producing an artificial pile-up at the boundary.

Because Khalatyan et al. (2024) and Ye et al. (2025) are anchored to APOGEE training labels, their MDFs decline rapidly below $[\mathrm{Fe/H}] \approx -2$, where APOGEE offers few measurements. In addition, Ye et al. (2025) disable their APOGEE-based NN flux correction for stars with $[\mathrm{Fe/H}] \leq -2.5$ and fit FERRE directly, which they release as a separate metal-poor subset. This treatment imprints a discontinuity near $[\mathrm{Fe/H}] = -2.5$, but the MDF tail slope remains consistent. Yang et al. (2025) train against metal-poor reference sets (PASTEL, SAGA) with a cost-sensitive objective to enhance recovery of VMP giants, similarly flattening their tail. Finally, Martin et al. (2024) as part of the Pristine effort targeting $[\mathrm{Fe/H}] < -1$ derive a comparatively larger fraction of metal-poor stars from synthetic CaHK, which also yields a flat density decline in their MDF.

The tiny bootstrap intervals do not include catalog-selection systematics.

The similarly flat MDF tail in the Hattori (2025) catalog is best explained by two design choices that boost recovery of VMP halo stars: (1) the training set was deliberately augmented with the VMP sample of Li et al. (2022), extending label coverage well below $[\mathrm{M/H}] \approx -2.5$ and increasing counts at the low-metallicity end; and (2) the recommended catalog is restricted to low-extinction sight lines, $E(B-V) < 0.1$, which preferentially probe high-latitude halo volume and improves completeness for distant halo giants. However, the same VMP set was also used for training by Zhang et al. (2023) and Ardern-Arentsen et al. (2025), yet both deliver noticeably steeper metal-poor tails. This suggests that simply augmenting training set with VMP stars does not necessarily determine the slope; sample selection and quality cuts play a more important role.

Andrae et al. (2023a) and Zhang et al. (2023) appear to under-represent the most metal-poor regime: their MDFs show the steepest slope for $[\mathrm{Fe/H}] < -2.0$ with $k > 1.4$, then flatten toward $k \approx 1.0$ as metallicity increases. Andrae et al. (2023a) selected sources in their vetted RGB sample to avoid spurious metal-poor stars, which could preferentially remove some VMP stars. Zhang et al. (2023) did not explicitly include VMP stars in their training set and their basic reliability cuts

Table 2. MDF tail slopes estimated over $-3 \leq [\mathrm{Fe/H}] \leq -2$ using an unbinned maximum-likelihood fit to a truncated exponential model. Reported values are the bootstrap median and 16th–84th percentile interval.

Color	Catalog	k
	Andrae et al. (2023)	$1.93^{+0.01}_{-0.01}$
	Zhang et al. (2023)	$2.09^{+0.00}_{-0.00}$
	Khalatyan et al. (2024)	$3.44^{+0.03}_{-0.03}$
	Hattori (2024)	$0.92^{+0.00}_{-0.00}$
	Fallows & Sanders (2024)	$1.65^{+0.01}_{-0.01}$
	Martin et al. (2024)	$0.79^{+0.00}_{-0.00}$
	Kane et al. (2024)	$1.40^{+0.02}_{-0.02}$
	Ardern-Arentsen et al. (2025)	$1.26^{+0.02}_{-0.02}$
	Ye et al. (2025)	$2.07^{+0.01}_{-0.01}$
	Yang et al. (2025)	$0.63^{+0.00}_{-0.00}$

individual [Fe/H] values → truncated exponential fit → bootstrap k

$-3 \leq [\mathrm{Fe/H}] \leq -2$ $k = \beta / \ln 10$ do not fit binned histogram

The tail rises toward higher metallicity; its steepness varies by catalog.

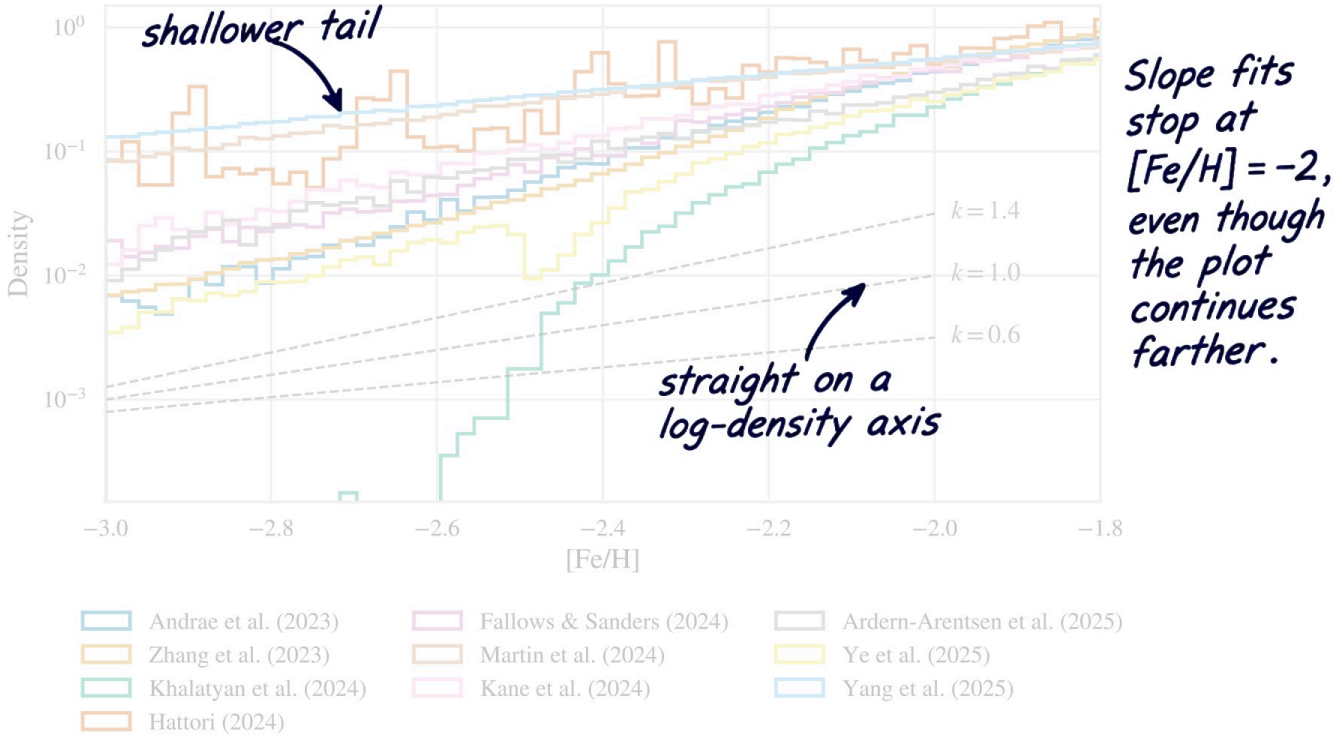


Figure 5. Logarithmic metal-poor MDFs from various Gaia XP catalogs after applying the quality cuts listed in Table 1. Bin width is for this and the panels in the following two figures is 0.02 dex. Dashed lines show exponential density profiles with reference slopes $k = 0.6$, 1.0, and 1.4 for comparison. Despite differences among the catalogs, most exhibit a near-exponential profile with $k \approx 1.0$. Catalogs from Khalatyan et al. (2024) and Ye et al. (2025), trained solely on APOGEE, under-represent VMP stars. Although this and the following figures could extend to [Fe/H] > -2.0, all slope estimates are restricted to $-3 <$ [Fe/H] < -2.0.

tend to cull low-SNR sources, disproportionately trimming the VMP tail. Overall, the metal-poor tails have logarithmic slopes centered near $k \approx 1$, and are well described by a single exponential over our metallicity window. We observe no bumps or peaks within $-3 <$ [Fe/H] < -1.8 after applying recommended quality cuts from the authors. Minor departures from linearity in log-density appear only at the boundaries-flattening in VMP-boosted catalogs and steepening where training or completeness drops below [Fe/H] ≈ -2.0-and are consistent with known selection and pipeline effects.

Figure 6 shows MDFs over $-3 <$ [Fe/H] < -1.8 split by Galactocentric radius R, while Figure 7 shows the same MDFs split by parallax-corrected magnitude, with the same reference slopes from Fig. 5. Because the y-axis is logarithmic, the curves tend to meet near [Fe/H] $\simeq -1.8$ at density of 1. A quick way to read an individual slope is to note the intercept at [Fe/H] $= -3$: for example, an intercept of 10^{-1} implies $k \approx 1.0/1.5 \simeq 0.67$, while 10^{-2} implies $k \simeq 1.33$. Most catalogs exhibit nearly identical slopes across the various R and parallax-corrected magnitude bins. The consistency across radius indicates that the MDF slope is not strongly sensitive to distance from the Galactic center, while the consistency across parallax-corrected magnitude suggests that the result is robust across different stellar types and luminosity classes despite the differing CMD footprints of the catalogs.

In Fig. 7, several samples contain fewer VMP stars at fainter parallax-corrected magnitudes. This is expected because some catalogs are dominated by giants and therefore contain relatively few intrinsically faint stars, while others have smaller overall samples in those bins. The MDF by Zhang et al. (2023) loses VMP stars more rapidly toward the faint end, consistent with their basic reliability cut, which preferentially culls lower-S/N sources with larger uncertainties. By contrast, the catalogs from Hattori (2025), Yang et al. (2025), and Martin et al. (2024) show nearly constant slopes across the parallax-corrected magnitude range. The first catalog supplemented its training set with VMP stars and uses a CMD-based quality filter that retains MSTO and main-sequence stars, while the latter two benefit from synthetic spectra in their pipelines, which help mitigate the scarcity of VMP dwarfs in the training data.

Changing radius mostly preserves each catalog's tail shape.

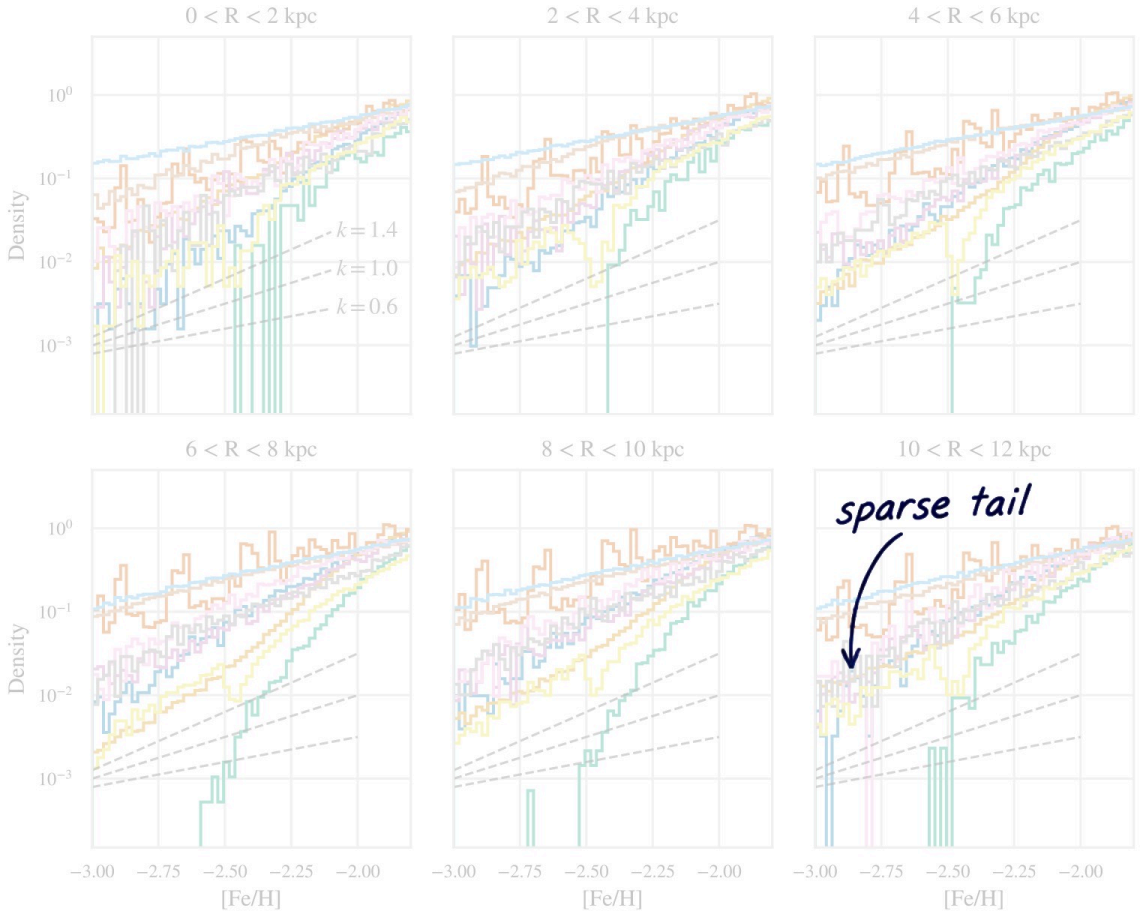


Figure 6. Logarithmic metal-poor MDFs from Gaia XP metallicity catalogs in bins of Galactocentric radius R. The panels show MDFs in 2 kpc intervals from 0 to 12 kpc. In each panel, dashed lines indicate exponential density profiles with three reference slopes ($k = 0.6, 1.0, 1.4$) for comparison, as in Fig. 5. The inferred MDF slopes remain broadly consistent across radial bins, except where the number of stars is too small to define the MDF robustly.

4.3. *MDF Slope Tells Early Assembly Tale*

We used the GCE model outlined in Sec. 3 from our first paper to explore the physical conditions behind different slopes of the metal-poor tail of MDF. The GCE models presented here are pre-selected in Paper I based on the following criteria: 1) [Fe/H] should reach -1.3 ± 0.05 dex after one Gyr; 2) [Mg/Fe] should reach 0.26 ± 0.05 dex after one Gyr; 3) [Mg/Fe] increases above 0.3 dex with a minimum increase of 0.05 dex while [Fe/H] surpasses -0.9 dex in the next Gyr.

We assigned each GCE model to one of three slope categories, $k = 0.6$, 1.0, or 1.4, by comparing its generated MDF to three reference MDFs representing these slope classes. For each model, we first computed a normalized metallicity histogram over the common [Fe/H] bins, and then evaluated a chi-square-like mismatch statistic with respect to each reference MDF,

$$\chi^2 = \sum_i \frac{(M_i - R_i)^2}{R_i + 10^{-10}}, \qquad (4)$$

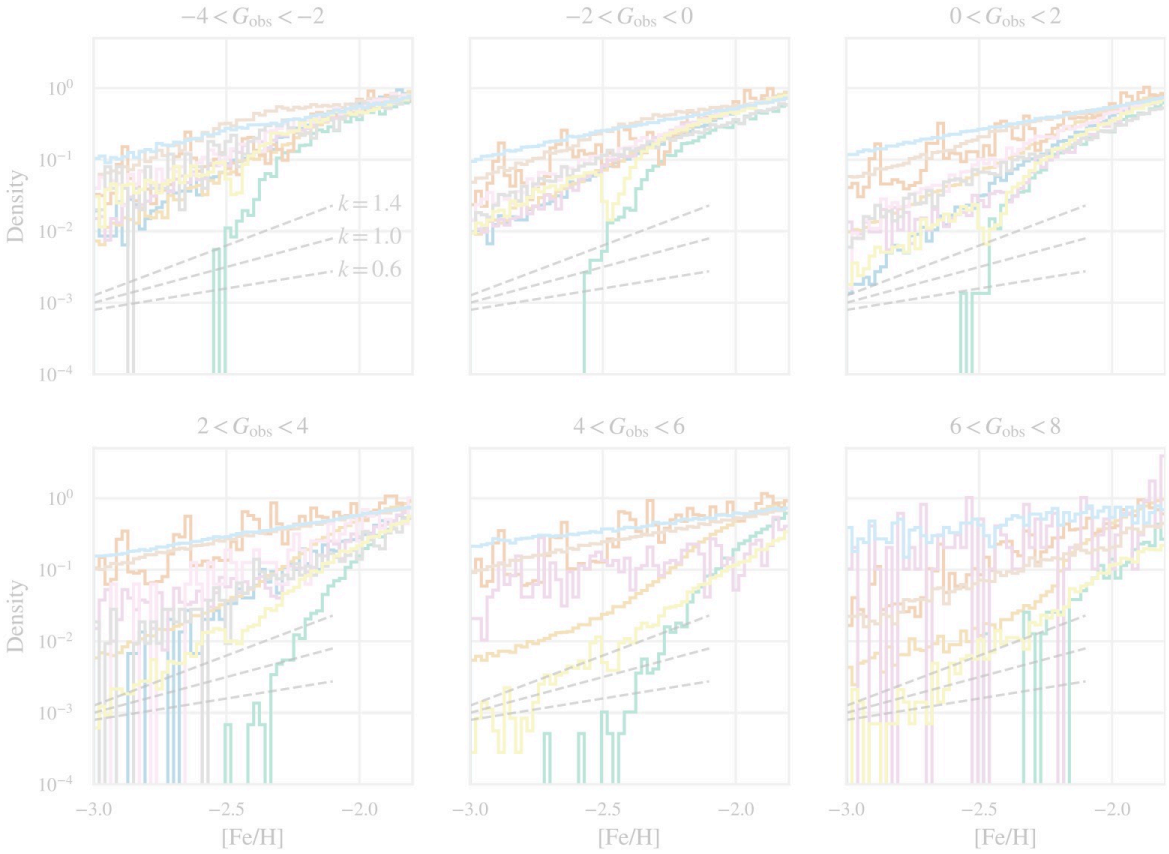
shape matching

where M_i is the model MDF in bin i and R_i is the corresponding reference MDF. Each model was assigned to the category that gave the smallest mismatch. We retained only models with a minimum mismatch below a threshold of $\chi^2 < 1.5$; models with larger deviations were discarded. This procedure therefore grouped models according to which reference MDF shape they most closely resembled, rather than by fitting each model slope independently.

Figure 8 shows the median MDFs of the subsets of GCE models that best match the three reference slopes from Fig. 5. The shaded bands indicate the inter-model spread, and the count of contributing models is annotated beside each k label. Except for $k = 1.0$, a single exponential profile is an imperfect description: for $k = 0.6$ the MDF steepens toward the metal-rich side once [Fe/H] $\gtrsim -2.0$, whereas for $k = 1.4$ it flattens at the metal-poor end below [Fe/H] $\lesssim -2.4$. We fitted the MDFs piecewise with sliding windows of 0.5 dex over

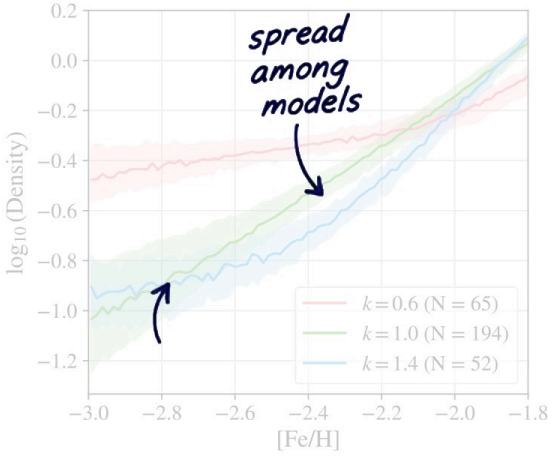
The model grid is grouped by nearest reference MDF, not by a separate slope fit for each model.

Nominal slope classes still contain curved model MDFs.



The faintest magnitude bins also become noisy and incomplete.

Figure 7. Logarithmic metal-poor MDFs from Gaia XP metallicity catalogs in bins of parallax-corrected magnitude, $G_{\rm obs}$. The panels show MDFs for subsamples divided by $G_{\rm obs}$ estimated from the observed G magnitude and parallax, without correcting for extinction. In each panel, dashed lines indicate exponential density profiles with three reference slopes ($k = 0.6, 1.0, 1.4$) for comparison, as in Fig. 5. The inferred MDF slopes remain broadly consistent across $G_{\rm obs}$ bins, except where the number of stars is too small to define the MDF robustly.



not one straight slope

Figure 8. Median MDFs from one-zone GCE models grouped by nearest reference MDF class corresponding to nominal slope $k = 0.6, 1.0, 1.4$. The number of models within each slope category is annotated. Lines and shaded regions show the median and 16th–84th percentile ranges, respectively.

$-3.0 < [\mathrm{Fe/H}] < -1.5$. We then retained the windows whose slopes agreed and showed small residuals in log-density, and took their common value as the slope of the metal-poor tail of each MDF. The main limitation is imposed by our GCE setup, specifically the assumption of a constant inflow rate, which we will discuss more in the following. In the resulting model ensemble, the $k = 1.0$ category is by far the most populated, with substantially fewer cases in the other two categories.

Figure 9 shows the parameter distributions of the models in three slope categories for our five free parameters: $m_{0,\rm cold}$, initial mass of the cold gas reservoir; $f_{\rm SNIa}$, fraction of white dwarfs with initial mass between $(3, 8)$ $M_\odot$ turing into Type Ia SNe; $t_{\rm cool}$, cooling timescale of warm ISM; $\epsilon_{\rm SF,before}$, SFE during the initial prototype phase before [α/Fe]-rise, and $\dot{m}_{\rm inflow,before}$. The diagonal panels show the one-dimensional distributions and the off-diagonal panels show the pairwise two-dimensional distributions. For the first column, models with $k = 1.0$ and $k = 1.4$ have similar values of $m_{0,\rm cold}$, while those with $k = 0.6$ start with a significantly more massive gas reservoir. $m_{*,cold}$ is positively

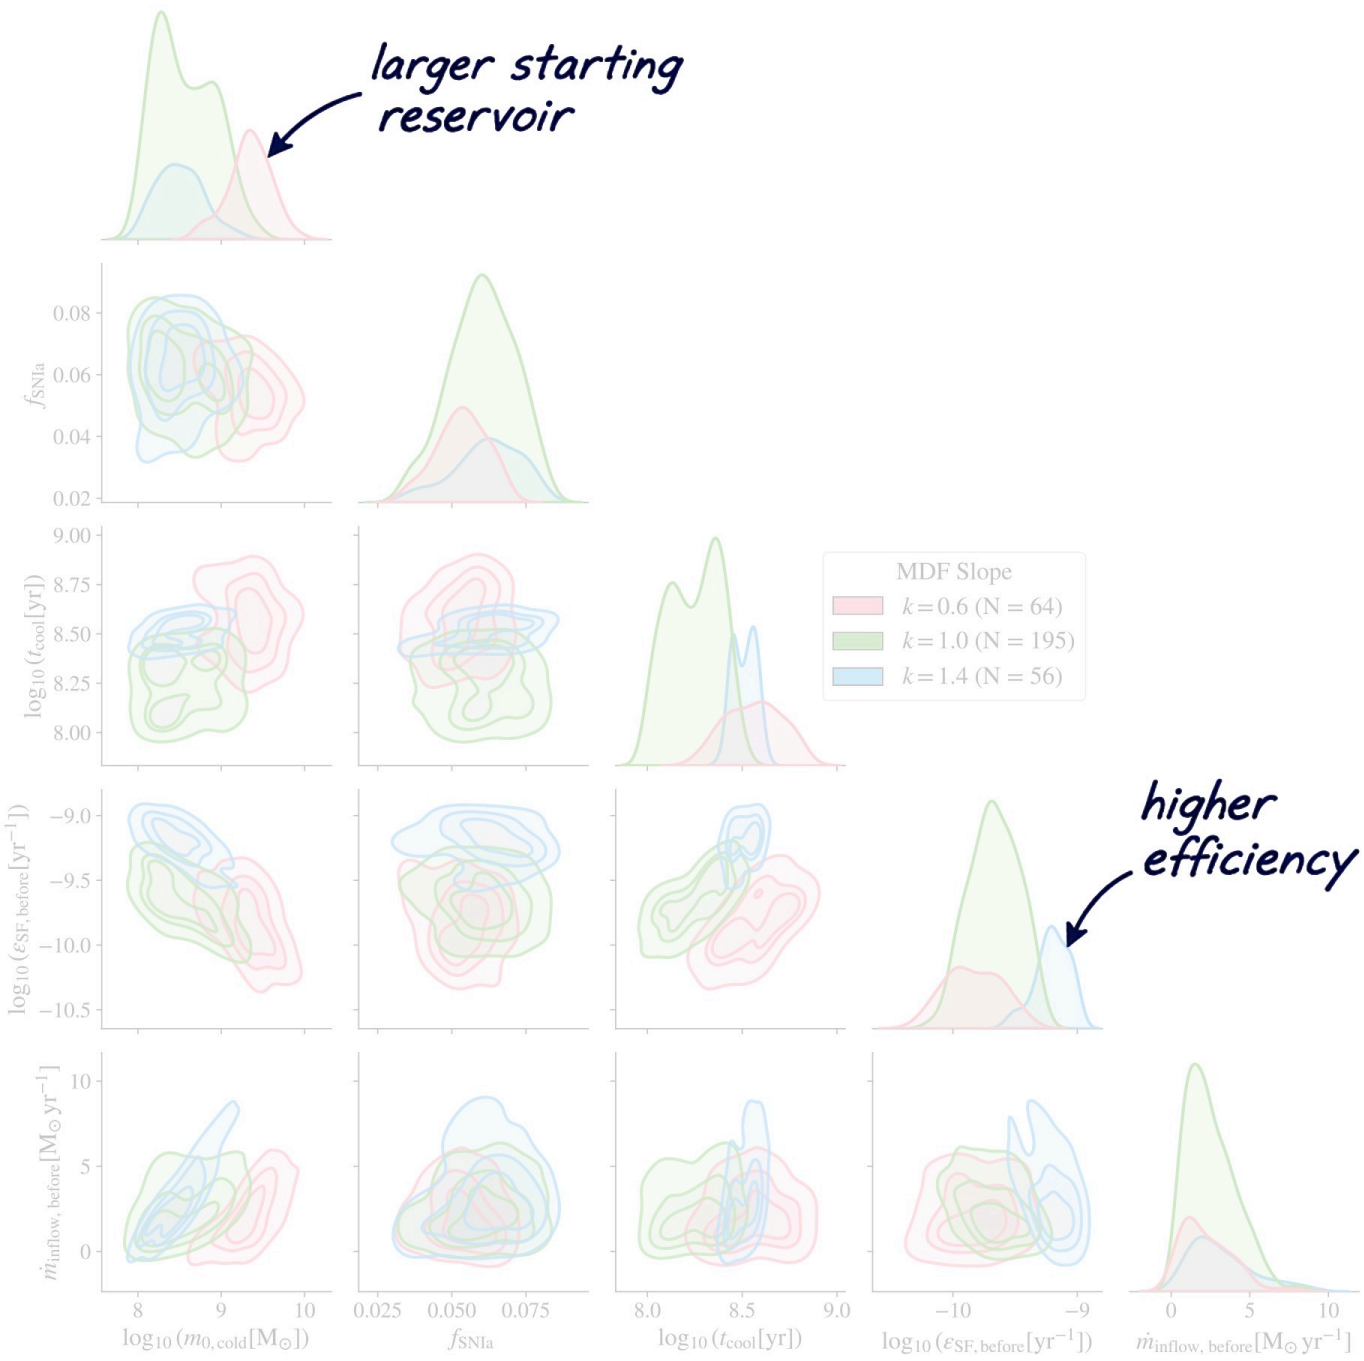


Figure 9. Two-dimensional marginal distributions of one-zone GCE model parameters, grouped by the metal-poor MDF exponential density slope k. The parameters, arranged left to right and bottom to top, are: the initial cold gas mass; the fraction of white dwarfs turning into SNe Ia; the warm ISM cooling timescale; the SFE during the initial proto-phase before the [α/Fe] rise; and the pre-[α/Fe]-rise gas inflow rate. One-dimensional marginal distributions are shown along the diagonal panels, while off-diagonal panels display kernel-density estimated contours enclosing 30%, 60%, and 90% of the models. This plot highlights how variations in k correlate with different GCE parameters.

An early burst leaves both a chemical tail and a stellar-mass signature.

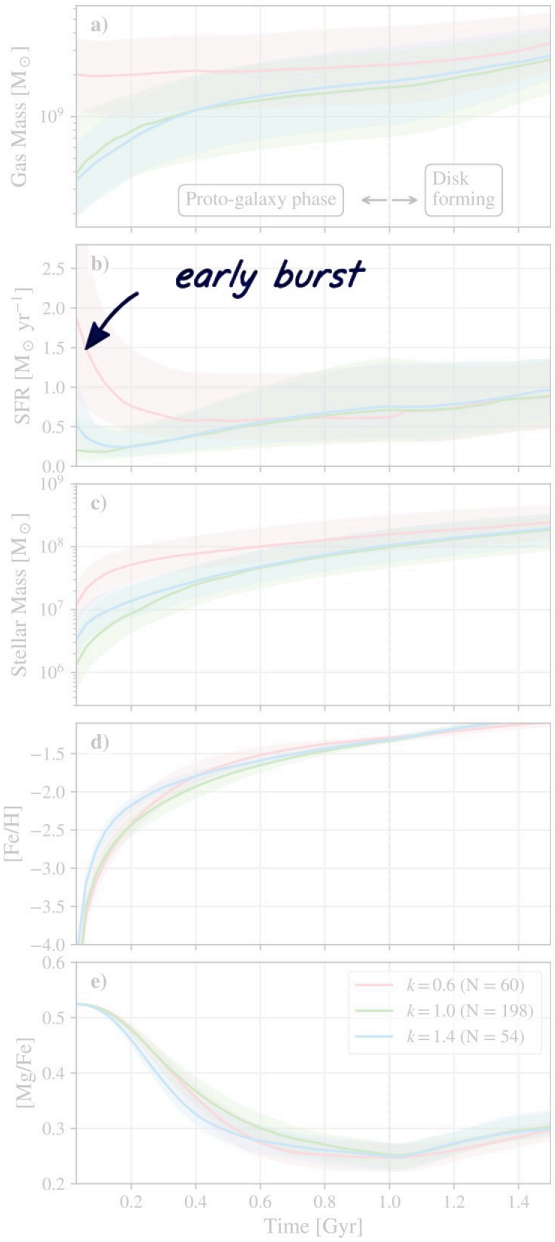


Figure 10. Time evolution of global quantities in one-zone GCE models, grouped by the metal-poor MDF slope k. Panels show (a) total mass of cold and warm ISM, (b) star formation rate (SFR), (c) cumulative stellar mass surviving to the present, (d) stellar [Fe/H], and (e) stellar [Mg/Fe]. Solid curves indicate the median of models in each k category, and shaded bands show the interquartile range (25-75th percentile). Models with $k = 0.6$ stand out with massive initial starbursts and near constant mass for their gas reservoirs, while models with $k = 1.4$ have moderate starbursts.

correlated with $\dot{m}_{\rm inflow,before}$ and negatively correlated with $\epsilon_{\rm SF,before}$. For the second column, the distributions of $f_{\rm SNIa}$ are indistinguishable among the three categories. As for $t_{\rm cool}$, models with $k = 1.0$ require more efficient cooling than the other categories. $\epsilon_{\rm SF,before}$ in-

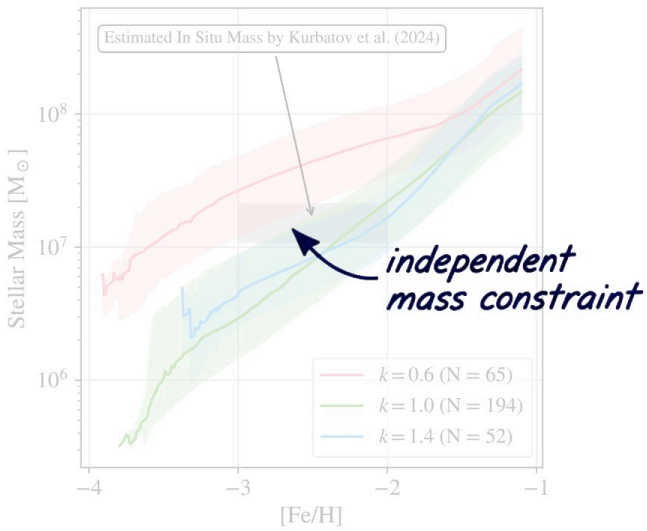


Figure 11. The cumulative stellar mass assembly history as a function of metallicity ([Fe/H]) for three slope classes in our GCE models. Lines and shaded regions show the median and 16th–84th percentile ranges, respectively. Kurbatov et al. (2024) estimated the mass of the in situ proto-Galaxy to be between 1.07×10^7 and 2.11×10^7 $M_\odot$ between [Fe/H] = -3.0 and -2.0. The models with $k = 0.6$ can be ruled out with this restriction.

The stellar-mass comparison also depends on the adopted IMF and model normalization.

creases with k and most models have low $\dot{m}_{\rm inflow,before}$ of less than 5 $M_\odot$ per year.

In our GCE framework, two factors control the shape of the resulting MDF. The first is the rate of chemical evolution, [Fe/H] in this particular case, and the second is the SFH. When chemical evolution stagnates or star formation intensifies, more stellar bins are born with a certain metallicity value, flattening the slope at low metallicities or steepening the slope at high metallicities. Therefore, when models start with a large amount of gas, they enjoy a faster rate of chemical evolution and require more pristine inflow gas to either sustain a relatively high SFR or slow down chemical evolution in order to maintain the same slope. The same logic holds for SFE, which explains the correlations among the parameters.

We translated the parameters in Fig. 9 into global properties over time under our GCE framework in Figure 10, to illustrate the conditions associated with each MDF slope. Panels (a)–(e) show: (a) total cold+warm ISM mass; (b) star-formation rate (SFR); (c) cumulative stellar mass surviving to the present; (d) stellar [Fe/H]; and (e) stellar [Mg/Fe]. Consistent with Fig. 9, models with $k = 0.6$ maintain an almost constant gas reservoir by starting with a large initial gas mass, receiving little subsequent inflow, and adopting a low SFE. An early star formation burst produces many stars with [Fe/H] < -2.0, inflating the low-metallicity tail. That burst also heats cold ISM into the warm phase and

Auriga's early gas buildup provides a comparison to the one-zone histories.

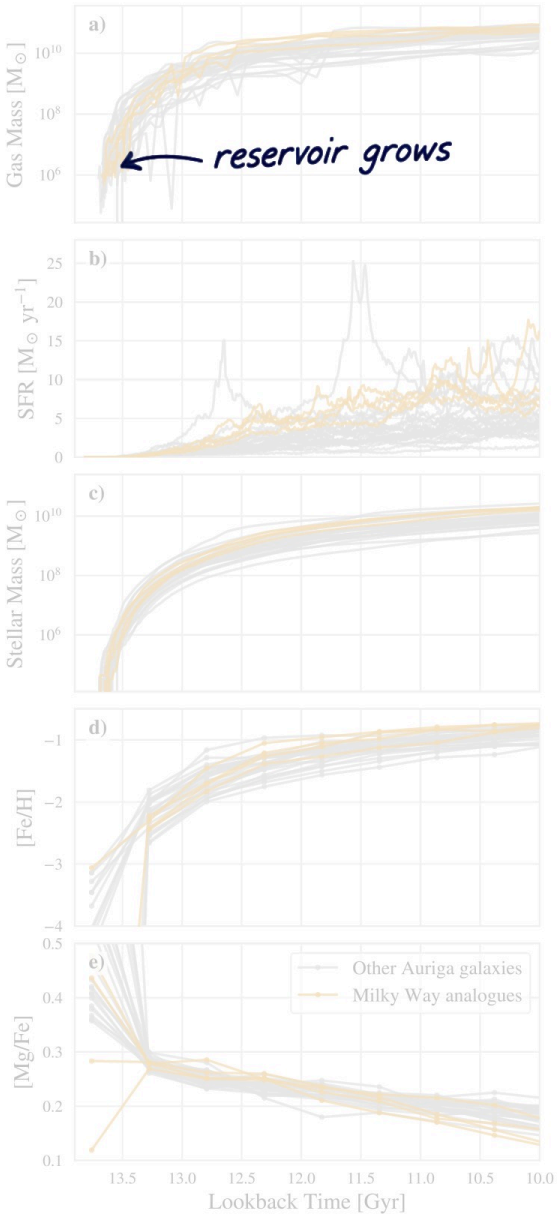


Figure 12. Time evolution of global properties for Auriga galaxies. Panels show: (a) total gas mass; (b) SFR; (c) cumulative mass of stars surviving to $z = 0$; (d) mass-weighted stellar [Fe/H]; and (e) mass-weighted stellar [Mg/Fe]. Tracks for the Milky Way analogues (Au-17, Au-18, Au-23, and Au-26) are highlighted in orange; the other Auriga galaxies are shown in gray. All of them show accumulation of gas and lack of starburst in the first Gyr, consistent with $k = 1.0$ models in Fig.10.

forward in time

quickly suppresses the SFR, and the delayed late-time enrichment keeps the high-metallicity end underpopulated, yielding a shallow overall slope.

The primary distinction between the $k = 1.0$ and $k = 1.4$ models is SFE. The $k = 1.4$ models have substantially higher SFE, achieve the fastest early chemical evolution, and exhibit a pronounced burst even with only a moderate initial gas reservoir. The combination of high SFR and rapid enrichment steepens the metal-poor tail, producing the large k. A limitation of our current GCE setup is the assumption of constant inflow beginning at $t = 0$ and the requirement that models reach [Fe/H] $= -1.3$ by one Gyr. For the steep class ($k = 1.4$), a truly constant slope would demand minimal star formation at the lowest metallicities; our target for [Fe/H] and early onset of gas inflow necessarily drive substantial early star formation and thus flattens the extreme metal-poor end. Conversely, a persistently shallow tail ($k = 0.6$) would prefer little late-time star formation, but the constant inflow unavoidably grows the gas reservoir and steepens the high-metallicity end. These behaviors point to the need for time-dependent inflow and/or SFE in future iterations to sustain single-slope behavior over the full $-3 <$ [Fe/H] < -1.8 range. Nevertheless, the constant inflow rate is sufficient for us to produce the desired slopes over much of our metallicity range.

Constant inflow limits which tail shapes the one-zone models can maintain.

Figure 11 shows the cumulative stellar mass as a function of metallicity for the three slope classes in our GCE models (Fig. 8). We construct this diagnostic because ages for the oldest stars are notoriously uncertain, whereas [Fe/H] is directly observable at scale from Gaia XP. As a reference point, we mark the estimated stellar mass of the in-situ proto-Galaxy (the Aurora population) from Kurbatov et al. (2024), who found a mass between 1.07×10^7 and $2.11 \times 10^7\,\mathrm{M}_\odot$ for stars in the range $-3.0 <$ [Fe/H] < -2.0. This mass was derived from modeling the spatial density distribution of metal-poor red giant branch stars from Gaia, accounting for the survey's selection function.

The curve shapes encode the same information highlighted in Fig. 10: in the burst-quenched case ($k = 0.6$), an early starburst builds $\sim 10^7\,\mathrm{M}_\odot$ of stars rapidly, but the large initial gas reservoir keeps the initial metallicity low, producing a fast early rise that soon plateaus. In contrast, the slow-assembly case ($k = 1.4$) begins with a milder burst into a less massive reservoir, reaching a higher initial metallicity and then growing steadily under sustained low-level inflow and higher SFE; the most common slopes in both GCE models and Auriga galaxies ($k \simeq 1.25$) lie between these limits. The absolute normalization of $M_\star(<$ [Fe/H]) depends on assumptions such as low-mass cutoff of the IMF. Lowering $m_{\rm min}$ from 0.5 to $0.1\,\mathrm{M}_\odot$ roughly doubles the surviving stellar mass, shifting the mass-metallicity relations in Fig. 11 upward at fixed MDF shape. However, $k = 1.0$ and $k = 1.4$ are still the preferred scenarios within the observed mass range. In short, the cumulative mass-

Two model families favor similar slopes, with very different spreads.

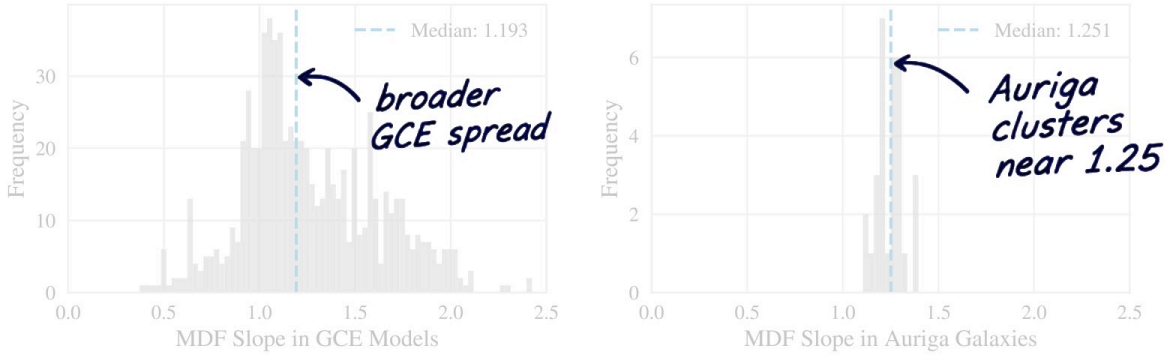


Figure 13. Distribution of exponential density slopes fitted for the metal-poor MDFs from GCE models exhibiting an [α/Fe]-rise and from Auriga galaxies measured between $-3.0 <$ [Fe/H] < -2.0. The histograms are computed using a uniform bin width of 0.025 in slope. The median slope for each group is marked with a vertical dashed line with its value annotated. The most common slopes among our GCE models with [α/Fe]-rise agree with those produced in Auriga. The [α/Fe]-rise is key to establishing the scenario of the proto-Galaxy and the observed MDF help us further narrow down.

metallicity plane provides an age-independent discriminator of proto-Galaxy histories and a direct bridge between XP-based MDFs, GCE models, and mass measurements from spatial modeling.

4.4. *Quiet Dawn Leads to Unified Slope*

We looked into the simulated galaxies in the Auriga project to verify the connection between the physical conditions and MDF slopes from our GCE models. The Auriga Project is a suite of high-resolution cosmological "zoom-in" simulations of Milky Way-mass galaxies (augmented recently to include a dwarf-galaxy sample), evolved with the moving-mesh code AREPO and ideal MHD to follow the coupled growth of dark matter, gas, stars, and black holes in a ΛCDM universe (Grand et al. 2017, 2024).

Auriga adopts AREPO's Galilean-invariant, Voronoi moving-mesh hydrodynamics–crucial for accurately capturing shocks, contact discontinuities, and rotating disks–and builds on a comprehensive galaxy-formation model (cooling, star formation, stellar/AGN feedback, metal enrichment) inherited from the Illustris framework (Springel 2010; Vogelsberger et al. 2013). The simulations self-consistently amplify magnetic fields and produce disk galaxies whose global properties (e.g., HI distributions, kinematics, scaling relations) compare well with observations, providing a physically grounded laboratory for disk formation and chemo-dynamical studies (Marinacci et al. 2017; Pakmor et al. 2017). A recent overview details the augmented suite and the first public data release–raw snapshots, group catalogs, merger trees, and analysis tools[8]–enabling community access to the methodology and outputs (Grand et al. 2024).

[8] https://wwwmpa.mpa-garching.mpg.de/auriga/data.html

Figure 12 shows the time evolution of global properties for the 30 Auriga galaxies in the original release, in the same style as Fig. 10. Panels display: (a) total gas mass; (b) SFR; (c) cumulative stellar mass; (d) mass-weighted stellar [Fe/H]; and (e) mass-weighted stellar [Mg/Fe]. The commonly used Milky Way analogues-Au-17, Au-18, Au-23, and Au-26-are highlighted in orange; all other galaxies are shown in gray. Au-17 and Au-18 best reproduce the Milky Way bulge/bar chemo-kinematics and exhibit quiescent merger histories since $z \sim 3.5$ (Fragkoudi et al. 2020). Au-18 also experiences an early Gaia-Sausage/Enceladus-like merger and has halo properties closest to the Milky Way (Merrow et al. 2024). Au-17 also undergoes a similar merger, but its properties are less consistent with those of GSE (Fattahi et al. 2019; Orkney & Laporte 2025).

Au-18, Au-23, and Au-26 possess Milky Way-like bars and have been used to time the disk/bar formation in our Galaxy (Zhang et al. 2024). To align the simulation with the stellar disk, the z-axis is defined by the net angular momentum vector of all stars within 0.1 times the virial radius ($< 0.1 \times r_{200c}$). We restricted the analysis to stars with $R < 10$ kpc and $|z| < 3$ kpc to match the spatial footprint of Gaia XP catalogs. Comparing Figs. 10 and 12, most Auriga systems are consistent with the $k = 1.0$ GCE category, characterized by negligible initial gas reservoirs and minimal initial star formation. For each galaxy, we constructed MDFs from the masses and [Fe/H] of surviving star particles at $z \sim 0$ and estimated the slope of the metal-poor tail over $-3 <$ [Fe/H] < -1.8.

Figure 13 compares the slope distributions: Auriga galaxies cluster around $k \approx 1.25$, whereas the GCE models span a broader range, $k \approx 0.5$–2.0. The MDFs in Auriga galaxies are impacted by young metal-poor stars because many galaxies have multiple mergers bringing

Auriga is a comparison sample, not an independent measurement of the Milky Way.

Selection checks support the broad tail shape without removing every bias.

in metal-poor gas over time. The slopes of many Auriga galaxies would be lower when these young metal-poor stars are removed. Nevertheless, the slopes in Auriga galaxies are consistent with what we expect from the GCE models based on the gas inflow history and SFH. The original Auriga suite does not explicitly sample initial conditions with a large, pre-assembled cold-gas reservoir that ignites an early star formation burst. Because the zoom-ins draw gas from cosmological inflow and assemble disks gradually, star formation peaks happen later typically from merger- or compaction-driven.

5. DISCUSSION

5.1. *Selection function considerations*

The study compares quality-controlled releases; it does not derive one common selection function.

The Gaia XP metallicity catalogs considered in this work do not share a single selection function. In addition to the native Gaia XP magnitude limit, each catalog has its own effective selection through the choice of training set, inference domain, and recommended quality cuts. Throughout this work, we therefore do not impose additional cuts of our own beyond the published recommended quality cuts; our goal is to compare the public catalogs within their validated regimes rather than construct a new common sample with a different and difficult-to-model selection function.

A related issue is evolutionary state. Because metal-poor giants are intrinsically brighter than dwarfs, giant-dominated catalogs probe larger distances and a different mixture of Galactic components than catalogs that include substantial numbers of dwarfs and subgiants. We do not apply a uniform giant/dwarf cut across catalogs, since several public releases are already restricted to vetted giant samples whereas others are intended to operate over broader CMD regions. Instead, we assess the effect of this heterogeneity empirically. As shown above, the logarithmic MDF slope remains broadly consistent across catalogs with different CMD footprints and also across bins of observed absolute magnitude and Galactocentric radius (Figs. 5, 7, and 6) indicating that luminosity-class and volume effects primarily alter the normalization and sampled volume rather than strongly changing the fitted MDF slope over the adopted metallicity interval.

Disk contamination is another concern because the Gaia XP samples span both low- and high-|z| regions. We do not apply a global kinematic cut to the catalogs because only a minority of stars (∼5–10%) have Gaia DR3 RVS radial velocities; instead, we use this subset only as a diagnostic. The available kinematics show that stars with [Fe/H] < -2.0 are predominantly slow-rotating, consistent with a proto-Galactic population, whereas the clearest fast-rotating contaminants become common closer to the metal-richer boundary, especially above [Fe/H] $\gtrsim -1.5$ (Fig. 2). These stars often deviate from the old, metal-poor isochrone locus in the CMD and occupy regions that are sparsely represented in the training sets of several XP-based pipelines (Fig. 4), indicating that they are more likely metal-rich disk contaminants with underestimated metallicities than genuine very metal-poor stars with large angular momentum. Catalogs with larger low-$|z|$ samples are therefore more susceptible to this effect (see the last panel of Fig. 2 and Fig. 3). To minimize such contamination, we restrict the MDF-tail analysis to [Fe/H] < -2.0 and treat the RVS-based kinematic distributions only as a validation check rather than as a selection criterion. The confinement of the clearest contaminants to the metal-richer boundary suggests that their impact on the fitted MDF slope over $-3 <$ [Fe/H] < -2.0 is small.

5.2. *A Unified Proto-Galactic Population*

We interpret the proto-Galaxy as the composite stellar population that formed up to the transition from dispersion-dominated assembly to the emergence of a rotation-supported disk. Across the Gaia XP metallicity catalogs, the metal-poor MDF tails over $-3 <$ [Fe/H] < -2.0 are generally well described by an approximately single-slope exponential with logarithmic slopes centered near $k \approx 1$. After applying each catalog's recommended quality cuts, we find no compelling evidence for secondary peaks or shoulders within this interval. Residual departures appear mainly near the boundaries and are consistent with known pipeline and selection effects: APOGEE-anchored products tend to steepen below [Fe/H] ≈ -2 because of sparse metal-poor training labels, whereas VMP-optimized or CaHK-based designs can flatten the extreme metal-poor end. These effects likely account for much of the inter-catalog variation in the fitted slopes, but they do not qualitatively change the main result that the metal-poor tail is broadly smooth and close to exponential over the adopted metallicity range.

smooth within this window

Taken together, these pipeline and selection effects explain the apparent differences among catalogs without requiring genuine departures from a single-exponential halo MDF over $-3 \leq$ [Fe/H] < -2.0. This result holds even when MDFs are reexamined in Galactocentric-radius or magnitude slices. Residual differences track known selection footprints-XP spectra exist only for $G < 17.65$, and some catalogs restrict inference to specific luminosity classes (e.g., vetted RGB), altering CMD occupancy and the inner/outer-halo mixture. Within these limits, the slope stability across R and magnitude supports a globally similar metal-poor tail.

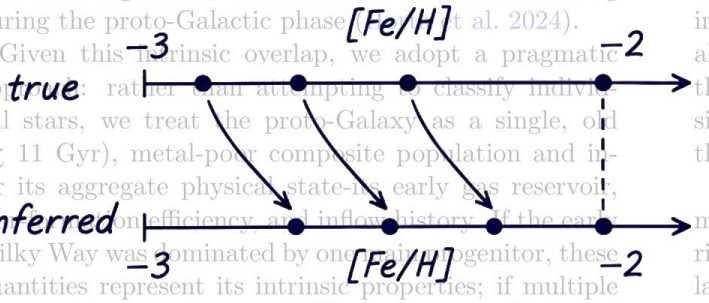
Catalog systematics can masquerade as a different gas-assembly history.

A smooth composite MDF does not identify each star's progenitor.

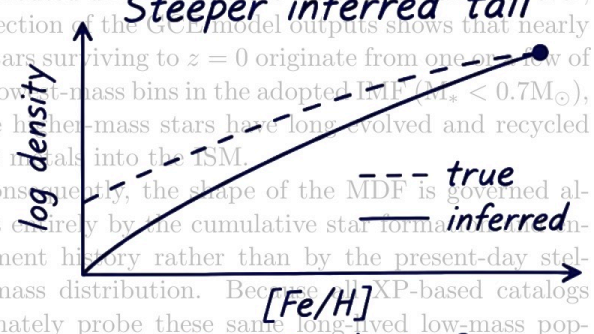
Both observations and simulations show that such old, metal-poor stars arise from a mixture of origins: an in-situ component heated by early mergers and accreted debris from one or two major progenitors (Tumlinson 2010; El-Badry et al. 2018; Horta et al. 2024). However, disentangling these channels is difficult even in simulations with perfect birth tags-kinematics and chemistry overlap strongly below [Fe/H] < -1, with mean abundance offsets $\lesssim 0.3$ dex and broad MDF overlap in both prograde and retrograde regions (Renaud et al. 2021; Khoperskov et al. 2023). Observational separations based on integrals of motion recover only a small subset of proto-Galactic structures (Brauer et al. 2022), and recent FIRE-2 analyses show significant kinematic and chemical blending during the proto-Galactic phase (Sarkar et al. 2024).

Given this intrinsic overlap, we adopt a pragmatic approach: rather than attempting to classify individual stars, we treat the proto-Galaxy as a single, old ($\gtrsim 11$ Gyr), metal-poor composite population and infer its aggregate physical state-its early gas reservoir, star-formation efficiency, and inflow history. If the early Milky Way was dominated by one main progenitor, these quantities represent its intrinsic properties; if multiple comparable contributors, they represent a luminosity-weighted mean. In either case, the chemical continuity of the metal-poor MDF supports a unified treatment of the dominant early star-forming system that can record the conditions of the proto-Galaxy.

Low-end metallicities overestimated



Same stars, different metallicity bins.
Positions are schematic, not four measured stars.



Schematic; one possible label-transfer bias. Other methods can flatten tails.

5.3. *Addressing Caveats on MDFs*

A systematic uncertainty arises from attenuation bias–the tendency of different inference methods to distort the intrinsic shape of the metal-poor MDF. Template-fitting and synthetic-spectra approaches that explicitly include very metal-poor stars (e.g., Ca H&K-based or VMP-augmented training) extend sensitivity into the lowest metallicities, but also introduce broader uncertainties that flatten the observed tail, yielding shallower slopes ($k_{\rm obs} < k_{\rm true}$). In contrast, catalogs derived by label transfer from APOGEE or other spectroscopic surveys suffer from a lack of true VMP training examples, even when augmented with a few LAMOST or Pristine stars, causing the model to compress the dynamic range and overpredict metallicities at the low end.

This truncation steepens the inferred tail ($k_{\rm obs} > k_{\rm true}$), mimicking the effect of an intrinsically smaller gas reservoir or higher star-formation efficiency in GCE interpretations. The net impact of attenuation therefore depends on the inference framework, the effective [Fe/H] prior, and the treatment of noise and regularization. Cross-catalog comparisons must thus account for this methodological bias before attributing differences in k to physical variations in the proto-Galaxy. A more detailed and careful intra-catalog examination is required in the future to further narrow down the slope of the MDF tail.

A related consideration concerns the implicit selection function of stars from the GCE models and its connection to the observed CMD coverage of Gaia XP catalogs. The MDF represents stars that have survived to the present day, effectively sampling the least massive, longest-lived stars. In principle, differences among catalogs in their CMD coverage–for example, whether they focus on main-sequence turnoff or red-giant branches–could bias the inferred MDF if stars of different masses contributed differently to the metal-poor tail. However, inspection of the GCE model outputs shows that nearly all stars surviving to $z = 0$ originate from one or a few of the lowest-mass bins in the adopted IMF ($M_* < 0.7 M_\odot$), since higher-mass stars have long evolved and recycled their metals into the ISM.

Consequently, the shape of the MDF is governed almost entirely by the cumulative star formation and enrichment history rather than by the present-day stellar mass distribution. Because all XP-based catalogs ultimately probe these same long-lived low-mass populations, their differing color footprints do not systematically affect the inferred slope k. In this regime, the MDF can be interpreted as a faithful record of the proto-Galaxy's early gas-regulation history, largely insensitive to catalog-specific evolutionary or photometric selections. Nevertheless, we implemented different selection functions based on Gaia magnitudes of the stellar component in our models and found that the resulting MDFs are identical. This agrees with Fig. 7 in which MDF slopes remain consistent across different parallax-corrected magnitude bins for catalogs with sufficient stars.

5.4. *Comparison with Previous Halo MDF Studies*

The slopes inferred from Gaia XP metallicity catalogs are broadly consistent with previous constraints on the Milky Way's metal-poor tail. In the Pristine survey, the inner-halo MDF between $-3.4 <$ [Fe/H] < -2.5 follows a single exponential with $k \equiv \Delta \log_{10} N / \Delta$[Fe/H] $= 1.0 \pm 0.1$ for main-sequence turnoff stars, establishing a canonical baseline of $k \simeq 1$ (Youakim et al. 2020). SkyMapper photometric metallicities for $\sim 10^5$ giants yield two regimes: $k = 1.5$ for $-3.0 <$ [Fe/H] < -2.3 and $k \simeq 1.0$ for $-2.3 <$ [Fe/H] < -1.5, with negligible dependence on Galactic height to $|z| \approx 5$–7 kpc (Da Costa et al. 2019; Chiti et al. 2021). These values bracket the range found here from Gaia XP, confirming

Compare slopes over matched metallicity windows and sampled populations.

that the inner-halo MDF tail typically follows an exponential decline with $k \approx 1$.

Other halo MDF studies, with differing spatial coverages and selection functions, have emphasized the metal-poor tail in the outer halo. The Hamburg/ESO Survey revealed a sharp decline near [Fe/H] ≈ -3.6, interpreted as evidence for a critical metallicity $Z_{\rm cr} \approx 10^{-3.4} Z_{\odot}$ required for low-mass star formation (Salvadori et al. 2007; Schörck et al. 2009). SDSS Stripe-82 and BOSS spectroscopy described a multi-modal halo MDF with peaks near -1.7 and -2.3, corresponding to inner- and outer-halo components with distinct kinematics (An et al. 2013; Allende Prieto et al. 2014). LAMOST giants at $|z| > 5$ kpc show a similar three-component structure-inner and outer halo plus a heated-disk population (Liu et al. 2018). The H3 Survey combined with Gaia data indicates that $\gtrsim 95\%$ of halo giants can be attributed to discrete merger remnants: beyond 15 kpc, the halo is dominated by GSE and Sagittarius debris, while the most metal-poor tail arises from lower-mass progenitors (Naidu et al. 2020; Bonifacio et al. 2021).

In contrast to these high-|z| or accretion-dominated samples, nearly all Gaia XP stars lie within $|z| < 2$ kpc (Fig. 3), probing the chemically mixed, in-situ inner halo. The exponential slopes derived here therefore trace the early enrichment history of the proto-Galactic gas reservoir rather than the outer-halo debris field. Methodologically, the all-sky XP spectroscopy provides a more uniform metallicity baseline and higher completeness at low |z| than prior spectroscopic surveys, reducing spatial biases in the inferred MDF shape. Physically, the near-unity slope $k \approx 1$ is consistent with a moderately efficient early star-formation episode within a well-mixed gas reservoir-close to the expectation of a one-zone, leaky-box model with limited metal retention. Steeper slopes ($k > 1$) observed in some photometric or high-|z| samples likely reflect either stronger outflows or incomplete sampling of the metal-poor tail.

Lowe et al. (2025) provided an independent check on the metal-poor MDF tail using follow-up spectroscopy of Gaia XP–selected candidates. They selected ∼6000 Gaia DR3 candidates from the Zhang et al. (2023) catalog, obtained 2dF+AAOmega spectroscopy, and then measured MDF tails for kinematically defined components (halo, prograde/retrograde disk, and GSE). In their Fig. 16, the fitted tail slopes over $-3 \leq$ [Fe/H] ≤ -2 are close to unity (for the halo they report $\Delta \log N / \Delta[\mathrm{Fe/H}] = 1.15 \pm 0.08$), consistent with the $k \approx 1$ slopes we found from Gaia XP MDF tails over the same metallicity range. In their appendix A, they compared their spectroscopic metallicities to several Gaia XP-based catalogs and found that the level of agreement is strongly pipeline-dependent, with substantial scatter and clear systematics at the metal-poor end. In particular, they show that XP products by Zhang et al. (2023), Andrae et al. (2023a), and Li et al. (2024) tend to assign systematically higher [Fe/H] to their most metal-poor stars, plausibly driven by reddening-related parameter degeneracies and limited metal-poor coverage in the training sets. This systematic over-estimation of [Fe/H] at the metal-poor end would shift intrinsically VMP/EMP stars toward higher inferred metallicities, thereby steepening the observed decline. The true metal-poor MDF tail would therefore be flatter (i.e., have a smaller k) than that measured from the affected catalogs.

Spectroscopic follow-up also finds a near-unity tail while exposing XP-dependent metallicity biases.

5.5. *Comparison with Previous GCE Studies*

The observed rise in [α/Fe] provides a key boundary condition for our GCE models introduced in Paper I, ruling out many implausible parameter combinations. The observed MDFs presented here further constrain the models by breaking degeneracies among parameters. As shown in Fig. 13, the median slope from our original model grid already agrees well with the median slope of the Auriga galaxies, and the overall slope distribution in the left panel matches the observed MDF slopes from the Gaia XP metallicity catalogs in Fig. 5. In Paper I, we found that the [α/Fe] rise alone sets only partial constraints: it defines a lower limit on the SFE and upper limits on the inflow rate and warm-ISM cooling timescale but does not meaningfully restrict the initial gas reservoir mass (see Fig. 7 in Paper I). It primarily constrains the SNe Ia fraction to $\sim 6\%$. The inclusion of the MDF tail from Gaia XP now allows us to tighten these limits substantially. For the most common slope category ($k = 1.0$), the feasible parameter ranges correspond to an initial cold-gas reservoir below $10^8\,\mathrm{M}_{\odot}$, a star-formation efficiency of a few $\times 10^{-10}\,\mathrm{yr}^{-1}$, and an inflow rate of approximately $2\,\mathrm{M}_{\odot}\,\mathrm{yr}^{-1}$.

Classical one-zone models show that the low-metallicity tail of the halo MDF should be exponential, with the power set by either lowering the effective yield (mass of freshly produced metals returned to the ISM per unit mass locked in long-lived stars/remnants) by a factor of 13 compared to the disk or adding outflow proportional to star formation (Hartwick 1976). This simple model successfully predicted the number of metal-poor stars down to [Fe/H] $= -3$ for early observations (Ryan & Norris 1991; Carney et al. 1996). However, Prantzos (2003) noted that this simple model over-predicts the number of stars below [Fe/H] < -4 if instantaneous recycling approximation (IRA) is relaxed and found that a non-IRA model with a moderate ini-

inner, low-|z| sample

outer, high-|z| sample

different population mixtures

Within this framework, sustained gas supply favors a quieter early buildup.

tial inflow and star-driven outflow recovers the observed exponential tail without significant reduction in the effective yield.

In a hierarchical extension, the halo MDF can be written as the sum of subhalo MDFs with mass-dependent yields (Prantzos 2008). In this picture, the low-metallicity bins ([Fe/H] < -2.0) are supplied primarily by numerous low-mass stellar subhaloes (10^{6-7} $M_\odot$). Independently, a merger-tree GCE model finds that the halo MDF requires strong feedback plus pre-enriched infall (Salvadori et al. 2007). In their fiducial solution, 60% of the metals that pre-enrich the Galactic medium are ejected by halos with a dark matter mass less than 6×10^9 $M_\odot$. Finally, a two-infall GCE model augmented with star-driven outflow reproduces the halo MDF for with a mass-loading factor of 14 and an early infall timescale less than 0.2 Gyr, implying rapid early enrichment and efficient metal loss (Brusadin et al. 2013). These studies depict an open, rapidly evolving proto-Milky Way regulated by gas flows and hierarchical assembly. An early infall episode coupled to outflows enables rapid enrichment while losing a large amount of metals. The lowest-metallicity stars chiefly arise in numerous low-mass progenitors whose ejecta pre-enriched incoming gas.

We employ a grid of one-zone, non-IRA GCE models that vary the initial gas reservoir, SNe Ia fraction, warm-ISM cooling timescale, SFE, and gas inflow rate over physically motivated ranges. We group outcomes by the exponential metal-poor-tail slope k. The models include outflows proportional to SFR and supernovae via mass-loading factors and were pre-selected in Paper I to achieve a characteristic turn in [Fe/H] -[Mg/Fe] within the first 1–2 Gyr, consistent with the observed proto-Galactic timeline. These models typically favors small gas reservoirs during the first Gyr. The models are capable of reproducing a single-slope exponential MDF with $k \simeq 1$, while departures reflect the simplifying assumption of a constant inflow rate.

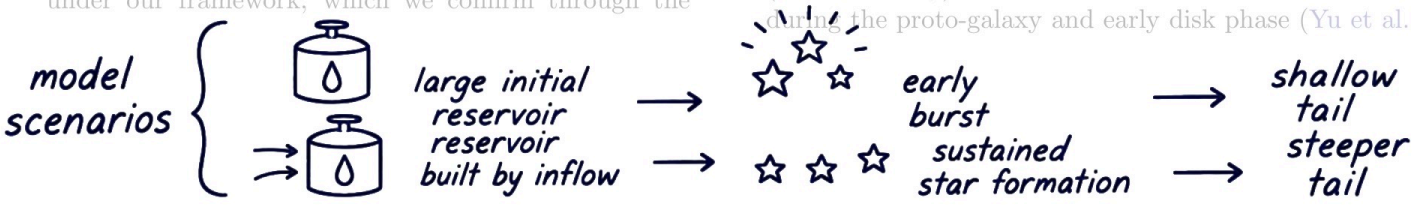
In our framework, k traces the early gas regulation: shallow tails ($k \simeq 0.6$) arise in models with large initial reservoirs and little subsequent inflow, whereas steeper tails ($k \simeq 1.0$–1.4) arise for smaller initial reservoirs plus sustained inflow that builds the gas reservoirs while delaying enrichment. The trade-off between SFE and gas cooling also modulate the mapping from the parameter space to MDF unless one parameter takes an extreme value. Fitting the same exponential form to local Auriga star-particle MDFs ($R < 10$ kpc, $|z| < 3$ kpc) yields a distribution clustered near $k \simeq 1.25$, consistent with a gas-regulated, infall-driven growth scenario under our framework, which we confirm through the global properties of Auriga galaxies. This scenario is also supported by more detailed examination of the gas accretion history of Auriga galaxies by Iza et al. (2022). This picture is compatible with classical one-zone interpretations in which the tail slope traces the effective yield. A moderate initial gas reservoir of $10^{8-9}M_\odot$ aligns with early-infall models (Prantzos 2003; Brusadin et al. 2013); cooling timescales need only be moderately short-no extreme suppression of the effective yield is required-and low-to-moderate inflow rates are consistent with pre-enrichment by numerous low-mass progenitors (Salvadori et al. 2007; Prantzos 2008).

5.6. *The Case for a Burstless Proto-Galaxy*

The MDF constraints from Gaia XP present a case for a largely burstless proto-Galaxy. While cosmological zoom-in simulations frequently show that Milky Way–mass progenitors at $z \gtrsim 2$ undergo bursty star formation, our GCE analysis shows that the observed, near-unity slope ($k \approx 1.0$) is inconsistent with a significant early starburst. A shallow MDF tail ($k \simeq 0.6$), the expected signature of a massive initial burst, is not observed. Instead, the exponential decline is best reproduced by models with sustained, low-level inflow and star formation.

Cosmological zoom-in simulations show that Milky Way-mass progenitors at $z \gtrsim 2$ undergo bursty star formation. In the cold-mode regime of Dekel & Birnboim (2006), efficient cooling prevents a stable virial shock for halos with $M_{\rm vir} \lesssim 10^{12}\,M_\odot$ at $z \gtrsim 2$, so dense, cold inflow drives recurrent bursts. In FIRE, these bursts are followed by energetic winds with mass loading $\eta \sim 10$ that expel a substantial fraction of the ISM and transiently depress the SFR (Muratov et al. 2015). In VINTERGATAN, high-z bursts arise from a turbulent, cold ISM reacting to frequent mergers (Segovia Otero et al. 2022). In FIRE-2, the bursty phase ends as a quasi-static, hot inner CGM develops when the halo crosses $M_{\rm vir} \sim 10^{12}\,M_\odot$ (Stern et al. 2021), accompanied by a rapid ($\lesssim 1$ Gyr) transition to time-steady, rotation-supported disk star formation (Gurvich et al. 2023). Throughout the proto-galaxy phase, the ISM remains turbulent and regulated by coupled inflow-outflow cycles; feedback-driven outflows throttle the net supply to the nascent disk until the accretion mode changes, but do not fully quench star formation (Semenov et al. 2025).

This early, burst-dominated mode plausibly overlaps the early epoch of disk formation (Yu et al. 2021). Kinematic diagnostics indicate that dispersion-dominated (low-circularity) orbits are characteristic of stars born during the proto-galaxy and early disk phase (Yu et al.



A smooth tail favors quiet assembly but cannot exclude every smaller burst.

2023; McCluskey et al. 2024), and stars formed in these early bursts can contribute to observed properties at $z \sim 0$ attributed to merger debris such as GSE (Rey et al. 2023). In our GCE framework, a shallow metal-poor tail ($k \approx 0.6$) arises from a large initial star-formation burst followed by ISM heating and weak subsequent inflow, whereas sustained gas accretion and weaker initial bursts generally yields steeper tails ($k \gtrsim 1$). MDFs inferred from Gaia XP catalogs appear smooth-lacking obvious multi-modal structure like the episodic features modeled by Ting & Ji (2025)-but measurement uncertainties can wash out burst-scale signatures, especially since the surviving stellar mass in an individual burst is less than 10^7 $M_\odot$. Kurbatov et al. (2024) estimated the stellar mass of the in situ "Aurora" and GSE to be $\sim 3\times10^7$ $M_\odot$ between $-3 <$ [Fe/H] < -2.0 which would prefer a steeper slope ($k \simeq 1.0$–1.4) based on Fig. 10 c) and 11.

Measurement errors can wash out burst-scale structure.

6. CONCLUSION

The chemical and kinematic structure of the oldest stars encodes the conditions under which the Milky Way first assembled. In our Paper I, we showed that the moderate rise in [α/Fe] near [Fe/H] ≈ -1.3 is best explained by the onset of α-enhanced inflow that supplied both fuel and angular momentum as the Galactic disk began to form. This feature identified the transition from the turbulent, dispersion-dominated Aurora phase to a rotation-supported system, but its amplitude alone could not uniquely constrain the proto-Galaxy's gas reservoir, inflow rate, or star-formation efficiency.

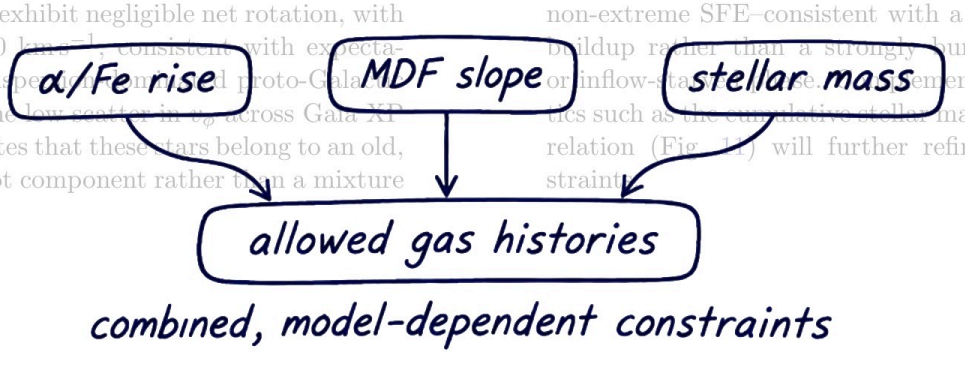
In this paper, we extend that analysis by turning to an independent diagnostic, the slope of the metal-poor tail of the Milky Way's MDF, which encodes the cumulative balance between metal production and gas supply during the earliest phase of Galactic evolution. Using metallicities from multiple Gaia DR3 XP-based catalogs, we quantify the logarithmic slopes of the MDF over $-3 <$ [Fe/H] < -2.0 and interpret them with a grid of one-zone, non-IRA GCE models calibrated to reproduce the [α/Fe] rise. This approach links the metal-poor MDF directly to the physical regulation of the proto-Galactic gas reservoir and provides a complementary constraint to abundance-ratio-based inferences.

Our main results are:

1. Kinematic diagnostics confirm that stars with [Fe/H] $\lesssim -1.5$ exhibit negligible net rotation, with median $v_\phi \approx 0$ km s^{-1}, consistent with expectations for the dispersion-supported proto-Galactic population. The low scatter in v_ϕ across Gaia XP catalogs indicates that these stars belong to an old, dynamically hot component rather than a mixture of rotating substructures. An apparent excess of stars with $v_\phi > 200$ km s^{-1} near [Fe/H] $\gtrsim -1.5$ (Fig. 2) is not a true halo feature but arises from thin-disk contaminants whose metallicities are underestimated in certain inference pipelines, particularly those trained primarily on red-giant stars or limited CMD domains.

2. MDFs from all Gaia XP catalogs, after the recommended quality cuts, are broadly consistent with single-slope exponentials over $-3 <$ [Fe/H] < -2.0, with slopes centered near $k \approx 1$ and an interquartile range of 0.6–2.0 (Figs. 5, 6, and 7). The slope remains largely invariant with Galactocentric radius ($R \lesssim 10$ kpc) and with parallax-corrected magnitude, indicating limited qualitative sensitivity to spatial incompleteness or luminosity-class selection. Modest inter-catalog differences follow known design choices: APOGEE-anchored label-transfer pipelines yield slightly steeper tails ($k \gtrsim 1.2$) owing to limited VMP training coverage, whereas VMP-optimized or CaHK-augmented models produce flatter tails ($k \lesssim 0.8$). These catalog-dependent effects plausibly explain much of the observed spread in k, but they do not alter the main conclusion that the Gaia XP metallicity catalogs consistently recover a broadly smooth, approximately exponential metal-poor tail.

3. Our grid of one-zone GCE models quantifies the direct link between the MDF tail slope and the proto-Galaxy's gas-regulation parameters. Shallow tails ($k \simeq 0.6$) occur only when the initial cold gas reservoir exceeds $\sim 10^9\,M_\odot$, the SFE is below $10^{-10}\,\mathrm{yr}^{-1}$, and inflow contributes less than 20% of the initial mass over the first Gyr. Such models produce an early starburst that rapidly heats and displaces the ISM, leading to a steep decline in the SFR. Conversely, steeper tails ($k \simeq$1.2–1.4) arise for smaller initial reservoirs ($\lesssim 10^8\,M_\odot$) with sustained, low-level inflow and SFE of order $10^{-9}\,\mathrm{yr}^{-1}$, indicative of gradual, infall-driven growth (Figs. 9 and 10). The Gaia XP-based slopes cluster near $k \simeq 1.0 \pm 0.2$, implying a proto-Galaxy with a moderate initial gas mass of $10^{8-9}\,M_\odot$, moderate inflow, and non-extreme SFE–consistent with a "quiet" early buildup rather than a strongly burst-dominated or inflow-starved system. Complementary diagnostics such as the cumulative stellar mass-metallicity relation (Fig. 11) will further refine these constraints.



The next step is to model survey selection alongside chemical evolution.

4. Among the GCE models that reproduce the observed [α/Fe]-rise associated with the Galactic spin-up, the modal metal-poor-tail slope is $k \simeq 1.2$. In the Auriga simulation suite, Milky Way-mass analogs exhibit nearly identical slopes ($k \simeq 1.25$) with little galaxy-to-galaxy scatter (Fig. 13). This further demonstrates the constraining power of the observed [α/Fe]-rise on the proto-Galaxy. These simulated systems assemble their baryons gradually over the first $\sim$1–2 Gyr without a large, isolated initial starburst, closely resembling our $k \simeq 1.0$ GCE models in which gas reservoirs build slowly under low-to-moderate inflow (Fig. 10 and 12). The concordance between the observed, modeled, and simulated slopes indicates a shared early phase in which a turbulent, feedback-regulated ISM throttles net accretion, naturally producing near-exponential metal-poor tails ($k \approx 1 - 1.2$) prior to the onset of α-enhanced refueling at the start of disk formation.

Our results, derived from the chemical imprint of the gas-regulation history, provide a diagnostic of the proto-Galaxy's initial conditions that is independent of—and more robust than—kinematic signatures. This distinction is crucial because the kinematic record of the earliest phase can be overwritten by later disruptive events. Orkney & Laporte (2025) demonstrates that massive radial mergers like the GSE can fundamentally scramble orbital information, making the kinematics of the proto-Galaxy's stars and stars kicked up from the disc (the Splash) indistinguishable at low metallicities. Their work shows that stellar kinematics at $z = 0$ rarely recover the true timing of the disc's initial spin-up due to this merger-induced scrambling. In contrast, the MDF is a product of the galaxy's internal gas supply and enrichment history, which is more resilient to such dynamical shuffling. Therefore, our method based on the MDF slope offers a more direct window into the conditions of the pre-merger, proto-Galactic gas reservoir, providing complementary constraints to purely kinematic approaches.

The Gaia XP metallicity catalogs differ in their training sets, flux calibrations, and CMD masks, and residual systematics-particularly near the Galactic plane-may still affect completeness and flatten the counts at [Fe/H] $\lesssim -2.5$. Our present one-zone framework, while flexible, assumes constant inflow within two one-Gyr phases and does not yet include explicitly time-variable, merger-driven accretion. Future work should incorporate a hierarchical forward model of the MDF that embeds each catalog's selection function within a multi-zone GCE framework, allowing inflow, star-formation efficiency, and feedback to vary continuously in time while marginalizing over Type Ia delay-time distributions and uncertain yields. The next generation of datasets-Gaia DR4/DR5, WEAVE (Dalton et al. 2012; Jin et al. 2024), and 4MOST (de Jong et al. 2019)-will provide higher-fidelity kinematics and XP-based [α/M] and [C/Fe] ratios, enabling cleaner separation of disk contaminants and improved recovery of the faintest metal-poor populations. Together, these advances will allow the exponential VMP tail to serve as a precise, physically interpretable observable linking the Milky Way's earliest gas supply, enrichment, and star-formation history to the emergence of its first long-lived disk.

This forward-modeling framework is proposed future work.

proposed

gas history → GCE model → survey selection → predicted catalog MDF

compare with observed MDFs

REFERENCES

Abohalima, A., & Frebel, A. 2018, ApJS, 238, 36, doi: 10.3847/1538-4365/aadfe9

Allende Prieto, C., Fernández-Alvar, E., Schlesinger, K. J., et al. 2014, A&A, 568, A7, doi: 10.1051/0004-6361/201424053

An, D., Beers, T. C., Johnson, J. A., et al. 2013, ApJ, 763, 65, doi: 10.1088/0004-637X/763/1/65

Andrae, R., Rix, H.-W., & Chandra, V. 2023a, ApJS, 267, 8, doi: 10.3847/1538-4365/acd53e

Andrae, R., Fouesneau, M., Sordo, R., et al. 2023b, A&A, 674, A27, doi: 10.1051/0004-6361/202243462

Ardern-Arentsen, A., Kane, S. G., Belokurov, V., et al. 2025, MNRAS, 537, 1984, doi: 10.1093/mnras/staf096

Ardern-Arentsen, A., Monari, G., Queiroz, A. B. A., et al. 2024, MNRAS, 530, 3391, doi: 10.1093/mnras/stae1049

Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, ApJ, 935, 167, doi: 10.3847/1538-4357/ac7c74

Beers, T. C., Preston, G. W., & Shectman, S. A. 1985, AJ, 90, 2089, doi: 10.1086/113917

—. 1992, AJ, 103, 1987, doi: 10.1086/116207

Bellazzini, M., Massari, D., Ceccarelli, E., et al. 2024, A&A, 683, A136, doi: 10.1051/0004-6361/202348106

Belokurov, V., & Kravtsov, A. 2022, MNRAS, 514, 689, doi: 10.1093/mnras/stac1267

Bessell, M. S., & Norris, J. 1984, ApJ, 285, 622, doi: 10.1086/162539

Bonifacio, P., Monaco, L., Salvadori, S., et al. 2021, A&A, 651, A79, doi: 10.1051/0004-6361/202140816

Borbolato, L., Rossi, S., Perottoni, H. D., et al. 2025, arXiv e-prints, arXiv:2504.00135, doi: 10.48550/arXiv.2504.00135

Brauer, K., Andales, H. D., Ji, A. P., et al. 2022, ApJ, 937, 14, doi: 10.3847/1538-4357/ac85b9
Bressan, A., Marigo, P., Girardi, L., et al. 2012, MNRAS, 427, 127, doi: 10.1111/j.1365-2966.2012.21948.x
Brusadin, G., Matteucci, F., & Romano, D. 2013, A&A, 554, A135, doi: 10.1051/0004-6361/201220884
Caffau, E., Bonifacio, P., François, P., et al. 2011, Nature, 477, 67, doi: 10.1038/nature10377
Carney, B. W., Laird, J. B., Latham, D. W., & Aguilar, L. A. 1996, AJ, 112, 668, doi: 10.1086/118042
Carollo, D., Beers, T. C., Lee, Y. S., et al. 2007, Nature, 450, 1020, doi: 10.1038/nature06460
Carter, C., Conroy, C., Zaritsky, D., et al. 2021, ApJ, 908, 208, doi: 10.3847/1538-4357/abcda4
Chandra, V., Semenov, V. A., Rix, H.-W., et al. 2024, ApJ, 972, 112, doi: 10.3847/1538-4357/ad5b60
Chen, B., Ting, Y.-S., & Hayden, M. 2024, PASA, 41, e063, doi: 10.1017/pasa.2024.56
Chiba, M., & Beers, T. C. 2000, AJ, 119, 2843, doi: 10.1086/301409
Chiti, A., Mardini, M. K., Frebel, A., & Daniel, T. 2021, ApJL, 911, L23, doi: 10.3847/2041-8213/abd629
Christlieb, N., Schörck, T., Frebel, A., et al. 2008, A&A, 484, 721, doi: 10.1051/0004-6361:20078748
Christlieb, N., Bessell, M. S., Beers, T. C., et al. 2002, Nature, 419, 904, doi: 10.1038/nature01142
Conroy, C., Bonaca, A., Cargile, P., et al. 2019, ApJ, 883, 107, doi: 10.3847/1538-4357/ab38b8
Conroy, C., Weinberg, D. H., Naidu, R. P., et al. 2022, arXiv e-prints, arXiv:2204.02989, doi: 10.48550/arXiv.2204.02989
Da Costa, G. S., Bessell, M. S., Mackey, A. D., et al. 2019, MNRAS, 489, 5900, doi: 10.1093/mnras/stz2550
Dalton, G., Trager, S. C., Abrams, D. C., et al. 2012, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 8446, Ground-based and Airborne Instrumentation for Astronomy IV, ed. I. S. McLean, S. K. Ramsay, & H. Takami, 84460P, doi: 10.1117/12.925950
De Angeli, F., Weiler, M., Montegriffo, P., et al. 2023, A&A, 674, A2, doi: 10.1051/0004-6361/202243680
de Jong, R. S., Agertz, O., Berbel, A. A., et al. 2019, The Messenger, 175, 3, doi: 10.18727/0722-6691/5117
Deason, A. J., Belokurov, V., & Evans, N. W. 2011, MNRAS, 411, 1480, doi: 10.1111/j.1365-2966.2010.17785.x
Dekel, A., & Birnboim, Y. 2006, MNRAS, 368, 2, doi: 10.1111/j.1365-2966.2006.10145.x
El-Badry, K., Wetzel, A., Geha, M., et al. 2016, ApJ, 820, 131, doi: 10.3847/0004-637X/820/2/131
El-Badry, K., Bland-Hawthorn, J., Wetzel, A., et al. 2018, MNRAS, 480, 652, doi: 10.1093/mnras/sty1864
Fallows, C. P., & Sanders, J. L. 2024, MNRAS, 531, 2126, doi: 10.1093/mnras/stae1303
Fattahi, A., Belokurov, V., Deason, A. J., et al. 2019, MNRAS, 484, 4471, doi: 10.1093/mnras/stz159
Fernández-Alvar, E., Kordopatis, G., Hill, V., et al. 2021, MNRAS, 508, 1509, doi: 10.1093/mnras/stab2617
Feuillet, D. K., Frankel, N., Lind, K., et al. 2019, MNRAS, 489, 1742, doi: 10.1093/mnras/stz2221
Fragkoudi, F., Grand, R. J. J., Pakmor, R., et al. 2020, MNRAS, 494, 5936, doi: 10.1093/mnras/staa1104
Frebel, A., Aoki, W., Christlieb, N., et al. 2005, Nature, 434, 871, doi: 10.1038/nature03455
Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, A&A, 595, A1, doi: 10.1051/0004-6361/201629272
Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, A&A, 674, A1, doi: 10.1051/0004-6361/202243940
Grand, R. J. J., Fragkoudi, F., Gómez, F. A., et al. 2024, Monthly Notices of the Royal Astronomical Society, 532, 1814, doi: 10.1093/mnras/stae1598
Grand, R. J. J., Gómez, F. A., Marinacci, F., et al. 2017, Monthly Notices of the Royal Astronomical Society, 467, 179, doi: 10.1093/mnras/stx071
Gurvich, A. B., Stern, J., Faucher-Giguère, C.-A., et al. 2023, MNRAS, 519, 2598, doi: 10.1093/mnras/stac3712
Hartwick, F. D. A. 1976, ApJ, 209, 418, doi: 10.1086/154735
Hattori, K. 2025, ApJ, 980, 90, doi: 10.3847/1538-4357/ad9686
Haywood, M., Di Matteo, P., Lehnert, M. D., Katz, D., & Gómez, A. 2013, A&A, 560, A109, doi: 10.1051/0004-6361/201321397
Hopkins, P. F., Kereš, D., Oñorbe, J., et al. 2014, MNRAS, 445, 581, doi: 10.1093/mnras/stu1738
Hopkins, P. F., Gurvich, A. B., Shen, X., et al. 2023, MNRAS, 525, 2241, doi: 10.1093/mnras/stad1902
Horta, D., & Schiavon, R. P. 2025, MNRAS, 537, 3730, doi: 10.1093/mnras/staf256
Horta, D., Cunningham, E. C., Sanderson, R., et al. 2024, MNRAS, 527, 9810, doi: 10.1093/mnras/stad3834
Huber, D., Slumstrup, D., Hon, M., et al. 2024, ApJ, 975, 19, doi: 10.3847/1538-4357/ad7110
Iza, F. G., Scannapieco, C., Nuza, S. E., et al. 2022, MNRAS, 517, 832, doi: 10.1093/mnras/stac2709
Jin, S., Trager, S. C., Dalton, G. B., et al. 2024, MNRAS, 530, 2688, doi: 10.1093/mnras/stad557
Kane, S. G., Belokurov, V., Cranmer, M., et al. 2025, MNRAS, 536, 2507, doi: 10.1093/mnras/stae2689
—. 2024, arXiv e-prints, arXiv:2409.00197, doi: 10.48550/arXiv.2409.00197

Keller, S. C., Schmidt, B. P., Bessell, M. S., et al. 2007, PASA, 24, 1, doi: 10.1071/AS07001
Khalatyan, A., Anders, F., Chiappini, C., et al. 2024, A&A, 691, A98, doi: 10.1051/0004-6361/202451427
Khoperskov, S., Minchev, I., Libeskind, N., et al. 2023, A&A, 677, A91, doi: 10.1051/0004-6361/202244234
Kurbatov, E. P., Belokurov, V., Koposov, S., et al. 2024, arXiv e-prints, arXiv:2410.22250, doi: 10.48550/arXiv.2410.22250
Li, H., Aoki, W., Matsuno, T., et al. 2022, ApJ, 931, 147, doi: 10.3847/1538-4357/ac6514
Li, J., Wong, K. W. K., Hogg, D. W., Rix, H.-W., & Chandra, V. 2024, ApJS, 272, 2, doi: 10.3847/1538-4365/ad2b4d
Liu, S., Du, C., Newberg, H. J., et al. 2018, ApJ, 862, 163, doi: 10.3847/1538-4357/aacf91
Lowe, B. D. C., Nordlander, T., Casagrande, L., et al. 2025, PASA, 42, e040, doi: 10.1017/pasa.2025.16
Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, AJ, 154, 94, doi: 10.3847/1538-3881/aa784d
Malhan, K., & Rix, H.-W. 2024, ApJ, 964, 104, doi: 10.3847/1538-4357/ad1885
Marinacci, F., Grand, R. J. J., Pakmor, R., et al. 2017, MNRAS, 466, 3859, doi: 10.1093/mnras/stw3366
Marocco, F., Eisenhardt, P. R. M., Fowler, J. W., et al. 2021, ApJS, 253, 8, doi: 10.3847/1538-4365/abd805
Martin, N. F., Starkenburg, E., Yuan, Z., et al. 2024, A&A, 692, A115, doi: 10.1051/0004-6361/202347633
Matijevič, G., Chiappini, C., Grebel, E. K., et al. 2017, A&A, 603, A19, doi: 10.1051/0004-6361/201730417
McCluskey, F., Wetzel, A., Loebman, S. R., et al. 2024, MNRAS, 527, 6926, doi: 10.1093/mnras/stad3547
Merrow, A., Grand, R. J. J., Fragkoudi, F., & Martig, M. 2024, MNRAS, 531, 1520, doi: 10.1093/mnras/stae1250
Muratov, A. L., Kereš, D., Faucher-Giguère, C.-A., et al. 2015, MNRAS, 454, 2691, doi: 10.1093/mnras/stv2126
Myeong, G. C., Evans, N. W., Belokurov, V., Sanders, J. L., & Koposov, S. E. 2018, ApJL, 856, L26, doi: 10.3847/2041-8213/aab613
Naidu, R. P., Conroy, C., Bonaca, A., et al. 2020, ApJ, 901, 48, doi: 10.3847/1538-4357/abaef4
Neeleman, M., Prochaska, J. X., Kanekar, N., & Rafelski, M. 2020, Nature, 581, 269, doi: 10.1038/s41586-020-2276-y
Nepal, S., Chiappini, C., Queiroz, A. B., et al. 2024, A&A, 688, A167, doi: 10.1051/0004-6361/202449445
Nomoto, K., Kobayashi, C., & Tominaga, N. 2013, ARA&A, 51, 457, doi: 10.1146/annurev-astro-082812-140956
Orkney, M. D. A., & Laporte, C. F. P. 2025, arXiv e-prints, arXiv:2509.09576, doi: 10.48550/arXiv.2509.09576
Pakmor, R., Gómez, F. A., Grand, R. J. J., et al. 2017, MNRAS, 469, 3185, doi: 10.1093/mnras/stx1074
Pillepich, A., Nelson, D., Springel, V., et al. 2019, MNRAS, 490, 3196, doi: 10.1093/mnras/stz2338
Prantzos, N. 2003, A&A, 404, 211, doi: 10.1051/0004-6361:20030445
—. 2008, A&A, 489, 525, doi: 10.1051/0004-6361:20079330
Queiroz, A. B. A., Anders, F., Santiago, B. X., et al. 2018, MNRAS, 476, 2556, doi: 10.1093/mnras/sty330
Renaud, F., Agertz, O., Read, J. I., et al. 2021, MNRAS, 503, 5846, doi: 10.1093/mnras/stab250
Rey, M. P., Agertz, O., Starkenburg, T. K., et al. 2023, MNRAS, 521, 995, doi: 10.1093/mnras/stad513
Rix, H.-W., Chandra, V., Andrae, R., et al. 2022, ApJ, 941, 45, doi: 10.3847/1538-4357/ac9e01
Ryan, S. G., & Norris, J. E. 1991, AJ, 101, 1865, doi: 10.1086/115812
Sahlholdt, C. L., Feltzing, S., & Feuillet, D. K. 2022, MNRAS, 510, 4669, doi: 10.1093/mnras/stab3681
Salvadori, S., Schneider, R., & Ferrara, A. 2007, MNRAS, 381, 647, doi: 10.1111/j.1365-2966.2007.12133.x
Santistevan, I. B., Wetzel, A., Sanderson, R. E., et al. 2021, MNRAS, 505, 921, doi: 10.1093/mnras/stab1345
Schörck, T., Christlieb, N., Cohen, J. G., et al. 2009, A&A, 507, 817, doi: 10.1051/0004-6361/200810925
Segovia Otero, Á., Renaud, F., & Agertz, O. 2022, MNRAS, 516, 2272, doi: 10.1093/mnras/stac2368
Semenov, V. A., Conroy, C., Smith, A., Puchwein, E., & Hernquist, L. 2025, ApJ, 990, 7, doi: 10.3847/1538-4357/addf48
Sestito, F., Longeard, N., Martin, N. F., et al. 2019, MNRAS, 484, 2166, doi: 10.1093/mnras/stz043
Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, AJ, 131, 1163, doi: 10.1086/498708
Smith, M. C., Ruchti, G. R., Helmi, A., et al. 2007, MNRAS, 379, 755, doi: 10.1111/j.1365-2966.2007.11964.x
Snaith, O., Haywood, M., Di Matteo, P., et al. 2015, A&A, 578, A87, doi: 10.1051/0004-6361/201424281
Springel, V. 2010, MNRAS, 401, 791, doi: 10.1111/j.1365-2966.2009.15715.x
Starkenburg, E., Martin, N., Youakim, K., et al. 2017, MNRAS, 471, 2587, doi: 10.1093/mnras/stx1068
Stern, J., Faucher-Giguère, C.-A., Fielding, D., et al. 2021, ApJ, 911, 88, doi: 10.3847/1538-4357/abd776
Suda, T., Katsuta, Y., Yamada, S., et al. 2008, PASJ, 60, 1159, doi: 10.1093/pasj/60.5.1159
Ting, Y.-S., & Ji, A. 2025, The Open Journal of Astrophysics, 8, 7, doi: 10.33232/001c.128414

Tumlinson, J. 2010, ApJ, 708, 1398, doi: 10.1088/0004-637X/708/2/1398

Vincenzo, F., Weinberg, D. H., Miglio, A., Lane, R. R., & Roman-Lopes, A. 2021, MNRAS, 508, 5903, doi: 10.1093/mnras/stab2899

Vogelsberger, M., Genel, S., Sijacki, D., et al. 2013, MNRAS, 436, 3031, doi: 10.1093/mnras/stt1789

Wisotzki, L., Christlieb, N., Bade, N., et al. 2000, A&A, 358, 77, doi: 10.48550/arXiv.astro-ph/0004162

Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, AJ, 140, 1868, doi: 10.1088/0004-6256/140/6/1868

Xiang, M., & Rix, H.-W. 2022, Nature, 603, 599, doi: 10.1038/s41586-022-04496-5

Xiang, M., Rix, H.-W., Yang, H., et al. 2025, Nature Astronomy, 9, 101, doi: 10.1038/s41550-024-02382-w

Yang, L., Yuan, H., Huang, B., et al. 2025, ApJS, 279, 7, doi: 10.3847/1538-4365/add5e3

Yanny, B., Rockosi, C., Newberg, H. J., et al. 2009, AJ, 137, 4377, doi: 10.1088/0004-6256/137/5/4377

Ye, X., Wu, W., Allende Prieto, C., et al. 2025, A&A, 695, A75, doi: 10.1051/0004-6361/202452871

Youakim, K., Starkenburg, E., Martin, N. F., et al. 2020, MNRAS, 492, 4986, doi: 10.1093/mnras/stz3619

Yu, S., Bullock, J. S., Klein, C., et al. 2021, MNRAS, 505, 889, doi: 10.1093/mnras/stab1339

Yu, S., Bullock, J. S., Gurvich, A. B., et al. 2023, MNRAS, 523, 6220, doi: 10.1093/mnras/stad1806

Yuan, Z., Li, C., Martin, N. F., et al. 2024, A&A, 691, L1, doi: 10.1051/0004-6361/202348593

Zhang, H., Belokurov, V., Evans, N. W., et al. 2024, MNRAS, 535, 2873, doi: 10.1093/mnras/stae2546

Zhang, X., Green, G. M., & Rix, H.-W. 2023, MNRAS, 524, 1855, doi: 10.1093/mnras/stad1941

Zhao, G., Zhao, Y.-H., Chu, Y.-Q., Jing, Y.-P., & Deng, L.-C. 2012, Research in Astronomy and Astrophysics, 12, 723, doi: 10.1088/1674-4527/12/7/002