



Reading the CARDS: The Imprint of Accretion History in the Chemical Abundances of the Milky Way’s Stellar Halo

Emily C. Cunningham¹ , Robyn E. Sanderson^{1,2} , Kathryn V. Johnston^{1,3} , Nondh Panithanpaisal² , Melissa K. Ness^{1,3} , Andrew Wetzel⁴ , Sarah R. Loebman⁵ , Ivanna Escala^{6,7,8,11} , Danny Horta⁹, and Claude-André Faucher-Giguère¹⁰

¹ Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA; ecunningham@flatironinstitute.org

² Department of Physics and Astronomy, University of Pennsylvania, 209 S 33rd Street, Philadelphia, PA 19104, USA

³ Department of Astronomy, Columbia University, 550 West 120th Street, New York, NY, 10027, USA

⁴ Department of Physics & Astronomy, University of California, Davis, CA 95616, USA

⁵ Department of Physics, University of California, Merced, 5200 Lake Road, Merced, CA 95343, USA

⁶ Department of Astronomy, California Institute of Technology, 1200 E California Boulevard, Pasadena, CA 91125, USA

⁷ The Observatories of the Carnegie Institution for Science, 813 Santa Barbara Street, Pasadena, CA 91101, USA

⁸ Department of Astrophysical Sciences, Princeton University, 4 Ivy Lane, Princeton, NJ 08544, USA

⁹ Astrophysics Research Institute, Brownlow Hill, Liverpool L3 5RF, UK

¹⁰ Department of Physics and Astronomy and CIERA, Northwestern University, 1800 Sherman Avenue, Evanston, IL 60201, USA

Received 2021 October 6; revised 2022 May 15; accepted 2022 June 13; published 2022 August 4

Abstract

In the era of large-scale spectroscopic surveys in the Local Group, we can explore using chemical abundances of halo stars to study the star formation and chemical enrichment histories of the dwarf galaxy progenitors of the Milky Way (MW) and M31 stellar halos. In this paper, we investigate using the chemical abundance ratio distributions (CARDS) of seven stellar halos from the Latte suite of FIRE-2 simulations. We attempt to infer galaxies’ assembly histories by modeling the CARDS of the stellar halos of the Latte galaxies as a linear combination of *template* CARDS from disrupted dwarfs, with different stellar masses M_* and quenching times t_{100} . We present a method for constructing these templates using present-day dwarf galaxies. For four of the seven Latte halos studied in this work, we recover the mass spectrum of accreted dwarfs to a precision of $<10\%$. For the fraction of mass accreted as a function of t_{100} , we find the residuals of 20%–30% for five of the seven simulations. We discuss the failure modes of this method, which arise from the diversity of star formation and chemical enrichment histories that dwarf galaxies can take. These failure cases can be robustly identified by the high model residuals. Although the CARDS modeling method does not successfully infer the assembly histories in these cases, the CARDS of these disrupted dwarfs contain signatures of their unusual formation histories. Our results are promising for using CARDS to learn more about the histories of the progenitors of the MW and M31 stellar halos.

Unified Astronomy Thesaurus concepts: Milky Way stellar halo (1060); Chemical abundances (224); Galaxy chemical evolution (580)

1. Introduction

In the hierarchical paradigm of galaxy evolution, the Milky Way (MW) is more than just one galaxy. As many as thousands of dwarf galaxies are predicted to have contributed to the assembly of the MW dark matter halo over cosmic time; the remains of these accreted systems make up the component of our Galaxy known as the stellar halo. The formation conditions of halo stars are imprinted in their chemical abundances. Therefore, through observations of the chemical properties of MW halo stars, we have the unique opportunity to study galaxy formation and evolution in the high-redshift universe: we can study the remains of galaxies that did not survive to present day, including populations that are too faint to observe directly at high redshift, on a star-by-star basis.

In the current era of MW studies, we have unprecedented knowledge of the motions and chemical properties of the stars in our Galaxy. This is in large part due to the data releases from the Gaia mission (Gaia Collaboration et al. 2016, 2018, 2021)

as well as complementary wide-field spectroscopic programs, such as the GALAH survey (De Silva et al. 2015), the H3 survey (Conroy et al. 2019b), APOGEE (Majewski et al. 2017), RAVE (Steinmetz et al. 2006), SEGUE (Yanny et al. 2009), and LAMOST (Cui et al. 2012). Our knowledge of the chemodynamical structure of our Galaxy will only continue to grow, with future data releases from the Gaia mission as well as the release of data from current ongoing and planned spectroscopic surveys, including WEAVE (Dalton et al. 2014), 4MOST (de Jong et al. 2019), DESI (Allende Prieto et al. 2020), and SDSS-V (Kollmeier et al. 2017). Thanks to these programs, we now have 6D phase space information as well as up to 30+ dimensions of chemical information (e.g., GALAH + DR3; Buder et al. 2021) for millions of stars in the Galaxy, a truly groundbreaking achievement.

Both kinematic and chemical abundance measurements can be pertinent to a star’s origin: the dynamical times of infalling and disrupting dwarfs are long compared to the age of the Galaxy (especially for stars in the outer regions of the halo). As a result, the kinematic properties of halo stars can retain a link to their initial conditions. Numerous studies have used Gaia data plus spectroscopic information to discover new substructures in the MW halo. The early Gaia data releases, cross-matched with spectroscopic surveys (SDSS and RAVE), enabled the discovery of the “Gaia-Sausage-Enceladus”

¹¹ Carnegie-Princeton Fellow.

(hereafter GSE; Belokurov et al. 2018; Helmi et al. 2018; Haywood et al. 2018; Koppelman et al. 2018; though we note that earlier work identified GSE members and discussed their possible accretion origin; such as Brook et al. 2003; and Nissen & Schuster 2010). The GSE substructure is a relatively metal-rich ($([\text{Fe}/\text{H}]) \sim -1.1$) population of stars on radially biased orbits, believed to be the debris from a massive, early accretion event and thought to be the dominant progenitor of the inner stellar halo. In addition, a population of stars on radial orbits, but with chemical abundances consistent with the MW thick disk, were also identified: it has been suggested that this population, referred to as the “in-situ halo” or the “Splash,” initially formed in the MW disk but was perturbed onto halo-like orbits by early mergers (Bonaca et al. 2017, 2020; Haywood et al. 2018; Di Matteo et al. 2019; Belokurov et al. 2020). While these two populations have been found to dominate local halo samples (see, e.g., Naidu et al. 2020), numerous additional substructures have been found through MW chemodynamical studies, including the Helmi stream (Helmi et al. 1999; Koppelman et al. 2019b), Sequoia (Myeong et al. 2019), Thamnos (Koppelman et al. 2019a), Aleph, Arjuna, I’itoi, and Wukong (Naidu et al. 2020). Through these chemodynamical studies, many of the major building blocks of the MW have been discovered and characterized (although we note that it remains challenging to link substructures to their dwarf galaxy progenitors, given that single progenitors can be responsible for multiple, overlapping substructures; see, e.g., Jean-Baptiste et al. 2017).

The remains of lower-mass galaxies accreted at early times are harder to find. Simulations have long predicted that massive dwarfs will contribute the most by mass to the stellar halo (e.g., Bullock & Johnston 2005; Santistevan et al. 2020), and are even predicted to contribute the vast majority of the halo’s metal-poor stars (e.g., Deason et al. 2016). This prediction has been borne out in the observations: debris from GSE has been found to dominate nearly every observational sample of halo stars that we have in the Gaia era, resulting in a stellar halo that is, on average, relatively metal-rich (Conroy et al. 2019a). In addition to the fact that they contribute small numbers of stars to the halo, many of the low-mass dwarf galaxies accreted by the MW at early times are difficult to identify because they are likely to be phase-mixed. This is particularly challenging when studying the inner halo (where we have the highest-quality spectroscopic data). While the dynamical times in the outer halo are long compared to the age of the Galaxy, enabling debris from common progenitors to be identified from kinematic measurements, the early low-mass progenitors of the inner halo are more likely to be phase-mixed (e.g., Johnston et al. 2008), having lost their kinematic memories to the evolving MW potential. We therefore cannot rely on using kinematics alone to detect the remains of these systems: chemical abundance information is required to find stars that formed in this elusive population of galaxies.

The payoff of detecting and characterizing these low-luminosity systems within our own halo could be tremendous. For example, the faintest, early galaxies are hypothesized to be significant contributors in driving the epoch of reionization at $z > 6$ (e.g., Kuhlen & Faucher-Giguère 2012; Robertson et al. 2015; Weisz & Boylan-Kolchin 2017; Faucher-Giguère 2020). However, even JWST will not be able to observe galaxies less luminous than the Large Magellanic Cloud directly at $z \sim 7$ (Boylan-Kolchin et al. 2015). It has long been known that

archaeological studies of dwarf galaxies in the Local Group (LG) provide an avenue for studying faint, high- z galaxies (near-field cosmology; e.g., Freeman & Bland-Hawthorn 2002; Ricotti & Gnedin 2005; Bovill & Ricotti 2011; Brown et al. 2012; Boylan-Kolchin et al. 2015). In principle, the disrupted dwarfs that make up the M31 and MW stellar halos can be used to increase our sample of local systems used to study the faint, early galaxies with resolved stars. However, the challenge remains identifying these low-mass systems within the stellar halo.

In this paper, we explore how the distributions of stellar chemical abundances can be used to infer the properties of stellar halo building blocks. We consider the modeling framework first proposed by Lee et al. (2015; hereafter L15): they proposed modeling the chemical abundance ratio distribution (CARD) of a stellar halo as a linear combination of “templates.” Here, the “templates” represent the abundance distributions of accretion events of varying stellar masses and accretion times. L15 developed and explored this method utilizing the Bullock & Johnston (2005) simulations, which are N -body simulations of accreted dwarfs onto a MW-like parent galaxy. In the context of a purely accreted stellar halo, the total halo CARD is simply a linear combination of the CARDS from all of the different accreted dwarf galaxies. Through the use of templates for different masses M_{sat} and accretion times t_{acc} , L15 posed the following question:

“How accurately can we determine the fraction of total stellar mass, A_j , contributed by satellites of various mass (M_{sat}) and accretion time (t_{acc}) to a stellar halo given a set of templates for the distribution $f_j(x_d, M_{\text{sat}}, t_{\text{acc}})$ of chemical abundances x_d found in those satellites, and observations of CARDS ($f(x_d)$) in the stellar halo?”

Using this technique on the Bullock & Johnston (2005) simulations, L15 found that they were generally able to recover the fractional mass contributions of accreted satellites to within a factor of two, and found that this method was particularly sensitive to low-mass satellites.

L15 had the key insight that there should be information about the Galaxy’s assembly history contained within the full distribution of chemical abundances in the MW halo. However, the study had its limitations in its ability to test how accurately we can constrain the assembly history from CARDS, both as a result of the chosen simulations and the available observations. In L15, because they worked with the Bullock & Johnston (2005) simulations, which are N -body simulations that include no hydrodynamics, the stellar properties of the infalling dwarfs (including chemical abundances) required prescriptions. The chemical properties of the Bullock & Johnston (2005) halos, presented in Font et al. (2006), were derived on an individual satellite-by-satellite basis, using the semianalytic chemical enrichment method from Robertson et al. (2005). While L15 showed that the assembly histories of the Bullock & Johnston (2005) halos could be recovered reasonably well, they were not working with simulations that modeled chemical enrichment fully self-consistently. In addition, when L15 first published this technique, there were very few stars with measured kinematic and chemical properties outside of the solar neighborhood; the data available for exploiting this technique were severely limited.

At the time of writing this publication, we are privileged to have spectroscopic metallicities measured for millions of stars in the MW. Furthermore, the deeper spectroscopic observations

of M31 have enabled α -abundance measurements and iron metallicities derived from full spectral synthesis for an ever-growing sample of individual stars in M31's halo (Vargas et al. 2014b; Escala et al. 2019, 2020a, 2020b, 2021; Gilbert et al. 2019, 2020) and M31's satellites (Vargas et al. 2014a; Kirby et al. 2020; Wojno et al. 2020). With more knowledge than ever about the chemodynamical structure of the LG, we are now in a position to explore more deeply the potential of utilizing CARDS to constrain the formation histories of LG galaxies, and address how this method might aid in identifying the stars that were born in high-redshift, low-mass galaxies that cannot be observed directly.

In this paper, we undertake a critical next step in exploring the feasibility of using CARDS to constrain the a galaxy's assembly history: we explore how CARDS can be leveraged to constrain the formation histories of the FIRE-2 zoom-in cosmological simulations of MW-mass galaxies (introduced in Wetzel et al. 2016). These simulations self-consistently model chemical enrichment, and contain many complications that are not included in the Bullock & Johnston (2005) models. Examples of such complications include dwarf-dwarf interactions, live dark matter (DM) halos and disks, the time variability of the host potential, and the role of feedback in shaping dwarf galaxy star formation histories.

In order to make the CARDS method into a practical method to apply to observations, we seek to address three key questions in this work:

1. Does the method proposed by L15 of modeling the stellar halo CARD as a linear combination of templates work in a more realistic setting?
2. Since we require a strategy for an observationally viable method for constructing these templates, can we use the CARDS of present-day dwarf galaxies to infer the properties of the disrupted dwarfs?
3. How can we assess the accuracy of this method?

It is crucial that we definitively answer this second question, as observationally, the construction of empirical templates for halo progenitors are likely to be limited to a sample of nearby present-day dwarf galaxies, where the member stars can be identified with high confidence. Abundance distributions have been measured already for a number of local dwarf galaxies (see, e.g., Kirby et al. 2009, 2010, 2017, 2020; Nidever et al. 2020; Hasselquist et al. 2021). The number of satellites with abundances measured over their full spatial extent, and to fainter magnitudes, will only increase with upcoming surveys; for example, mapping these distributions is one of the main objectives of the Subaru Prime Focus Spectrograph Galactic Archaeology program (Takada et al. 2014). Therefore, from an observational perspective, local dwarf galaxies would be the ideal choice as the basis for empirical template CARDS for stellar halo building blocks.

However, it is well known that MW halo stars from disrupted dwarfs have, on average, higher α -abundances relative to stars in their surviving dwarf galaxy counterparts at the same metallicity (e.g., Venn et al. 2004; Naidu et al. 2022). This is a natural consequence of the fact that halo stars formed in dwarfs that did not survive until the present day (e.g., Robertson et al. 2005): the dwarf galaxies that survive until the present day are more likely to have extended star formation histories than halo progenitors. Therefore, if using stars in dwarf galaxies to constrain the properties of halo progenitors,

the potentially different timescales in star formation must be taken into account. In addition, given that the present-day dwarfs are likely to be in more isolated environments at early times than the halo progenitors, these different environments could have an effect on their chemical enrichment and evolution (e.g., Corlies et al. 2013).

In this work, we demonstrate a method for creating templates for the progenitors of stellar halos using exclusively present-day dwarf galaxies (accounting for the fact that the two populations are likely to be forming stars on different timescales), and assess the performance of these templates relative to the templates constructed from streams and phase-mixed debris. While an in-depth exploration and comparison of the chemical enrichment histories of present-day dwarfs versus halo progenitors is beyond the scope of this work, we can investigate whether halo progenitor CARD templates can be reliably constructed from a sample of dwarf galaxies.

This paper is organized as follows. In Section 2, we introduce the simulations used in this work. We summarize the key relevant details about the FIRE-2 zoom-in cosmological simulations in Section 2.1 and introduce the catalog of halo progenitors and present-day dwarf galaxies in Section 2.2. In Section 3, we explore the CARDS in the simulation and demonstrate their connection to the formation history of the host halo (in Section 3.1) and the dynamical state of the accreted satellite (i.e., classification as phase-mixed debris, coherent stream, or dwarf galaxy; in Section 3.2). In Section 4, we describe the CARDS technique (in Section 4.1) and our method for constructing CARD templates (in Section 4.2). In Section 5, we present the results of our modeling procedure for all the simulations used in this work, and discuss in detail a successful case as well as a less successful case. In Section 6, we discuss some of the limitations of this technique, and discuss potential improvements and extensions of the method presented here. We conclude in Section 7.

2. Simulations

In order to test the CARDS method in a more realistic setting, in this work, we make use of two suites of cosmological zoom-in simulations from the Feedback in Realistic Environments (FIRE) project.¹² We begin in Section 2.1 by summarizing the relevant properties of the simulations; in Section 2.2, we introduce the subset of stellar halo and dwarf galaxy star particles from the simulation that we use in this analysis.

2.1. FIRE-2 Simulations

All of the simulations used in this work were run using the GIZMO¹³ gravity plus hydrodynamics code in meshless finite-mass mode (Hopkins 2015) with the FIRE-2 physics model (Hopkins et al. 2018). While we refer the reader to the above papers for details of the FIRE-2 implementation, we summarize some of the key details here, especially those related to star formation and chemical enrichment.

The FIRE-2 gas model includes metallicity-dependent treatment of radiative heating and cooling processes, over a temperature range of 10–10¹⁰ K (Hopkins et al. 2018). The FIRE-2 simulations were evolved with the Faucher-Giguère et al. (2009) UV background model, in which reionization

¹² FIRE project website: <http://fire.northwestern.edu>.

¹³ <http://www.tapir.caltech.edu/~phopkins/Site/GIZMO.html>

occurs early ($z_{\text{reion}} \sim 10$). Star particles form from self-gravitating, cold, dense, molecular, Jeans-unstable gas clouds (following Krumholz & Gnedin 2011). Star particles are treated as individual stellar populations with a Kroupa (2002) initial mass function. Type II supernovae (SNe) rates, as well as mass- and metal-loss rates, are derived from stellar population models (STARBURST99; Leitherer et al. 1999). The metal yields for Type Ia SNe are taken from Iwamoto et al. (1999), and the yields for Type II SNe are from Nomoto et al. (2006).

The three sources of chemical enrichment in the FIRE-2 simulations are Type Ia SNe, Type II SNe, and stellar winds (produced predominantly by asymptotic giant branch stars and O-type stars). The simulations self-consistently track eleven chemical abundances: H, He, C, N, O, Ne, Mg, Si, S, Ca, and Fe. We utilize the runs implementing subgrid turbulent metal diffusion (Hopkins 2017), which Escala et al. (2018) demonstrated results in realistic spreads in the abundance distributions in dwarf galaxies, while not resulting in significant changes to star formation (Su et al. 2017).

In this work, our primary focus is a set of seven galaxies from the Latte suite of FIRE-2 simulations (introduced in Wetzel et al. 2016). The Latte suite are simulations of isolated galaxies with present-day masses in the range of $M_{200m} = 1\text{--}2 \times 10^{12} M_{\odot}$, similar to the MW and M31. These simulations are high resolution, with initial masses of star particles with $\sim 7000 M_{\odot}$. While the $z=0$ particle masses depend on age, the average particle mass is $\sim 5000 M_{\odot}$, as a result of stellar mass loss.

While we focus on the stellar halos of seven of the isolated MW analogs from the Latte suite in this paper, we also make use of the satellites from the ELVIS on FIRE suite (Garrison-Kimmel et al. 2019a, 2019b) to increase our sample of dwarf galaxies. The ELVIS on FIRE suite consists of three simulations, each containing a M31–MW analog pair; the DM halos selected for the zoom-in simulation were chosen to have the approximate distance and relative velocities of the MW and M31 (Garrison-Kimmel et al. 2019a). While the physics in the Latte runs and the ELVIS on FIRE runs we use in this work are the same, we note that the mass particle resolution in the ELVIS on FIRE suite is approximately twice as high, with initial star particle masses of $\sim 3500\text{--}4000 M_{\odot}$; however, for the analysis presented here, the difference in resolution does not affect the outcome.

The properties of the host galaxies in these simulations have been demonstrated to show broad agreement with the MW and M31, including the stellar–halo-mass relation (Hopkins et al. 2018), stellar halos (Bonaca et al. 2017; Sanderson et al. 2018), and the radial and vertical structure of their disks (Ma et al. 2017; Sanderson et al. 2020; Bellardini et al. 2021). Critically for this work, the satellite populations of both of these suites of simulations have also been shown to agree with many observed properties, such as their masses and velocity dispersions (Wetzel et al. 2016; Garrison-Kimmel et al. 2019a), star formation histories (Garrison-Kimmel et al. 2019b), and radial distributions (Samuel et al. 2020).

While the satellite galaxies in these simulations have been shown to have many properties in common with the observed satellites around the MW and M31, they are generally found to be too metal-poor compared to the observed dwarf galaxies (Escala et al. 2018; Wheeler et al. 2019; Panithanpaisal et al. 2021). This is likely at least in part because of the assumed delay-time distributions for SNe Ia. These simulations adopt a

SNe Ia rate with prompt and delayed components (Mannucci et al. 2006); adopting a SNe Ia delay-time distribution that results in a larger total number of SNe Ia events (e.g., a power-law rate; Maoz & Graur 2017) helps to resolve the discrepancy between the simulated and observed iron abundances (Gandhi et al. 2022). For this work, we emphasize that we do not require quantitative agreement with simulated and observed abundances. The empirical CARD templates we construct in this work from the simulations are intended for use within the simulations only. While the method presented here of using the empirical templates to infer the halo progenitor properties can be applied to observational data, we emphasize that the template CARDS from the simulations should not be applied to the observations.

2.2. Halo Progenitors and Dwarf Galaxies

In investigating how CARDS of stellar halos can be used to infer assembly histories, and in comparing the abundance distributions of disrupted dwarfs with the present-day dwarfs, we must first define our stellar halo and dwarf galaxy samples within the simulation. In this section, we introduce the catalog of halo star particles and dwarf galaxies that we use for this experiment. We limit our study to the catalog of star particles belonging to dwarf galaxies, stellar streams, and phase-mixed debris constructed in Panithanpaisal et al. (2021; hereafter P21).¹⁴ By using this subset of particles, we can explore the efficacy of the CARDS framework for a purely accreted stellar halo in a more realistic setting than that in L15 while retaining control of the precise composition of the accreted stellar halo, and make comparisons between the halo progenitor and present-day dwarf galaxy populations. We explore these questions in the limit of no contamination from disk or in situ halo stars, and leave the strategies for modeling potential contamination to future work.

While the primary goal of P21 was to identify present-day streams, as a result of their stream candidate identification method, they also found present-day dwarf galaxies and phase-mixed debris. P21 implemented the following criteria to identify the stream candidates:

1. Particles in stream candidates must be within the virial radius of the main host at present day, although bound in a different subhalo 2.7–6.5 Gyr ago.
2. Stream candidates must have $120 < N < 10^5$ star particles; this results in a stellar mass limit for the progenitors of $10^{5.5} M_{\odot} < M_{\star} < 10^9 M_{\odot}$ for the isolated MW simulations.¹⁵

With stream candidates in hand, each candidate is classified as either phase-mixed debris, a coherent stream, or a dwarf galaxy based on its local velocity dispersion and maximum pairwise distance between the star particles at present day. To be classified as a coherent stream or as phase-mixed debris, the star particles must have a maximum pairwise distance greater than 120 kpc. This pairwise distance threshold, on the order of

¹⁴ Stream catalog publicly available: <https://flathub.flatironinstitute.org/sapfire>.

¹⁵ This corresponds to an upper mass limit of $\sim M_{\star} < 10^{8.5} M_{\odot}$ for the paired simulations, which have typical particle masses of $\sim 3000 M_{\odot}$. While the stellar mass limits of the catalogs are slightly different, this does not affect the outcome of our study, as we limit our templates to the range of $10^6\text{--}10^9 M_{\odot}$ (see Section 4.2), and there are no surviving satellites above the P21 mass limit in the paired simulations (Samuel et al. 2022).

Table 1

Each of the Simulations Discussed in This Work, with the Fraction of Accreted Stars (Excluding Those Bound in Satellites) Included in the P21 Catalog

Simulation	$M_{*,\text{acc}} (10^9 M_{\odot})$	% Stellar Halo Mass in P21
m12b	11.1	2.4
m12c	1.8	26.4
m12f	6.6	9.3
m12i	2.9	18.0
m12m	3.8	9.3
m12r	7.8	0.4
m12w	2.9	14.6

Note. To compute the total accreted contribution, we calculate the total mass in star particles that are located within the virial radius at present day, but that formed at a distance larger than 30 kpc from the center of main halo. We exclude star particles in bound subhalos at present day from this estimate. Halos with the largest total accreted stellar masses have the lowest fraction of their halos included in P21, as these halos experience at least one major accretion event above the mass limit imposed by P21.

the size of the host galaxy halo, effectively distinguishes between the disrupted systems (phase-mixed debris and coherent streams) and present-day bound dwarf galaxies, which remain compact in the position space. P21 used a linear kernel support vector machine (SVM) in order to derive an average, local velocity dispersion threshold as a function of the stellar mass to distinguish the phase-mixed debris from the coherent streams (see their Equation (2)). Using disrupted systems from four of the simulations, P21 classified the coherent streams and phase-mixed debris by eye to use as the training set for the SVM. In practice, the coherent streams have lower local velocity dispersions on average than the phase-mixed debris (at fixed M_*). P21 determines the local velocity dispersion for a star particle using its nearest neighbors in phase space (using 20 nearest neighbors if the stream candidate has more than 300 star particles; otherwise, 7 nearest neighbors are used); they then average over all star particles in the stream candidate to get the average local velocity dispersion. This local velocity dispersion threshold is $\sim 20 \text{ km s}^{-1}$ for a stream candidate with $M_* \sim 10^7 M_{\odot}$ (see their Figure 1).

In summary, the present-day dwarf galaxies are compact in velocity and position space; the coherent streams have low local velocity dispersions but are extended in position space; and the phase-mixed debris have larger local velocity dispersions and can be extended in position space. Each accretion event is assigned a single classification, even if portions of the debris might be classified differently (e.g., a dwarf remnant with long tidal tails will likely be classified as a coherent stream). We define the stellar halo samples analyzed in this work to include the phase-mixed debris and streams from P21, and exclude the present-day dwarf galaxies from our stellar halo samples.

We note that not every accreted star particle in the simulation is included in the P21 catalog. In particular, the accreted dwarfs that exceed the particle limit from P21 (with $M_* > 10^9 M_{\odot}$, bordering on major merger classification) are not included, as well as the low-mass systems containing fewer than 120 star particles ($\sim 10^5 M_{\odot}$). In addition, a number of early mergers (that are unbound prior to 6.5 Gyr ago) are also excluded (these earlier mergers are discussed in more detail in D. Horta et al., 2022, in preparation). Table 1 summarizes the total fraction of the accreted stellar halo that is included in the P21 catalog for the seven halos analyzed in this work. To estimate the total

mass of the accreted stellar component, we compute the mass of all star particles that are located within the virial radius of the main host at present day, but that formed at least 30 kpc from the main host (excluding stars bound in satellites). The total accreted stellar masses range from $1\text{--}11 \times 10^9 M_{\odot}$. Table 1 shows that the fraction of accreted star particles included in P21 varies widely from simulation to simulation. Halos that have a relatively large accreted mass usually have experienced at least one massive accretion event above the P21 limit; the two halos with the largest accreted stellar masses (m12b and m12r) have the lowest fraction of their accreted stars included in the P21 catalog. For all simulations, we can see that the material that is accreted early does contribute significantly to the stellar mass of the overall halo; these accretion events are explored in more detail in D. Horta et al. (2022, in preparation).

While there are many accreted halo star particles that are not included in our analysis, we note that we do not require a fully complete halo sample to test the validity of this method. In order to address whether or not the CARDS method works in a cosmological simulation, we require a sample of the accretion events whose progenitors have been well characterized. We therefore limit the scope of this study to using the accretion events and dwarf galaxies from the P21 catalog.

3. Simulation CARDS

We seek to use the chemical abundance distribution of halo stars to infer properties of their dwarf galaxy progenitors. In this section, we demonstrate how halos with different formation histories come to have different CARDS (Section 3.1), in order to motivate why we can use the halo star CARDS to make inferences about the dwarf galaxy progenitors of stellar halos. In Section 3.2, we demonstrate the relationship between a halo progenitor CARD and its dynamical state (e.g., whether or not they are phase-mixed, coherent streams, or present-day dwarfs).

3.1. The Link Between CARDS and Assembly History

In this subsection, we demonstrate the relationship between halo CARDS and host halo formation history, by discussing in detail three hosts from the Latte suite of simulations. The top panels of Figure 1 show the 2D histograms of formation time versus formation distance for every star particle in simulations m12m, m12f, and m12r (with density plotted on a logarithmic scale). We note that formation time is defined such that $t_{\text{form}} = 0$ at the beginning of the simulation, and the formation distance is physical distance (as opposed to comoving distance). Looking in this plane, we can see where the star particles are forming in these different simulated galaxies over time. The high-density regions at the bottom of each figure, showing the star particles with formation distances $0 < d_{\text{form}} < 20 \text{ kpc}$, are particles that form in the disk of the host galaxy. The streaks of star particles that start at larger d_{form} and move toward the disk over time show dwarf galaxies falling into the main host halo. Figure 1 demonstrates that these three simulations have very different formation histories: for example, m12m's evolutionary history (shown in the top left panel of Figure 1) is characterized by many early mergers, whereas m12r (top right panel of Figure 1) experiences two massive, recent accretion events ($M_* > 10^9 M_{\odot}$, within the last $\sim 2 \text{ Gyr}$). We note that m12r was one of the Latte hosts chosen

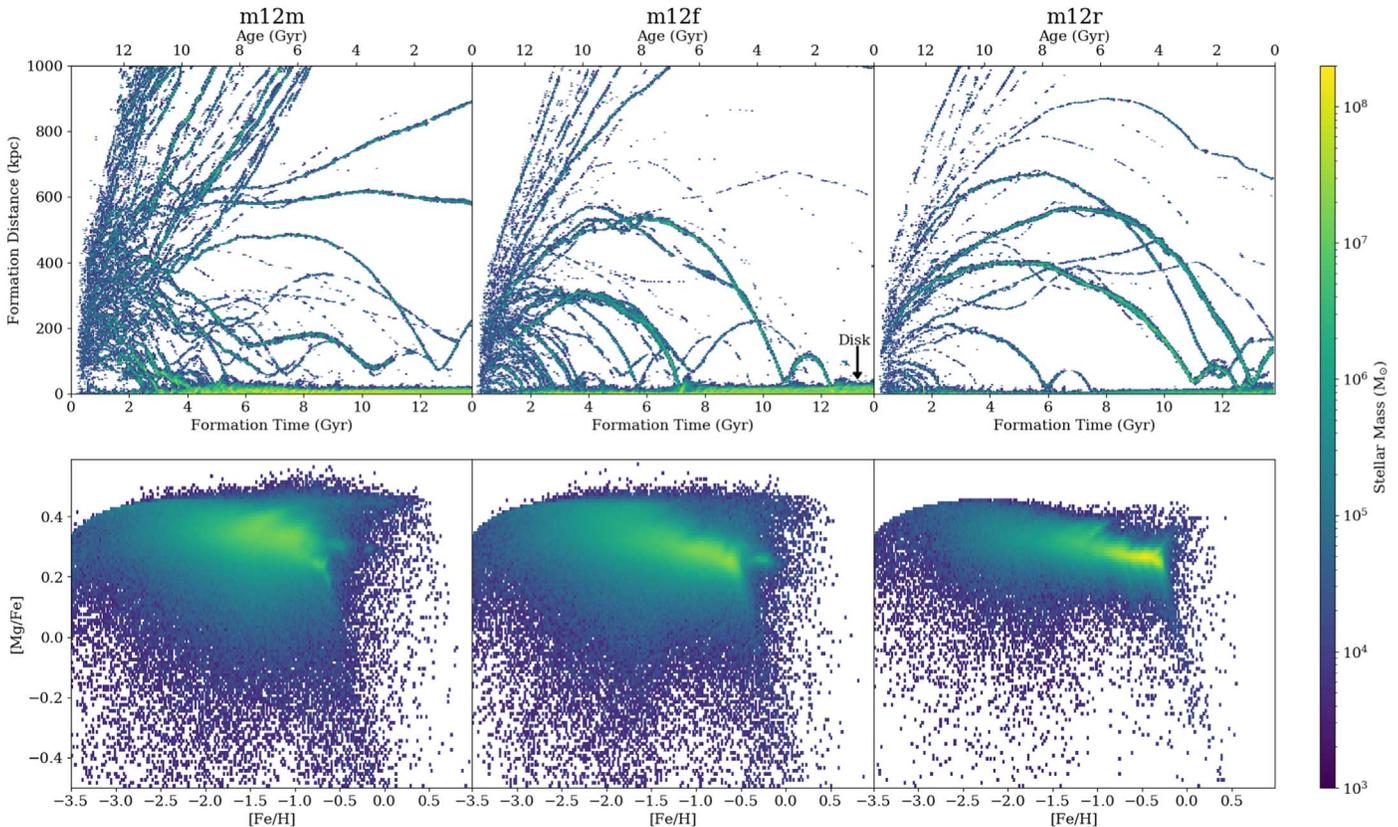


Figure 1. Top panels: formation histories for three of the Latte simulations (m12m, m12f, and m12r), as shown by 2D histograms of formation distance vs. formation time (defined such that $t_{\text{form}} = 0$ at the beginning of the simulation), for all star particles. Density is normalized on a logarithmic scale. In this plane, the thick regions of high density at small values of formation distance show star particles forming in the disks of the main host galaxies. The lower density streaks that start at larger formation distance and move toward lower formation distances indicate star particles forming in dwarf galaxies and falling into the host. From these figures, we see that m12m’s assembly history (left) is characterized by early accretion; though it also has several massive dwarfs in the halo at present day. In contrast, m12r (right) experiences several recent, massive accretion events, with less activity at early times. The intermediate case is shown by m12f (center), which experiences a relatively recent massive accretion (~ 2 Gyr ago) event as well as one ~ 6 Gyr ago. Lower panels: the corresponding CARD for the stellar halo of each for galaxy, taking star particles that have formation distances greater than 30 kpc from the main host, but that lie within 100 kpc of the main host at present day. These visibly different distributions reflect the formation histories of each of their hosts. For example, m12m (left), characterized by early accretion, has higher density at lower [Fe/H] and higher [Mg/Fe]. In contrast, the recent accretion experienced by m12r (left) results in an abundance distribution that has high density at high [Fe/H] and lower [Mg/Fe].

with the intention of simulating a MW host with an LMC-analog; see Samuel et al. (2020).

The lower panels of Figure 1 show the chemical abundance ratio distributions (CARDs) for these same galaxies, in [Mg/Fe] versus [Fe/H]; the abundances are shown for all stars that formed outside of 30 kpc but are within 100 kpc of the host galaxy at present day. These CARDs are visibly quite different, with apparent links to the formation histories of their host halos.

The galaxy m12m (left panels of Figure 1) experiences a number of massive accretion events early on, and while there are many dwarfs in the vicinity of m12m at present day, most have large pericenters (as seen by the fact that the wiggles do not intersect with the disk). Most of the density in the CARD for m12m is at lower [Fe/H] ($\sim -2 < [\text{Fe}/\text{H}] < -1$) and higher [Mg/Fe] ([Mg/Fe] > 0.25). In contrast, m12r (right panels of Figure 1) has an assembly history characterized by recent, massive accretion events, indicated in the top panel of Figure 1 by the thick streaks of high density that approach the disk at young ages/recent times. The resulting CARD shows high density at higher [Fe/H] and low [Mg/Fe] relative to m12m, consistent with the picture that the halo is dominated by dwarfs that were accreted recently and had a prolonged star formation

and chemical enrichment history. Finally, m12f has experienced several massive accretion events over the course of its formation, with a major event recently (~ 2 Gyr ago) and another major event at intermediate times (~ 6 Gyr ago); the resulting CARD shows density at both low [Fe/H] and high [Mg/Fe] as well as higher [Fe/H] and lower [Mg/Fe].

These different distributions arise due to the different timescales of chemical enrichment. At early times, the galaxies are enriched by Type II SNe (beginning around ~ 4 Myr in the simulation), and producing iron and the α -elements at a relatively constant fraction. As a result, early in a galaxy’s lifetime, the mean $[\alpha/\text{Fe}]$ is roughly constant. Once the Type Ia SNe turn on (at first around 40 Myr, with the rate dependent on the star particle age), the $[\alpha/\text{Fe}]$ ratio of stars born in the enriched interstellar matter (ISM) starts to decrease, because Type Ia SNe produce more iron than they do α -elements. Therefore, the CARD of the stellar halo in a galaxy contains key information about the star formation histories of the halo progenitors.

However, what remains an open question is whether or not we can use CARDs from present-day dwarf galaxies to infer the properties of halo progenitors. In the next section, we

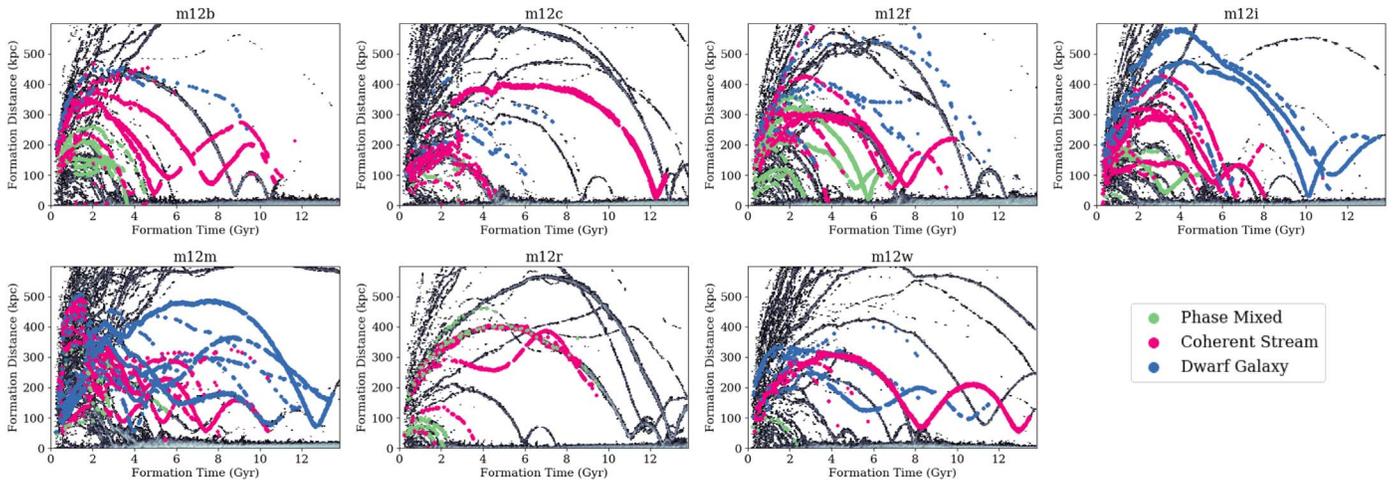


Figure 2. Same as the top panels of Figure 1, but for the set of seven Latte simulations studied in this work. Star particles from the P21 stream candidates catalog (used in this work) are indicated by the colored points. Star particles belonging to surviving dwarf galaxies (within virial radius of the host galaxy at present day in the simulations) are shown in blue; particles classified as members of coherent stellar streams are in pink; and phase-mixed debris are shown in green. Not every star particle that is accreted in the simulation is included in the P21 catalog (and therefore in this analysis); see Section 2.2 for details.

introduce the chemical properties for the catalog of dwarf galaxies and halo progenitors used in this work.

3.2. Linking Dynamical State and Chemical Abundances

Chemical abundance distributions can be used to *tell time* within a dwarf galaxy or halo progenitor: because the different chemical enrichment processes (e.g., Type Ia SNe, Type II SNe, and stellar winds) occur on different timescales; the distribution of a dwarf or progenitor’s chemical abundances is informative for determining how much and when stellar mass formed in the system’s evolution. An additional clue constraining the amount of time a halo progenitor had to form stars is its *dynamical state* at present day, which we take here to mean whether or not a halo progenitor is phase-mixed, a coherent stream, or remains a dwarf galaxy at the end of the simulation. Because these different classifications imply different times for truncating star formation (with phase-mixed debris, on average, implying an earlier accretion time than a surviving dwarf), their chemical properties are in turn expected to be different. In this section, we compare the chemical properties of phase-mixed debris, coherent streams, and dwarf galaxies from the P21 catalog. This has implications for using the CARDS technique in an observational context, where we will likely need to rely on a sample of present-day dwarf galaxies for constructing an empirical template set to model the MW’s halo progenitors.

Figure 2 shows the 2D histogram of age versus formation distance for the seven of the isolated MW simulations used in this work in gray, with the star particles included in the P21 catalog overlotted in color. Points are color-coded by classification: dwarf galaxies are in blue, coherent streams are in pink, and phase-mixed debris are in green.

Figure 2 highlights some systematic differences between the dwarf galaxy, stream, and phase-mixed populations. For example, most of the star particles associated with the phase-mixed debris (colored in green) are primarily composed of older star particles, consistent with the picture that they must have been accreted relatively early on in order to have had sufficient time to become phase-mixed at present day. The differences between the populations are highlighted in more detail in Figure 3. The left-hand panel of Figure 3 shows t_{100}

versus the total stellar masses for all objects from the P21 catalog, where we define t_{100} to be the time for a dwarf or halo progenitor to form 100% of its stars (i.e., the quenching time). As previously, the points are color-coded by classification (dwarf galaxies in blue, coherent streams in pink, and phase-mixed objects in green) while the marker shape indicates whether or not the object is in an isolated MW simulation (circles) or in a paired galaxy simulation (diamonds; the ELVIS on FIRE suite). While there is substantial overlap in this plane among the three classifications, at fixed stellar mass, the phase-mixed objects have smaller t_{100} values (on average) than the surviving dwarf galaxies, as a result of the fact that they are disrupted earlier and have less time to form more stellar mass. This has implications for the abundance distributions (which we can directly observe), as shown by the right panel of Figure 3: at fixed mean metallicity ($\langle[\text{Fe}/\text{H}]\rangle$), the phase-mixed debris have higher mean magnesium abundances compared to the surviving dwarf galaxies. We note that, while there is a correlation between mass and metallicity in the simulations (e.g., Wetzell et al. 2016; Escala et al. 2018; P21), there is not a direct correlation between mean $[\text{Mg}/\text{Fe}]$ and t_{100} . As the left panel highlights, t_{100} is also correlated with stellar mass, and therefore mean $[\text{Fe}/\text{H}]$. Therefore, the quenching time t_{100} is only correlated with mean $[\text{Mg}/\text{Fe}]$ at fixed the stellar mass.

We note that the right panel of Figure 3 shows a trend of increasing $\langle[\text{Mg}/\text{Fe}]\rangle$ with increasing $\langle[\text{Fe}/\text{H}]\rangle$ at lower metallicities, contrary to the expected trend. This trend is due to the particle initialization in the FIRE-2 physics model: all star particles are initialized with a metallicity of $[\text{Fe}/\text{H}] = -4$ and $[\text{Mg}/\text{Fe}] = 0$ (this is not the case for the FIRE-3 physics model, which initializes star particles from primordial gas; Hopkins et al. 2022). This results in a maximum magnesium abundance for a given iron abundance, which causes the ridge feature seen in the simulation CARDS at high $[\text{Mg}/\text{Fe}]$ and low $[\text{Fe}/\text{H}]$ (e.g., lower panels of Figure 1). At low metallicity, the majority of the particles will have magnesium abundances at or near the maximum allowed $[\text{Mg}/\text{Fe}]$ value available to them, as they are enriching primarily through Type II SNe. However, as the mean metallicity increases, we see the expected trend of decreasing mean $[\text{Mg}/\text{Fe}]$ with increasing mean $[\text{Fe}/\text{H}]$, as star particles enrich due to Type Ia SNe.

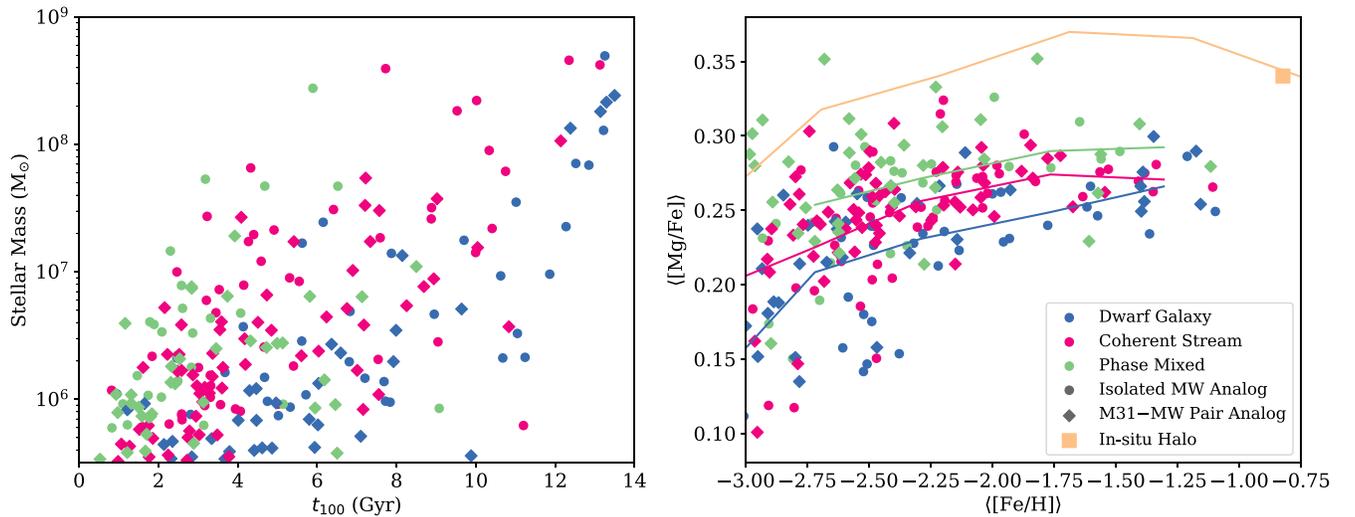


Figure 3. Left: stellar mass vs. quenching time (t_{100}) for the full P21 catalog. Points are color-coded based on their classification as dwarf galaxies, coherent streams, or phase-mixed debris. Diamond-shaped points are for objects in the paired M31–MW analog simulations; circular points are objects from the isolated MW analog simulations. While there is substantial overlap in this plane across the three classifications, at a given stellar mass, dwarf galaxies tend to have the latest quenching times, and phase-mixed debris tends to have the earliest quenching times. Right: mean [Mg/Fe] vs. [Fe/H] for the full P21 catalog. At fixed metallicity, phase-mixed debris is, on average, more enhanced in Mg than dwarf galaxies. The orange square indicates the mean [Fe/H] and [Mg/Fe] for the in situ halo, which is, on average, more metal-rich and enhanced in magnesium compared to the accreted populations; see Appendix D for more details. Solid lines indicate the means of each population as a function of metallicity.

Figure 3 also shows the mean magnesium and iron abundances for the in situ halo, both the overall mean (orange square) as well as the mean magnesium abundance in bins of metallicity. For the purposes of this illustration, we define the in situ halo as the particles with formation distances of less than 40 kpc and present-day distances that are greater than 40 kpc. We exclude the particles that are within 10 kpc of the Galactic midplane as well as particles that are assigned to a known accretion event in either P21 or D. Horta et al. (2022, in preparation; this catalog contains the star particles from the most massive accretion events that are excluded from the P21 catalog). We find that in situ halo stars are, on average, enhanced in magnesium and more metal-rich than the accreted populations studied in this work (see Appendix D for more details), and that they are enhanced in magnesium relative to the accreted populations at all metallicities.

The fact that the phase-mixed debris is, on average, α -enhanced (at fixed metallicity) with respect to present-day dwarf galaxies is consistent with a long-known observational result: the accreted halo stars are observed, on average, to be α -enhanced relative to stars of the same metallicity that formed in LG dwarf galaxies (see, e.g., Venn et al. 2004; Naidu et al. 2022). Figure 3 demonstrates that this arises because the accreted halo stars form in dwarfs that disrupt before $z = 0$, and do not have as much time to enrich their ISMs in iron through Type Ia SNe compared to present-day dwarfs (though we note that this discussion here excludes in situ halo stars; see Appendix D for additional discussion). Comparing the present-day dwarf populations with halo star populations therefore results in a “temporal bias” (e.g., Robertson et al. 2005; Font et al. 2006; Johnston et al. 2008): at a fixed stellar mass, the dwarf galaxies have, on average, spent a longer period of time forming their stars. Therefore, if we are to use dwarf galaxies to recover the properties of accreted debris, we must correct for this effect in constructing our templates. In the subsequent section, we describe our method for using the CARDS for present-day dwarf galaxies as templates for halo progenitors.

4. Methodology

With the CARDS method, first proposed by L15, we model the abundance distributions in the stellar halo as a linear combination of template CARDS. Two of the primary goals of this paper are to test the CARDS modeling procedure in the more realistic setting of a cosmological simulation, and to explore how to construct template CARDS in a potentially observationally viable manner. In this section, we describe the CARDS modeling method, and present our method for constructing empirical templates.

We begin by describing the CARDS modeling method in Section 4.1; for the reader’s convenience, we have summarized a number of relevant definitions and notations in Table 2. In Section 4.2, we describe our procedure for creating empirical templates, and introduce the template sets used for analysis in this work.

4.1. Modeling Framework

In this subsection, we demonstrate our procedure for modeling the CARDS of the Latte halos using a grid of template abundance distributions. For the purposes of keeping our notation general, we define the vector of chemical abundance ratios as \mathbf{x}_d : in this work, we limit ourselves to the cases where $\mathbf{x}_d = ([\text{Fe}/\text{H}], [\text{Mg}/\text{Fe}])$ or $\mathbf{x}_d = ([\text{Fe}/\text{H}], [\text{Mg}/\text{C}])$ (though, in principle, this method could make use of many more elements and dimensions). While the simulations track 11 elements, we use the elements most dominated by each of the three channels of chemical enrichment (Fe for Type Ia SNe, Mg for Type II SNe, and C for stellar winds). We define our grid of templates over a range of stellar mass and quenching time. The bins in stellar mass are denoted with the index i , and the bins in quenching time are denoted with the index j (we define the specifics of the grid used in this work in Section 4.2.1).

As in L15, we express the model for the (normalized) full abundance distribution of the accreted halo as a linear

Table 2
Useful Definitions

Abbreviations and Definitions Referenced in the Text	
CARD	Chemical abundance ratio distribution; the density of stars in the chemical plane, for two or more elements (e.g., [Fe/H] versus [Mg/Fe]).
Timescales	
t_{form}	Formation time of a star particle; time a star particle formed in the simulation, with $t_{\text{form}} = 0$ defined as the beginning of the simulation
t_{100}	Quenching time of a halo progenitor or dwarf galaxy. This is equivalent to the formation time of the youngest star particle associated with a given halo progenitor or dwarf galaxy.
Templates	
Grid Dimensions	$10^6 M_{\odot}$ to $10^9 M_{\odot}$ (in intervals of $10^{0.5} M_{\odot}$) in stellar mass; 0 to 14 Gyr in $t_{100,\text{temp}}$ (in intervals of 2 Gyr).
i	Index for the stellar mass grid dimension.
j	Index for the quenching time (t_{100}).
Template Sources	Objects with empirical CARDS used for constructing templates (e.g., present-day dwarfs, or streams phase-mixed debris).
CARD _{temp,ij}	Template abundance distribution for a halo progenitor with $(M_{\text{sat},i}, t_{100,j})$.
$t_{100,\text{source}}$	The quenching time for a halo progenitor or dwarf galaxy used to construct a set of templates.
$t_{100,\text{temp},j}$	The mean quenching time corresponding to a specific CARD _{temp,ij} .
$t_{100,\text{temp},\text{max},j}$ ($t_{100,\text{temp},\text{min},j}$)	The maximum (minimum) quenching time corresponding to a specific CARD _{temp,ij} .
Template Sets	
S/PM Templates	Templates constructed using streams and phase-mixed debris as sources. Default template set is for [Mg/Fe] vs. [Fe/H]. Each halo has its own set of S/PM templates, constructed, excluding its own halo progenitors as sources.
DG Templates	Templates constructed using present-day dwarf galaxies as sources.

combination of individual templates. We define our accreted halo to be all star particles that are identified as belonging to phase-mixed debris or a coherent stream in P21, and exclude other particles from analysis. The CARDS for each simulation including only these particles are shown in Section 5 and Appendix B. The abundance distribution of our accreted stellar halo, which we define as $\text{CARD}_{\text{halo}}$, can thus be expressed as follows:

$$\text{CARD}_{\text{halo,model}}(\mathbf{x}_d) = \sum_i \sum_j A_{ij} \text{CARD}_{\text{temp},ij}(\mathbf{x}_d | M_{\text{sat},i}, t_{100,j}), \quad (1)$$

where A_{ij} is the coefficient for the template for an accretion event with stellar mass $M_{\text{sat},i}$ and quenching time $t_{100,j}$ (that has normalized density $\text{CARD}_{\text{temp},ij}(\mathbf{x}_d | M_{\text{sat},i}, t_{100,j})$). In this context, the template coefficients A_{ij} represent the fraction of mass contributed to the halo from accretion events in a bin $M_{\text{sat},i}, t_{100,j}$. We also include an additional template that has uniform density across the full abundance plane; this serves as our *noise* template.

We can solve for our linear coefficients by specifying a loss function and optimizing it under a simple set of constraints. If we define a vector of templates as $\overrightarrow{\text{CARD}}_{\text{temp}}(\mathbf{x}_d)$ and a coefficient vector \vec{A} , we can rewrite the above expression:

$$\text{CARD}_{\text{halo,model}}(\mathbf{x}_d) = \overrightarrow{\text{CARD}}_{\text{temp}}(\mathbf{x}_d) \cdot \vec{A}. \quad (2)$$

We define our loss function to be the absolute value of the difference between our model CARD density and the true density from the simulations:

$$\text{Loss Function} = |\overrightarrow{\text{CARD}}_{\text{temp}} \cdot \vec{A} - \text{CARD}_{\text{halo,true}}|. \quad (3)$$

We constrain the problem solely by requiring that each coefficient $A_{ij} \in [0, 1]$ and that the coefficients sum to 1 ($\sum A = \sum_i \sum_j A_{ij} = 1$). We solve for the template coefficients

using the `scipy.optimize.minimize` routine. We emphasize that this model as written contains no knowledge of physics, and could be expanded in a Bayesian framework to incorporate additional information (e.g., the mass of the halo). However, as we show in the next section, even this simple model can be used to recover information about the simulation accretion histories, in some cases with surprising accuracy.

While the modeling method is simple, it relies on a set of template abundance distributions: how best to define and create these templates is less simple. In the next subsection, we introduce our method for creating templates.

4.2. Templates

The CARDS modeling method, as described in the previous subsection, relies on a set of template abundance distributions for halo progenitors with different stellar masses and quenching times. In this subsection, we present our method for constructing CARD templates. In our efforts to test the CARDS method in a realistic setting, as well as make the CARDS method ultimately practical for observations, we require a framework for creating templates that is observationally viable. While there are several valid approaches for constructing the template CARDS, including from purely theoretical models, here we explore an empirical approach, relying on the abundance distributions of present-day dwarf galaxies within the simulations. We assess this empirical approach within simulations, motivated by the ultimate goal of using the observed dwarf galaxy CARDS to recover the properties of the progenitors of the MW and M31 stellar halos.

Because we are constructing empirical templates, we rely on a catalog of *sources*: the sources for our templates refer to the abundance distributions used to create the templates. We consider two different populations of template sources: (1) halo progenitors (e.g., streams and phase-mixed debris) that contribute to other halos in the simulation suite, and (2)

present-day dwarf galaxies (using all available dwarf galaxies from the P21 catalog). By comparing these two catalogs of template sources, we can test the expectation that the disrupted dwarfs may be *better* template sources than the present-day dwarfs, and determine if using the observed dwarf galaxy CARDS will be viable template sources for the MW and M31 stellar halo progenitors.

We begin in Section 4.2.1 by discussing how we define the axes of our grid of templates; i.e., the properties of halo progenitors that we are attempting to constrain with this method. In Section 4.2.2, we then describe our method for creating templates, and show how we can use phase-mixed debris, streams, or dwarf galaxies as template sources. In Section 4.2.3, we present our CARD templates sets in [Mg/Fe] versus [Fe/H], constructed from streams and phase-mixed debris as well as an additional template set constructed only from present-day dwarf galaxies. Additional templates used in this work (for the [Mg/C] versus [Fe/H] distributions) can be found in Appendix A.

4.2.1. Defining the Grid of Templates

Given that we want to use these templates to extract information about the evolution of the halo over time, we first address the question of which timescale we should be using in our template grid. In L15, the two axes of their template grids were stellar mass and accretion time. However, as discussed in P21, assigning an individual accretion time to these events in the context of a cosmological simulation is not straightforward. For example, a dwarf may cross the virial radius of the host galaxy multiple times before becoming tidally disrupted. Furthermore, as can be seen in Figure 2, many of the halo progenitors and dwarf galaxies continue to form stars after falling into the host and completing one or more pericentric passages (for a more detailed discussion of satellite quenching, see Samuel et al. 2022). As a result, in these simulations, the accretion time is not equivalent to the quenching time (as it is in the Bullock & Johnston 2005 simulations). Because the quenching time is the more relevant timescale for star formation and chemical enrichment, in this work, we use the quenching time t_{100} as the time dimension for our templates, which we define to be the time a dwarf takes to form 100% of its stars (defined such that $t=0$ at the beginning of the simulation). This timescale is also equivalent to the formation time of the youngest star particle associated with a given progenitor or dwarf galaxy, and is approximately equivalent to the age range of the associated star particles (all sources in the P21 catalog begin forming stars very near the beginning of the simulation).

In Figure 4, we introduce our format for our grid of halo progenitor CARD templates. Each row in Figure 4 represents a fixed stellar mass range, with the stellar mass increasing from the top row ($M_* \sim 10^6 M_\odot$) to the bottom row ($M_* \sim 10^9 M_\odot$). Each column represents a range for t_{100} : the far left-hand panels show the templates for halo progenitors with $0 < t_{100} < 2$ Gyr (i.e., templates for halo progenitors that quench within the first 2 Gyr of the simulation). The far right panels show the templates for halo progenitors with $12 < t_{100} < 14$ Gyr (this can theoretically include halo progenitors that are still forming stars at the present day). The particular templates shown in Figure 4 are used to demonstrate our procedure for creating templates, which we discuss in detail in the subsequent subsection.

Our grid of templates extends from $10^6 M_\odot$ to $10^9 M_\odot$ (in intervals of $10^{0.5} M_\odot$) in stellar mass, and 0 to 14 Gyr in quenching time, t_{100} (in intervals of 2 Gyr). We denote i as the index for stellar mass and j as the index over quenching time. Each grid cell (i, j) shows the template CARD $\text{CARD}_{\text{temp},ij}$ for a halo progenitor with the stellar mass $M_{*,i}$ and quenching time $t_{100,j}$. We note that each grid cell (i, j) is actually representing a range of stellar mass and t_{100} , and as such has a corresponding $M_{*,\text{min},i}$, $M_{*,\text{max},i}$ as well as a $t_{100,\text{min},j}$, $t_{100,\text{max},j}$. However, for convenience in notation, we generally refer to a given grid cell with indices (i, j) as containing the template CARD $\text{CARD}_{\text{temp},ij}$ for a halo progenitor with properties $(M_{*,i}, t_{100,j})$.

4.2.2. Creating Templates

The observed CARDS from present-day dwarf galaxies in the LG provide one avenue for constructing the empirical CARD templates for MW and M31 halo progenitors. One challenge, as discussed in Section 2.2 and shown in Figure 3, is that the dwarf galaxies occupy a different region in the (M_*, t_{100}) plane from halo progenitors (though there is substantial overlap). On average, at a fixed stellar mass, the dwarf galaxies will have longer quenching times than the phase-mixed debris and stellar streams, as a result of the fact that the surviving dwarf galaxies are not quenched, as a result of being accreted by the host. If we are to construct empirical template sets from the observations, we will want to use the abundance distributions from local low-mass dwarfs in creating our templates. Given that dwarf galaxies tend to have larger t_{100} than streams and phase-mixed debris (as seen in Figure 3), how can we use the dwarf galaxy CARD templates to recover the properties (including t_{100}) of the halo progenitors?

In order to compare the abundance distributions of dwarf galaxies and halo progenitors, we have to account for the fact that, at fixed stellar mass, the t_{100} distributions for these two populations are not the same. Therefore, we define two versions of t_{100} :

1. $t_{100,\text{source}}$ is the time a template source (either a halo progenitor or present-day dwarf galaxy) from the P21 catalog takes to form 100% of its stars. This is the quantity plotted on the y-axis in Figure 3.
2. $t_{100,\text{temp}}$ is the time a halo progenitor represented by specific CARD $\text{CARD}_{\text{temp},ij}$ template takes to form 100% of its stars. This is the quantity shown in Figure 4 (as well as subsequent figures containing CARD templates).

Figure 4 demonstrates the structure of our template grid, with the stellar mass increasing from the top row to the bottom row and $t_{100,\text{temp}}$ increasing from left to right. To construct our templates, we begin by specifying a column j , and work from left to right. Given a specified column j , we then work with one source at a time. For each object in the source catalog:

1. We select all star particles that formed before $t_{100,\text{temp},\text{max},j}$. For example, to make the leftmost column of templates, which are for the accretion events with $0 < t_{100,\text{temp}} < 2$ Gyr, we use all the star particles from a given accretion event that formed in the first 2 Gyr of the simulation.
2. We compute the stellar mass of the subset of star particles with $t_{\text{form}} < t_{100,\text{temp},\text{max},j}$. Based on this stellar mass, the accretion event is assigned to a stellar mass bin i (or row in Figure 4).

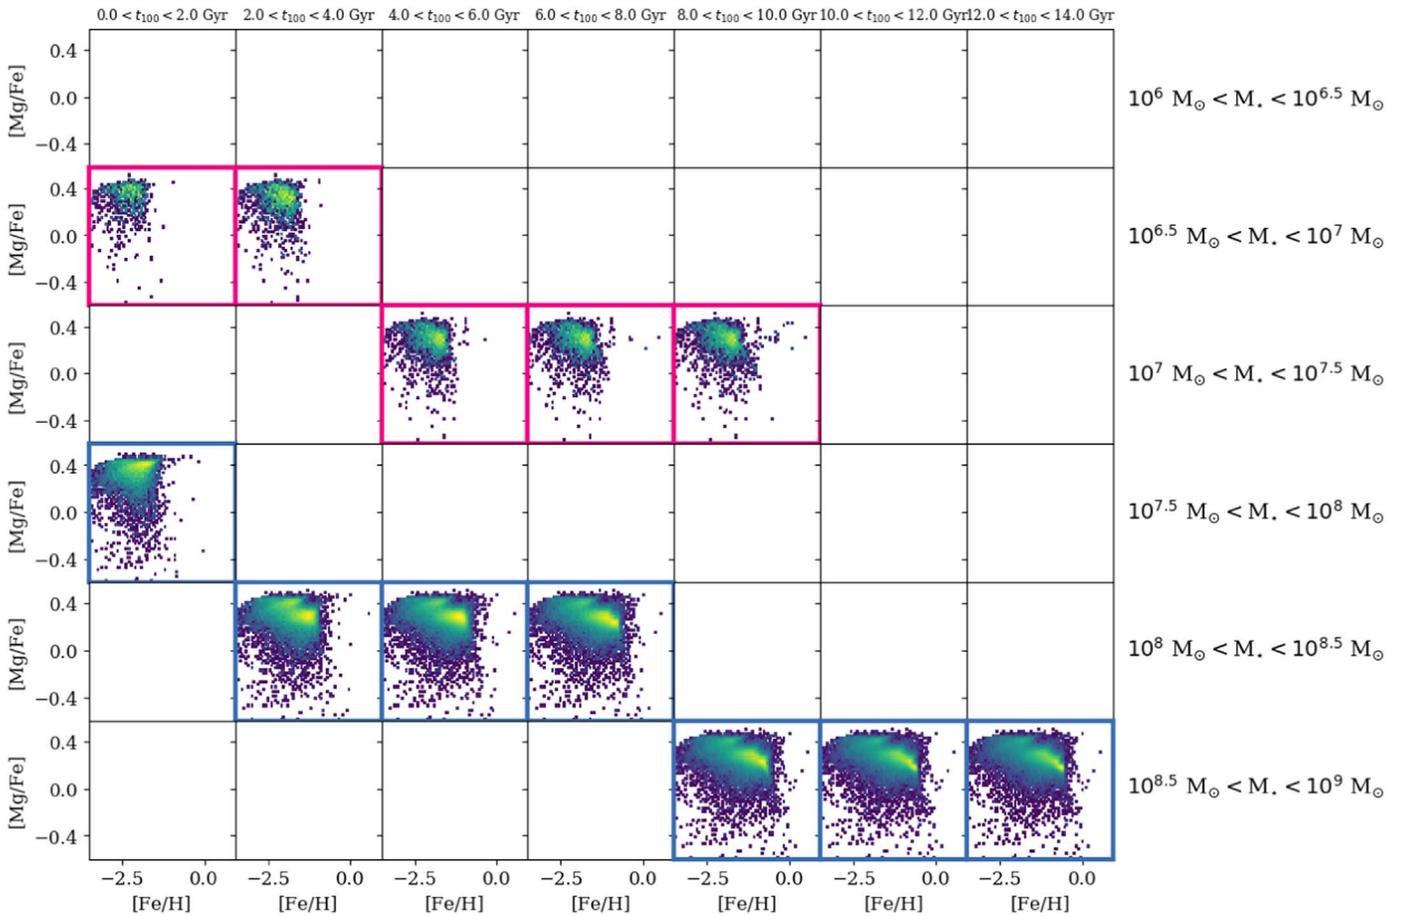


Figure 4. A grid of templates, including only two sources from the P21 catalog in the m12m simulation, to introduce our grid axes and demonstrate our procedure for creating templates. All of the templates shown in the bottom three rows (outlined in blue) are created from a single present-day dwarf galaxy, with $M_* = 10^{8.7} M_\odot$, $t_{100, \text{source}} = 13.25$ Gyr. To create the templates with $0 < t_{100, \text{temp}} < 2$ Gyr, we consider only the CARDS for the accretion events for star particles that formed within the first 2 Gyr of the simulation. Because the massive dwarf is still forming stars at present day in the simulation, it contributes to each bin of $t_{100, \text{temp}}$; it moves to progressively larger mass bins as it forms stars throughout the simulation. The templates shown in the second and third rows (outlined in pink) are created from a present-day coherent stream ($M_* = 10^{7.15} M_\odot$, $t_{100, \text{source}} = 10.0$ Gyr). Because the stream stops forming stars after 10 Gyr, it does not contribute to the templates in the far two left columns (where $t_{100, \text{temp}} > t_{100, \text{source}}$).

3. We compute the CARD for the subset of star particles with $t_{\text{form}} < t_{100, \text{temp}, \text{max}, j}$. This CARD is then averaged with the CARDS from other template sources that contribute to the same bin (i, j), to create the CARD temp, ij .

Once we have assigned each source to a bin (i, j), we compute the average normalized density of all accretion events that contribute to each bin to create the template CARD temp, ij . The distributions from individual events are normalized before averaging so that the most massive events do not dominate the average densities. We then repeat this procedure for each column (i.e., the next interval for $t_{100, \text{temp}}$). We note that, once the template source is no longer forming stars (due to tidal disruption, quenching, etc.), it no longer contributes to templates. Accretion events only contribute to the templates with $t_{100, \text{temp}} \leq t_{100, \text{source}}$.

Figure 4 demonstrates our procedure for creating templates, using two sources from the P21 catalog: a present-day coherent stream ($M_* = 10^{7.15} M_\odot$, $t_{100, \text{source}} = 10.0$ Gyr; outlined in pink) and a present-day dwarf galaxy ($M_* = 10^{8.7} M_\odot$, $t_{100, \text{source}} = 13.25$ Gyr; outlined in blue). While we only used two sources, as a result of our procedure for creating templates, each source contributes to multiple templates on the grid in

Figure 4. The present-day dwarf galaxy contributes to seven different templates: one for each time interval (as the dwarf galaxy is forming stars over the full duration of the simulation) and all with stellar mass $M_* > 10^{7.5} M_\odot$. As the dwarf’s stellar mass increases over time, it contributes to the templates in higher stellar mass bins. In contrast, the present-day stream contributes to only five templates, all with $M_* < 10^{7.5} M_\odot$: because the stream has $t_{100} = 10.0$ Gyr, it no longer contributes to the templates beyond the $8 < t_{100, \text{temp}} < 10$ Gyr time interval (i.e., it does not contribute to the two rightmost columns in Figure 4).

The template set shown in Figure 4 is purely for illustrative purposes: a template set created from only two sources is obviously of limited usefulness. In the subsequent subsection, we introduce the various template sets used for analysis in this paper.

4.2.3. Template Sets

Using the procedure outlined in the previous section, we use the P21 catalog to create four primary sets of templates: two template sets for the distributions in $[\text{Mg}/\text{Fe}]$ versus $[\text{Fe}/\text{H}]$ and two template sets for the distributions in $[\text{Mg}/\text{C}]$ versus $[\text{Fe}/\text{H}]$. We choose to include carbon in addition to magnesium

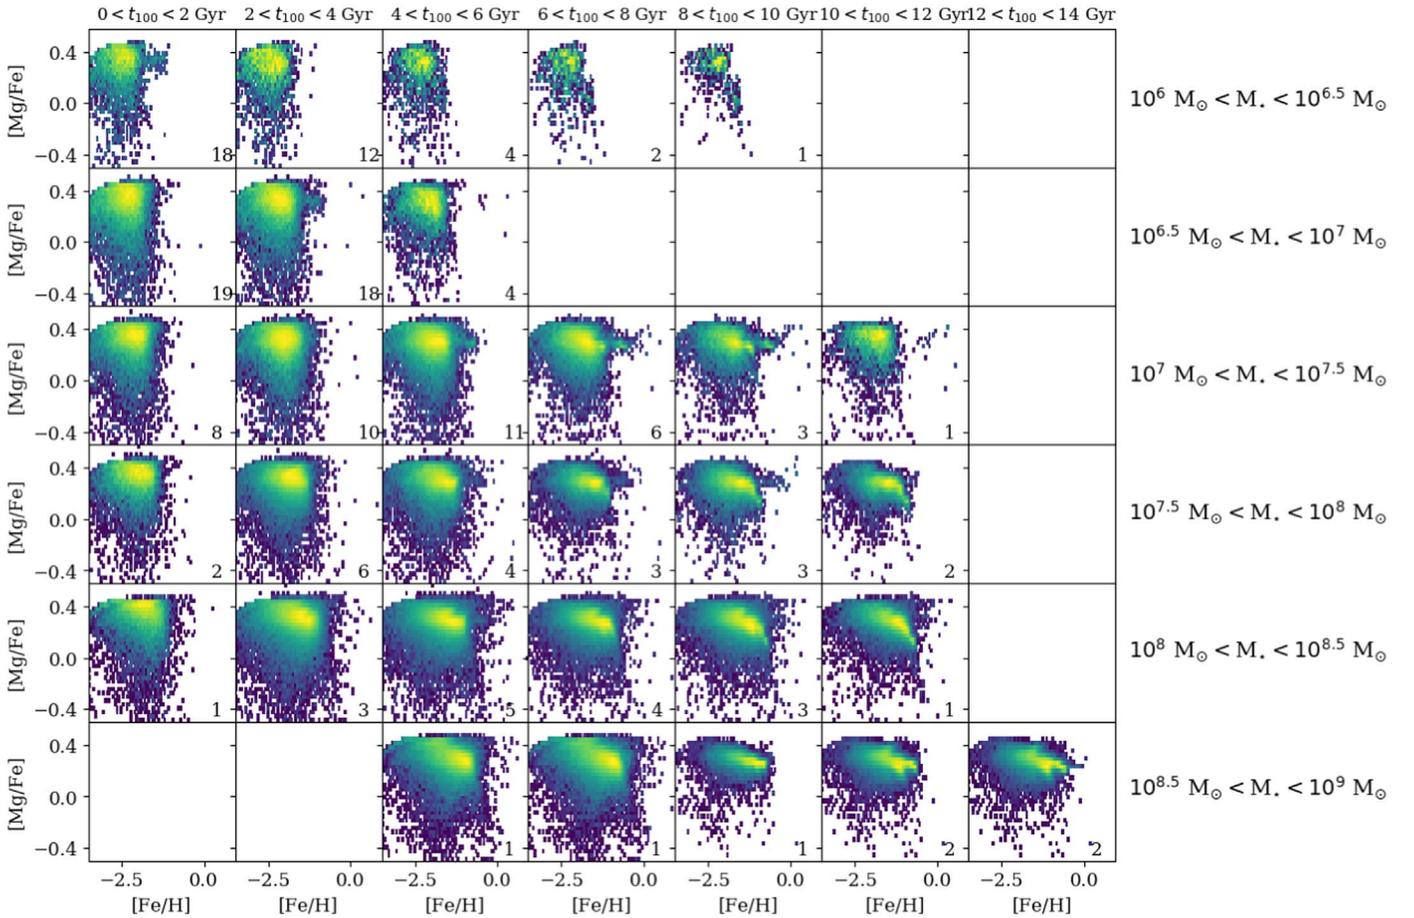


Figure 5. Templates for accretion events (constructed from the streams and phase-mixed debris, used to define the stellar halos in this work) for different stellar masses and age ranges. Stellar mass of the event increases from the top row ($10^6 M_\odot < M_* < 10^{6.5} M_\odot$) to the bottom row ($10^{8.5} M_\odot < M_* < 10^9 M_\odot$). Age range, or duration of star formation, increases from left to right. The number in the lower right of each panel indicates the number of dwarf galaxies that were used to construct a given template. We see the general expected trends in the CARDs in these templates; more massive dwarf galaxies have CARDs extending to higher metallicities. At fixed stellar mass, galaxies that assemble more quickly (i.e., have shorter age ranges) have more density in their CARDs at higher [Mg/Fe] values than their counterparts with more extended star formation histories.

because carbon predominantly traces the enrichment due to stellar winds, therefore incorporating information about an additional source of chemical enrichment (on a different timescale). For each combination of elements, we create one template set from the material that has been phase-mixed or has formed a stream in the isolated MW simulations (i.e., the accretion events whose properties we are trying to recover) as well as the dwarf galaxies at present day.

Figure 5 shows the resulting CARD templates in [Mg/Fe] versus [Fe/H] created from all of the phase-mixed debris and coherent streams from the P21 catalog in the Latte suite. These templates are hereafter referred to as the stream/phase-mixed (S/PM) templates. Figure 6 shows the [Mg/Fe] versus [Fe/H] templates created from the surviving dwarf galaxies (from both the Latte and ELVIS on FIRE suites), hereafter referred to as the DG templates. In both Figures 5 and 6, as in Figure 4, the templates increase in stellar mass from the top row ($M_* \sim 10^6 M_\odot$) to the bottom row ($M_* \sim 10^9 M_\odot$). In constructing the templates, we average over as many accretion events with the desired properties that we have available; the number of events contributing to each template is indicated by the number in the lower right corner of each panel in Figures 5 and 6. In this work, we also make use of template sets constructed from [Mg/C] versus [Fe/H]; these template sets

are shown in Appendix A. In Appendix A, we also discuss several additional template sets that we tested, including “master” templates (utilizing halo progenitors and dwarf galaxies as sources) as well as “early” and “late” templates (dividing template sources by t_{100}). We refer the reader to Appendix A for further details.

As noted in Section 4.2, we create a distinct S/PM template set for each individual halo, which excludes all of that halo’s progenitors (where each simulation’s “stellar halo” is defined as the ensemble of star particles belonging to streams and phase-mixed debris identified in P21). Therefore, we note that the template set shown in Figure 5, which includes all of the P21 Latte halo progenitors, is not used for analysis. We choose to exclude a given halo’s progenitors from its template set in order to better directly compare the performance of S/PM templates with the templates created from present-day dwarf galaxies. If a given halo’s progenitors are included in the S/PM templates, the S/PM templates will always perform better than the dwarf galaxy templates by design.

Both template sets show the same general trends as expected from the chemical enrichment models. As the stellar mass increases, the CARDs extend to higher and higher iron metallicities. At a fixed stellar mass, the templates with shorter age ranges (i.e., smaller t_{100} , on the left-hand side of the grid of

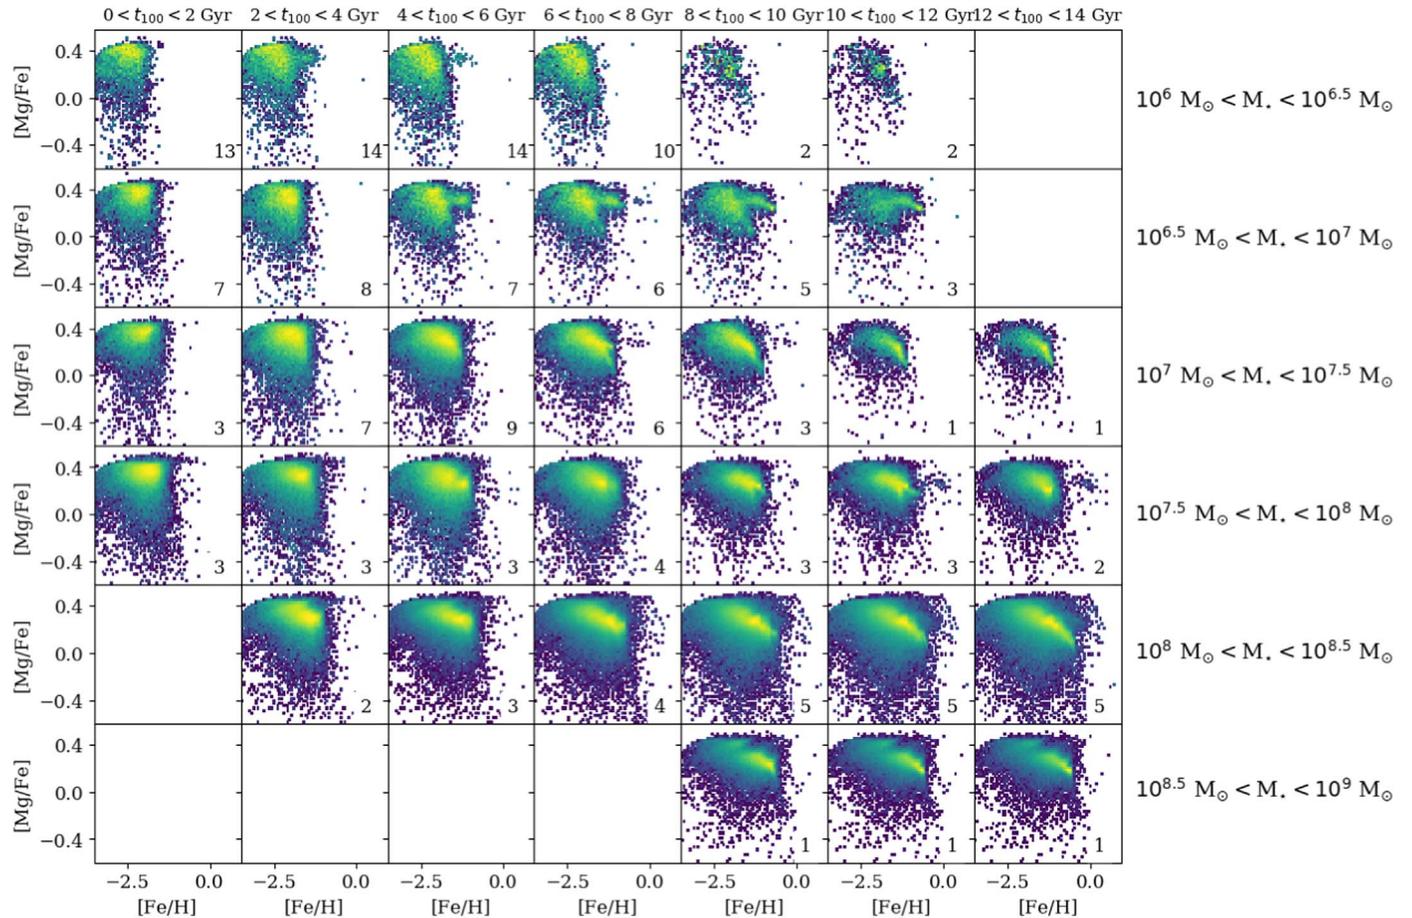


Figure 6. Same as Figure 5, but for templates constructed from dwarf galaxies.

templates) have a higher density in their CARs at higher $[\text{Mg}/\text{Fe}]$, whereas the templates with larger age ranges (i.e., high t_{100} , on the right-hand side of the grid) have a higher density at lower $[\text{Mg}/\text{Fe}]$. This results from the longer timescales associated with Type Ia SNe (the primary source of iron enrichment) with respect to Type II SNe (the primary source of α -element enrichment).

While the expected general trends are found in both template sets, there are a few salient differences between the two template sets worth noting. First of all, the fact that the present-day dwarf galaxies are *temporally biased* with respect to their disrupted counterparts can be identified based on the low number of S/PM templates with long age ranges relative to the DG templates (i.e., the grid in Figure 5 is relatively sparsely populated on the right-hand side compared to the grid in Figure 6). Therefore, had we constructed the template sets based on the present-day properties alone, we would have found that many of the S/PM accretion events would not have been found within the parameter space of the present-day DG templates.

In addition, we note that this technique relies upon the assumption that the dwarf galaxies and halo progenitors with similar stellar masses and age ranges will have similar CARs; we assume that the CARs within our template set will be representative of the halo progenitor CARs. Comparing the different CARs on Figures 5 and 6, we can immediately see that not all of the distributions within the same gridpoints are

identical. We explore some of these differences in more detail in Section 6.

We emphasize that these templates are empirical templates derived from the FIRE simulations and are designed to be used on the FIRE simulations. The templates from the FIRE simulations should not be applied to observational data. While the overall properties of the simulations make them excellent testbeds for studying galaxy formation phenomena, it remains important to keep in mind that the abundances are not literal abundances, but rather labels in a simulation. As discussed in Section 2.1, the average $[\text{Fe}/\text{H}]$ of the stars in the satellites in these these simulations are too low compared to the observed dwarf galaxies (Escala et al. 2018; Wheeler et al. 2019; P21), likely as a result of the assumed delay-time distributions for SNe Ia. In addition, the harsh ridge seen in all the CARs at high $[\text{Mg}/\text{Fe}]$ results because of the initialization: star particles are initialized with $[\text{Fe}/\text{H}] = -4$ and $[\alpha/\text{Fe}] = 0$, resulting in a hard upper limit for the α abundances at a given iron abundance. Both of these particular issues are addressed with the new FIRE-3 physics model (Hopkins et al. 2022). Therefore, while these templates can be used in estimating the assembly histories of the Latte galaxies, we emphasize that an empirical template set for use on observational data should be constructed from observational data (or from theoretical models calibrated to reproduce the observed abundance distributions).

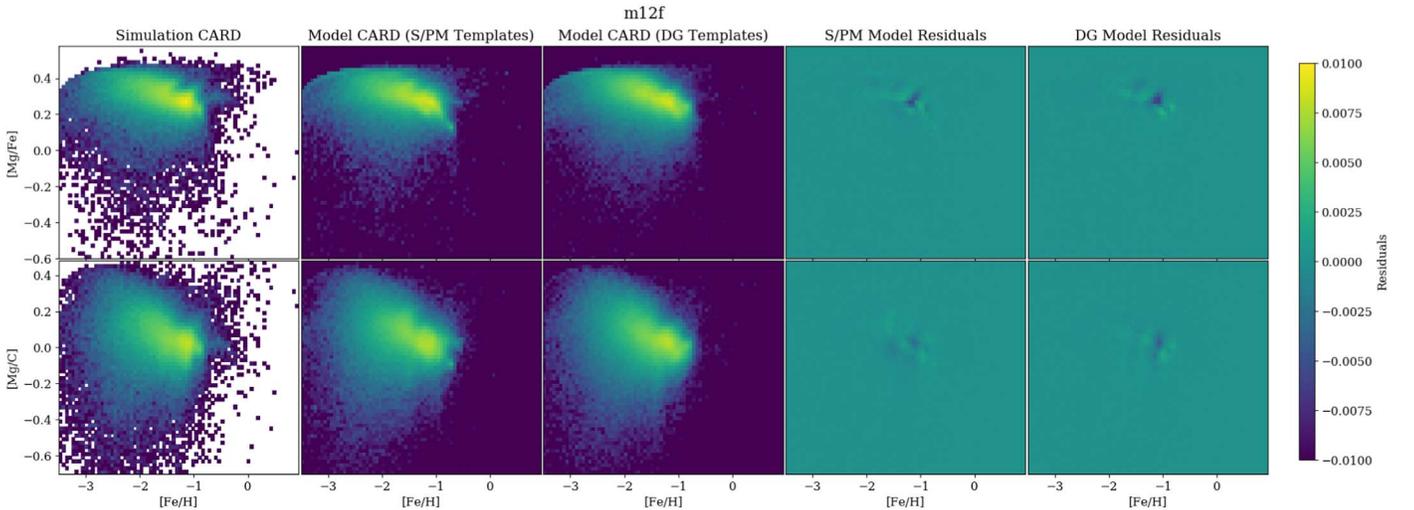


Figure 7. CARDs results for the simulation m12f. The true simulation CARDs (using the streams and phase-mixed debris identified in P21) are shown in the far left panels, with [Mg/Fe] vs. [Fe/H] plotted in the top row and [Mg/C] vs. [Fe/H] shown in the bottom row. The resulting model CARDs using the S/PM templates are shown in the second from the left panels, and the CARDs using the DG templates are in the center. Residuals from all models are shown in the right-hand panels; the magnitude of the residuals is shown by the colorbar. For the case of m12f, the models utilizing all of the different template sets produce a good fit to the data, as demonstrated by the very low residuals.

5. Results

In this section, we discuss the results of using CARDs to estimate the assembly histories of the stellar halos from the Latte suite of simulations. We discuss the results from two simulations in detail: one in which the technique performs very well (m12f) and one for which the technique does not accurately recover the assembly history (m12c).

For our two example cases, in Section 5.1, we first explore how well our four template sets are able to reproduce the stellar halo CARDs. In Section 5.2, we discuss how well the CARDs technique recovers the mass spectrum of accreted dwarfs. In Section 5.3, we demonstrate how well we can recover the mass accreted as a function of progenitor quenching time. In Section 5.4, we summarize the results from our analysis for all the simulations. For figures and detailed discussion of the remaining five simulations not discussed in this section, we refer the reader to Appendix B.

5.1. Model CARDs

In assessing the performance of this technique, we begin by examining the results from what we are modeling directly: the simulation abundance distributions.

In this section, we compare the CARDs derived from our modeling procedure with the true simulation CARDs, for the two simulations we discuss in detail. We first focus on our success case: m12f. The CARDs from the accreted star particles (as defined in Section 2.2) are shown in the far left-hand panels of Figure 7; we model both [Mg/Fe] versus [Fe/H] (top panels) as well as [Mg/C] versus [Fe/H] (lower panels).¹⁶ The resulting model CARDs, using our S/PM templates (second from the left) and dwarf galaxy templates (middle panels), appear to be a good representation of the data. This is quantified by the model residuals, as shown in the far right panels; the value of the residuals (defined here as model

–data) are indicated by the colorbar. We see that, for this simulated galaxy, the residuals are quite low, and the CARDs constructed from all four sets of templates match the halo CARDs very well.

In contrast, Figure 8 shows the CARDs results for m12c. Even by eye, the fact that the model CARDs are not good representations for the simulation CARDs is immediately apparent. The CARD for m12c shows a high density near [Fe/H] ~ -1 and [Mg/Fe] ~ 0.2 . This feature, which arises from the most massive accretion event in m12c ($M_* = 10^{8.7}$ and $t_{100} = 12.35$ Gyr), is not captured by the templates; as a result, the model strongly underpredicts the density in this region of the CARD. Unsurprisingly, given the poor fit to the data, the CARDs technique does not recover the assembly history correctly.

The inability of the method to accurately reproduce the CARD of m12c highlights the limitations of a fundamental assumption of the model: that accretion events with similar M_* , t_{100} will have similar CARDs. We discuss the diversity in the formation histories (and therefore the CARDs) of the most massive dwarfs in our sample in detail in Section 6.

5.2. Mass Spectra of Accreted Dwarfs

Having assessed the fit of our model CARDs, we turn now to the ability of the technique to recover the assembly histories. We emphasize again that the model presented in this paper contains no information about mass or time: the model only uses the simulation CARDs and the CARDs from the templates. However, using the stellar mass and t_{100} labels for each of our templates, we can assess how well the technique recovers the assembly histories of the simulated galaxies (which are modeled indirectly). We consider the mass spectrum of accreted dwarfs (i.e., the fraction of mass contributed to the halo as a function of the mass of the progenitor) and the stellar mass accreted as a function of time (where we again use t_{100} instead of *accretion time*).

Figure 9 shows the mass spectra for the simulations m12f (left panels) and m12c (right panels), with the true mass spectra shown as the gray squares connected by the gray lines. The

¹⁶ The reader may note that CARD in the top left panel of Figure 7 is different from the m12f CARD shown in Figure 1; this is due to the fact that the massive accretion event in m12f is above the mass limit for the P21 catalog and excluded from this analysis.

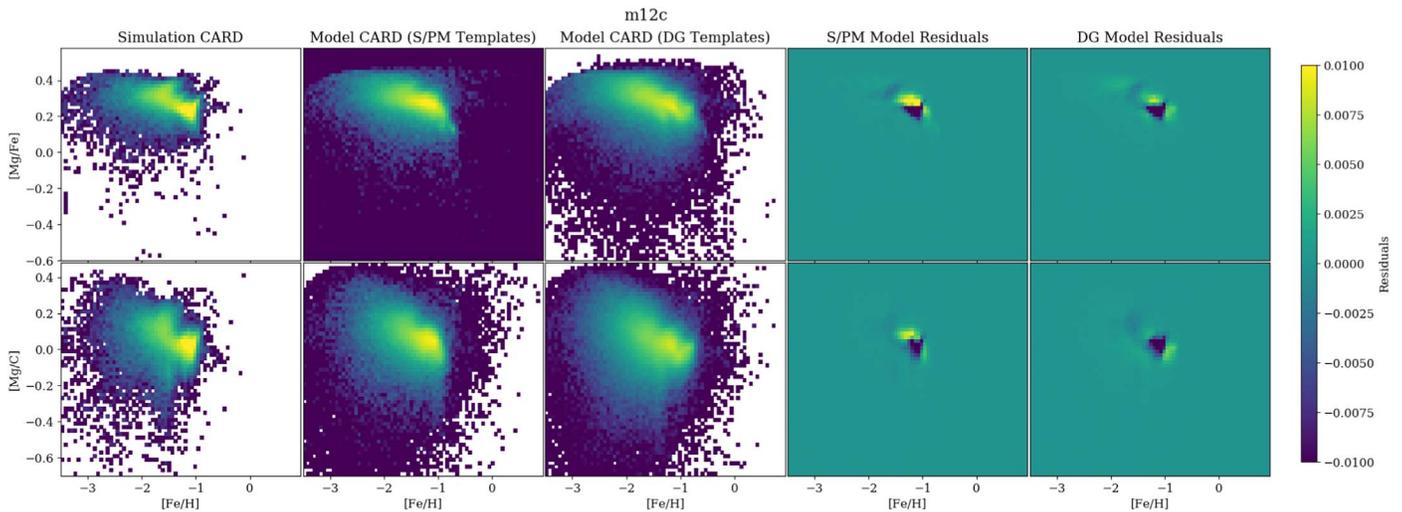


Figure 8. Same as Figure 7, but for the simulation m12c. In contrast to the m12f results, none of the models shown here are a good fit to the data: the high-density feature in the simulation CARD at $[\text{Fe}/\text{H}] \sim -1$ is not well represented by the templates. As a result, all templates result in model CARDS with very high residuals.

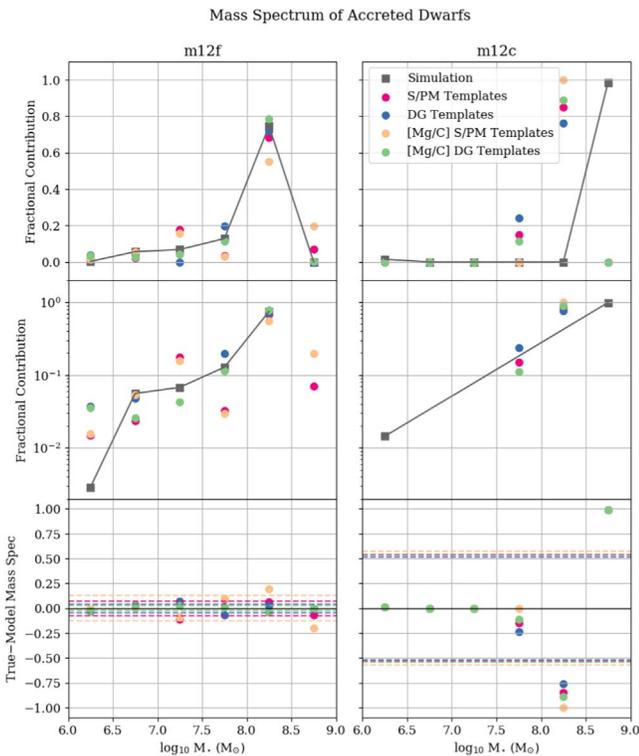


Figure 9. Mass spectra of accreted dwarfs as derived from the different template sets, for the best case (m12f) and a failure case (m12c). Top panels show the mass fraction as a function of progenitor stellar mass on a linear scale; middle panels show the mass fractions on a logarithmic scale. The true mass spectrum from the simulations are shown by the gray squares. The resulting mass spectra derived using the stream/phase-mixed templates, dwarf galaxy templates, the S/PM templates for the $[\text{Mg}/\text{C}]$, and the dwarf galaxy templates for the $[\text{Mg}/\text{C}]$ distributions are shown by the pink, blue, peach, and green points, respectively. Lower panels show the residuals from the results derived from each set of templates, with the rms dispersion for each fit indicated by the dashed lines. The results for m12f are shown on the left, and the results for m12c are shown on the right.

derived mass spectra from the models using the S/PM templates, the DG templates, the $[\text{Mg}/\text{C}]$ S/PM templates, the $[\text{Mg}/\text{C}]$ DG templates are shown in pink, blue, peach, and green, respectively. The top panels show the mass

spectrum plotted linearly in a mass fraction; the middle left panel shows the results logarithmically in a mass fraction. The lowest panel shows the residuals from each model, as well as the rms dispersion of the residuals as the dashed lines.

It is immediately apparent that the mass spectrum of accreted dwarfs is recovered extremely accurately for m12f, using all template sets. The rms dispersion of the residuals is 0.05 for the S/PM templates, 0.04 for the dwarf galaxy templates in $[\text{Mg}/\text{Fe}]$ versus $[\text{Fe}/\text{H}]$; 0.09 for the $[\text{Mg}/\text{C}]$ versus $[\text{Fe}/\text{H}]$ S/PM templates; and 0.03 for the templates for $[\text{Mg}/\text{C}]$ versus $[\text{Fe}/\text{H}]$. While the logarithmic panel highlights that the dwarf galaxy templates overestimate the contribution of the lowest-mass bin (finding $\sim 3\%$ – 5% contribution), overall, the CARDS modeling technique recovers the mass spectrum very accurately, with the $[\text{Mg}/\text{C}]$ DG templates recovering the true mass spectrum to the highest precision. We also note that a very low-mass contribution from high mass accretion events is predicted for both sets of S/PM templates; given that this is unphysical (as a low contribution from a $10^9 M_\odot$ accretion event implies a very large halo mass), this could be addressed by constructing a more complex model incorporating prior information on the total mass of the halo and modeling the number of accretion events instead of the mass fractions.

In contrast, the right-hand panels of Figure 9 show the mass spectrum results for the m12c analysis. Unsurprisingly (based on the poor fit to the simulation CARD), the true mass spectrum is not well recovered. All template sets perform equally poorly at recovering the mass spectrum. Because the dwarf progenitor of the most massive accretion event in m12c has high density at $[\text{Fe}/\text{H}] \sim -1$, and does not continue to enrich to higher $[\text{Fe}/\text{H}]$ and lower $[\alpha/\text{Fe}]$ as would be expected, the modeling procedure derives a mass spectrum where the dominant event occurs in a lower-mass bin.

We see from the m12f results that this technique is capable of accurately recovering the mass spectra of accreted dwarf galaxies. For 4/7 of the Latte suite halos (m12b, m12f, m12m, and m12w), we are able to recover the mass spectrum to a precision of 10% or less, using the DG and/or the $[\text{Mg}/\text{C}]$ DG templates. However, for the remaining 3/7 galaxies (m12c, m12i, and m12r), the mass spectra are not as accurately recovered: this occurs, as in m12c, when the CARDS from the

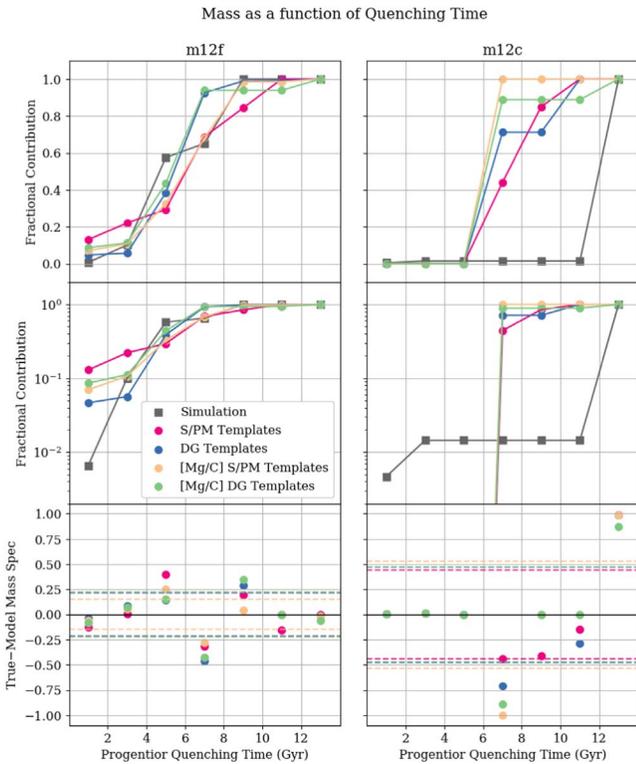


Figure 10. Cumulative mass fraction as a function of t_{100} , for the success case (m12f, left) and failure case (m12c, right).

halo accretion events are not well represented by the template set.

5.3. Mass Assembly versus Quenching Time

In the previous subsection, we discussed how well the CARDs method recovered the mass accreted as a function of the mass of the progenitor (i.e., the mass spectrum of accreted dwarfs). In this section, we look at the fraction of mass accreted as a function of the progenitor’s quenching time. Figure 10 shows the accreted stellar mass fraction as a function of progenitor t_{100} for m12f (left) and m12c (right). Because we are using t_{100} as a proxy for what is traditionally thought of as accretion time, we show the cumulative distributions; though the residuals in the lower panels of Figure 10 are shown for the mass fractions and not the cumulative mass fractions, as in Figure 9.

Based on the lower panels of Figure 10, it is immediately apparent that the residuals for the mass fraction as a function of t_{100} are significantly larger than those for the mass spectra. Even for m12f, the residuals have an rms dispersion of $\sim 0.2\text{--}0.3$; while the distribution of mass accreted as a function of t_{100} is not recovered as accurately as the mass spectrum, the model still does recover the fact that most of the mass is from accretion events with intermediate t_{100} (i.e., $3\text{ Gyr} < t_{100} < 7\text{ Gyr}$).

However, as with the mass spectrum, the inferred mass assembly over time for m12c is in poor agreement with its actual assembly history. Because of the anomalous CARD for the most massive accretion event in this halo, which does not enrich beyond $[\text{Fe}/\text{H}] \sim -1$ and $[\text{Mg}/\text{Fe}] \sim 0.2$, the model created from the templates favors an earlier assembly history than the galaxy actually experienced.

While generally the mass spectrum is recovered with higher accuracy than the mass assembled as a function of time, there are a number of prospects for potential improvements to the methods that could be explored in future work. For example, in this work, we have focused exclusively on t_{100} in creating templates and quantifying the properties of halo progenitors; however, it is possible that other timescales, such as t_{90} or the interquartile age range, could improve the inferences. Exploring the moments of the abundance distributions, as well as incorporating additional elements tracking different sources of chemical enrichment, could also prove useful in constraining the properties of the progenitors.

However, as with m12c in this work, if the template CARDs are not a good representation of the halo progenitor CARDs, the inferences for the assembly histories derived using this method will not be correct. Fortunately, our first clue that there is a mismatch between the halo progenitor CARDs and the template CARDs are the high residuals in comparing the model CARDs with the data. In the next section, we summarize the results of modeling all seven of the halos from the Latte suite used in this work, and demonstrate the connection between the residuals in the CARDs models and the residuals in the inferred assembly histories.

5.4. Summary of All Simulation Results

While we show the detailed results for the assembly histories of the individual simulations in Appendix B, Figure 11 summarizes the residuals from modeling the CARDs of all of the Latte halos. The top panels demonstrate the results from inferring the mass spectrum of accreted dwarfs, and the lower panels show the results from inferring the fraction of mass accreted as a function of progenitor t_{100} . In all panels, each color represents the results from a different simulation. The circular points show the results from using the $[\text{Mg}/\text{Fe}]$ versus $[\text{Fe}/\text{H}]$ template sets, while the square points indicate the results from using the $[\text{Mg}/\text{C}]$ versus $[\text{Fe}/\text{H}]$ template sets. The open symbols show the results from using the S/PM templates, and the filled points show the results from the DG templates.

The x -axes on the left panels of Figure 11 indicate the rms dispersion of the residuals from the CARDs models (i.e., the rms dispersion of the residuals plotted in the right-hand panels of Figures 7 and 8). The y -axes of the top panels show the rms dispersion of the residuals for the mass spectra, and the y -axes of the lower panels show the rms dispersion for the mass assembly histories.

This figure demonstrates the following key results:

1. The CARDs method infers the mass spectrum of accreted dwarfs to a precision of 10% or better using our DG templates for four out of the seven simulations (m12b, m12f, m12m, and m12w). The top left panel of Figure 11 shows a clustering of points in the lower left-hand corner: for these simulations, the CARD has been modeled accurately (i.e., with low residuals), and the mass spectrum has also been recovered accurately. While the m12w CARD model has high residuals, its mass spectrum is still inferred relatively accurately for the DG templates and the $[\text{Mg}/\text{C}]$ templates; see Appendix B for a detailed discussion of the individual simulation.
2. The fraction of mass as a function of t_{100} is generally inferred less accurately than the mass spectrum, with a

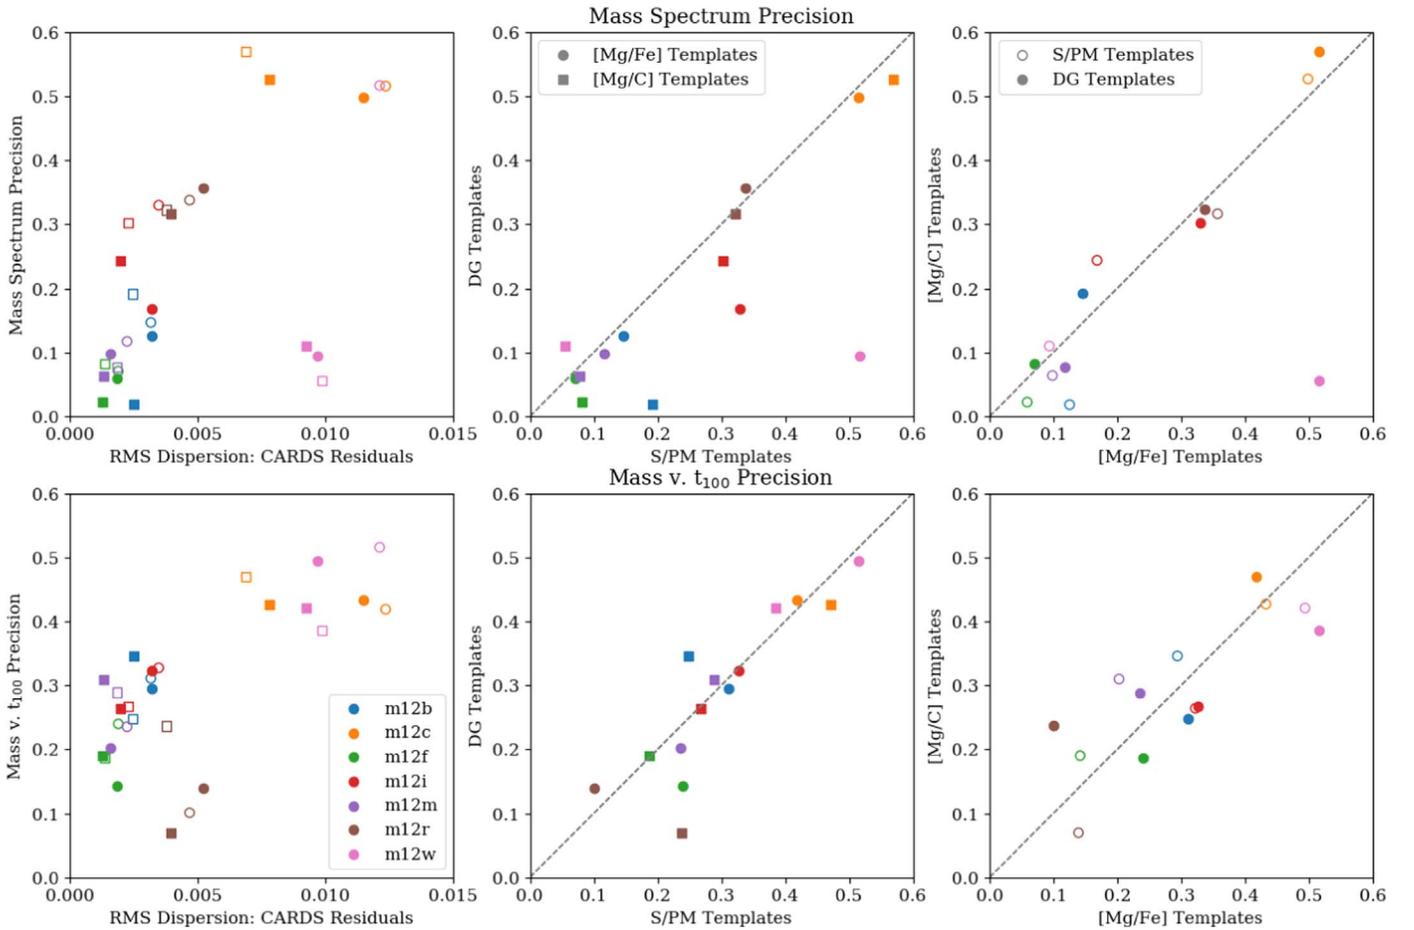


Figure 11. Summary of our results for all simulations. Colors across all panels correspond to results from specific simulations. Top panels show the resulting precision from inferring the mass spectrum of accreted dwarfs, and bottom panels show resulting precision from inferring mass as a function of t_{100} . Left panels: inferred precision of mass spectrum and t_{100} vs. the rms dispersion of the model CARDS. Middle panels: comparison of the results from the DG templates vs. the S/PM templates. Results from the [Mg/C] templates are shown as squares while the results from the [Mg/Fe] templates are shown as circles. The DG templates perform comparably to or better than the S/PM templates. Right panels: comparison of the results from the [Mg/C] templates vs. the [Mg/Fe] templates. S/PM template results are shown as open symbols, while DG template results are filled circles. We see some marginal improvement in the inferred mass spectra by including [Mg/C], although no real improvement in the inferred quenching times.

precision on the order of $\sim 20\%$ – 30% for five out of the seven simulations (m12b, m12f, m12i, m12m, and m12r). Comparing the top left panel with the bottom left panel of Figure 11 shows that the residuals are higher for the mass accreted as a function of quenching time than for the mass spectrum. The only exception to this is m12r; see Appendix B for further discussion.

- Our first indication that the assembly history will not be inferred accurately is the quality of the fit to the halo CARD. The x -axis on the left panels of Figure 11 indicates the rms dispersion of the residuals from the CARDS models (i.e., the rms dispersion of the residuals plotted in the right-hand panels of Figures 7 and 8). We see a correlation between the rms dispersion in the CARDS residuals and the rms dispersion of the parameters of the accretion history (the exception being the mass spectrum inference for m12w; see Appendix B for discussion). As discussed in the previous section, if the templates are not representative of the halo accretion events, the model CARDS will be a poor fit to the true CARDS, and the derived accretion histories will not be correct.

- The DG templates perform comparably to, if not better than, the S/PM templates. In the middle panel of Figure 11, we directly compare the results of our modeling procedure using DG templates (y -axis) versus S/PM templates (x -axis). The fact that most of the points lie below the 1:1 line in the top middle panel of Figure 11 demonstrates that the DG templates infer the mass spectrum better than the S/PM templates for the majority of the simulations, for both [Mg/Fe] versus [Fe/H] template sets as well as [Mg/C] versus [Fe/H]. This is likely a result of the fact that the DG templates contain more sources than the S/PM templates; we have more CARDS to average over in creating the DG templates than for the S/PM templates. We test this assertion by creating three sets of *subsampled* DG templates, and find that the templates created from the smaller sample of dwarfs are more likely to infer the mass spectrum less accurately than the full DG template set. See Appendix A for more details about this analysis.
- The inclusion of information from carbon (which traces stellar wind enrichment) in the [Mg/C] templates marginally improves the inferences of the mass spectrum, but does not appear to improve the inferences of the

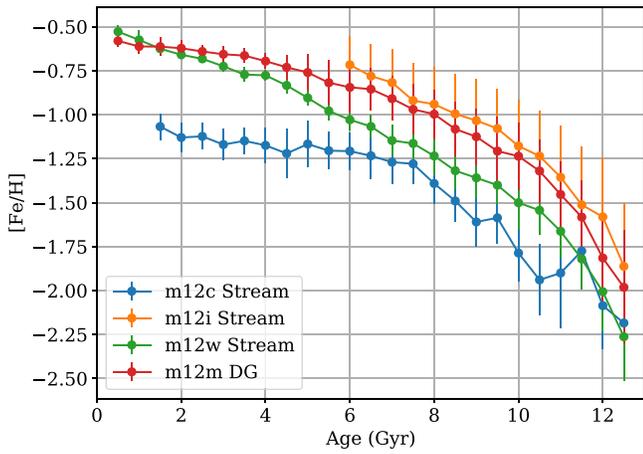


Figure 12. Mean iron abundance as a function of stellar age for the massive streams and dwarf galaxy in the P21 catalog. Errorbars indicate the sample variance in each age bin. The iron abundance of m12c is lower than the other dwarfs/halo progenitors in its mass range at all stellar ages.

quenching times. The right panels of Figure 11 compare the results from using the [Mg/C] templates versus the [Mg/Fe] templates. The top panel demonstrates that the inferences of the mass spectrum for both template sets are comparable, with most results lying near the 1:1 line. However, the [Mg/C] templates do infer the mass spectrum slightly more accurately on average. The results for mass accreted versus t_{100} are comparable for the two template sets, with the [Mg/Fe] templates performing slightly better on average.

What gives rise to our less accurate inferences? If the template CARDS are not representative of the halo progenitors, our inferences about the assembly history will not be accurate. In particular, Figure 11 highlights that the inferences for m12c and m12w are usually the worst: while both of these simulations experience a massive, late time accretion event, their CARDS are very different. In the next section, we look more closely at the four most massive sources in the P21 catalog in detail. For the results of modeling the CARDS excluding accretion events with $M_* > 10^8 M_\odot$, we refer the reader to Appendix C.

6. Discussion

6.1. Evolutionary Histories of Massive Dwarfs

In the previous section, we discussed at length that the simulation m12c’s CARD and assembly history were poorly recovered, largely due to the fact that m12c’s most massive accretion event ($M_* = 10^{8.7} M_\odot$, $t_{100} = 12.35$ Gyr) has a CARD that is not well represented by the templates. A fundamental assumption of the CARDS modeling method presented here and in L15 is that halo progenitors with the same (M_* , t_{100}) should have similar CARDS. Why does the most massive accretion event in m12c have such an odd abundance distribution for a stream with its mass and t_{100} ? Here, we discuss in more detail the chemical enrichment histories of the most massive streams and dwarf galaxies in our sample, to better understand the diversity of the observed CARDS.

The P21 catalog contains four objects that have stellar masses $M_* > 10^{8.5} M_\odot$; massive streams in m12c, m12w, and m12i; as well as a present-day dwarf galaxy in m12m (the only source for the most massive template for the DG templates).

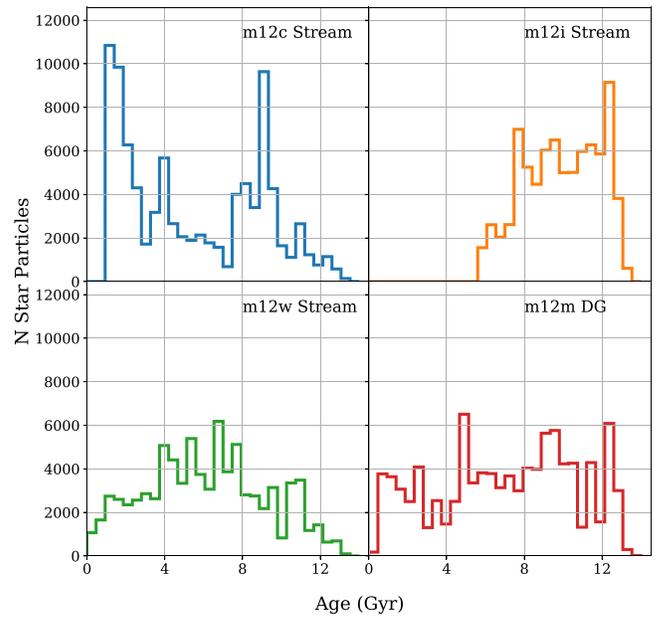


Figure 13. Age distributions of star particles for the massive streams and dwarf galaxy from the P21 catalog. The peaks in the age distributions for m12c illustrate that star formation history of the m12c stream is significantly burstier than the other massive streams and the massive dwarf galaxy, with a burst at ~ 9 and a burst around ~ 2 (following its first pericentric passage). The stream in m12c forms approximately 25% of its mass after its first pericentric passage. As a result of its bursty star formation history, it does not enrich as steadily as the other massive systems.

Figure 12 shows the mean metallicity of particles for each of these massive objects as a function of age; the stream in m12c (shown in blue) is significantly more metal-poor at all times than the other massive streams and dwarf galaxy.

To get insight as to why m12c is more metal-poor than the other systems with comparable stellar mass, Figure 13 shows the age distributions (i.e., the star formation histories) of the star particles in each system. The stream in m12c has a very peaky age distribution, indicating two major bursts of star formation. The first burst occurs at around ~ 9 Gyr, while the second burst occurs around ~ 2 Gyr. The stream progenitor first crosses the virial radius of the halo 3 Gyr ago, and experiences its first pericentric passage 1.5 Gyrs ago; its last burst of star formation, which is responsible for approximately 25% of the stellar mass in the system, begins after the progenitor has entered the halo and continues beyond its first pericentric passage. This is in stark contrast to the framework in, for example, the Bullock & Johnston (2005) simulations, in which all dwarfs are assumed to lose their gas to ram pressure stripping upon falling into the main host halo.

However, this is consistent with a recent study Di Cintio et al. (2021), who find that close pericentric passages of relatively gas-rich systems can result in pericentric passage triggered starbursts, arising due to the compression of cold gas. Di Cintio et al. (2021) find that 25% of the dwarfs in their simulations experience these starbursts, requiring that the gas mass fraction be $M_{\text{gas,inf}}/M_{\text{vir,inf}} \geq 10^{-2}$ and pericenter be greater than 10 kpc. An example of gas compression at the pericenter triggering star formation in satellites in the Auriga simulations is also discussed in Simpson et al. (2018). While a detailed exploration of the gas properties of the halo progenitors in the Latte simulations is beyond the scope of this work (see Samuel et al. 2022), we emphasize that, in these

simulations, star formation after infall (and after first pericentric passage) can significantly contribute to the stellar mass of halo progenitors, and affects its overall CARD.

Figures 12 and 13 highlight both a limitation of the CARDS technique as well as the potential insights we can glean from studying the abundance patterns of halo progenitors. The diversity in the SFHs and the abundance distributions at high masses mean that their properties are not necessarily well recovered from the templates; if the abundance distribution of the halo progenitor is not well represented by the template abundance distribution, the technique will not derive the correct parameters for the assembly history. However, the signatures of these interesting formation histories are encoded in the abundance patterns. We therefore emphasize the fact that measuring the stellar parameters and abundances within the MW stellar halo provides a unique window into galaxy formation in the high-redshift universe.

6.2. Future Work Within the Simulations

While there are limitations to using CARDS to constrain assembly histories, as discussed in the previous section, there are also prospects for further improving and expanding this method. First of all, a natural extension of the work presented here is to supplement the P21 catalog with the earlier accretion events identified and presented in D. Horta et al. (2022, in preparation). The Horta catalog includes accreted galaxies down to a stellar mass limit of $10^7 M_{\odot}$; extending the catalog of known accretion events in the Latte suite down to lower-mass limits at earlier times is an ongoing effort. Once we have complete catalogs in hand for the progenitors of all the Latte halos, it will be possible to test the effectiveness of the CARDS framework at recovering the full assembly histories of the Latte suite.

In addition, improvements in the simulation may improve our ability to constrain the quenching time of our halo progenitors. The upcoming FIRE-3 physics model (Hopkins et al. 2022) includes many updates and improvements in modeling stellar evolution. In addition to the updated physics, there will also be a flexible new framework for exploring nucleosynthesis and metal enrichment (A. Wetzel et al. 2022, in preparation), enabling the user to implement different feedback and chemical enrichment models in post-processing. With this model (known as the “age-tracer” model), revised yields (as well as additional sources of enrichment, including r- and s-process) can be implemented to compute revised abundances at $z = 0$. This has particularly exciting implications for this work: creating CARD templates that include information about more elements, including those that trace chemical enrichment on different timescales, is likely to improve the ability of the CARDS method to recover the quenching times of halo progenitors. Furthermore, the option for implementing flexible yields enables the option for tuning the $z = 0$ abundance distributions to match the observed dwarf galaxies. After matching the simulated CARDS to observed dwarf galaxy CARDS, we could create a template set from the simulations that could be used on the observations.

7. Conclusions

The halo of the MW contains the remnants from as many as thousands of low-mass galaxies that formed at high redshift and did not survive to the present day. However, we currently lack

quantitative frameworks for quantifying the contributions from these systems and identifying the stars that were born within them. In this paper, we have experimented with modeling the CARDS of the stellar halos of the Latte simulations, as a linear combination of empirical templates for halo progenitors with different stellar masses and accretion times.

In the Introduction, we pose three questions we seek to address in this work. We return to these questions here, and summarize our main results.

1. *Does the method proposed by L15 of modeling the stellar halo CARD as a linear combination of templates work in a more realistic setting?* We find that CARDS can be effectively used to explore histories in a realistic setting.
 - (a) We demonstrate the link between formation history and CARDS in the Latte simulations. Each of the Latte halos has distinct CARDS, as a result of their diverse formation histories.
 - (b) The CARDS method recovers the mass spectrum very accurately (to within 10%) for four out of the seven presented here. The CARDS method shows promise for accurately inferring the mass spectrum of accreted dwarfs in LG galaxies.
 - (c) The fraction of mass accreted as a function of halo progenitor quenching time is generally inferred less accurately than the mass spectrum. We infer the accreted mass fraction as a function of progenitor quenching time with a precision of 20%–30% for five out of the seven halos studied in this work. We discuss the prospects for improving the performance of this technique, such as including information from additional elements, as well as upcoming improvements to the chemical evolution models in the simulations.
2. *What is an observationally viable method for constructing templates? Can we use the CARDS of present-day dwarf galaxies to infer the properties of the disrupted dwarfs?* We present a method for using dwarf galaxies to construct templates for halo progenitors, and find that the templates from the dwarf galaxies successfully recover the halo assembly histories.
 - (a) Building off of the discussion of P21, we present the differences between the present-day dwarf galaxies, coherent streams, and phase-mixed debris. At a fixed stellar mass, dwarf galaxies tend to have longer t_{100} 's than coherent streams and phase-mixed debris. At a fixed t_{100} , phase-mixed debris tends to have the largest stellar mass. This in turn has implications for the CARDS of the different populations: we show how phase-mixed debris is, on average, enhanced in magnesium relative to the dwarf galaxies with the same mean metallicity. Therefore, a framework for using dwarf galaxy CARDS to model streams and phase-mixed debris must take this effect into account.
 - (b) We present a method for creating templates for halo progenitors (i.e., streams or phase-mixed debris) using exclusively present-day dwarf galaxies, which accounts for the fact that present-day dwarf galaxy populations generally form their stars over a longer period of time. Our S/PM and DG templates are generally in good agreement.
3. *How can we assess the accuracy of this method?* We identify the failure cases to this method when the best-fit

models are poor fits to the data, as indicated by high model residuals.

- (a) For two out of the seven halos studied in this work, the CARDS from the templates (using the best-fit model) are not good fits to the data, as demonstrated by the high rms dispersion of the residuals. We find that a poor fit to the data generally indicates that the derived properties of the accretion histories are untrustworthy. This arises as a result of the fact that the abundance distributions for all dwarf galaxies and halo progenitors exhibit diversity in their star formation and chemical enrichment histories. We highlight an example due to pericentric passage induced starburst in m12c. Therefore, in interpreting the distributions of halo progenitors, it remains important to recall that we are studying the remains of galaxies, which may have experienced their own unusual histories before being disrupted in the MW halo.

This work marks only a first step in exploring the full capabilities of using detailed studies of dwarf galaxy populations to unravel the MW assembly history using CARDS. One complication not addressed in this work includes observational uncertainties; the precision required for spectroscopic abundance measurements (for both halo stars and dwarf galaxy stars) remains to be explored. Furthermore, we have neglected the presence of an in situ halo population. In addition, we have constructed our templates in the context of simulations, where we have had access to all information about our dwarf galaxies, including ages and star formation histories; the optimal methods for constructing dwarf galaxy templates from observations remains to be explored.

However, in this work, we have undertaken a critical first step of testing the CARDS modeling framework in the context of a cosmological simulation. We emphasize again that the model presented here contains no physics or prior information, and there are many ways in which more complexity and information could be incorporated into this model. For example, prior information about the mass of the halo, the mass function of DM halos from simulations, or kinematic information are all likely to improve the quality of the inferences. However, for the purposes of this paper, we demonstrate the wealth of information that can be extracted from chemical abundances alone. The power of chemical abundances alone is especially important for characterizing the origin of the phase-mixed halo.

Through the current and next generation of spectroscopic surveys, we are well positioned to use halo star chemical abundances to learn more about the evolutionary histories of the halo progenitors of the MW and M31. Only in the LG can we test our understanding of star formation and chemical enrichment in low-mass galaxies in the high-redshift universe, on a star-by-star basis.

E.C.C. is supported by a Flatiron Research Fellowship at the Flatiron Institute. The Flatiron Institute is supported by the Simons Foundation. The data used in this work were, in part, hosted on facilities supported by the Scientific Computing Core at the Flatiron Institute, a division of the Simons Foundation. E. C.C. would like to thank Duane Lee, Alis Deason, Erin Kado-Fong, Ted Mackereth, and the members of the Center for Computational Astrophysics Dynamics Group and Astronomical Data Group for useful scientific discussions. R.E.S. and N.

P. acknowledge support from NASA grant 19-ATP19-0068. R. E.S. further acknowledges support from NSF grant AST-2009828 and grant HST-AR-15809 from the Space Telescope Science Institute (STScI), which is operated by AURA, Inc., under NASA contract NAS5-26555. K.V.J.'s contributions were supported by NSF grant AST-1715582. A.W. received support from the following: NSF grants CAREER 2045928 and 2107772; NASA Astrophysics Theory Program grants 80NSSC18K1097 and 80NSSC20K0513; Hubble Space Telescope (HST) grants GO-14734, AR-15057, AR-15809, GO-15902 from STScI; a Scialog Award from the Heising-Simons Foundation; and a Hellman Fellowship. I.E. is supported by a Carnegie-Princeton Fellowship through the Carnegie Observatories. C.-A.F.-G. was supported by NSF through grants AST-1715216, AST-2108230, and CAREER award AST-1652522; by NASA through grant 17-ATP17-0067; by STScI through grant HST-AR-16124.001-A; and by the Research Corporation for Science Advancement through a Cottrell Scholar Award. We ran simulations using XSEDE, supported by NSF grant ACI-1548562; Blue Waters, supported by the NSF; Pleiades, via the NASA High-End Computing program through the NASA Advanced Supercomputing Division at Ames Research Center.

Appendix

In this appendix, we present additional figures and analysis for the interested reader. In Section A, we present example template sets in [Mg/C] versus [Fe/H] used in this analysis. In Section, we present detailed results for all the remaining Latte halos studied in this work. As in Section 5, we present the model CARDS for each halo as well as the inferred mass spectrum and fraction of mass accreted as a function of progenitor quenching time. In Section C, we present results for modeling the Latte halos excluding the halo progenitors with stellar masses greater than $10^8 M_{\odot}$.

Appendix A Additional Templates

Figure 14 shows the [Mg/C] versus [Fe/H] templates (referred to as the [Mg/C] templates throughout the text), constructed from streams and phase-mixed debris (top) and present-day dwarf galaxies (bottom). In this work, we found that including information from carbon marginally improved the inferred mass spectrum (see Section 7); however, we have dealt exclusively with templates constructed in 2D. We could further experiment with constructing templates with additional dimensions.

While we do not show them here, we also mention that we experimented with constructing *master* template sets, constructed from streams, phase-mixed debris, and dwarf galaxies. However, we did not include the results from this analysis here, as the results were essentially identical to the results from the DG templates. With the exception of the most massive row of templates, combining the S/PM sources with the DG sources did not significantly change the average template densities, resulting in very similar model CARDS to those produced by the DG templates.

For a subset of the halos, we also experiment modeling their early accretion events (accreted prior to 6 Gyr ago) using only an *early* template set (constructed from sources with $t_{100} < 6$

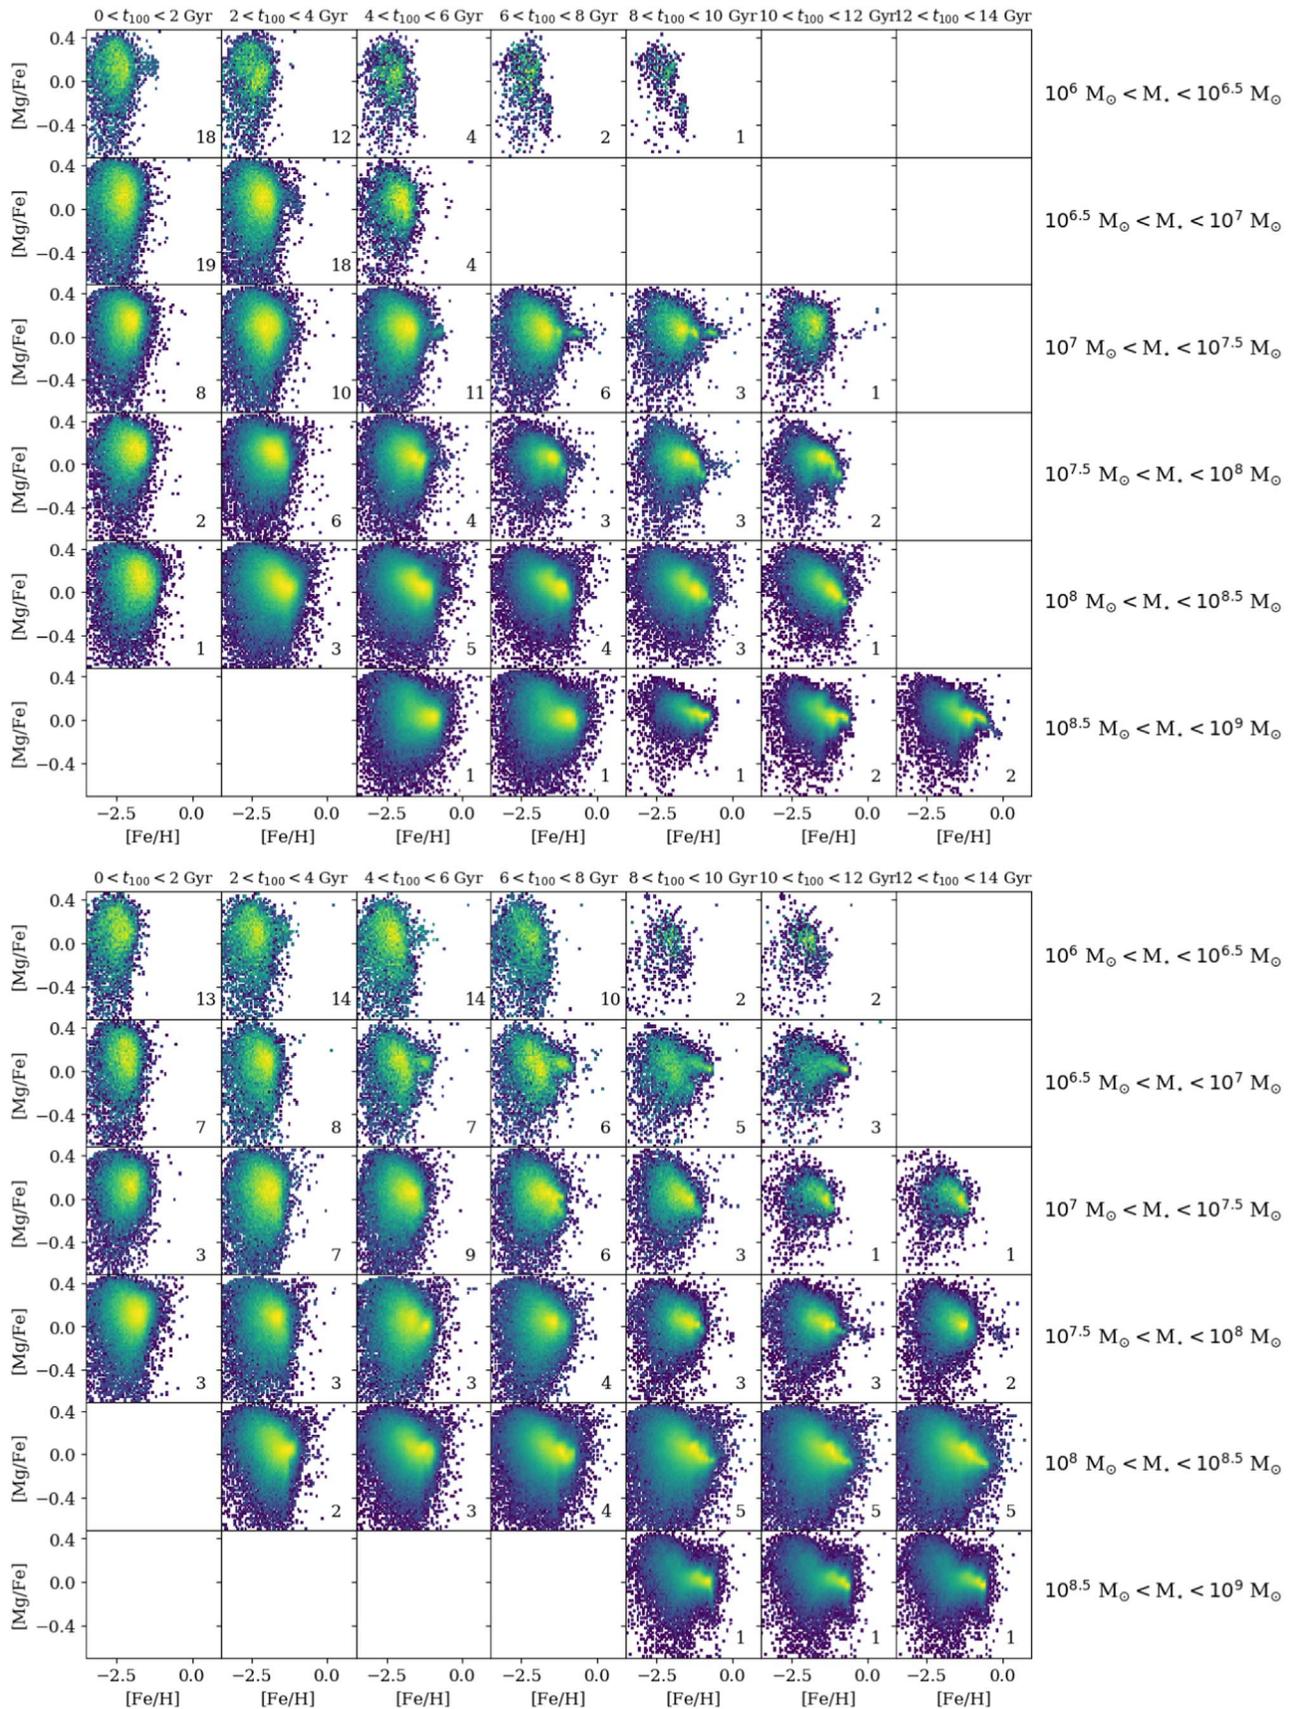


Figure 14. CARD templates in [Mg/C] vs. [Fe/H]. Top panels show the templates constructed from the streams and phase-mixed debris, while the lower panels show the templates constructed from present-day dwarf galaxies.

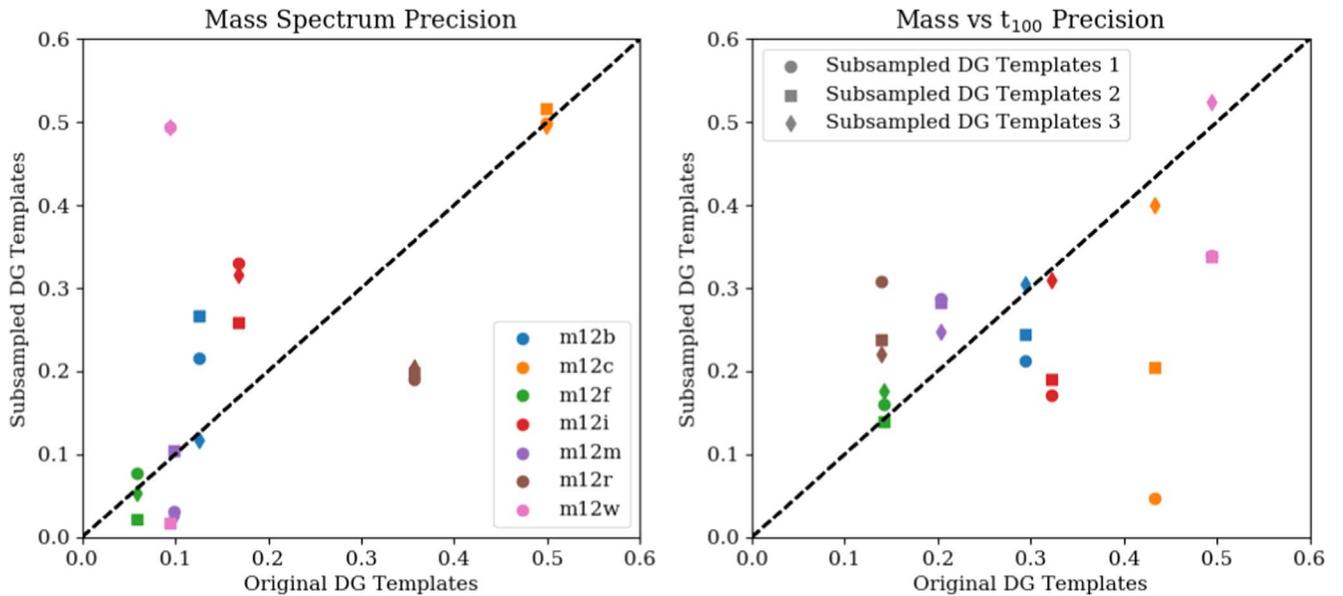


Figure 15. Results for inferring the mass spectra (left) and mass accreted as a function of t_{100} (right) for our original dwarf galaxy template set (x -axis) and our three sets of *subsamped* DG template sets. On average, the DG template sets infer the mass spectrum of accreted dwarfs more accurately than the subsampled template sets. The accuracy in inferring the mass accreted as a function of t_{100} , which is generally inferred less accurately than the mass spectrum, varies for each of the halos.

Gyr) and a *late* template set (from sources with $t_{100} > 6$ Gyr). For the halos we tested, we find that they perform comparably, as with the S/PM and DG templates. For one of the halos we tested (m12f), the *late* templates did perform much better than the early templates. Because m12f experiences the most massive accretion event prior to 6 Gyr, its CARD is better matched by the older stellar populations of massive streams and dwarf galaxies with later t_{100} 's. Therefore, as with the other template sets, if the stellar populations of the halo are well represented by the templates, we will do a reasonable job of inferring the assembly history.

Finally, motivated by the performance of the DG templates relative to the S/PM templates (see Section 5.4), we construct three sets of *subsamped* DG templates, including only dwarf galaxies from a different subset of the simulations (though we note that we include m12m in each set of subsampled templates, as it contains the only dwarf galaxy with $M_* > 10^{8.5} M_\odot$ in our sample). The results of this analysis are summarized in Figure 15. We find that the subsampled DG templates are more likely to infer the mass spectrum less accurately than the DG templates constructed from the full sample. The subsampled templates do occasionally perform just as well or better than the full DG template set, which will happen if a particular template is a very good match to a specific accretion event halo, but overall, the full DG templates infer the mass spectra more accurately than the subsampled templates.

For the results of the mass versus t_{100} , the differences in the performance for the subsampled templates are not as clear, with an improved inferred accuracy for some halos (e.g., m12i, m12w, and m12c), while the original DG templates perform better for m12f, m12m, and m12r. We therefore conclude that the number of template sources is not the limiting factor in improving the inferred mass accreted as a function of quenching time. We note that the S/PM templates perform equally well to the DG templates in inferring the mass accreted as a function of t_{100} (middle lower panel of Figure 11), consistent with this result.

Appendix B Detailed Results for All Simulations

In this appendix, we present and describe the detailed results from the five other simulations studied in this work. In Figures 16–19 we show the results from all the remaining simulations; we discuss the failure cases of this method in more detail below.

1. *m12i*. The second columns of the panels in Figures 18 and 19 highlight that none of the models do particularly well at recovering the assembly history of m12i. The reason for this is again the fact that the most massive accretion event in m12i ($M_*, t_{100} =$) is not represented by the templates. We can see this if we look back to Figure 5: the most massive accretion event from m12i is the only accretion event that contributes to the third and fourth panels along the bottom row of the S/PM templates in Figure 5. Furthermore, we can see these squares are left blank in Figure 6, as there are no dwarf galaxies included in P21 whose stellar populations can be used to populate those grid squares. This event is the only one within the entire P21 catalog to have $M_* > 10^{8.5}$ and a quenching time $t_{100} < 8$ Gyr. Therefore, the templates constructed here do not accurately recover their properties. Expanding the templates to include more massive systems could help improve the performance of the method here.
2. *m12r*. In Figure 1, we highlighted the formation history of m12r characterized by massive, recent accretion. From Figure 2, it becomes clear that these massive, recent accretion events are too massive to make it into the P21 catalog; the top left panels of Figure 17 show the CARD for the m12r particles modeled in this work. This abundance distribution looks different from the other halos, and has a feature at high metallicity and relatively high magnesium that is inconsistent with the other halos and the templates. This feature results in a poorly inferred mass spectrum (Figure 18), making m12r one of the

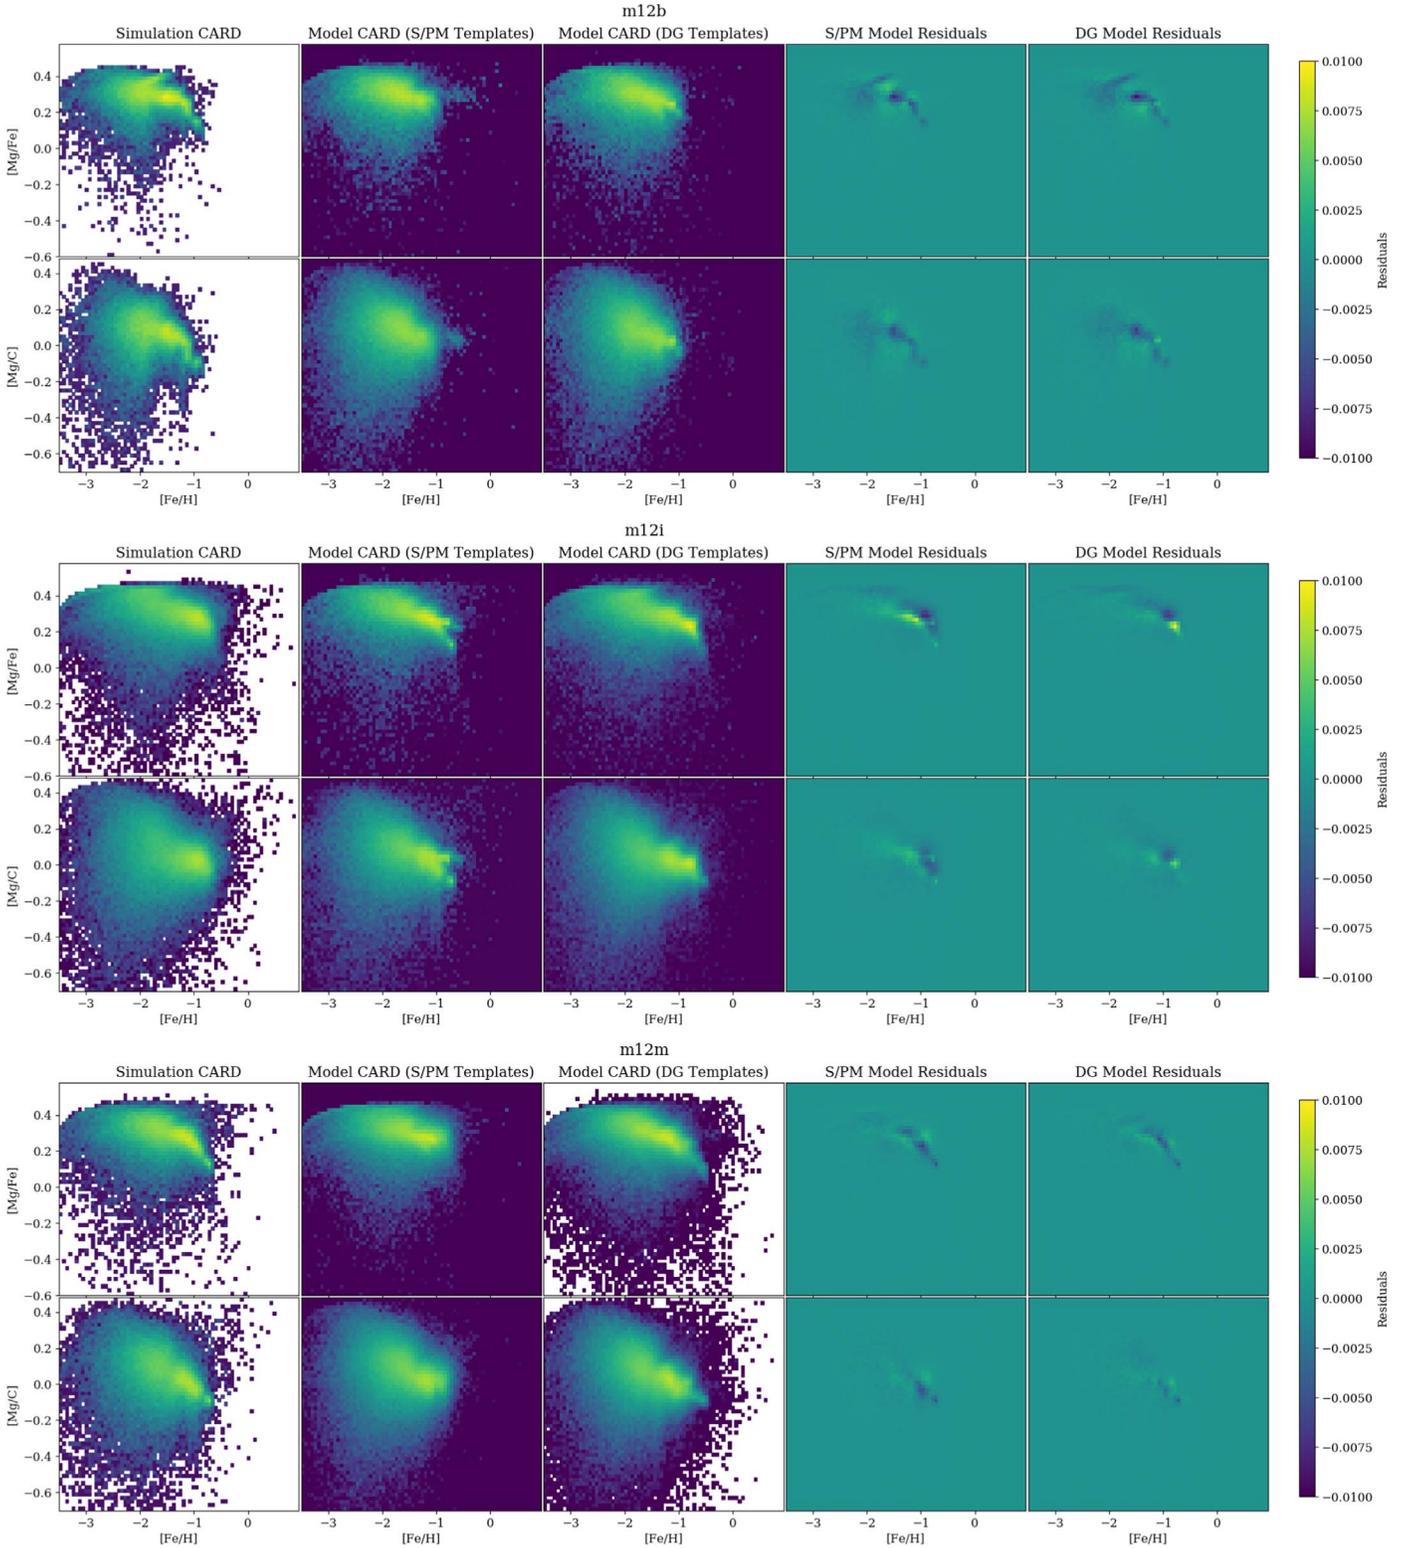


Figure 16. Same as Figures 7 and 8, but for the simulations m12b, m12i, and m12m.

failure cases for this method.

The reason that m12r has a strange abundance distribution is a result of group infall. While the most massive infalling satellites are excluded from this analysis as a result of the P21 selection criteria, the dominant accretion event included in this analysis in m12r ($M_{\star} = 1.9 \times 10^7 M_{\odot}$, $t_{100} = 9.1$ Gyr) has a CARD that

extends to higher metallicity than its counterparts in the template sets. This halo progenitor falls in with one of the massive infalling satellites (M , t_{100}), experiencing several pericentric passages during their collective infall and bursts of star formation during this interaction. Because it has a higher metallicity than expected based on its stellar mass and t_{100} , the CARDS method infers a stellar mass

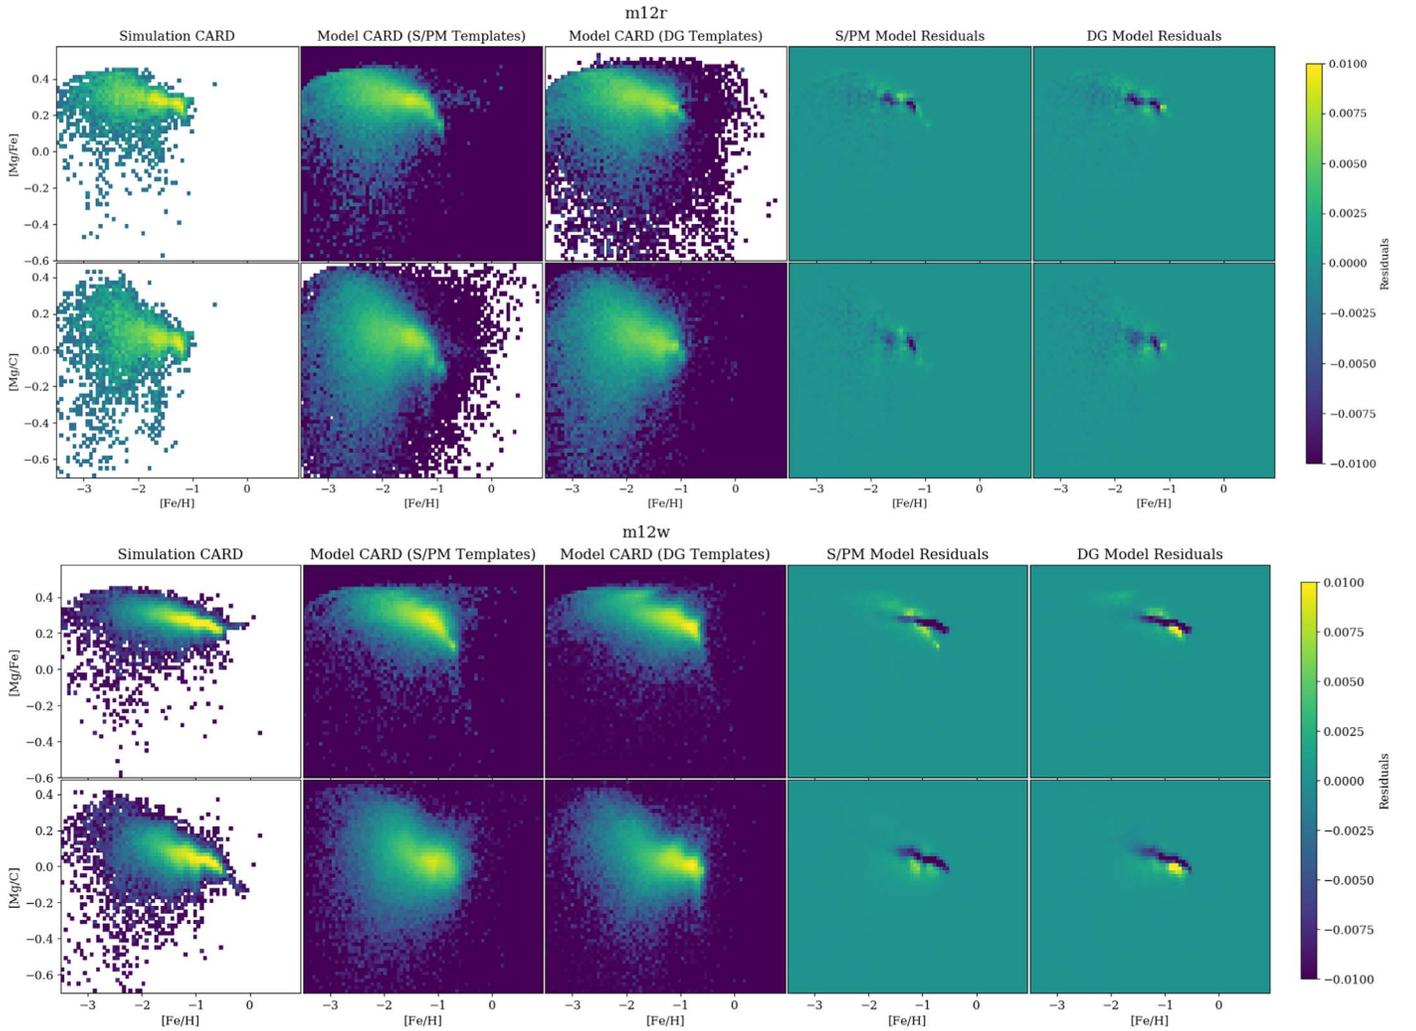


Figure 17. Same as Figures 7, 8, and 16 but for the simulations m12r and m12w.

that is too high for this progenitor. While this is technically a *failure case* for the method, it is also another example of how clues as to the interesting evolutionary histories of the halo progenitors are encoded in the halo CARDS.

3. *m12w*. Results from modeling the stellar halo of m12w are shown in the lower panels of 17. The simulation m12w is an interesting case, as its mass spectrum is inferred extremely accurately, but the timing of its major accretion event is inferred inaccurately, and the CARD models show relatively high residuals for all models. From the residuals, we see that the shape of the turnover in the [Mg/Fe] versus [Fe/H] plane is not a good match to the templates; however, as can be seen in Figure 12, the iron distribution as a function of age for the m12w massive stream is not too dissimilar from that of the massive dwarf in m12m. Therefore, while there are high residuals as a result of the different Mg distributions, the iron distributions are similar enough such that the mass of the accretion event is recovered accurately with the DG templates (though not its quenching time).

Appendix C Modeling Only Lower-mass Events

In Section 6, we discussed how diverse evolutionary histories, particularly for the massive halo progenitors, resulted in failure modes for the CARDS method. In this appendix, we repeat our analysis, but excluding all accretion events above $M_* > 10^8 M_\odot$. Given we can often use kinematic information in conjunction with chemical information to identify stars belonging to the most massive accretion events (e.g., Sagittarius, GSE), we test here the performance of the CARDS method once the most massive progenitors are subtracted off.

Figure 20 shows the resulting inferred mass spectra, over the stellar mass range of $10^6 M_\odot < M_* < 10^8 M_\odot$. With the exception of m12r, whose results are unchanged (there is no accretion event with $M_* > 10^8 M_\odot$ included in m12r in the P21 catalog), we see improvements in the residuals for the inferred mass spectra for all halos. The results of the inferred mass accreted as a function of quenching time are shown in Figure 21.

Appendix D The In Situ Halo

While we have focused our analysis on the accreted populations in the Latte halos, there is also a contribution to

Mass Spectrum of Accreted Dwarfs

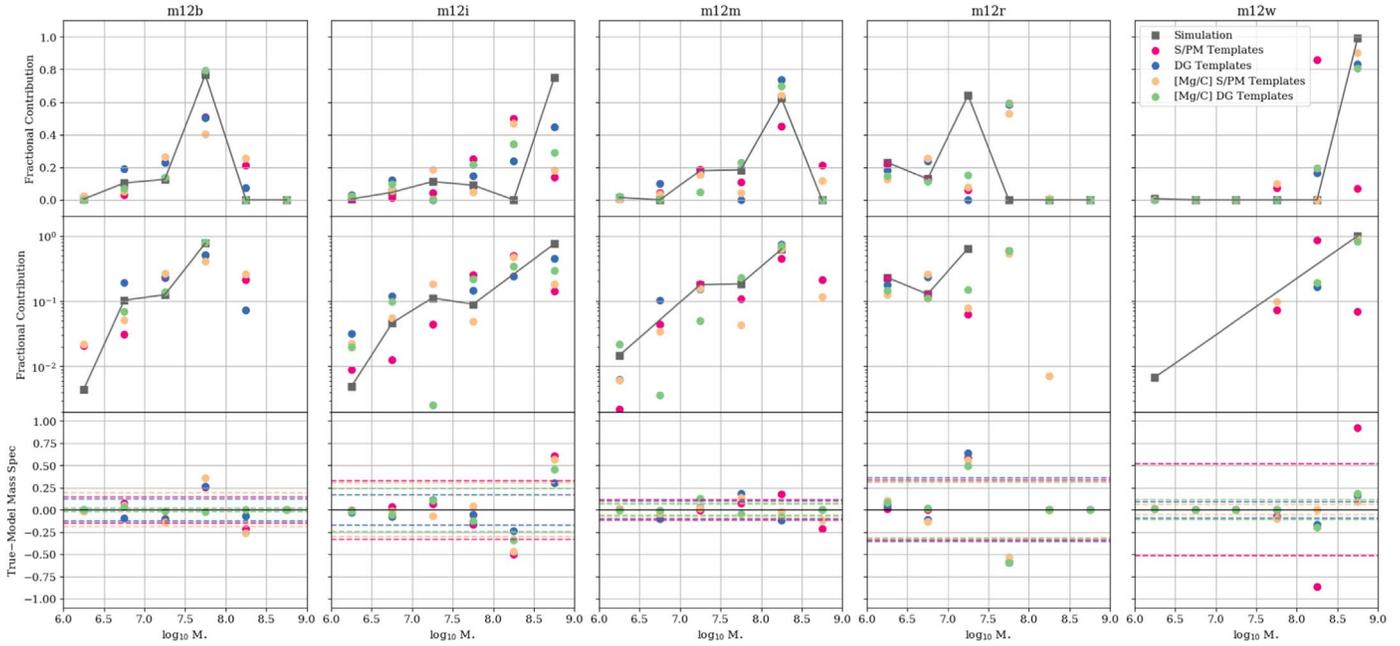


Figure 18. Same as Figure 9, but for the remaining Latte stellar halos not discussed in detail in Section 5. As in previous figures, gray squares mark the true values from the simulations, while pink, blue, peach, and green circles show the results from the S/PM templates, the DG templates, the [Mg/C] S/PM templates, and the [Mg/C] DG templates, respectively.

Mass Fraction vs Quenching Time

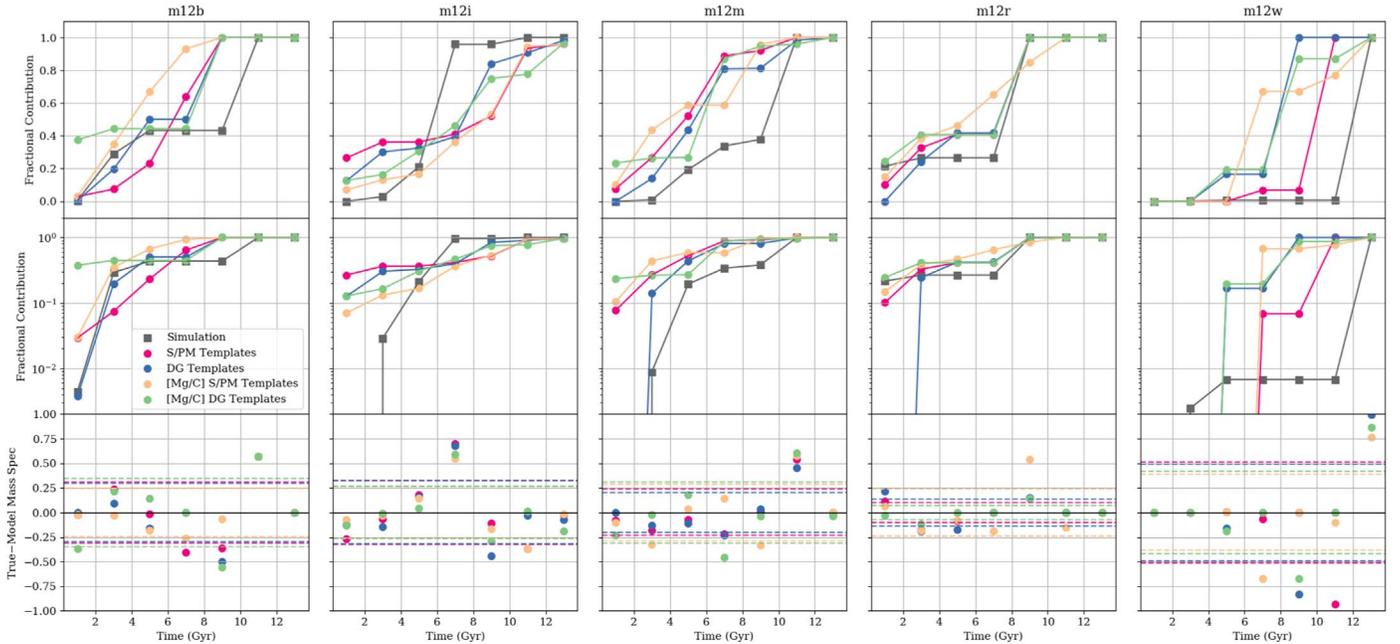


Figure 19. Same as Figure 10, but for the remaining Latte stellar halos not discussed in detail in Section 5. As in previous figures, gray squares mark the true values from the simulations, while pink, blue, peach, and green circles show the results from the S/PM templates, the DG templates, the [Mg/C] S/PM templates, and the [Mg/C] DG templates, respectively.

the stellar halos of these galaxies from the stars that formed in situ in the MW-like host disk, which are then heated onto halo-like orbits. In this appendix, we discuss how in situ halo stars could be incorporated into the proposed analysis method.

In situ halo stars are known to contribute to the stellar halo in the Latte suite of simulations (Yu et al. 2020; Ostdiek et al. 2020; Santistevan et al. 2021; Kado-Fong et al. 2022), as well as other simulation suites (e.g., Zolotov et al. 2009, 2010). In

the observations, halo stars are often interpreted as having potentially in situ origins if they have either (1) high- α abundances and relatively high metallicities, consistent with the thick disk, but eccentric orbits (e.g., Bonaca et al. 2017; Belokurov et al. 2020; Naidu et al. 2020); or (2) high- α , metal-poor stars with more disk-like kinematics (i.e., the metal-poor, dynamically hot tail of the disk; e.g., Norris et al. 1985; Chiba & Beers 2000; Beers et al. 2002; Sestito et al. 2019; Di Matteo

Mass Spectrum of Accreted Dwarfs

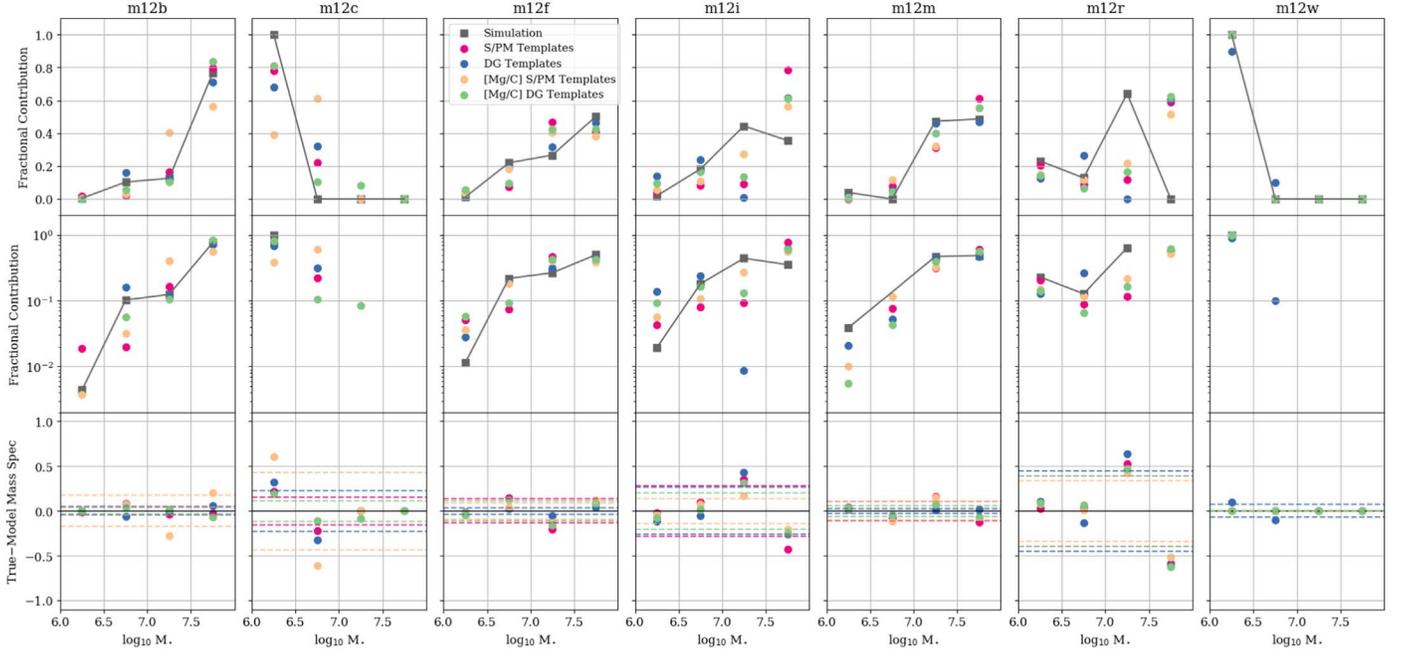


Figure 20. Same as Figure 18, but excluding accretion events above $M_* > 10^8 M_\odot$. As in previous figures, gray squares mark the true values from the simulations, while pink, blue, peach, and green circles show the results from the S/PM templates, the DG templates, the [Mg/C] S/PM templates, and the [Mg/C] DG templates, respectively.

Mass Fraction vs Quenching Time

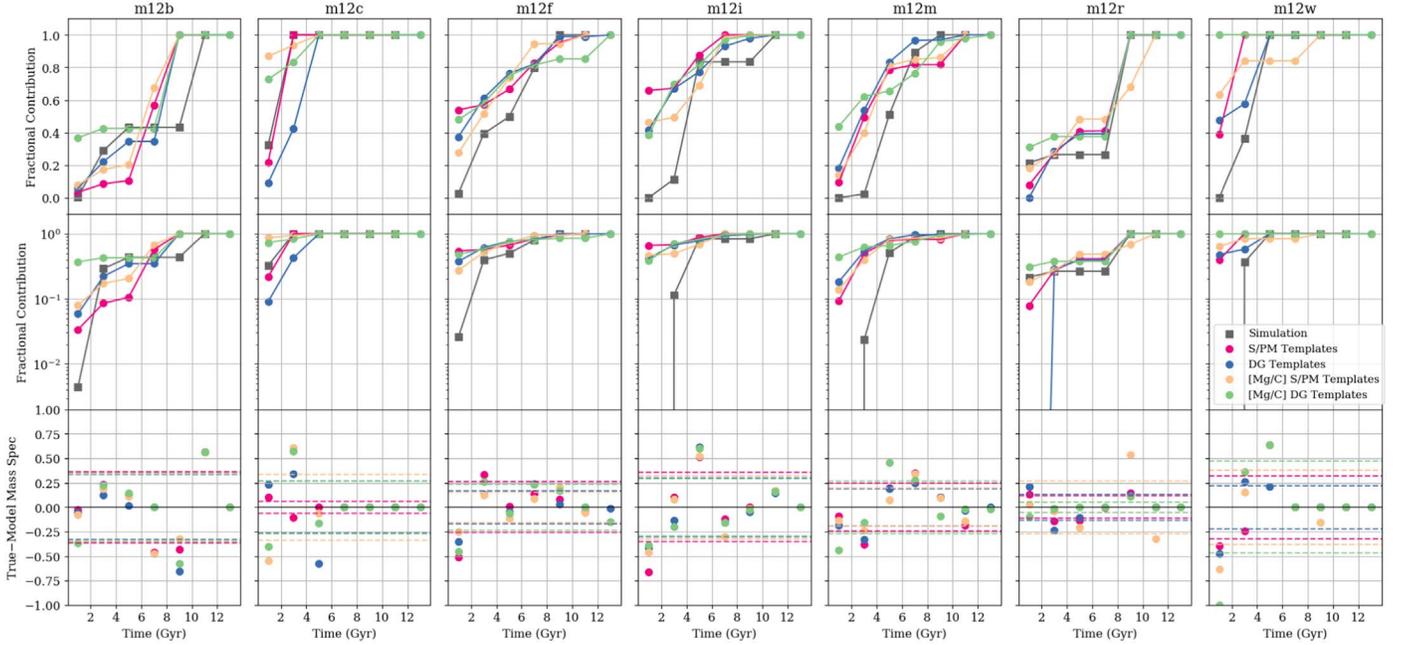


Figure 21. Same as Figure 19, but excluding accretion events above $M_* > 10^8 M_\odot$. As in previous figures, gray squares mark the true values from the simulations, while pink, blue, peach, and green circles show the results from the S/PM templates, the DG templates, the [Mg/C] S/PM templates, and the [Mg/C] DG templates, respectively.

et al. 2020; Venn et al. 2020; Cordoni et al. 2021). The fact that metal-poor, high- α stars form in galaxies at all masses poses a challenge for definitively identifying the origins of individual stars. However, by making use of the CARDs framework, we can model the contribution to the low-metallicity tail of the MW from galaxies of all masses, and constrain the origin of metal-poor stars statistically.

To select in situ halo stars, we take all star particles that have a formation distance of less than 40 kpc but are now at distances greater than 40 kpc from the Galactic center. We also exclude the particles that are within 10 kpc of the Galactic midplane as well as the star particles that are classified as belonging to an accretion event from the P21 catalog or the D. Horta et al. (2022, in preparation) catalog (this catalog contains the star particles from the most massive accretion events that

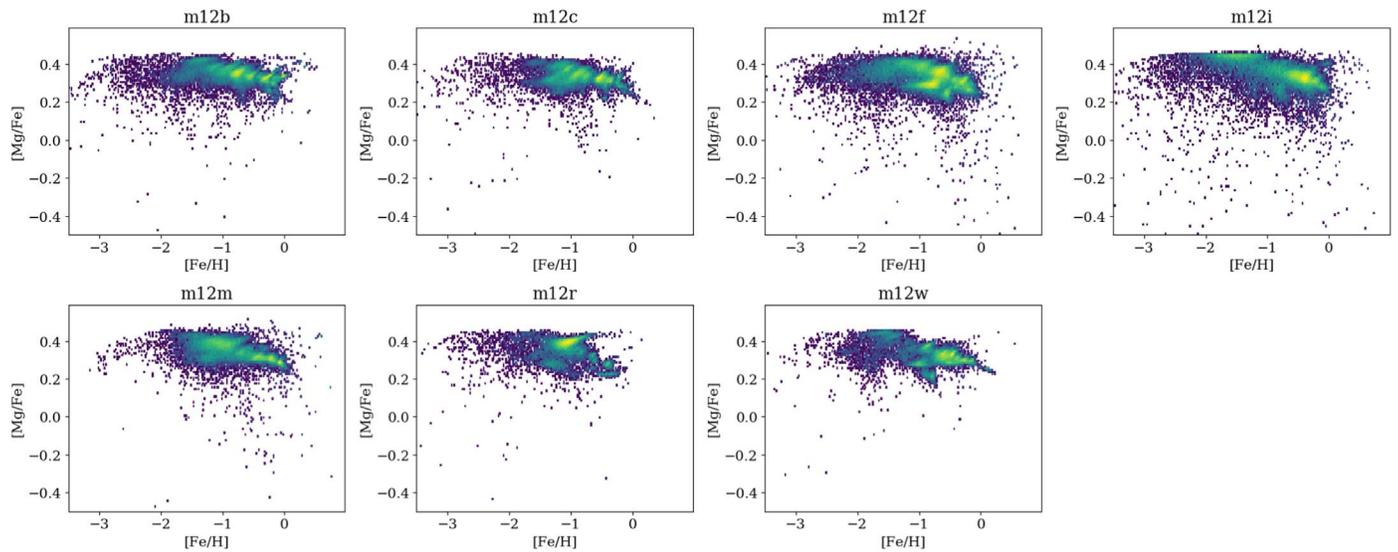


Figure 22. CARDs for the in situ halo from the seven Latte galaxies defined in this work. Relative to the accreted populations, the in situ halo stars are, on average, enhanced in magnesium and more metal-rich compared to the accreted populations. While the in situ and accreted halo star CARDs do overlap, this can be accounted for in the CARDs modeling framework by including a template for the in situ stars in the analysis.

are excluded from the P21 catalog). While this selection does not encompass all star particles that could be classified as in situ halo (in particular those very close to the disk), it does safely exclude the star particles that should be classified as belonging to the disk.

Figure 22 shows the CARDs for the in situ stellar halos for each Latte simulation discussed in this work. While there is some diversity in the in situ halo CARDs across the simulations (related to the diversity of formation histories of their disks), overall, the CARDs for the in situ stellar halos are more metal-rich and enhanced in magnesium compared to the accreted halo CARDs (see also Figure 3). While the distributions for the accreted and in situ halos overlap, the contribution to the in situ halo could be modeled by incorporating a template for in situ halo stars. While there are in situ metal-poor high- α stars, they make up a small fraction of the overall in situ halos in these simulations (of our in situ sample here, only 2% have $[\text{Fe}/\text{H}] < -2$). This is consistent with the recent finding from Santistevan et al. (2021), who studied the origin of metal-poor prograde stars in the Latte suite (analogous to the observed populations discussed in, e.g., Sestito et al. 2019; Venn et al. 2020; Di Matteo et al. 2020), and found their origin to be primarily from large, gas-rich mergers at early times (with a subdominant contribution from in situ stars). We leave a more comprehensive discussion of the properties of the in situ halo stars in these simulations, including how these properties depend on the definition of in situ, to future work.

ORCID iDs

Emily C. Cunningham <https://orcid.org/0000-0002-6993-0826>

Robyn E. Sanderson <https://orcid.org/0000-0003-3939-3297>

Kathryn V. Johnston <https://orcid.org/0000-0001-6244-6727>

Nondh Panithanpaisal <https://orcid.org/0000-0001-5214-8822>

Melissa K. Ness <https://orcid.org/0000-0001-5082-6693>

Andrew Wetzel <https://orcid.org/0000-0003-0603-8942>

Sarah R. Loebman <https://orcid.org/0000-0003-3217-5967>
 Ivanna Escala <https://orcid.org/0000-0002-9933-9551>
 Claude-André Faucher-Giguère <https://orcid.org/0000-0002-4900-6628>

References

- Allende Prieto, C., Cooper, A. P., Dey, A., et al. 2020, *RNAAS*, 4, 188
 Beers, T. C., Drilling, J. S., Rossi, S., et al. 2002, *AJ*, 124, 931
 Bellardini, M. A., Wetzel, A., Loebman, S. R., et al. 2021, *MNRAS*, 505, 4586
 Belokurov, V., Erkal, D., Evans, N. W., Koposov, S. E., & Deason, A. J. 2018, *MNRAS*, 478, 611
 Belokurov, V., Sanders, J. L., Fattahi, A., et al. 2020, *MNRAS*, 494, 3880
 Bonaca, A., Conroy, C., Wetzel, A., Hopkins, P. F., & Kereš, D. 2017, *ApJ*, 845, 101
 Bonaca, A., Conroy, C., Cargile, P. A., et al. 2020, *ApJL*, 897, L18
 Bovill, M. S., & Ricotti, M. 2011, *ApJ*, 741, 17
 Boylan-Kolchin, M., Weisz, D. R., Johnson, B. D., et al. 2015, *MNRAS*, 453, 1503
 Brook, C. B., Kawata, D., Gibson, B. K., & Flynn, C. 2003, *ApJL*, 585, L125
 Brown, T. M., Tumlinson, J., Geha, M., et al. 2012, *ApJL*, 753, L21
 Buder, S., Sharma, S., Kos, J., et al. 2021, *MNRAS*, 506, 150
 Bullock, J. S., & Johnston, K. V. 2005, *ApJ*, 635, 931
 Chiba, M., & Beers, T. C. 2000, *AJ*, 119, 2843
 Conroy, C., Naidu, R. P., Zaritsky, D., et al. 2019a, *ApJ*, 887, 237
 Conroy, C., Bonaca, A., Cargile, P., et al. 2019b, *ApJ*, 883, 107
 Cordoni, G., Da Costa, G. S., Yong, D., et al. 2021, *MNRAS*, 503, 2539
 Corlies, L., Johnston, K. V., Tumlinson, J., & Bryan, G. 2013, *ApJ*, 773, 105
 Cui, X.-Q., Zhao, Y.-H., Chu, Y.-Q., et al. 2012, *RAA*, 12, 1197
 Dalton, G., Trager, S., Abrams, D. C., et al. 2014, *Proc. SPIE*, 9147, 91470L
 de Jong, R. S., Agertz, O., Berbel, A. A., et al. 2019, *Msngr*, 175, 3
 De Silva, G. M., Freeman, K. C., Bland-Hawthorn, J., et al. 2015, *MNRAS*, 449, 2604
 Deason, A. J., Mao, Y.-Y., & Wechsler, R. H. 2016, *ApJ*, 821, 5
 Di Cintio, A., Mostoghiu, R., Knebe, A., & Navarro, J. 2021, *MNRAS*, 506, 531
 Di Matteo, P., Haywood, M., Lehnert, M. D., et al. 2019, *A&A*, 632, A4
 Di Matteo, P., Spite, M., Haywood, M., et al. 2020, *A&A*, 636, A115
 Escala, I., Gilbert, K. M., Kirby, E. N., et al. 2020a, *ApJ*, 889, 177
 Escala, I., Kirby, E. N., Gilbert, K. M., et al. 2020b, *ApJ*, 902, 51
 Escala, I., Gilbert, K. M., Wojno, J., Kirby, E. N., & Guhathakurta, P. 2021, *AJ*, 162, 45
 Escala, I., Kirby, E. N., Gilbert, K. M., Cunningham, E. C., & Wojno, J. 2019, *ApJ*, 878, 42
 Escala, I., Wetzel, A., Kirby, E. N., et al. 2018, *MNRAS*, 474, 2194
 Faucher-Giguère, C.-A. 2020, *MNRAS*, 493, 1614

- Faucher-Giguère, C.-A., Lidz, A., Zaldarriaga, M., & Hernquist, L. 2009, *ApJ*, **703**, 1416
- Font, A. S., Johnston, K. V., Bullock, J. S., & Robertson, B. E. 2006, *ApJ*, **646**, 886
- Freeman, K., & Bland-Hawthorn, J. 2002, *ARA&A*, **40**, 487
- Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, *A&A*, **595**, A1
- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *A&A*, **616**, A1
- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2021, *A&A*, **649**, A1
- Garrison-Kimmel, S., Hopkins, P. F., Wetzel, A., et al. 2019a, *MNRAS*, **487**, 1380
- Garrison-Kimmel, S., Wetzel, A., Hopkins, P. F., et al. 2019b, *MNRAS*, **489**, 4574
- Gandhi, P. J., Wetzel, A., Hopkins, P. F., et al. 2022, arXiv:2202.10477
- Gilbert, K. M., Kirby, E. N., Escala, I., et al. 2019, *ApJ*, **883**, 128
- Gilbert, K. M., Wojno, J., Kirby, E. N., et al. 2020, *AJ*, **160**, 41
- Hasselquist, S., Hayes, C. R., Lian, J., et al. 2021, *ApJ*, **923**, 172
- Haywood, M., Di Matteo, P., Lehnert, M. D., et al. 2018, *ApJ*, **863**, 113
- Helmi, A., Babusiaux, C., Koppelman, H. H., et al. 2018, *Natur*, **563**, 85
- Helmi, A., White, S. D. M., de Zeeuw, P. T., & Zhao, H. 1999, *Natur*, **402**, 53
- Hopkins, P. F. 2015, *MNRAS*, **450**, 53
- Hopkins, P. F. 2017, *MNRAS*, **466**, 3387
- Hopkins, P. F., Wetzel, A., Kereš, D., et al. 2018, *MNRAS*, **480**, 800
- Hopkins, P. F., Wetzel, A., Wheeler, C., et al. 2022, arXiv:2203.00040
- Iwamoto, K., Brachwitz, F., Nomoto, K., et al. 1999, *ApJS*, **125**, 439
- Jean-Baptiste, I., Di Matteo, P., Haywood, M., et al. 2017, *A&A*, **604**, A106
- Johnston, K. V., Bullock, J. S., Sharma, S., et al. 2008, *ApJ*, **689**, 936
- Kado-Fong, E., Sanderson, R. E., Greene, J. E., et al. 2022, *ApJ*, **931**, 152
- Kirby, E. N., Gilbert, K. M., Escala, I., et al. 2020, *AJ*, **159**, 46
- Kirby, E. N., Guhathakurta, P., Bolte, M., Sneden, C., & Geha, M. C. 2009, *ApJ*, **705**, 328
- Kirby, E. N., Rizzi, L., Held, E. V., et al. 2017, *ApJ*, **834**, 9
- Kirby, E. N., Guhathakurta, P., Simon, J. D., et al. 2010, *ApJS*, **191**, 352
- Kollmeier, J. A., Zasowski, G., Rix, H.-W., et al. 2017, arXiv:1711.03234
- Koppelman, H., Helmi, A., & Veljanoski, J. 2018, *ApJL*, **860**, L11
- Koppelman, H. H., Helmi, A., Massari, D., Price-Whelan, A. M., & Starkenburg, T. K. 2019a, *A&A*, **631**, L9
- Koppelman, H. H., Helmi, A., Massari, D., Roelenga, S., & Bastian, U. 2019b, *A&A*, **625**, A5
- Kroupa, P. 2002, *Sci*, **295**, 82
- Krumholz, M. R., & Gnedin, N. Y. 2011, *ApJ*, **729**, 36
- Kuhlen, M., & Faucher-Giguère, C.-A. 2012, *MNRAS*, **423**, 862
- Lee, D. M., Johnston, K. V., Sen, B., & Jessop, W. 2015, *ApJ*, **802**, 48
- Leitherer, C., Schaerer, D., Goldader, J. D., et al. 1999, *ApJS*, **123**, 3
- Ma, X., Hopkins, P. F., Wetzel, A. R., et al. 2017, *MNRAS*, **467**, 2430
- Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, *AJ*, **154**, 94
- Mannucci, F., Della Valle, M., & Panagia, N. 2006, *MNRAS*, **370**, 773
- Maoz, D., & Graur, O. 2017, *ApJ*, **848**, 25
- Myeong, G. C., Vasiliev, E., Iorio, G., Evans, N. W., & Belokurov, V. 2019, *MNRAS*, **488**, 1235
- Naidu, R. P., Conroy, C., Bonaca, A., et al. 2020, *ApJ*, **901**, 48
- Naidu, R. P., Conroy, C., Bonaca, A., et al. 2022, arXiv:2204.09057
- Nidever, D. L., Hasselquist, S., Hayes, C. R., et al. 2020, *ApJ*, **895**, 88
- Nissen, P. E., & Schuster, W. J. 2010, *A&A*, **511**, L10
- Nomoto, K., Tominaga, N., Umeda, H., Kobayashi, C., & Maeda, K. 2006, *NuPhA*, **777**, 424
- Norris, J., Bessell, M. S., & Pickles, A. J. 1985, *ApJS*, **58**, 463
- Ostdiek, B., Necib, L., Cohen, T., et al. 2020, *A&A*, **636**, A75
- Panithanpaisal, N., Sanderson, R. E., Wetzel, A., et al. 2021, *ApJ*, **920**, 10
- Ricotti, M., & Gnedin, N. Y. 2005, *ApJ*, **629**, 259
- Robertson, B., Bullock, J. S., Font, A. S., Johnston, K. V., & Hernquist, L. 2005, *ApJ*, **632**, 872
- Robertson, B. E., Ellis, R. S., Furlanetto, S. R., & Dunlop, J. S. 2015, *ApJL*, **802**, L19
- Samuel, J., Wetzel, A., Santistevan, I., et al. 2022, *MNRAS*, **514**, 5276
- Samuel, J., Wetzel, A., Tollerud, E., et al. 2020, *MNRAS*, **491**, 1471
- Sanderson, R. E., Garrison-Kimmel, S., Wetzel, A., et al. 2018, *ApJ*, **869**, 12
- Sanderson, R. E., Wetzel, A., Loebman, S., et al. 2020, *ApJS*, **246**, 6
- Santistevan, I. B., Wetzel, A., El-Badry, K., et al. 2020, *MNRAS*, **497**, 747
- Santistevan, I. B., Wetzel, A., Sanderson, R. E., et al. 2021, *MNRAS*, **505**, 921
- Sestito, F., Longeard, N., Martin, N. F., et al. 2019, *MNRAS*, **484**, 2166
- Simpson, C. M., Grand, R. J. J., Gómez, F. A., et al. 2018, *MNRAS*, **478**, 548
- Steinmetz, M., Zwitter, T., Siebert, A., et al. 2006, *AJ*, **132**, 1645
- Su, K.-Y., Hopkins, P. F., Hayward, C. C., et al. 2017, *MNRAS*, **471**, 144
- Takada, M., Ellis, R. S., Chiba, M., et al. 2014, *PASJ*, **66**, R1
- Vargas, L. C., Geha, M. C., & Tollerud, E. J. 2014a, *ApJ*, **790**, 73
- Vargas, L. C., Gilbert, K. M., Geha, M., et al. 2014b, *ApJL*, **797**, L2
- Venn, K. A., Irwin, M., Shetrone, M. D., et al. 2004, *AJ*, **128**, 1177
- Venn, K. A., KIELTY, C. L., Sestito, F., et al. 2020, *MNRAS*, **492**, 3241
- Weisz, D. R., & Boylan-Kolchin, M. 2017, *MNRAS*, **469**, L83
- Wetzel, A. R., Hopkins, P. F., Kim, J.-h., et al. 2016, *ApJL*, **827**, L23
- Wheeler, C., Hopkins, P. F., Pace, A. B., et al. 2019, *MNRAS*, **490**, 4447
- Wojno, J., Gilbert, K. M., Kirby, E. N., et al. 2020, *ApJ*, **895**, 78
- Yanny, B., Rockosi, C., Newberg, H. J., et al. 2009, *AJ*, **137**, 4377
- Yu, S., Bullock, J. S., Wetzel, A., et al. 2020, *MNRAS*, **494**, 1539
- Zolotov, A., Willman, B., Brooks, A. M., et al. 2009, *ApJ*, **702**, 1058
- Zolotov, A., Willman, B., Brooks, A. M., et al. 2010, *ApJ*, **721**, 738