Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **492,** 8–21 (2020)
Advance Access publication 2019 November 28

doi:10.1093/mnras/stz3331

A stochastically sampled IMF alters the stellar content of simulated dwarf galaxies

Elaad Applebaum [1]★ Alyson M. Brooks,[1] Thomas R. Quinn[2] and Charlotte R. Christensen[3]

[1]*Department of Physics and Astronomy, Rutgers, The State University of New Jersey, 136 Frelinghuysen Road, Piscataway, NJ 08854, USA*
[2]*Department of Astronomy, University of Washington, Box 351580, Seattle, WA 98195, USA*
[3]*Physics Department, Grinnell College, 1116 Eighth Avenue, Crinnell, IA 50112, USA*

Accepted 2019 November 25. Received 2019 September 20; in original form 2018 October 27

ABSTRACT

Cosmological simulations are reaching the resolution necessary to study ultra-faint dwarf galaxies. Observations indicate that in small populations, the stellar initial mass function (IMF) is not fully populated; rather, stars are sampled in a way that can be approximated as coming from an underlying probability density function. To ensure the accuracy of cosmological simulations in the ultra-faint regime, we present an improved treatment of the IMF. We implement a self-consistent, stochastically populated IMF in cosmological hydrodynamic simulations. We test our method using high-resolution simulations of a Milky Way halo, run to $z = 6$, yielding a sample of nearly 100 galaxies. We also use an isolated dwarf galaxy to investigate the resulting systematic differences in galaxy properties. We find that a stochastic IMF in simulations makes feedback burstier, strengthening feedback, and quenching star formation earlier in small dwarf galaxies. For galaxies in haloes with mass $\lesssim 10^{8.5}$ $M_\odot$, a stochastic IMF typically leads to lower stellar mass compared to a continuous IMF, sometimes by more than an order of magnitude. We show that existing methods of ensuring discrete supernovae incorrectly determine the mass of the star particle and its associated feedback. This leads to overcooling of surrounding gas, with at least ∼10 per cent higher star formation and ∼30 per cent higher cold gas content. Going forwards, to accurately model dwarf galaxies and compare to observations, it will be necessary to incorporate a stochastically populated IMF that samples the full spectrum of stellar masses.

Key words: methods: numerical – supernovae: general – galaxies: dwarf – galaxies: formation – galaxies: star formation.

1 INTRODUCTION

Prior to the last decade, dwarf galaxies posed a long-standing challenge to galaxy formation models within the context of the cold dark matter (CDM) paradigm. In recent years, however, enormous progress has been made in simulating these low-mass systems, owing to increasing resolution and careful modelling of the baryonic processes involved. For example, repeated fluctuations of the gravitational potential well of a halo due to supernova-induced gas outflows have been shown to flatten the central dark matter profile (e.g. Read & Gilmore 2005; Mashchenko, Wadsley & Couchman 2008; Governato et al. 2010; Pontzen & Governato 2012; Di Cintio et al. 2014; Read, Agertz & Collins 2016; Fitts et al. 2017), alleviating tension with observations (e.g. Simon et al. 2005; de Blok et al. 2008; Kuzio de Naray, McGaugh & de Blok 2008). Supernova-driven outflows also remove low-angular momentum gas, leading to the creation of bulgeless dwarf galaxies (Governato et al. 2010; Brook et al. 2011; Teyssier et al. 2013). Similar effects together with enhanced tidal stripping within the parent halo successfully reduce the central densities of the most massive Milky Way satellites (e.g. Zolotov et al. 2012; Brooks & Zolotov 2014; Sawala et al. 2016; Tomozeiu, Mayer & Quinn 2016; Wetzel et al. 2016; Garrison-Kimmel et al. 2019), solving the well-known 'too big to fail' problem (Boylan-Kolchin, Bullock & Kaplinghat 2011, 2012). Destruction of satellite haloes from interactions with the central disc potential reduces the theoretically expected number of MW satellites (e.g. Brooks et al. 2013; Sawala et al. 2016; Wetzel et al. 2016; Garrison-Kimmel et al. 2019), solving the 'missing satellites' problem (Klypin et al. 1999; Moore et al. 1999). Proper consideration of the biases introduced when comparing

★ E-mail: applebaum@physics.rutgers.edu

observations to theoretical results also brings theory into agreement with observations (Brooks et al. 2017).

With the recent successes in modelling dwarf galaxies, various groups are pushing the limits of resolution even further into the ultra-faint dwarf (UFD) galaxy range ($M_{\rm star} \lesssim 10^5$ M$_\odot$; e.g. Munshi et al. 2013; Wheeler et al. 2015; Munshi et al. 2017, 2019).

Despite significant differences in feedback models and simulation details, many groups have succeeded in reproducing realistic dwarf galaxies. This is largely a result of galaxy self-regulation; i.e. global galaxy properties are robust to the details of star formation and feedback, since increased feedback suppresses future star formation (Saitoh et al. 2008; Shetty & Ostriker 2008; Hopkins, Quataert & Murray 2011; Christensen et al. 2014; Agertz & Kravtsov 2015; Benincasa et al. 2016; Hopkins et al. 2018). As simulations approach the UFD regime, however, the self-regulation of galaxies breaks down. Munshi et al. (2019) demonstrated this conclusively, finding that different star formation prescriptions predict different numbers of UFD galaxies. Further, the interplay of supernova feedback and reionization with the gas in these low-mass haloes leads to diverging star formation behaviour at early times. Thus, without a concerted effort to study subgrid models in the low-mass galaxy regime, the predictive power of cosmological simulations will diminish.

In this work, we investigate another physical prescription that can alter UFD simulation results: the realistic sampling of stars and its impact on subsequent feedback. Since the high-mass stars providing the bulk of stellar feedback are relatively uncommon, nuances in choice of sampling may not 'average out' in UFDs. It is therefore important to correctly determine the stellar masses within a population; this is done through the use of the stellar initial mass function (IMF), which describes the number distribution of stars as a function of their birth mass.

Estimates of the IMF in large stellar populations have shown remarkable consistency. Parametrizations generally find a steep power-law slope for more massive stars consistent with the original Salpeter (1955) estimate, with a 'knee' at $m_{\rm star} \sim$ M$_\odot$, below which the distribution involves a shallower decline (e.g. Kroupa, Tout & Gilmore 1993; Kroupa 2001; Chabrier 2003). While theoretical expectations predict systematic variation in the IMF with environment (Kroupa et al. 2013) in large resolved stellar populations there is limited evidence of deviations from the universal IMF (Bastian, Covey & Meyer 2010). While several recent observations suggest systematic variations (e.g. van Dokkum & Conroy 2010; Cappellari et al. 2012; Conroy & van Dokkum 2012; Geha et al. 2013; Kalirai et al. 2013; Gennaro et al. 2018), there is neither consensus on their significance nor a clear physical driver for their variation, with dominant candidates including metallicity (e.g. Martín-Navarro et al. 2015; Gennaro et al. 2018) and velocity dispersion (e.g. La Barbera et al. 2013; Spiniello et al. 2014; Rosani et al. 2018).

Despite the general success of the IMF formalism in describing galaxies and large stellar populations, in small populations it is clear that the current IMF formalism is insufficient. The inherently discrete nature of stars makes a continuous description unrealistic. To find a better description, a variety of observations can be used. For example, there exists an average relationship between the mass of an embedded star cluster and the mass of the most massive star residing in the cluster, such that more massive stars tend to reside in more massive clusters (Weidner, Kroupa & Bonnell 2010; Weidner, Kroupa & Pflamm-Altenburg 2013a; see also Cerviño et al. 2013 for a detailed discussion). Other observations have shown lower values of Hα- or Hβ-to-FUV luminosity ratios in galaxies with low star formation rates (SFRs), indicating a relative dearth of very high-mass stars (Lee et al. 2009, 2016; Meurer et al. 2009).

It follows from the above observations that there is a tendency for fewer high-mass stars to form in small, low-SFR populations. Observations of SFR indicators can be explained via bursty star formation histories (Weisz et al. 2012; Guo et al. 2016; Emami et al. 2019), but IMF sampling effects may also contribute (e.g. Pflamm-Altenburg, Weidner & Kroupa 2009; Fumagalli, da Silva & Krumholz 2011; Eldridge 2012).

There are two broad theories of how the IMF should be sampled to explain these observations. The first is the integrated galactic IMF (IGIMF), presented in Kroupa & Weidner (2003), Weidner & Kroupa (2006), Weidner et al. (2010), Weidner et al. (2013b) and Yan, Jerabkova & Kroupa (2017), which assumes a deterministic relationship between the mass of a star cluster and the stellar mass distribution within it. The other predominant explanation for the observations is that the IMF is sampled randomly (e.g. Elmegreen 2006; Corbelli et al. 2009; Calzetti et al. 2010; Fumagalli et al. 2011; Andrews et al. 2013, 2014). In every star formation event, stars form in a way approximated as being drawn from an underlying probability density function – the universal IMF. Since the probability of forming high-mass stars is rare, in small populations there is an *average* tendency for massive stars to form in massive clusters, mimicking the proposed IGIMF.

The only restriction in stochastically sampling the IMF is that a star cannot form with a greater mass than the cluster in which it resides (i.e. stars cannot form with more mass than their available gas reservoirs). Until recently, star particles in cosmological simulations were large enough to ignore all the nuances of IMF sampling, and stellar feedback was calculated by treating star particles as a simple stellar population with a uniform, continuous IMF. With sufficiently massive star particles, this was a reasonably accurate approximation. At lower particle masses, however, the above model has proven increasingly unrealistic. Further, star particles are small enough such that not only is a uniform IMF no longer consistent with observations, but a naive calculation of supernova counts per time-step yields fractions of supernovae exploding (Revaz et al. 2016).

In small galaxies, the credibility of simulated results depends upon proper treatment of the IMF. Carigi & Hernandez (2008) demonstrated that a stochastically sampled IMF does not converge to the underlying continuous IMF until $M_{\rm star} \sim 10^5$ M$_\odot$. Cosmological galaxy simulations are now pushing to high enough resolution to study stellar populations in the ultra-faint regime; at these scales, not only individual star particles but also entire galaxies will have incompletely sampled IMFs.

Limited work has been done in investigating the impacts of IMF sampling within cosmological simulations. In post-processing, Sparre et al. (2017) used the SLUG code (da Silva, Fumagalli & Krumholz 2012, 2014; Krumholz et al. 2015) to show how a stochastic IMF increases the scatter of dwarf galaxies' Hα-to-FUV ratios, but the simulation itself assumed a fully populated IMF. Revaz et al. (2016) studied the effects on IMF sampling on stellar chemical abundances in isolated dwarf galaxies and found that a continuous IMF becomes unrealistic at star particle masses below $\sim 10^5$ M$_\odot$. They further found that regardless of sampling method, the combined IMF of multiple-star particles together will be undersampled below particle masses of $\sim 10^3$ M$_\odot$. Bracketing the case of a stochastic IMF, Hensler, Steyrleithner & Recchi (2017) compared a truncated and a filled IMF in simulations of dwarf galaxies, and found that truncation suppresses the self-regulation of star formation. Several cosmological simulations have

Table 1. Properties of the cosmological simulations used in this work, including particle masses, the expectation value of core collapse supernovae per star particle, and the softening length.

Simulation	$M_{\rm dark}$ ($M_\odot$)	$M_{\rm gas}$ ($M_\odot$)	$M_{\rm star}$ ($M_\odot$)	$\langle M_{\rm SNII}\rangle$	Softening Length (pc)
Elena (Milky Way)	1.8×10^4	3.3×10^3	990	9.9	87
Isolated 10^9 $M_\odot$	1.0×10^4	1.4×10^3	420	4.2	21

discussed or incorporated methods of discretizing stellar feedback from Type-II supernovae (e.g. Stinson et al. 2010; Agertz et al. 2013; Hopkins et al. 2014, 2018; Rosdahl et al. 2018). The most common method has been to decide whether or not a star explodes by drawing from a binomial or Poisson distribution derived from an average measure of supernova rates. The drawbacks of this method are discussed in Section 2.3. Su et al. (2018) took first steps in investigating IMF sampling and stochastic effects more closely; they found a dramatic decrease in star formation when discretizing their supernovae compared to continuous energy injection. However, their model does not sample the full range of masses in the IMF and still calculates feedback by drawing from a Poisson distribution.

More work has been done in high-resolution simulations that do not model cosmological contexts. Grudić & Hopkins (2019) used the same methodology as Su et al. (2018) on molecular cloud scales and found similar results. Sormani et al. (2017) introduced a method based on discretizing stars into mass bins that are then Poisson sampled. Other groups (e.g. Gatto et al. 2017; Geen et al. 2018) separate their IMF into high- and low-mass regimes, and stochastically sample only within the high-mass regime. We note that some recent simulations of very small, isolated dwarf galaxies now track the evolution of individually sampled stars (Hu et al. 2017; Emerick, Bryan & Mac Low 2019) and recent work with isolated Milky Way–mass galaxies includes a stochastically populated IMF within star particles (Fujimoto, Krumholz & Tachibana 2018), but this has never previously been attempted in cosmological simulations.

In this paper, we present a new prescription for star formation that stochastically samples the full spectrum of masses in the IMF and individually tracks the evolution of high-mass stars within them. This methodology ensures conservation of mass and self-consistency of radiative and supernova feedback. We discuss the simulations and the sampling method in Section 2, and compare to existing discretization methods. In Section 3, we demonstrate the effects of improved IMF sampling. In Section 4, we discuss implications of this sampling method for future observational predictions. We conclude in Section 5.

2 METHODS

We implement a new stochastic IMF treatment for star particles in our simulations. The updated sampling changes the stellar mass distribution from a smooth IMF to a set of discrete stars; these stars are then used to calculate supernova explosions, metal production, and high-energy radiation output. We emphasize that the actual feedback mechanisms remain unaltered; what changes in the new recipe is the timing and quantity of feedback each star particle produces.

2.1 Simulations

To test the new IMF prescription, we ran cosmological zoom-in simulations of a Milky Way–mass galaxy with and without the stochastic IMF. These were run until immediately after reionization, by which time UFD galaxies – those with $M_{\rm star} \lesssim 10^5$ $M_\odot$, and therefore unconverged IMFs – have formed most or all of their stars (Brown et al. 2014; Weisz et al. 2014). Since the same dark matter particles form the same haloes in simulations with and without a stochastic IMF, we were able to match galaxies that formed in corresponding haloes between runs. To ensure our results are independent of any intrinsic scatter in the cosmological runs (and allow a finer time resolution of outputs), we also used an isolated dwarf galaxy. For both treatments of the IMF, we ran the same 10^9 $M_\odot$ halo 50 times, and compared the ensemble behaviour.

The relevant properties of each simulation can be found in Table 1. The cosmological simulations used in this work were selected from a uniform resolution, dark matter-only simulation of 50 Mpc per side, run with Planck cosmological parameters (Planck Collaboration XIII 2016). A region around the selected halo in these simulations was rerun at higher resolution using the 'zoom-in' technique (Katz & White 1993). The halo used in this work was selected to resemble the Milky Way at $z = 0$, and is one of the haloes of the DC Justice League suite of simulations (Bellovary et al. 2019), nicknamed 'Elena.' We use a gravitational force softening length of 87 pc and equivalent resolution to a 6144^3 grid. The dark matter particle masses are 1.8×10^4 $M_\odot$, the gas particles begin with 3.3×10^3 $M_\odot$, and the star particles form with a mass of 994 $M_\odot$. Versions of the DC Justice League suite at this resolution are being run to the present day, and will constitute the highest resolution Milky Way simulations to date.

The initial conditions for the 10^9 $M_\odot$ isolated dwarf galaxy have been described previously (Kaufmann et al. 2007; Stinson et al. 2007; Christensen et al. 2010). In short, the initial conditions consist of an equilibrium halo with a Navarro–Frenk–White concentration of $c = 8$. Dark matter velocities were determined via the Eddington inversion method of Kazantzidis, Magorrian & Moore (2004). Gas particles were assigned temperatures to ensure hydrostatic equilibrium before cooling, and were given a uniform rotational velocity corresponding to a spin parameter $\sim$0.04. Dark matter particles within the virial radius have a mass of 1.0×10^4 $M_\odot$, gas particles have a mass of 1.4×10^3 $M_\odot$, and star particles from with 425 $M_\odot$. The force softening length is 0.1 per cent the virial radius, or 21 pc.

The stochastically populated IMF is incorporated into the *N*-body + smoothed particle hydrodynamics (SPH) code CHANGA (Menon et al. 2015), a fully cosmological simulation code that includes physics from the GASOLINE2 code (Wadsley, Keller & Quinn 2017), but utilizes the CHARM++ runtime system for dynamic load balancing to efficiently scale up to thousands of cores. All simulations discussed in this work smooth over 32 nearest-neighbour gas particles.

As discussed above, feedback from high-mass stars is crucial for modelling realistic galaxies. In this work, we use the 'blastwave' supernova feedback mechanism described in detail in Stinson et al. (2006), whereby mass, energy, and chemically enriched material are deposited into neighbouring gas when a massive star dies as a

Type-II supernova. With the existing continuous IMF, the minimum and maximum stars that explode in a given time-step are calculated based on the stellar lifetime parametrizations of Raiteri, Villata & Navarro (1996). The number and mass in supernovae are then determined by integrating along the IMF between these stellar masses. We note that the default time-step for calculating feedback and star formation in CHANGA is 1 Myr. We deposit 1.5×10^{51} erg per supernova[1] among gas particles within the blast radius calculated using McKee & Ostriker (1977), then gas cooling is shut off until the end of the blastwave's snowplow phase. We assume stars between 8 and 40 $M_\odot$ explode as supernovae, while more massive stars collapse into black holes. Future work will incorporate a stochastic IMF using the 'superbubble' feedback mechanism (Keller et al. 2014). Results from CHANGA and GASOLINE with blastwave feedback have been used to reproduce a variety of observations, including the stellar mass–halo mass relation (Munshi et al. 2013, 2017), the mass–metallicity relation (Brooks et al. 2007), the baryonic Tully–Fisher relation (Christensen et al. 2016; Brooks et al. 2017), the abundance of Damped Lyman α systems (Pontzen et al. 2008), and the properties of dwarf Spheroidal Milky Way satellites (Brooks & Zolotov 2014). These models also produced the first simulated cored dark matter density profiles and bulgeless disc galaxies (Governato et al. 2010; Brook et al. 2011; Governato et al. 2012).

In addition to feedback from core collapse supernovae, CHANGA incorporates metal cooling and diffusion in the ISM (Shen, Wadsley & Stinson 2010), a time-dependent UV background (Haardt & Madau 2012), Type-Ia supernovae, mass-loss in stellar winds, and metal enrichment (Stinson et al. 2006), and supermassive black hole formation, growth, and feedback (Tremmel et al. 2015, 2017).

CHANGA includes a star formation recipe based on the local abundance of molecular hydrogen (H_2; Christensen et al. 2012). This scheme includes calculations for the formation of H_2, shielding from dissociative Lyman Werner (LW) radiation, and production of LW photons from high-mass stars. LW photon production from star particles is calculated using STARBURST99 (Leitherer et al. 1999); for star particles represented by a uniform IMF, as done until now, the calculation is based on a single-age, simple stellar population.

Haloes in the cosmological simulations are identified using AMIGA'S HALO FINDER (Gill, Knebe & Gibson 2004; Knollmann & Knebe 2009). Haloes are defined as the radius within which the density reaches a redshift-dependent overdensity criterion using the approximation of Bryan & Norman (1998).[2] Virial radius and halo mass are defined according to this overdensity.

2.2 Stochastically populated IMF

Ideally, no approximations would be necessary in cosmological simulations and every individual star would be tracked within a given stellar population. However, tracking hundreds of individual stars within each of millions of star particles would be computationally prohibitive. Any approximation, then, should be guided by two considerations: (1) we wish to preserve the highest accuracy for the individual stars whose feedback has the greatest impact on galaxy evolution, and (2) we wish to preserve the highest accuracy for individual stars that are rarest, and therefore are most altered by the approximations of a continuous, universal IMF. Therefore, we strive to maintain highest accuracy in the high-mass component; specifically, stars that eventually release energy as Type-II supernovae.

Due to the nature of random sampling, the mass of a simulated stellar population cannot be predetermined if using a stochastically populated IMF; rather, only an estimate can be known a priori (see e.g. Cerviño et al. 2013). Since in our simulations we form star particles of a given mass, we must use an algorithm to stochastically populate stars in a way that reaches the desired mass of our population.

When stochastically sampling from the IMF, the IMF is treated as a probability density function, so that its area is normalized to one but its form is otherwise unchanged. The algorithm we adopt is the stop-nearest method (e.g. da Silva et al. 2012; Eldridge 2012). With this method, stars are drawn from the IMF until the desired mass is first exceeded. Then, the last star drawn is either kept or discarded, depending on whether the total mass is closer with or without its inclusion.

The methodology, then, for our stochastically populated IMF is as follows:

(i) If we determine a star particle forms, the formation mass is the target mass.

(ii) Following the stop-nearest method, we draw stars from the IMF until we pass our target mass threshold, then either keep or discard the last star based on which brings us closer to the target mass.

(iii) We discard all stars below a cut-off mass[3] and reapproximate the low-mass stars as a continuous IMF, normalized such that the total mass of the star particle is the target mass.

(iv) For feedback dependent on stellar masses above the cut-off, we use the individual high-mass stars to calculate the timing and quantity. For a cut-off of 8 $M_\odot$, all energy and metals from Type-II supernovae and LW photon production are calculated discretely; Type-Ia supernovae and stellar winds are calculated as in the case of a continuous IMF, since they come from the stellar mass range approximated as a continuous distribution.

This methodology shares similarities with that of, e.g. Gatto et al. (2017); however, we sample stars over the entire mass range of the IMF rather than just among the high-mass end to avoid imposing a restriction on either the mass in high-mass stars or the number. We thus allow for substantially more variation in the mass and number of high-mass stars per star particle. The numerical implications of this method are discussed briefly in Appendix B.

An example of the methodology can be seen in Fig. 1, where we show three different realizations of a stochastically populated star particle with a total mass of 500 $M_\odot$. In a population of this mass, we expect approximately 5.5 stars with 8 $M_\odot < m_{\rm star} <$100 $M_\odot$, and 105 $M_\odot$ in mass for the same range. For each realization, we show the change in number and mass of high-mass stars from these expected values, as well as the per cent changes these correspond to. Since only about 20 per cent of the mass in these star particles is in the high-mass range, even relatively large deviations in the mass content of the discrete portion result in apparently small changes in the normalization of the low-mass portion. To see how the sampling affects the resulting stellar feedback, Fig. 2 shows the supernova rate in 1 Myr time-steps for the same three star particles.

[1] Rates of supernova energy deposition were determined using the parameter optimization technique described in Tremmel et al. (2017) and Anderson et al. (2017).

[2] At $z = 0$, the overdensity compared to the critical density is $\rho/\rho_c \approx 100$.

[3] Our fiducial cut-off mass is 8 $M_\odot$. Appendix A discusses what would happen if we raised the value of the cut-off.

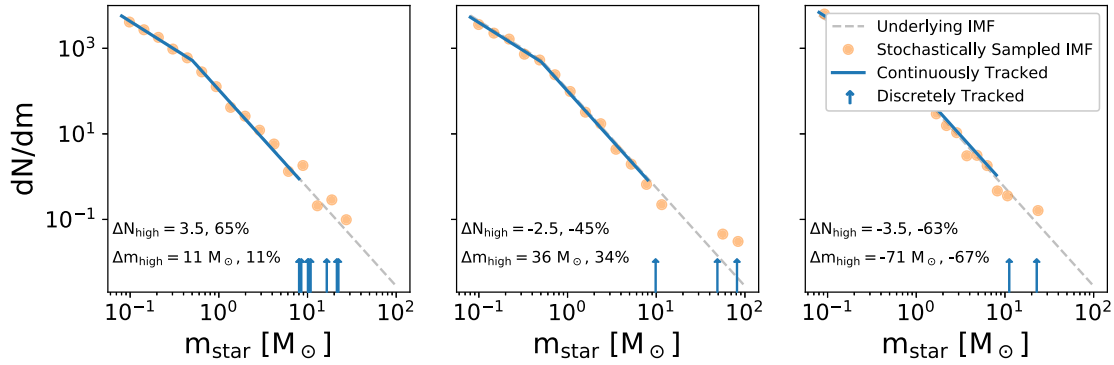

Figure 1. Three realizations of a stochastically populated 500 $M_\odot$ star particle, following the methodology of Section 2.2. The grey dashed line shows the universal underlying IMF of Kroupa (2001), scaled to the population mass. The orange circles represent the IMF recovered from stochastically drawing from the entire IMF to which we then apply a continuous/discrete cut-off. The solid blue line shows the portion of the IMF that is approximated as continuous after sampling. The blue arrows represent individual, discrete high-mass stars tracked within the star particle. The difference in number of high-mass stars from the 5.5 stars expected above 8 $M_\odot$ using a continuous IMF, as well as the difference in mass of high-mass stars from the expected 105 $M_\odot$, is given for each realization. A stochastically populated IMF leads to large variation in mass and number of high-mass stars in each star particle.

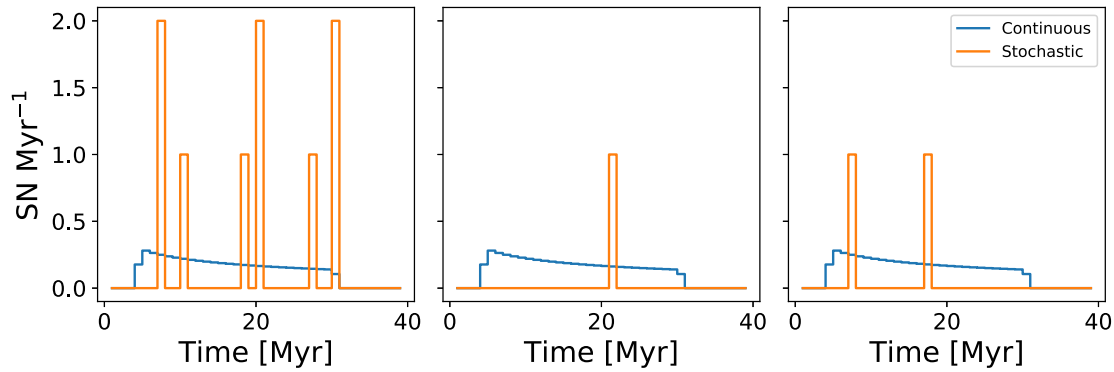

Figure 2. The supernova rate in 1 Myr intervals for the same three star particles as Fig. 1, with a stochastic IMF and a continuous IMF. The stochastically populated IMF leads to long periods of no feedback, with several intervals of stronger feedback. There are also different total numbers of supernovae and greatly varied timings. Though the middle star particle has three discretely tracked high-mass stars, only one star is below 40 $M_\odot$ and therefore in the mass range of Type-II supernovae.

These three realizations demonstrate several important features introduced by a stochastically populated IMF. The first, and most fundamental, is that high-mass stars are discretized, and therefore Type-II supernovae are discretized as well. We discuss in Section 2.3 how proper discretization of stars can *only* be done with a comprehensive consideration of IMF sampling methodology. Another unsurprising change is that the actual number of high-mass stars within a star particle can vary greatly. For example, there are eight stars with $m_{star} > 8$ $M_\odot$ in the left-most star particle of Fig. 1, and only two in the right-most star particle. Less obvious, the number and mass in high-mass stars is only loosely correlated. The middle star particle, for example, has fewer high-mass stars than the average, but has more mass in this range because the few high-mass stars that did form tended to be more massive. As seen in the middle panel of Fig. 2, we also note that since Type-II supernovae are restricted to stars with 8 $M_\odot < m_{star} <$40 $M_\odot$, only one of the high-mass stars in this star particle would actually explode as a supernova.

2.3 Existing discretization methods

Currently, there is one dominant method of discretizing supernova feedback in cosmological simulations (e.g. Stinson et al. 2010; Hopkins et al. 2014; Rosdahl et al. 2018; Smith, Sijacki & Shen 2018; Su et al. 2018), which we will henceforth refer to as 'quantized feedback.' While this is not the only method used in the simulation community, it seems to be the prevailing method among high-resolution cosmological simulations, and so we focus on it here. We emphasize that the issues raised in the following discussion are true for any method that does not sample the distribution of stars at birth, but rather calculates supernova explosions 'on the fly', while leaving the remainder of the IMF unchanged.

In quantized feedback, the number of supernovae in a given time-step is drawn from either a binomial distribution or Poisson distribution.[4] The supernova mean rate may be taken from rate tables (as in Hopkins et al. 2014) or from the expectation number of supernovae in that time-step (e.g. Stinson et al. 2010 or the RIMFS method of Revaz et al. 2016). While this method guarantees that only integer numbers of stars explode in a given time-step, we note that at high resolution there are several internal inconsistencies upon closer inspection. The importance of these inconsistencies

[4]A binomial distribution with many trials and a small probability of success converges to a Poisson distribution, so these two formulations are approximately equivalent.

on simulation results depends on the size of the ensemble of star particles considered. Here, we start with the case of a single-star particle.

First, this method cannot guarantee that the mass used in calculating feedback is the same as the dynamical mass of the particle. The mass of stars below the quantized regime is fixed, while the mass of stars within the quantized regime can vary by more than a factor of 2. In terms of the number of stars within the star particle, a fixed number of stars (those below the quantization limit) are added to a Poisson distribution of stars (those that are quantized). This combination will rarely yield the assumed initial particle mass. For smaller star particles, this effect will be more severe. The star particle mass and the mass removed from any individual parent gas particle will therefore be inconsistent. If, on the other hand, mass conservation is enforced (for example, by adding back missing mass or subtracting excess mass in the low-mass end after supernovae explode), then this method will require the mass in low-mass stars to be amended in real time in an unphysical way, leading to more internal inconsistencies. Additionally, imposing a limit (minimum) on the number of supernovae that can explode will artificially concentrate supernovae to go off in early (late) times.

To see this in more detail, we consider many trials of a 250 $M_\odot$ star particle with a Kroupa (2001) IMF and demonstrate that without an a priori knowledge of the stellar mass distribution, the stated star particle mass is inconsistent with the initial mass implied by the number of supernovae explosions. We consider a slightly simplified version of what occurs in simulations, assuming for each trial that we begin with the same metallicity. The mass below 8 $M_\odot$ in a star particle of this mass and IMF is 198 $M_\odot$, with the remainder falling in the high-mass portion of the IMF. For each particle, we iterate forwards in time-steps of 10^5 yr, and at each time-step determine whether a supernova explodes by the method of Stinson et al. (2010) (i.e. drawing from a binomial distribution). To make this comparison more applicable to other simulators without upper limits on core collapse supernovae, we do not place a Type II upper limit of 40 $M_\odot$. If a supernova explodes, we add the initial mass of the star that exploded to the initial low-mass total of the star particle. Since the range outside of Type-II supernovae is not quantized, its mass is unaffected by this procedure. The results of this process are shown in Fig. 3. While we use the specific methodology of Stinson et al. (2010), the results of Fig. 3 are broadly true for any scheme that uses a binomial or Poisson distribution to determine whether a supernova explodes in real time at each time-step.

A second drawback of this method, similar to the first, is that other forms of feedback are incorrectly calculated for every individual star particle. Much as mass is not conserved because the number and mass of high-mass stars is not known a priori, neither is any form of feedback that relies on high-mass stars. Photoionizing radiation and photoelectric feedback in simulations are calculated from the distribution of stars within a population, often relying on stellar population synthesis codes (e.g. Christensen et al. 2012; Hopkins, Quataert & Murray 2012; Agertz et al. 2013; Rosdahl et al. 2013; Stinson et al. 2013; Ceverino et al. 2014). If the actual distribution is not known ahead of time (and, as shown previously, is almost certainly different from the assumed distribution), then the photon estimates will be inconsistent. Since various feedback effects add non-linearly (Hopkins et al. 2014), any simulation that includes different feedback mechanisms must ensure that all are consistent. The only way to ensure consistency is to sample the IMF at the formation time of the particle.

To see how the results diverge between quantized feedback and a stochastic IMF, we run an ensemble of dwarf galaxies using the quantized feedback of Stinson et al. (2010). We compare the results in Section 3.4.

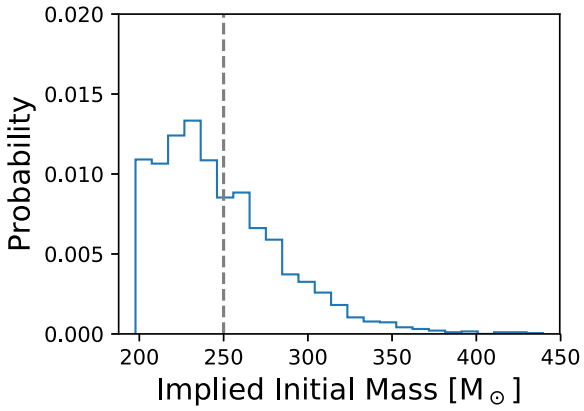

Figure 3. A Monte Carlo simulation of quantized feedback in which supernovae explode each time-step according to a binomial distribution given by their expected number. Thus, the mass in stars less massive than Type-II progenitors is fixed while the mass in high-mass stars can vary widely. The probability distribution of the resultant total initial mass of the star particle is shown. The vertical dashed line shows the assumed initial mass. Quantized feedback can result in almost 40 per cent difference between the assumed initial mass of the star particle and the mass implied by the number and mass of supernova explosions. To enforce mass conservation, the number and mass in low-mass stars would have to be constantly updated in real time as supernovae explode, leading to further inconsistencies.

We note one more advantage of the stochastic IMF over quantized feedback: any other sampling method can be easily incorporated. As discussed above, there is still ongoing debate over the fundamental way in which high-mass stars form, and therefore the proper way to sample from the IMF. If we wanted to use, for example, the sorted sampling of Weidner & Kroupa (2006), we could do so provided we implement a treatment for clusters. Choosing sampling methods is impossible in methods such as quantized feedback where the mass distribution is not known from birth.

3 RESULTS

3.1 Cosmological star formation

For this paper, we restrict our attention to galaxies residing in well-resolved haloes with $M_{\rm vir} > 10^7$ $M_\odot$. We focus exclusively on galaxies that exist in both runs to ensure our results are converged. To avoid Poisson noise in star formation between matching haloes, we only consider galaxies where at least one of the two runs formed at least four star particles. At $z = 6$, Elena forms 86 galaxies satisfying the above criteria that are either in the field or satellites of the main halo.

The most direct way to measure the impact of any prescription is through its effect on the stellar mass of galaxies. There is active research in the literature regarding the abundance of dwarf galaxies, which will soon be discovered in unprecedented numbers with the Large Synoptic Survey Telescope. There is also growing consensus that there is large scatter in the stellar mass–halo mass relationship at the low-mass end (Lin & Ishak 2016; Garrison-Kimmel et al. 2017; Munshi et al. 2017; Kulier et al. 2019). While we do not expect a stochastic IMF to significantly alter these relationships in massive galaxies, we do expect a resulting change in feedback to

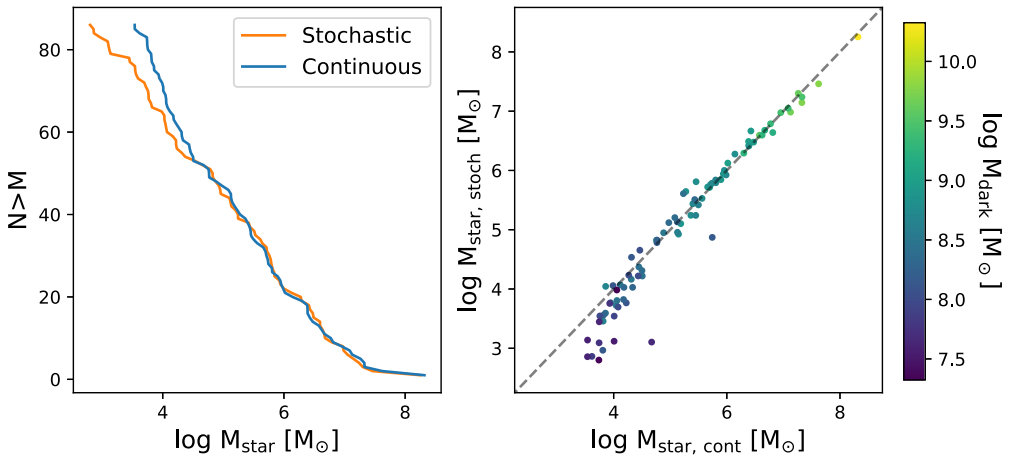

Figure 4. *Left*: The cumulative stellar mass function of galaxies in our sample for runs with a stochastic and continuous IMF. While converged above $\sim 10^5$ M$_\odot$, the run with a stochastic IMF is shifted towards smaller stellar masses below this mass. *Right*: Stellar masses for galaxies residing in matching haloes between the two simulations. Points are coloured according to the mass of the dark matter halo hosting the galaxy. The dashed grey line shows equal masses between runs. The smallest galaxies see a systematic reduction in stellar mass with a stochastic IMF compared to the run without a stochastic IMF.

affect dwarf galaxies, particularly those small enough such that their IMF is not converged.

The left-hand panel of Fig. 4 shows the cumulative stellar mass functions at $z = 6$ of the cosmological runs. The runs are identical other than the treatment of the IMF. Galaxies with $M_{\rm star} \gtrsim 10^5$ M$_\odot$ are fully converged, with identical stellar mass functions above this point. Below this mass, however, there is a clear shift towards smaller stellar masses in the run with a stochastic IMF.

To further quantify the differences in stellar mass, the right-hand panel of Fig. 4 shows the stellar masses of all galaxies in the left-hand panel, with corresponding galaxies matched between the two runs. The galaxies are coloured according to the mass of the dark matter halo in which they reside.[5] The figure shows that in the smallest galaxies, the stellar mass is generally lower with the stochastic prescription than with the continuous prescription. In the low-mass range, galaxies with a stochastic IMF see a reduction in stellar mass of up to an order of magnitude compared to the equivalent galaxy with a continuous IMF.

The trend of star formation suppression with a stochastic IMF becomes clearer if we instead arrange the stellar masses according to their dark matter masses. Fig. 5 shows the change in stellar mass as a function of halo mass for both sets of galaxies. Below a halo mass of $\sim 10^{8.5}$ M$_\odot$, galaxies formed with a stochastic IMF tend to have their star formation suppressed compared to a continuous IMF. Above this mass there is little change. This trend appears in addition to the generally greater scatter at lower halo mass which results from the breakdown of halo self-regulation. This dependence clarifies one of the outlier results of Fig. 4, where a somewhat more massive galaxy ($\sim 10^6$ M$_\odot$ in the run with a continuous IMF) still sees a reduction of an order of magnitude in stellar mass in the run with a stochastic IMF. The cause for this is that while its stellar mass is greater, its dark matter halo is more typical of the type that hosts ultra-faint galaxies.

[5]More precisely, they are coloured by the mean of matching halo masses in the two runs; however, the dark matter halo masses of the stochastic and non-stochastic runs are typically not more than 5 per cent different.

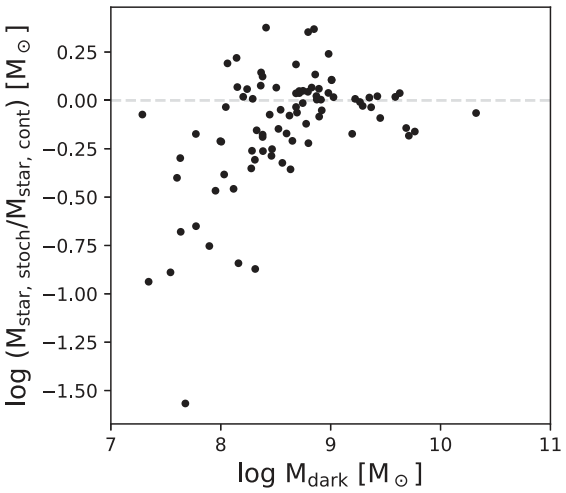

Figure 5. Difference in stellar mass as a function of dark matter halo mass for matched galaxies across runs with and without a stochastic IMF. The dashed grey line indicates equal masses between runs. Lower halo masses have increased scatter, and starting at $M_{\rm dark} \sim 10^{8.5}$ M$_\odot$, galaxies in the run with a stochastic IMF are generally lower than with a continuous IMF.

3.2 Bursty feedback

3.2.1 Supernova timing

To explain the reduced star formation with a stochastically sampled IMF as compared to the continuous IMF, a simple guess would be that since we now allow the total number of supernovae per star particle to vary, we now have varying total levels of feedback. At low SFR, however, a stochastically sampled IMF is expected to under-fill the high-mass end of the IMF, which would lead to fewer supernovae and less energy in feedback. If total supernova energy were the dominant factor, then, we would expect to see higher stellar masses in the run with a stochastic IMF, which is the opposite of the results shown in Fig. 4. In fact, among galaxies with suppressed star formation, roughly equal numbers have above and below average supernova total feedback.

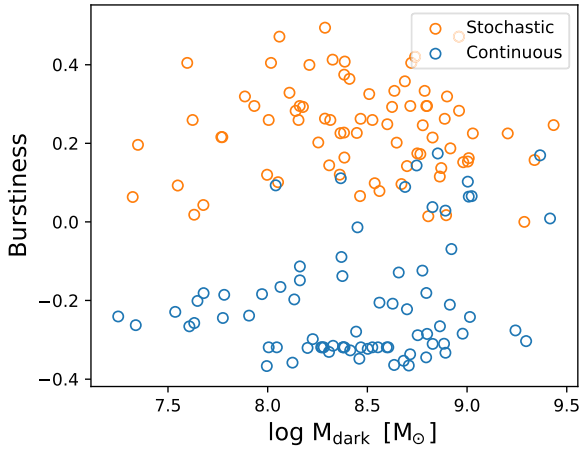

Figure 6. Burstiness of the supernova rate as a function of dark matter halo mass for runs with and without a stochastic IMF. Burstiness is calculated according to equation (1) using 1 Myr intervals during the first 30 Myr of star formation. Independent of galaxy properties, the run with a stochastic IMF leads to burstier supernova feedback, which leads to the suppression of star formation in dwarf galaxies that are too low-mass to self-regulate.

More important than the absolute number of supernovae, then, is the timing of the explosions. As is clear in Fig. 2, the supernova feedback with a stochastic IMF becomes much more temporally clustered. To see this in more detail, we can define a burstiness parameter (Goh & Barabási 2008 and similar to equation 1 of Mistani et al. 2016) as

$$B = \frac{\sigma/\mu - 1}{\sigma/\mu + 1}, \tag{1}$$

where σ is the standard deviation of the supernova rate, and μ is the mean supernova rate. Using this definition, the burstiness ranges from -1 to 1; a uniform distribution has a burstiness $B = -1$, an exponential distribution has a burstiness $B = 0$, and the burstiness approaches 1 as $\sigma/\mu \to \infty$. We calculate the rate in 1 Myr intervals[6] for the first 30 Myr of every galaxy's star formation (coinciding roughly with the longest-lived supernova from the first star particle to form in the galaxy). The results are plotted in Fig. 6.

As seen in the figure, the supernova rate is significantly burstier in runs with a stochastic IMF than runs with a continuous IMF. Crucially, the increase in burstiness applies to all galaxies, not only galaxies with low SFR. Since feedback with a stochastic IMF is universally more effective, it may seem surprising that we only see an impact in small galaxies. As was shown in Fig. 5, this stems from the stronger dependence on halo mass than stellar mass. The bursty supernova feedback leads to more effective heating of gas, as will be shown explicitly in the next section. However, only in small haloes less able to self-regulate does this more effective feedback completely shut off future star formation. The deeper potential wells of the higher mass haloes minimize the effects of burstier feedback.

[6]The burstiness parameter will be somewhat dependent on the binning chosen for the supernova rate, but we confirmed that the increased burstiness of the stochastic IMF over the continuous IMF is independent of bin size. We choose binning in 1 Myr intervals to be consistent with the timing of feedback in our simulations.

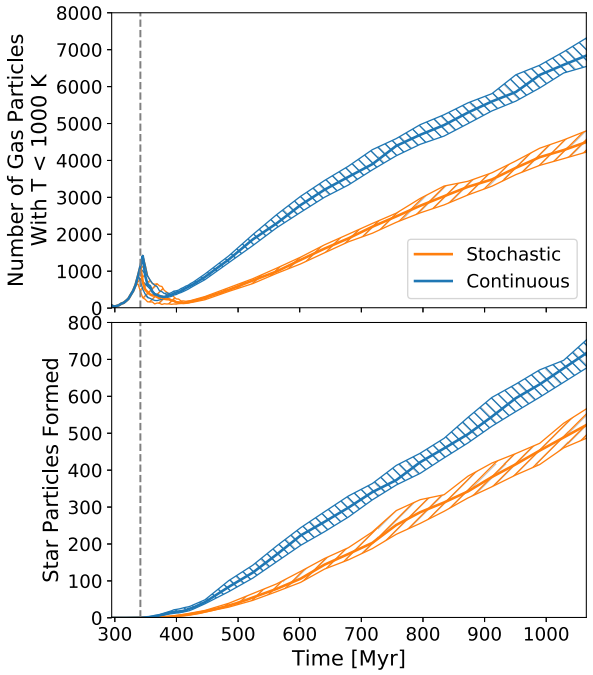

Figure 7. Star formation in the isolated dwarf galaxy runs, with and without a stochastic IMF. Each IMF treatment was run 50 times. Central lines show the medians, and bands represent the interquartile range. The grey line represents the median time of the first star formation. *Top*: The available cold gas as a function of time. With a stochastic IMF, about half as much gas is available compared to runs with a continuous IMF. *Bottom*: Cumulative star formation as a function of time. With a stochastic IMF, star formation is suppressed compared to runs with a continuous IMF, resulting from supernova feedback more effectively heating surrounding gas (see top panel).

3.2.2 *Isolated runs*

The same simulation run multiple times can have differing galaxy properties, owing to stochastic variations in numerical codes (Keller et al. 2019). To further investigate the results of the previous section and to ensure our results are independent from the intrinsic scatter in galaxy stellar mass, we consider an isolated simulation of a 10^9 M$_\odot$ halo, with initial star particle masses of 420 M$_\odot$.

To quantify the significance of differences between IMF treatments, we simulate the isolated halo 50 times with a stochastic IMF and 50 times without. The results are shown in Fig. 7, where we focus on the first Gyr. The bottom panel of the figure shows the cumulative number of star particles that have formed as a function of time, displaying both the median of all runs and the interquartile range. Clearly, the majority of the time, the run with a stochastic IMF forms fewer stars throughout the entire duration of the simulation.

To see why star formation is suppressed, the top panel of Fig. 7 shows the number of gas particles with a temperature below 1000 K. While the conditions for star formation are based on H_2 abundance (Christensen et al. 2012), this serves as a proxy for the number of gas particles that could potentially form stars. We see that with a stochastic IMF, the gas is more effectively heated – at any given time, up to half as many gas particles have cooled to below 1000 K, in line with the reduction in star formation. We have verified that in our cosmological simulation, the galaxies tend to have similar gas masses, with no systematic change as with stellar mass. This

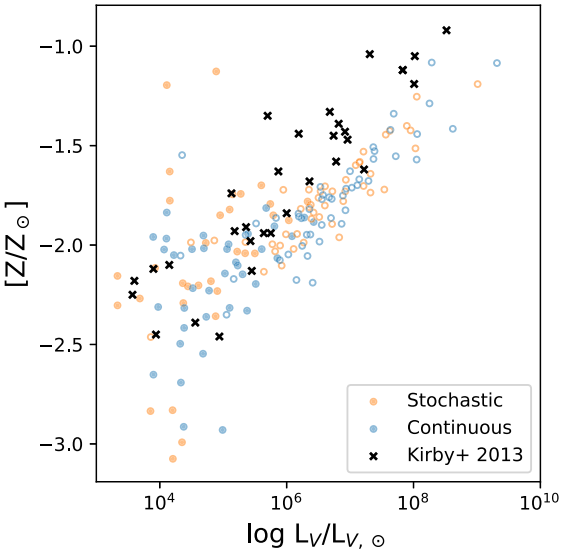

Figure 8. Mean stellar metallicity as a function of luminosity for galaxies in our sample, with and without a stochastic IMF. Filled circles show galaxies residing in dark matter haloes below $10^{8.5}$ $M_\odot$, approximately corresponding to galaxies we expect to be quenched by reionization (Tollerud & Peek 2018). Since the sample is at $z = 6$, only low-mass galaxies that have stopped forming stars match the $z = 0$ data. More massive galaxies are expected to increase in metallicity with time. There are no clear differences between the IMF treatments, though there may be slightly greater scatter in metallicity with a stochastic IMF.

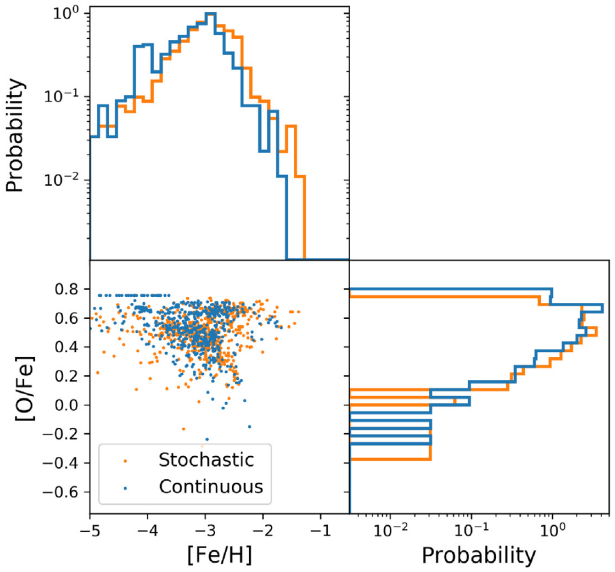

Figure 9. Chemical abundances of all-star particles in galaxies with $z = 6$ luminosities $L_V \leq 10^{5.5}$ $L_{V,\odot}$ and $M_{\rm dark} \leq 10^{8.5}$ $M_\odot$, corresponding roughly to $z = 0$ UFDs. The abundance ratios between the two models are largely consistent, likely as a result of metal diffusion in the ISM, which reduces scatter introduced by IMF sampling effects (Revaz et al. 2016). The high number of stars at [O/Fe] $\sim$ 0.75 using a continuous IMF result from the stars formed from the metal content of the very first time-step in which supernovae explode. There may be a slight shift to higher [Fe/H] with a stochastic IMF, but it is well below observational precision.

indicates that more effective gas heating dominates over more efficient gas expulsion.

In a cosmological setting, these effects can be significantly exaggerated. In the cosmological runs compared to the isolated runs, the differences between the stellar masses of galaxies can increase from a factor of $\sim$1.5 to $\sim$10, as in Fig. 4.

3.3 Metallicity

Beyond energy deposition, high-mass stars return processed material to the ISM. As previously emphasized, a stochastic IMF significantly alters the distribution of high-mass stars, including their masses and numbers. Since the metal production depends non-linearly on the mass of the exploding star (Raiteri et al. 1996), one might expect greater variation in the chemical enrichment of galaxies with a stochastic IMF.

Fig. 8 shows the luminosity–metallicity relationship of all galaxies in the sample, along with data from Kirby et al. (2013). To calculate the galaxy metallicities, we apply a floor for individual star particles of $Z > 10^{-5}$. Luminosities are calculated using PARSEC[7] isochrones (Bressan et al. 2012). We note that our sample is at $z = 6$; for faint galaxies that have likely stopped forming stars indefinitely, we match the data well, thopaugh our scatter below $L_V \sim 10^5$ $L_{V,\odot}$ is higher than the observations. Some of the excess scatter is likely reduced due to mergers at later times, when some of the faint galaxies with extremely low metallicities are incorporated into larger, higher metallicity galaxies. For larger galaxies that lie below the $z = 0$ observations, ongoing star formation will increase the metallicities with time. A forthcoming paper (Munshi et al. in prep) will discuss these relationships and demonstrate consistency with the data in the present day.

Surprisingly, runs with the stochastic and continuous IMF are consistent at all luminosities, though the scatter may increase for faint galaxies with a stochastic IMF. It is likely that metal diffusion in the ISM (Shen et al. 2010) quickly obscures any systematic differences in future generations of stars that form, consistent with the findings of Revaz et al. (2016), who found that introducing metal diffusion in dwarf galaxies reduced scatter introduced by IMF sampling effects. The overall impression is that the stochastic IMF has little impact on the chemical evolution of galaxies.

We also investigate whether the stellar chemical abundances change with a stochastically populated IMF. Fig. 9 shows the abundance ratios for galaxies with ($z = 6$) luminosities $L_V \leq 10^{5.5}$ $L_{V,\odot}$ and dark matter masses below $10^{8.5}$ $M_\odot$, which after passive stellar evolution would correspond roughly to today's UFDs. As above, there is almost no difference between the models; the [O/Fe] versus [Fe/H] distributions overlap completely. With the continuous IMF there is a small overabundance of stars at [O/Fe] $\sim$ 0.75, which corresponds to the maximum abundance ratio possible that results from the very first time-step in which supernovae explode when integrating along the continuous IMF. Additionally, there is a mild indication that stars with a stochastic IMF are shifted to slightly higher [Fe/H]. This shift makes sense given that requiring whole supernovae to explode means that more processed material is ejected at once into the ISM. On the whole, however, it seems that even when looking at the chemical composition of stars there is little to distinguish the models, especially given observational uncertainties.

[7] http://stev.oapd.inaf.it/cgi-bin/cmd

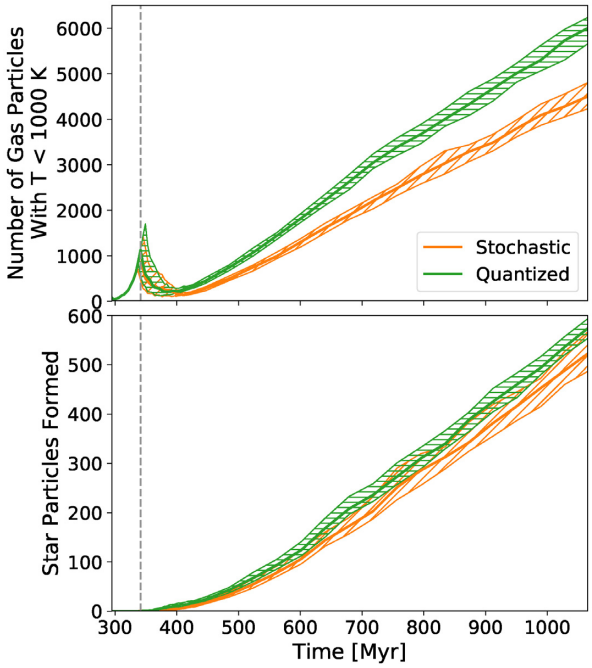

Figure 10. Star formation in the isolated dwarf galaxy runs, comparing runs with a stochastic IMF versus runs with quantized feedback. Each IMF treatment was run 50 times. Central lines show the medians, and bands represent the interquartile range. The grey line represents the median time of the first star formation. *Top*: The available cold gas as a function of time. Quantized feedback leaves ∼30 per cent more cold gas available to form stars compared to runs with a stochastic IMF. *Bottom*: Cumulative star formation as a function of time. Quantized feedback results in ∼10 per cent more star formation than a stochastic IMF.

3.4 Stochastic versus quantized feedback

As discussed in Section 2.3, quantized feedback is an existing method that discretizes supernovae. To test how it compares to a stochastic IMF, we ran the isolated dwarf galaxy 50 times with quantized Type II supernovae (but otherwise the same feedback implementation). Fig. 10 shows the quantity of cold gas and cumulative star formation as a function of time for the two IMF treatments. Interestingly, there is consistently ∼30 per cent more available cold gas and ∼10 per cent more star formation with the quantized feedback as compared to the stochastic IMF.

Stellar feedback, therefore, is less effective with quantized feedback than with a fully self-consistent stochastic IMF. There are multiple differences that could contribute to this behaviour. First, as discussed in Section 2.3, the LW photon production with quantized feedback cannot accurately reflect the internal distribution of high-mass stars. Different feedback mechanisms interact non-linearly (Hopkins et al. 2014), such that one can reinforce the effect of the others. The stochastic IMF ensures that LW radiation and supernovae come from the same stars, which may result in stronger disruption of gas. Additionally, since UV luminosities depend highly non-linearly on stellar mass, estimating the ionizing photon counts from fractions of massive stars may result in less initial LW radiation than for individual (whole) massive stars, though averaged over many particles the LW outputs converge. Another difference may come from the timing and distribution of supernovae; if the stochastic IMF leads to more temporally clustered explosions, then the cumulative heating of gas may be more effective. Monte Carlo simulations of a stochastically populated IMF show that approximately 25 per cent of 420 $M_\odot$ star particles with feedback calculated every 1 Myr will have a single time-step with multiple supernovae, as opposed to none using quantized feedback. Thus, some of the difference could result from our use of a binomial rather than Poisson distribution in calculating the quantized feedback, and we caution that a different implementation may yield closer results. We note that the average total number of supernovae per star particle is the same between the prescriptions.

We caution that these results come from an isolated dwarf galaxy. In a cosmological context, the presence of tidal interactions, ionizing radiation, and gas outflows may increase the differences in these results. In particular, the presence of a larger cold gas reservoir when using quantized feedback could lead to more extended star formation histories in UFD galaxies during and after the epoch of reionization.

In our case, because the primary source of energetic feedback comes from Type-II supernovae, the differences between a stochastic IMF and quantized feedback are conspicuous but limited. This demonstrates the general robustness of the supernova feedback physics implemented in our simulations. Some simulations incorporate significant quantities of radiative feedback from high-mass stars, beyond merely a Lyman–Werner prescription for H_2 destruction. While capturing more subgrid processes, they may also suffer from an even greater internal inconsistency in the use of quantized feedback, due to the non-linear nature of different feedback mechanisms, which may result in even greater differences between quantized feedback and a stochastically sampled IMF.

4 IMPLICATIONS FOR FUTURE PREDICTIONS

4.1 Star formation quenching

The most significant difference between the stochastic IMF and the continuous IMF is the lower stellar mass of a large fraction of dwarf galaxies. This is not unexpected, and is similar to what was found by Su et al. (2018). The effect we find, however, is less extreme, owing mostly to the differences in our feedback implementations. Our feedback is calculated in 1 Myr intervals, compared to the smaller time-scales in the FIRE simulations, which can be shorter than 10^4 yr. Thus, continuous injection of supernova energy results in relatively large per-step feedback in our simulations as compared to continuous injection in FIRE. Additionally, in CHANGA instantaneous energy deposition is incorporated as a 1 Myr heating rate to avoid numerical instabilities (see footnote 63 in Kim et al. 2016), meaning that even with a stochastic IMF supernova energy is effectively continuously injected when considering sufficiently small time intervals.

The cause of this suppression in star formation is clear from its effects on the gas in the isolated dwarf runs: supernova feedback with a stochastic IMF is more effectively preventing gas from cooling. In the first Gyr of star formation, there can be more than a factor of 2 difference in the amount of gas available to form stars, with a corresponding difference in the amount of star formation. The reason for the suppression of star formation is that supernova energy is deposited in a shorter time frame; this is not only because stars are discrete, but because stochasticity in the stellar masses allows supernovae to cluster temporally, as was shown in Fig. 2.

We also find a strong halo mass dependence in our results, as was shown in Fig. 5. Owing to our large sample, we are able to explore these trends. The continuous IMF treatment yields largely the same results for galaxies residing in haloes more massive than $\sim 10^{8.5}$ $M_\odot$

at $z = 6$. These more massive haloes are able to self-regulate even with the stronger supernova feedback. Smaller haloes, however, have too small a potential well to prevent supernovae from driving out gas from star-forming regions. This is also approximately the (high-redshift) mass scale where reionization is thought to suppress star formation (Quinn, Katz & Efstathiou 1996; Thoul & Weinberg 1996; Barkana & Loeb 1999; Bullock, Kravtsov & Weinberg 2000; Gnedin 2000; Okamoto, Gao & Theuns 2008; Tollerud & Peek 2018); most galaxies in haloes of these masses will quench, if not from supernova feedback, then from the combined effects of feedback and reionization (Benson et al. 2002; Somerville 2002; Hoeft et al. 2006; Nickerson et al. 2011). For a given reionization model, our results show that a stochastic IMF suppresses star formation even further in low-mass haloes.

4.2 Reionization

Though in some cases galaxies quenched at high redshift can restart star formation in later times (Wright et al. 2019), even a temporary suppression of star formation would have significant implications for the epoch of reionization. Early galaxies are thought to be the primary source of ionizing radiation (Stark 2016), with significant contributions from dwarf galaxies. In fact, though not directly observable, inferences from the Local Group imply that very small dwarf galaxies (as faint as $M_V \sim -3$) may have contributed to reionization (Weisz & Boylan-Kolchin 2017). What is still unknown, however, is the fraction of ionizing radiation they provided, and similarly, their luminosity function at high redshifts. There have been many simulation predictions of the reionization era, reaching down to dwarf galaxy scales (e.g. O'Shea et al. 2015; Finlator et al. 2016; Gnedin 2016; Ocvirk et al. 2016; Xu et al. 2016; Anderson et al. 2017; Ma et al. 2018), and simulations continue to push to higher resolution.

When properly accounting for IMF sampling effects, we have seen that star formation is often suppressed earlier, indicating that by $z \sim 6$, the faintest galaxies will constitute a reduced fraction of the ionizing photon budget. This may be countered by higher escape fractions resulting from hotter bubbles of gas around these small galaxies. Detailed explorations of the implications of IMF sampling on reionization will be pursued in future work, but what is already clear is that accurate predictions will require a stochastically populated IMF.

4.3 Tracking high-mass stars

One of the key new features of this prescription is the tracking of individual high-mass star data. For every star particle, we now have a list of the masses of every star above 8 $M_\odot$ residing within it. This opens up new science avenues that were not available before.

For example, cosmological simulations may be used to estimate the size and evolution of H II regions, as in Anderson et al. (2017), based on the ionizing photon output from star particles. Rather than assume ionizing photon production from an SSP, this new prescription allows us to use the specific stars to determine the photon rate, which will add variability to our predictions.

Predictions for stellar remnants, such as pulsar counts, can now be directly inferred from star particles in our simulations. Further, now that we have entered the era of gravitational wave astronomy (Abbott et al. 2016), simulations can be used to predict the merger rates of binary compact objects. This has so far been accomplished by pairing simulation outputs with population synthesis models (e.g. Chakrabarti et al. 2017; Mapelli et al. 2017; O'Shaughnessy et al. 2017; Schneider et al. 2017). Having precise high-mass star information can allow us to refine such predictions by using the actual high-mass star counts from the simulations. Additionally, since the Milky Way has accreted many small galaxies over its lifetime, we can still see the imprint of these single rare events. For example, in the Milky Way stellar halo, accreted UFD-like galaxies may have contributed over half of *r*-process enhanced metal-poor stars (Brauer et al. 2019). Such predictions, however, are highly sensitive to the actual numbers of massive binary stars in small populations.

5 SUMMARY

Motivated by the inability of low-mass haloes to self-regulate, and the resulting divergence of different prescriptions in UFDs, we investigated the treatment of stellar feedback in simulations. In this work, we presented a new treatment of the IMF in cosmological simulations. Informed by observations, we stochastically sample stars from the IMF within each star particle. As a compromise with computational reality, once we have stochastically populated a star particle, we only track individual stars above 8 $M_\odot$, so that feedback dependent on high-mass stars is calculated for discrete stars, and feedback dependent on low-mass stars is calculated as before for continuous populations.

To investigate the effects of our new stochastic IMF, we used cosmological zoom-in simulations run to $z = 6$ to compare the stellar masses of the resulting ∼100 galaxies. We found that while galaxies residing in dark matter haloes greater than $\sim 10^{8.5}$ $M_\odot$ remain unchanged by the new sampling method, galaxies in smaller haloes typically have lower stellar masses, by up to an order of magnitude. When comparing the supernova rate via a burstiness parameter, we found that a stochastic IMF leads to significantly burstier feedback because of the greater temporal clustering of supernovae.

To see how the burstier supernova feedback impacts the gas in a galaxy, we used a simulated isolated dwarf galaxy. We ran many versions of both the continuous and stochastically sampled IMF treatment to study systematic differences in star formation, while bracketing scatter between runs. We found that during the first billion years, the galaxy with the stochastic IMF formed as few as half as many stars. The suppression in star formation results from a reduction in the number of gas particles available to form stars. Feedback with a stochastic IMF is more effective at heating surrounding gas and preventing gas from cooling and condensing into stars.

While another method ('quantized feedback') is typically used to ensure supernovae are discrete, we found that this method has several inconsistencies, since the actual distribution of high-mass stars is unknown until all supernovae have exploded. We found that the star formation results of quantized feedback were intermediate between a continuous and a stochastic IMF. Quantized feedback leads to ∼10 and ∼30 per cent more star formation and available cold gas, respectively. If high-energy radiation contributed energy or momentum in our simulations, it is possible the results would be even more dissimilar.

To test the new IMF prescription, this work focused on galaxies at high redshift. Since the galaxies in the mass range where this prescription is most impactful stop forming stars shortly after reionization, this was sufficient to draw conclusions for faint galaxies. Future work will investigate the effects of a stochastic IMF in simulations run to the present day.

This IMF prescription is ideal for high-resolution simulations; as star particle masses decrease, we can lower the cut-off mass to discretely track lower mass stars, and incorporate discrete treatments for Type-Ia explosions and mass-loss due to stellar winds. Further, our unique ability to track the evolution of individual stars in cosmological simulations will allow us to make more specific predictions for any observations dependent on the number and distribution of high-mass stars. Future work, for example, can constrain the rates of compact object binary mergers detectable by gravitational wave experiments.

As we explore smaller stellar populations in simulations, we are now afforded the opportunity to investigate astrophysics on smaller scales. With the future predictive power of our simulations in mind, we have implemented a novel stochastic IMF in our cosmological simulations. In future studies of faint galaxies, including their stellar populations and radiative contributions to the epoch of reionization, it will be necessary to use such a stochastically populated IMF to accurately model these phenomena. Otherwise, observational predictions will systematically overestimate the star formation in UFD galaxies.

ACKNOWLEDGEMENTS

The authors would like to thank the anonymous referee for useful comments that improved the manuscript. We would also like to thank Ferah Munshi and Michael Tremmel for useful discussions relating to this work. EA and AMB acknowledge support from NSF grant AST-1813871. Resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center. EA acknowledges support from the National Science Foundation (NSF) Blue Waters Graduate Fellowship. This research is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (awards OCI-0725070 and ACI-1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications.

REFERENCES

Abbott B. P. et al., 2016, Phys. Rev. Lett., 116, 061102
Agertz O., Kravtsov A. V., 2015, ApJ, 804, 18
Agertz O., Kravtsov A. V., Leitner S. N., Gnedin N. Y., 2013, ApJ, 770, 25
Anderson L., Governato F., Karcher M., Quinn T., Wadsley J., 2017, MNRAS, 468, 4077
Andrews J. E. et al., 2013, ApJ, 767, 51
Andrews J. E. et al., 2014, ApJ, 793, 4
Barkana R., Loeb A., 1999, ApJ, 523, 54
Bastian N., Covey K. R., Meyer M. R., 2010, ARA&A, 48, 339
Bellovary J. M., Cleary C. E., Munshi F., Tremmel M., Christensen C. R., Brooks A., Quinn T. R., 2019, MNRAS, 482, 2913
Benincasa S. M., Wadsley J., Couchman H. M. P., Keller B. W., 2016, MNRAS, 462, 3053
Benson A. J., Lacey C. G., Baugh C. M., Cole S., Frenk C. S., 2002, MNRAS, 333, 156
Boylan-Kolchin M., Bullock J. S., Kaplinghat M., 2011, MNRAS, 415, L40
Boylan-Kolchin M., Bullock J. S., Kaplinghat M., 2012, MNRAS, 422, 1203
Brauer K., Ji A. P., Frebel A., Dooley G. A., Gomez F. A., O'Shea B. W., 2019, ApJ, 871, 247
Bressan A., Marigo P., Girardi L., Salasnich B., Dal Cero C., Rubele S., Nanni A., 2012, MNRAS, 427, 127
Brooks A. M., Zolotov A., 2014, ApJ, 786, 87
Brooks A. M., Governato F., Booth C. M., Willman B., Gardner J. P., Wadsley J., Stinson G., Quinn T., 2007, ApJ, 655, L17
Brooks A. M., Kuhlen M., Zolotov A., Hooper D., 2013, ApJ, 765, 22
Brooks A. M., Papastergis E., Christensen C. R., Governato F., Stilp A., Quinn T. R., Wadsley J., 2017, ApJ, 850, 97
Brook C. B. et al., 2011, MNRAS, 415, 1051
Brown T. M. et al., 2014, ApJ, 796, 91
Bryan G. L., Norman M. L., 1998, ApJ, 495, 80
Bullock J. S., Kravtsov A. V., Weinberg D. H., 2000, ApJ, 539, 517
Calzetti D., Chandar R., Lee J. C., Elmegreen B. G., Kennicutt R. C., Whitmore B., 2010, ApJ, 719, L158
Cappellari M. et al., 2012, Nature, 484, 485
Carigi L., Hernandez X., 2008, MNRAS, 390, 582
Cerviño M., Román-Zúñiga C., Luridiana V., Bayo A., Sánchez N., Pérez E., 2013, A&A, 553, A31
Ceverino D., Klypin A., Klimek E. S., Trujillo-Gomez S., Churchill C. W., Primack J., Dekel A., 2014, MNRAS, 442, 1545
Chabrier G., 2003, PASP, 115, 763
Chakrabarti S., Chang P., O'Shaughnessy R., Brooks A. M., Shen S., Bellovary J., Gladysz W., Belczynski C., 2017, ApJ, 850, L4
Christensen C., Quinn T., Governato F., Stilp A., Shen S., Wadsley J., 2012, MNRAS, 425, 3058
Christensen C. R., Quinn T., Stinson G., Bellovary J., Wadsley J., 2010, ApJ, 717, 121
Christensen C. R., Governato F., Quinn T., Brooks A. M., Shen S., McCleary J., Fisher D. B., Wadsley J., 2014, MNRAS, 440, 2843
Christensen C. R., Davé R., Governato F., Pontzen A., Brooks A., Munshi F., Quinn T., Wadsley J., 2016, ApJ, 824, 57
Conroy C., van Dokkum P. G., 2012, ApJ, 760, 71
Corbelli E., Verley S., Elmegreen B. G., Giovanardi C., 2009, A&A, 495, 479
da Silva R. L., Fumagalli M., Krumholz M., 2012, ApJ, 745, 145
da Silva R. L., Fumagalli M., Krumholz M. R., 2014, MNRAS, 444, 3275
de Blok W. J. G., Walter F., Brinks E., Trachternach C., Oh S.-H., Kennicutt R. C., 2008, AJ, 136, 2648
Di Cintio A., Brook C. B., Macciò A. V., Stinson G. S., Knebe A., Dutton A. A., Wadsley J., 2014, MNRAS, 437, 415
Eldridge J. J., 2012, MNRAS, 422, 794
Elmegreen B. G., 2006, ApJ, 648, 572
Emami N., Siana B., Weisz D. R., Johnson B. D., Ma X, El-Badry K., 2019, ApJ, 881, 71
Emerick A., Bryan G. L., Mac Low M.-M., 2019, MNRAS, 482, 1304
Finlator K., Oppenheimer B. D., Davé R., Zackrisson E., Thompson R., Huang S., 2016, MNRAS, 459, 2299
Fitts A. et al., 2017, MNRAS, 471, 3547
Fujimoto Y., Krumholz M. R., Tachibana S., 2018, MNRAS, 480, 4025
Fumagalli M., da Silva R. L., Krumholz M. R., 2011, ApJ, 741, L26ga
Garrison-Kimmel S., Bullock J. S., Boylan-Kolchin M., Bardwell E., 2017, MNRAS, 464, 3108
Garrison-Kimmel S. et al., 2019, MNRAS, 487, 1380
Gatto A. et al., 2017, MNRAS, 466, 1903
Geen S., Watson S. K., Rosdahl J., Bieri R., Klessen R. S., Hennebelle P., 2018, MNRAS, 481, 2548
Geha M. et al., 2013, ApJ, 771, 29
Gennaro M. et al., 2018, ApJ, 855, 20
Gill S. P. D., Knebe A., Gibson B. K., 2004, MNRAS, 351, 399
Gnedin N. Y., 2000, ApJ, 542, 535
Gnedin N. Y., 2016, ApJ, 825, L17
Goh K. I., Barabási A. L., 2008, EPL, 81, 48002
Governato F. et al., 2010, Nature, 463, 203
Governato F. et al., 2012, MNRAS, 422, 1231
Grudić M. Y., Hopkins P. F., 2019, MNRAS, 488, 2970
Guo Y. et al., 2016, ApJ, 833, 37
Haardt F., Madau P., 2012, ApJ, 746, 125
Hensler G., Steyrleithner P., Recchi S., 2017, in Gil de Paz A., Knapen J. H., Lee J. C., eds, Proc. IAU Symp. 321, Formation and Evolution of Galaxy Outskirts. Cambridge Univ. Press, Cambridge, England, p. 99
Hoeft M., Yepes G., Gottlöber S., Springel V., 2006, MNRAS, 371, 401

Hopkins P. F., Quataert E., Murray N., 2011, MNRAS, 417, 950
Hopkins P. F., Quataert E., Murray N., 2012, MNRAS, 421, 3488
Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, MNRAS, 445, 581
Hopkins P. F. et al., 2018, MNRAS, 480, 800
Hu C.-Y., Naab T., Glover S. C. O., Walch S., Clark P. C., 2017, MNRAS, 471, 2151
Kalirai J. S. et al., 2013, ApJ, 763, 110
Katz N., White S. D. M., 1993, ApJ, 412, 455
Kaufmann T., Mayer L., Wadsley J., Stadel J., Moore B., 2007, MNRAS, 375, 53
Kazantzidis S., Magorrian J., Moore B., 2004, ApJ, 601, 37
Keller B. W., Wadsley J., Benincasa S. M., Couchman H. M. P., 2014, MNRAS, 442, 3013
Keller B. W., Wadsley J. W., Wang L., Kruijssen J. M. D., 2019, MNRAS, 482, 2244
Kim J.-h. et al., 2016, ApJ, 833, 202
Kirby E. N., Cohen J. G., Guhathakurta P., Cheng L., Bullock J. S., Gallazzi A., 2013, ApJ, 779, 102
Klypin A., Kravtsov A. V., Valenzuela O., Prada F., 1999, ApJ, 522, 82
Knollmann S. R., Knebe A., 2009, ApJS, 182, 608
Kroupa P., 2001, MNRAS, 322, 231
Kroupa P., Weidner C., 2003, ApJ, 598, 1076
Kroupa P., Tout C. A., Gilmore G., 1993, MNRAS, 262, 545
Kroupa P., Weidner C., Pflamm-Altenburg J., Thies I., Dabringhausen J., Marks M., Maschberger T., 2013, Planets, Stars, and Stellar Systems, Vol. 5. Springer, Dordrecht, p. 115
Krumholz M. R., Fumagalli M., da Silva R. L., Rendahl T., Parra J., 2015, MNRAS, 452, 1447
Kulier A., Padilla N., Schaye J., Crain R. A., Schaller M., Bower R. G., Theuns T., Paillas E., 2019, MNRAS, 482, 3261
Kuzio de Naray R., McGaugh S. S., de Blok W. J. G., 2008, ApJ, 676, 920
La Barbera F., Ferreras I., Vazdekis A., de la Rosa I. G., de Carvalho R. R., Trevisan M., Falcón-Barroso J., Ricciardelli E., 2013, MNRAS, 433, 3017
Lee J. C. et al., 2009, ApJ, 706, 599
Lee J. C., Veilleux S., McDonald M., Hilbert B., 2016, ApJ, 817, 177
Leitherer C. et al., 1999, ApJS, 123, 3
Lin W., Ishak M., 2016, J. Cosmol. Astropart. Phys., 10, 025
Ma X. et al., 2018, MNRAS, 477, 219
Mapelli M., Giacobbo N., Ripamonti E., Spera M., 2017, MNRAS, 472, 2422
Martín-Navarro I. et al., 2015, ApJ, 806, L31
Mashchenko S., Wadsley J., Couchman H. M. P., 2008, Science, 319, 174
McKee C. F., Ostriker J. P., 1977, ApJ, 218, 148
Menon H., Wesolowski L., Zheng G., Jetley P., Kale L., Quinn T., Governato F., 2015, Comput. Astrophys. Cosmol., 2, 1
Meurer G. R. et al., 2009, ApJ, 695, 765
Mistani P. A. et al., 2016, MNRAS, 455, 2323
Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, ApJ, 524, L19
Munshi F. et al., 2013, ApJ, 766, 56
Munshi F., Brooks A. M., Applebaum E., Weisz D. R., Governato F., Quinn T. R., 2017, preprint (arXiv:1705.06286)
Munshi F., Brooks A. M., Christensen C., Applebaum E., Holley-Bockelmann K., Quinn T. R., Wadsley J., 2019, ApJ, 874, 40
Nickerson S., Stinson G., Couchman H. M. P., Bailin J., Wadsley J., 2011, MNRAS, 415, 257
O'Shaughnessy R., Bellovary J. M., Brooks A., Shen S., Governato F., Christensen C. R., 2017, MNRAS, 464, 2831
O'Shea B. W., Wise J. H., Xu H., Norman M. L., 2015, ApJ, 807, L12
Ocvirk P. et al., 2016, MNRAS, 463, 1462
Okamoto T., Gao L., Theuns T., 2008, MNRAS, 390, 920
Pflamm-Altenburg J., Weidner C., Kroupa P., 2009, MNRAS, 395, 394
Planck Collaboration et al., 2016, A&A, 594, A13
Pontzen A., Governato F., 2012, MNRAS, 421, 3464
Pontzen A. et al., 2008, MNRAS, 390, 1349
Quinn T., Katz N., Efstathiou G., 1996, MNRAS, 278, L49
Raiteri C. M., Villata M., Navarro J. F., 1996, A&A, 315, 105
Read J. I., Gilmore G., 2005, MNRAS, 356, 107
Read J. I., Agertz O., Collins M. L. M., 2016, MNRAS, 459, 2573
Revaz Y., Arnaudon A., Nichols M., Bonvin V., Jablonka P., 2016, A&A, 588, A21
Rosani G., Pasquali A., La Barbera F., Ferreras I., Vazdekis A., 2018, MNRAS, 476, 5233
Rosdahl J., Blaizot J., Aubert D., Stranex T., Teyssier R., 2013, MNRAS, 436, 2188
Rosdahl J. et al., 2018, MNRAS, 479, 994
Saitoh T. R., Daisaka H., Kokubo E., Makino J., Okamoto T., Tomisaka K., Wada K., Yoshida N., 2008, Publ. Astron. Soc. Japan, 60, 667
Salpeter E. E., 1955, ApJ, 121, 161
Sawala T. et al., 2016, MNRAS, 457, 1931
Schneider R., Graziani L., Marassi S., Spera M., Mapelli M., Alparone M., Bennassuti M. d., 2017, MNRAS, 471, L105
Shen S., Wadsley J., Stinson G., 2010, MNRAS, 407, 1581
Shetty R., Ostriker E. C., 2008, ApJ, 684, 978
Simon J. D., Bolatto A. D., Leroy A., Blitz L., Gates E. L., 2005, ApJ, 621, 757
Smith M. C., Sijacki D., Shen S., 2018, MNRAS, 478, 302
Somerville R. S., 2002, ApJ, 572, L23
Sormani M. C., Treß R. G., Klessen R. S., Glover S. C. O., 2017, MNRAS, 466, 407
Sparre M., Hayward C. C., Feldmann R., Faucher-Giguère C.-A., Muratov A. L., Kereš D., Hopkins P. F., 2017, MNRAS, 466, 88
Spiniello C., Trager S., Koopmans L. V. E., Conroy C., 2014, MNRAS, 438, 1483
Stark D. P., 2016, ARA&A, 54, 761
Stinson G., Seth A., Katz N., Wadsley J., Governato F., Quinn T., 2006, MNRAS, 373, 1074
Stinson G. S., Dalcanton J. J., Quinn T., Kaufmann T., Wadsley J., 2007, ApJ, 667, 170
Stinson G. S., Bailin J., Couchman H., Wadsley J., Shen S., Nickerson S., Brook C., Quinn T., 2010, MNRAS, 408, 812
Stinson G. S., Brook C., Macciò A. V., Wadsley J., Quinn T. R., Couchman H. M. P., 2013, MNRAS, 428, 129
Su K.-Y. et al., 2018, MNRAS, 480, 1666
Teyssier R., Pontzen A., Dubois Y., Read J. I., 2013, MNRAS, 429, 3068
Thoul A. A., Weinberg D. H., 1996, ApJ, 465, 608
Tollerud E. J., Peek J. E. G., 2018, ApJ, 857, 45
Tomozeiu M., Mayer L., Quinn T., 2016, ApJ, 827, L15
Tremmel M., Governato F., Volonteri M., Quinn T. R., 2015, MNRAS, 451, 1868
Tremmel M., Karcher M., Governato F., Volonteri M., Quinn T. R., Pontzen A., Anderson L., Bellovary J., 2017, MNRAS, 470, 1121
van Dokkum P. G., Conroy C., 2010, Nature, 468, 940
Wadsley J. W., Keller B. W., Quinn T. R., 2017, MNRAS, 471, 2357
Weidner C., Kroupa P., 2006, MNRAS, 365, 1333
Weidner C., Kroupa P., Bonnell I. A. D., 2010, MNRAS, 401, 275
Weidner C., Kroupa P., Pflamm-Altenburg J., 2013a, MNRAS, 434, 84
Weidner C., Kroupa P., Pflamm-Altenburg J., Vazdekis A., 2013b, MNRAS, 436, 3309
Weisz D. R., Boylan-Kolchin M., 2017, MNRAS, 469, L83
Weisz D. R. et al., 2012, ApJ, 744, 44
Weisz D. R., Dolphin A. E., Skillman E. D., Holtzman J., Gilbert K. M., Dalcanton J. J., Williams B. F., 2014, ApJ, 789, 148
Wetzel A. R., Hopkins P. F., Kim J.-h., Faucher-Giguère C.-A., Kereš D., Quataert E., 2016, ApJ, 827, L23
Wheeler C., Oñorbe J., Bullock J. S., Boylan-Kolchin M., Elbert O. D., Garrison- Kimmel S., Hopkins P. F., Kereš D., 2015, MNRAS, 453, 1305
Wright A. C., Brooks A. M., Weisz D. R., Christensen C. R., 2019, MNRAS, 482, 1176
Xu H., Wise J. H., Norman M. L., Ahn K., O'Shea B. W., 2016, ApJ, 833, 84
Yan Z., Jerabkova T., Kroupa P., 2017, A&A, 607, A126
Zolotov A. et al., 2012, ApJ, 761, 71

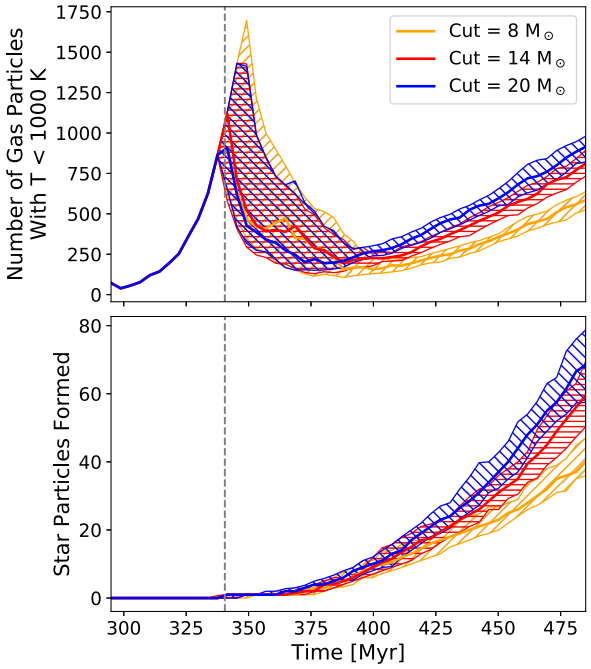

Figure A1. The same as Figs 7 and 10, but instead comparing three versions of the stochastic IMF with different cut-off masses. The higher the cut-off, and therefore the more supernovae are approximated as continuous, the more cold gas is available to form stars and the more stars form. It is therefore important to capture all core collapse supernovae discretely in any IMF sampling method.

APPENDIX A: CUT-OFF MASS

It is important to capture all Type-II supernovae when employing a discretization scheme. That is, shifting the cut-off mass above the minimum core-collapse supernova mass lessens the impact of the discretization. Fig. A1 shows this, where we have plotted the first ∼100 Myr of star formation of 50 runs each of the isolated dwarf galaxy. As we shift the cut-off higher, thereby approximating more of the Type II supernovae as continuous rather than discrete, we approach the fully continuous case of greater star formation and more cold gas. Thus, in order to capture the full impact of the energy deposition, *all* supernovae must be described within the framework of a stochastic IMF.

APPENDIX B: NUMERICAL CONSIDERATIONS

Here, we briefly discuss the computational considerations of applying the IMF sampling method presented in this work. Compared to quantized sampling, this method clearly requires more memory allocated per star particle (specifically, an array containing the list of masses of discrete massive stars tracked). For this reason, this method is best applied only in very high-resolution cosmological simulations. For most codes that already associate many pieces of information with their star particles (e.g. temperature, chemical composition, formation time, mass, etc.), a reasonable constraint is to no more than double the memory required per star particle, or equivalently (for typical cosmological codes) to limit the maximum number of individual stars expected during sampling to be of order 10. For this reason, this method is intended for use in simulations with star particle masses $\lesssim$1000 $M_\odot$.

In terms of computation time, sampling in this way is more intensive than quantized feedback or a method akin to Gatto et al. (2017) in which the IMF is sampled only in the high-mass regime. However, as discussed in Section 2.2, sampling instead over the entire regime imposes no restriction on the mass or number of stars in the high-mass regime, and so allows for greater variation in both. Fortunately, though the sampling takes time, it is negligible compared to the computation time involved in other aspects of the simulation. Timing comparisons of simulations suggest differences between the two methods are at no more than the per cent level.

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