Evidence of a Non-universal Stellar Initial Mass Function. Insights from *HST* Optical Imaging of Six Ultra-faint Dwarf Milky Way Satellites*

Mario Gennaro¹ , Kirill Tchernyshyov², Thomas M. Brown¹ , Marla Geha³, Roberto J. Avila¹, Puragra Guhathakurta⁴ , Jason S. Kalirai^{1,2} , Evan N. Kirby⁵ , Alvio Renzini⁶, Joshua D. Simon⁷, Jason Tumlinson^{1,2} , and Luis C. Vargas^{3,8}

¹ Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA; gennaro@stsci.edu

² Department of Physics and Astronomy, The Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218, USA

³ Department of Astronomy, Yale University, New Haven, CT 06520, USA

⁴ UCO/Lick Observatory and Department of Astronomy and Astrophysics, University of California, Santa Cruz, CA 95064, USA

⁵ California Institute of Technology, 1200 East California Boulevard, MC 249-17, Pasadena, CA 91125, USA

⁶ Osservatorio Astronomico, Vicolo Dell'Osservatorio 5, I-35122 Padova, Italy

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

⁸ SecurityScorecard Inc., 214 West 29th St, New York, NY 10001, USA

Received 2017 December 12; revised 2018 January 18; accepted 2018 January 18; published 2018 February 28

Abstract

Using deep observations obtained with the Advanced Camera for Surveys (ACS) on board the *Hubble Space Telescope* (*HST*), we demonstrate that the sub-solar stellar initial mass function (IMF) of six ultra-faint dwarf Milky Way satellites (UFDs) is more bottom light than the IMF of the Milky Way disk. Our data have a lower-mass limit of $\sim 0.45 M_{\odot}$, while the upper limit is $\sim 0.8 M_{\odot}$, set by the turnoff mass of these old, metal-poor systems. If formulated as a single power law, we obtain a shallower IMF slope than the Salpeter value of -2.3 , ranging from -1.01 for Leo IV to -1.87 for Boötes I. The significance of these deviations depends on the galaxy and is typically 95% or more. When modeled as a log-normal, the IMF fit results in a higher peak mass than in the Milky Way disk, but a Milky Way disk value for the characteristic system mass ($\sim 0.22 M_{\odot}$) is excluded at only 68% significance, and only for some UFDs in the sample. We find that the IMF slope correlates well with the galaxy mean metallicity, and to a lesser degree, with the velocity dispersion and the total mass. The strength of the observed correlations is limited by shot noise in the number of observed stars, but future space-based missions like the *James Webb Space Telescope* (*JWST*) and the *Wide-Field Infrared Survey Telescope* (*WFIRST*) will enhance both the number of dwarf Milky Way satellites that can be studied in such detail and the observation depth for individual galaxies.

Key words: galaxies: dwarf – galaxies: stellar content – Local Group – methods: statistical – stars: luminosity function, mass function

1. Introduction

Since the pioneering work of Salpeter (1955), increasing evidence has been collected that the mass distribution of newly formed stars, the initial mass function (IMF), has a constant behavior through space and time, at least in our Galaxy (Kroupa 2002; Bochanski et al. 2010). Studies of the Milky Way field population and stellar clusters suggest that the vast majority of stars were drawn from a universal IMF (Bastian et al. 2010).

However, the IMF universality has recently been questioned by studies of extragalactic environments. For example, van Dokkum & Conroy (2010, 2011, 2012) and Conroy & van Dokkum (2012) detect signatures of a possible overabundance of low-mass stars in the spectra of giant elliptical galaxies. They interpret these signatures as evidence of a steeper IMF at low stellar masses. In addition, Cappellari et al. (2012) measure significant variations of the IMF within a large sample of early-type galaxies. They infer this by disentangling the dark matter and stellar contributions to the galaxies' gravitational fields using integral-field stellar kinematics and dynamical modeling.

The variations in the stellar mass-to-light ratios for their galaxies seem to rule out the universality of the IMF.

On the other hand, Smith (2014) critically compare the results of the spectroscopic method of van Dokkum & Conroy (2010, 2011, 2012) and Conroy & van Dokkum (2012) and the dynamics-based one by Cappellari et al. (2012). They find that while both methods agree in suggesting a global bottom-heaviness of the IMF over their respective observational samples, a detailed comparison for the giant ellipticals in common between the samples shows large systematics and lack of agreement, however, thus casting doubts on the individual results. Furthermore, Newman et al. (2017) perform a similar comparison between dynamical, spectroscopic, and, in addition, lensing-based mass estimates for a sample of three nearby giant ellipticals. They show that if the underlying IMF is parametrized as a single or broken power law, the results for the three methods disagree for two out of the three galaxies in the sample. The tension between the approaches is partly relieved when considering more flexible, including non-parametric, IMF forms.

While the implications of a possible, even though not yet completely proven, non-universality of the IMF are intriguing, they are based on indirect evidence from integrated measurements. As such, a complementary approach, using direct counts of resolved low-mass stellar populations, is certainly necessary. Thanks to, and only with, the *Hubble Space Telescope* (*HST*), such a study is possible within the system

* Based on observations made with the NASA/ESA *Hubble Space Telescope*, obtained at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-26555. These observations are associated with program GO-12549.

of Milky Way satellites, including the Magellanic Clouds and several smaller dwarfs. The Milky Way satellites strongly differ from the Galaxy in a multitude of aspects (e.g., morphology, metallicity, and star formation history) and therefore are the ideal targets for a study of the possible IMF variations among different types of galaxies. Several recent studies have attempted to reach the faint end of the IMF in different Milky Way satellites. These studies usually adopt a single power-law model for the IMF, parametrized by a slope, and compare their best-fit results to an average value for the Milky Way, typically -2.3 to -2.35 (Salpeter 1955; Kroupa 2001). The IMF in our Galaxy is observed to experience a flattening, or turnover, below $\sim 0.5 M_{\odot}$ (Kroupa 2001; Chabrier 2003), so results of these studies need to be taken with great care when attempting a comparison with the Milky Way. The single power-law parametrization is convenient, however, and facilitates comparison with, e.g., the extragalactic studies, and we adopt it throughout this paper, together with a log-normal parameterization.

For example, Kalirai et al. (2013), using the Advanced Camera for Surveys (ACS) on board the *HST*, found that the IMF in the Small Magellanic Cloud (SMC) outskirts has a slope of $-1.9^{+0.15}_{-0.10}$ (3σ uncertainty), shallower than the typical Salpeter value of -2.35 . Moreover, their analysis suggests that a single power law can be used to fit their observations down to the completeness limit of $0.37 M_{\odot}$ (i.e., no turnover or flattening). However, Offner et al. (2014) claim that this result is not in contrast with an underlying log-normal IMF (Chabrier 2003), which cannot be ruled out at a 2σ level. Recently, Geha et al. (2013), also using *HST*/ACS, found that the sub-solar IMF slope in the ultra-faint dwarfs (UFDs) Hercules and Leo IV (slopes and 1σ uncertainties of $1.2^{+0.4}_{-0.5}$ and $1.3^{+0.8}_{-0.8}$, respectively) is much shallower than the slope for the Galaxy, which seems to confirm a trend in mass function slope versus total host galaxy mass (with steeper slopes at higher galaxy mass). However, as with the SMC, while the single power-law model yields very different results between the UFDs and the Milky Way, a log-normal IMF with the same parameters as for the Galaxy cannot be completely ruled out for the UFDs.

In this paper we expand the study by Geha et al. (2013) by adding four more UFDs to the sample of Local Group galaxies with a well-characterized IMF down to the low-mass stellar regime. The four additional UFDs are Boötes I (Boo I), Canes Venatici II (CVn II), Coma Berenices (Com Ber), and Ursa Major I (UMA I). These four UFDs are observed within the same *HST*/ACS program (*HST* GO Program ID: 12549, PI: T. Brown) as Hercules and Leo IV, and for consistency, we include these two previously studied UFDs in our new analysis. Other than just expanding the sample, we have improved our analysis by developing a new fitting technique. Our results for Hercules and Leo IV are consistent with those of Geha et al. (2013), although our refined analysis produces more robust uncertainty estimates, which are also smaller in Leo IV.

2. The Data

The data used in the current analysis have already been used in Brown et al. (2014) and is also available online⁹. We refer the reader to that paper for a description of the observations and data reduction steps. In brief, we use deep optical *HST*/ACS

images taken in the F606W and F814W filters and perform aperture (at the bright end, stars with signal-to-noise ratios $\gtrsim 100$) and point-spread function (PSF) fitting photometry (for all other stars). Artificial star tests are also performed to recover the magnitude-error distributions. These test are used in the rest of the paper for synthetic color-magnitude diagram (CMD) realizations. With respect to the Brown et al. (2014) paper, we further select the data to work with a sample that excludes most contaminants (background galaxies and field stars). The CMDs of all six UFDs are shown in Figure 1.

The selection process consists of several steps:

1. We first identify regions of the CMD that are obviously far from the galaxy's main sequence; these regions are visible in the figure. Stars in these regions are excluded.
2. We identify the base of the red giant branch (RGB) and exclude stars brighter than that point. The reason for doing this is to avoid problems with small color mismatches between the models and data. Such mismatches would systematically affect the fitting procedure. Excluding RGB stars has no impact on the IMF fitting for two reasons: first, these stars are few; second, these stars all have very similar masses, and these are very similar to the turnoff mass (not excluded), such that excluding them reduces the dynamical mass range for IMF fitting by less than a percent. The base of the RGB is identified by eye, using a comparison between observations and synthetic CMDs to identify the departure in color between the observed and predicted RGB locus.
3. We identify a faint limit below which we exclude all stars. This limit has a twofold motivation: it accounts for the increase in scatter from photometric errors, and it accounts for increasing contamination on the blue side of the main sequence at fainter magnitudes. After experimenting with the Besançon galactic model¹⁰ (Robin et al. 2003), we can attribute this contamination to Galactic halo white dwarfs. Given that we stop considering stars well above where the "plume" of contaminants approaches the galaxies' main sequences (in color), we do not attempt to better model such contaminants. We prefer to lose some of the faint UFD stars as the purity of the sample is deemed here more important than its completeness. At the same time, the data 50% completeness limit is typically within 0.1 mag of the adopted faint limit, and with a mass-magnitude relation of $0.1 M_{\odot} \text{ mag}^{-1}$ at $0.5 M_{\odot}$ (approximately the mass limit for all the six UFDs), the corresponding loss in mass range depth is less than $0.01 M_{\odot}$, thus negligible.
4. After applying all the above selection criteria, the data are binned in 0.1 mag wide bins in m_{814} and a clipping in $m_{606}-m_{814}$ color is applied, using the `scipy.stats.sigmaclip` module. The mean and standard deviation, σ , of the color of the stars in each bin is computed and stars are clipped if their colors differ more than 3.5σ from the mean. The process is iterated until no more stars fall outside the 3.5σ interval. The stars excluded in this way appear as blue dots outside the shaded area in Figure 1.
5. The same criteria are applied to the artificial star tests.

Table 1 reports the total numbers of selected stars and the brightness limits we imposed in the selection.

⁹ <http://archive.stsci.edu/doi/resolve/resolve.html?doi=10.17909/T99H5G>.

¹⁰ Available at <http://model.obs-besancon.fr/>.

Figure 1. Color–magnitude diagrams for the six ultra-faint dwarfs. The dashed lines represent the bright limit at the base of the red giant branch and the faint limit where both contamination and photometric errors become important. The yellow shading represents regions where no main-sequence member stars are expected. The stars that are used in the analysis are shown in black, while those that are excluded are plotted in blue. Stars are excluded if they are within the orange areas, above the bright limit and below the faint limit. Additionally, for the stars between the limits and outside the orange regions, a 3σ clipping in color—at fixed magnitude—is used to exclude outliers.

Table 1
Number of Selected Stars for Each Ultra-faint Dwarf Galaxy

Galaxy	Selected Stars [number]	Bright Limit [m_{814}]	Faint Limit [m_{814}]
Boo I	2698	22.5	27.5
CVn II	1727	24.5	29.2
Com Ber	2011	22.0	27.25
Hercules	2457	24.5	28.5
Leo IV	1016	24.5	28.75
UMa I	1678	23.7	27.6

3. The Stellar Models

We adopt the same α -enhanced stellar models as in Brown et al. (2014), computed using the Victoria-Regina evolutionary code (see VandenBerg et al. (2014), and references therein), in the metallicity range $-4.0 < [\text{Fe}/\text{H}] < -1.0$ dex. These models have $[\alpha/\text{Fe}] = +0.4$ dex, and a further oxygen enhancement, $\Delta[\text{O}/\text{Fe}]$ that increases at decreasing $[\text{Fe}/\text{H}]$ values (details in Section 3.2 of Brown et al. 2014). In the current work, however, we are very close to the limit where such models with enhanced abundances have been calculated, namely $0.4 M_{\odot}$. Having models for masses below $0.4 M_{\odot}$ is important for the simulation of binary stars that have a component less massive than this limit. At the same time, due

to photometry errors, some stars with $M < 0.4 M_{\odot}$ could be up-scattered above the detection threshold, thus, even for single stars it is important to have models that reach slightly below the average mass detection limit. We extend the grid of models by computing α -enhanced models without extra $\Delta[\text{O}/\text{Fe}]$. These models are then “patched” by forcing them to agree with the oxygen-enhanced ones at $0.4 M_{\odot}$, and by assuming that a rigid shift applies to the model for all masses below $0.4 M_{\odot}$. This is not a perfect solution, but the shift amounts to at most a few hundredths of magnitudes in a region of the CMD that is affected by a much larger photometric scatter. Indeed, this solution is only an approximation for those stars that are scattering from the edge of the detection limit into our fitting region.

4. Fitting Technique

The fitting technique used in this work is based on proposing possible stellar population parameter sets, simulating individual CMDs according to those parameter sets, comparing the simulated CMDs to the observed CMD, and keeping or discarding the proposed stellar population parameter sets according to the result of the comparisons. These comparisons are made in a way that ensures that the distribution of the kept parameter sets is close to the probability distribution over stellar population parameters defined in Gennaro et al. (2015).

We use this new technique instead of the one described in the above work or CMD template fitting (see e.g., Harris & Zaritsky 2001; Dolphin 2002; Ng et al. 2002; Vergely et al. 2002; Cignoni et al. 2006; Aparicio & Hidalgo 2009) because for our problem, repeatedly simulating individual CMDs is less computationally intensive than computing the tables of likelihood functions required for the Gennaro et al. (2015) approach or the CMD templates required for the template fitting approach.

The mechanics of the fitting procedure can be broken into two largely independent parts: a function for quantitatively comparing CMDs, and an algorithm for using that comparison to explore a probability distribution. An observation of one star or unresolved stellar system is a point in $\mathbb{R}^d$, where $\mathbb{R}$ is the set of real numbers, and d is the number of bands used in the CMD. Comparing CMDs requires a definition of a distance between unordered collections of points in $\mathbb{R}^d$. There are multiple ways of defining such a distance, each with its own advantages and disadvantages. We use the so-called *kernel distance* (Phillips & Venkatasubramanian 2011). The kernel distance can be computed in closed form and fulfills the assumptions of some useful theorems on the convergence of our algorithm for interpreting comparisons between CMDs. We describe the kernel distance in more detail in Section 4.2.

The algorithm we use to interpret comparisons between simulated CMDs and the observed CMD is called approximate Bayesian computation Markov chain Monte Carlo (ABC-MCMC, Marjoram et al. 2003). ABC-MCMC is a specific form of approximate Bayesian computation, or ABC. ABC is an approach to Bayesian inference designed for problems where evaluating the actual posterior probability distribution is computationally infeasible or actually impossible. While astronomers have invented many ABC-like techniques, the conscious use of ABC in astronomy started with Cameron & Pettitt (2012) and Weyant et al. (2013). Knowing that a technique is an example of ABC is useful because there is a large technical literature on the statistical properties of many types of ABC. For example, in Section 4.2 we use a theorem from the ABC literature to affirm our choice of comparison function. We describe ABC in more detail in Section 4.1.

Appendix A shows a series of tests we have performed to empirically validate our fitting technique. In these tests, we simulate catalogs that are meant to resemble our Hercules galaxy data set and attempt to recover the input stellar population parameters. The tests demonstrate that our technique has the ability to accurately recover IMF parameters for both functional forms used in this paper (single power-law IMF and log-normal IMF). The recovered IMF parameters are on average unbiased, even though we are only working with a very limited mass range of $0.5\text{--}0.8 M_{\odot}$. The limited mass range does mean, however, that the variance of our IMF parameter estimates is fairly high; it also means that we cannot select between the two possible functional forms.

4.1. ABC and ABC-MCMC

ABC is a technique for doing inference when generating realizations of a model given input parameters is simple, but evaluating the model likelihood is inconvenient or impossible. In the CMD fitting problem, we are dealing with the inhomogeneous Poisson point process (PPP) model of a resolved stellar population, which was described in Gennaro et al. (2015). In that work, the likelihood of an observed flux

given an intrinsic flux was assumed to be a simple, closed-form function. The assumption of a closed-form likelihood allowed us to numerically compute the integrals involved in the PPP likelihood function (see, e.g., Equations (11)–(14) of Gennaro et al. 2015) in a reasonable amount of time. In this work, we directly use precomputed artificial star tests instead of a closed-form function and incorporate more stellar parameters than were included in our previous work. These two changes make our previous exact approach too time-consuming on our immediately available computational resources.

ABC is usually introduced via the following thought experiment from Rubin (1984). Suppose you have an observation, y , and a model with parameters θ for that observation from which you can generate simulated observations. Randomly select a possible parameter set, θ_i , and then use the model to generate a simulated observation y_i . If the simulated observation and actual observation are exactly equal, keep θ_i ; otherwise, discard it. By repeating this procedure, one can generate draws from the probability distribution of θ given the observations y .

If y is a continuous variable, this procedure is only possible as a thought experiment—the probability of generating an exact match to a continuous variable, such as the flux of a star, is zero. If instead of requiring an exact match we are willing to accept a y_i that is within some distance threshold, or tolerance, τ of the actual measured y , the probability of actually generating accepted θ_i draws is no longer zero. The cost of this computational tractability is that the θ_i draws will now follow an *approximation* to the original probability distribution. This procedure was originally proposed in Pritchard et al. (1999) in the context of human population genetics.

ABC as described so far does not have guidelines on how to efficiently propose θ_i values. To propose θ_i values more efficiently, we use ABC within an MCMC approach (see e.g., Marjoram et al. 2003). ABC can be dropped into an existing MCMC approach by treating the acceptance or rejection of a proposed value within ABC as a noisy estimate of the likelihood function. When deciding whether to accept or reject an MCMC step, the likelihood term in the numerator of the MCMC acceptance ratio is 1 if the proposed y_i is close enough to y and 0 if it is not. The likelihood term in the denominator is 1 by definition, since the ABC likelihood of an accepted value is 1. The rest of the MCMC acceptance ratio is unchanged.

We adopt the Foreman-Mackey et al. (2013) python implementation of the Goodman & Weare (2010) affine-invariant sampler. We start our MCMC runs with high values of the tolerance, τ , and typically run 80 chains for 30 steps. We then use the 90% quantile of the simulated-observed CMD distance values across the 80 walkers to define a new tolerance level. This means that at the beginning of a new iteration of 30 runs, 10% of the walkers will find themselves outside of the acceptable parameters zone. They will be slowly drawn in by mixing with the now 90% of good walkers. As soon as all of the walkers are in an acceptable state, the iteration is completed, a new tolerance computed, and a new iteration begins. We continue to adapt the tolerance until the new value differs from the old by less than 0.2%. At this point we run the chains for 2500 steps. We further throw away the first 500 steps and thin the chains by a factor of 50, a value that seemed appropriate for our MCMC experiments once the results were obtained and the autocorrelation times estimated. In this way,

we are left with independent draws from a good approximation of the true posterior.

4.2. Kernel Distances

In order to decide whether to accept or reject a proposed parameter set in the ABC procedure described above, we need to compute the distance between a simulated data set and the actual data set; in the terms of our problem, we need to compute a distance between a simulated and observed CMD in order to decide whether to accept a stellar population parameter set. We have chosen the kernel distance for this purpose (Phillips & Venkatasubramanian 2011). If $\mathcal{X}$ and $\mathcal{Y}$ are collections of points in $\mathbb{R}^d$ and $K(\mathbf{x}, \mathbf{y})$ is a *kernel function*, i.e., a function that maps pairs of inputs to the interval $[0, 1]$, then the square of the kernel distance between $\mathcal{X}$ and $\mathcal{Y}$ is

$$\rho_K^2(\mathcal{X}, \mathcal{Y}) \equiv \sum_{\mathbf{x} \in \mathcal{X}} \sum_{\mathbf{x}' \in \mathcal{X}} K(\mathbf{x}, \mathbf{x}') + \sum_{\mathbf{y} \in \mathcal{Y}} \sum_{\mathbf{y}' \in \mathcal{Y}} K(\mathbf{y}, \mathbf{y}') - 2 \sum_{\mathbf{x} \in \mathcal{X}} \sum_{\mathbf{y} \in \mathcal{Y}} K(\mathbf{x}, \mathbf{y}). \quad (1)$$

If the kernel function is positive definite, the corresponding kernel distance is a metric, meaning that it is non-negative, symmetric, obeys the triangle inequality, and evaluates to 0 if and only if $\mathcal{X}$ is equal to $\mathcal{Y}$. We have chosen the Gaussian kernel, which is positive definite.

In our implementation the choice of sigma depends on the data. We first standardize our CMD into the $(0, 1) \times (0, 1)$ rectangle. Working in this system ensures that similar weight is given to the magnitude and the color dimension, which are naturally quite inhomogeneous (the magnitude range of our observations being about 10 times larger than the color range). We then compute the average minimum distance between the points (i.e., the average of the minimum Euclidean distance between each point and all other points). In the case of a uniform distribution this number is approximately $1/N^2$, but our data are not uniform. We then set sigma to three times this value. The idea is that σ is chosen so that within each kernel there will be a few neighbors. If σ is too large, the kernel distance essentially returns the square root of the square difference in the number of simulated and observed stars; if σ is too small, it returns the square root of the square sum of the number of simulated and observed stars that do not exactly line up.

While we are not aware of any prior examples of ABC being done using the kernel distance, Bernton et al. (2017) have performed an empirical and theoretical study of ABC with the Wasserstein, or Earth mover's, distance. The Wasserstein distance is a metric on collections of points that has good theoretical properties but is non-trivial to actually compute. Bernton et al. (2017) establish conditions under which the Wasserstein distance-ABC probability distribution converges to the actual probability distribution as the tolerance τ is reduced (Proposition 5.1). The kernel distance meets the two assumptions required by that proposition, meaning that the kernel distance-ABC probability distribution also converges to the actual probability distribution. We have not determined whether the other properties of Wasserstein distance-ABC also carry over to kernel distance-ABC; this is left to future work.

4.3. IMF Parametrization

We assume two possible forms for the IMF: a single power law, and a log-normal distribution. Our parametrization describes the system mass function, not the individual stars' mass function. We only consider stellar systems made of one or two stars—no higher order multiples are included. We thus have to specify a model for binarity. This model is very simple: one parameter is used to characterize the binary fraction, $\text{bf} = \frac{\# \text{ binaries}}{\# \text{ singles} + \# \text{ binaries}}$. We also define a fixed distribution (i.e., one we do not fit) for the binary mass ratio, i.e., $q = \frac{m_i}{m_p}$, assuming a uniform distribution for q between 0 and 1.

Another unknown in the problem is the number of formed stars, which correspond to the intensity of the underlying inhomogeneous PPP (Gennaro et al. 2015). The intensity is the mean number of formed stars, with the actual number of formed stars being drawn from a Poisson distribution with mean equal to the intensity. The intensity can be seen as an overall, unknown normalization and thus depends on the range of masses over which the IMF model is specified; in our case, this range is $(0.35, 8) M_\odot$.

We thus have three parameters to fit in the case of the single power-law (SPL) model: intensity, power-law slope, and binary fraction; for the log-normal (LN) model, there are four parameters: intensity, binary fraction, m_c , and σ . The mass distributions are specified as

$$p_{\text{SPL}}(m|\alpha) \propto m^\alpha \quad (2)$$

$$p_{\text{LN}}(m|m_c, \sigma) \propto \frac{1}{m} e^{-\frac{1}{2} \left(\frac{\log m - \log(m_c)}{\sigma} \right)^2}. \quad (3)$$

Note that the log-normal formulation above is non-standard because we use the base-10 logarithm (indicated by log) instead of the base- e one (indicated by ln). We keep this notation in order to allow for a direct comparison with the system mass function of Chabrier (2003), where m_c is $0.22 M_\odot$, and σ is 0.57 (see their Equation (18)). To obtain a standard log-normal, the base-10 logarithms in Equation (3) can be replaced by base- e ones, while $\sigma \rightarrow \sigma' = \sigma \times \ln(10)$.

4.4. Assumptions in Generating Synthetic CMDs

The position of a stellar system in the CMD does not only depend on its mass and binary properties, but also on its age, metallicity, distance, and extinction in each band. Rather than simultaneously fit for all these properties, we adopt priors on their distribution.

We use the star formation histories (SFHs) derived for the six UFDs in Brown et al. (2014) as priors on age for our simulations. The SFHs are described as a two-burst model, with each burst being a normal distribution with $\sigma = 0.1$ Gyr, and mean value given by Table 2 of Brown et al. (2014). All the SFHs are dominated by one of the two normal distributions, except for UMa I, for which we have two almost identical bursts of star formation at 14.1 and 11.6 Gyr. The mean ages are in the interval 12.7–13.9 Gyr over the sample of UFDs.

Similarly to the SFHs, we adopt the metallicity distribution functions (MDFs) derived in Brown et al. (2014), based on spectroscopic measurements of $[\text{Fe}/\text{H}]$. Brown et al. (2014) did not assume a functional form for the MDFs, but performed a non-parametric fit to tabulated MDFs sampled at 0.2 dex interval in the $[-4.0, -1.0]$ dex range (see the shaded areas in Figure 3 of Brown et al. 2014).

We also adopt the same A_V values and apparent distance moduli from that paper, together with the assumption that the reddening law can be described as in Cardelli et al. (1989), to derive values of the true distance modulus, and of the extinction in each band, A_{606} and A_{814} . We note that Brown et al. (2014) assume the extinction law by Fitzpatrick (1999); this law and the one by Cardelli et al. (1989) are nearly identical in the wavelength regime bracketed by our photometric bands. Given their distance, it is in practice reasonable to assume that all of the stars in each of the six UFDs are at the same distance, and thus we adopt a single value for the prior on distance. We also assume that all stars are subject to the same reddening, so fixed global values for each UFD are used as priors on the individual photometric bands' extinction.

4.5. Summary of the Method

For each ABC-MCMC step, we have a tuple of parameters describing the underlying PPP: the intensity, Λ_i , one or two parameters describing the IMF distribution, $\theta_{\text{IMF},i}$, and one parameter for the binary fraction, $\theta_{\text{Bin},i}$. We extract an integer number, n_i , of “born” stellar systems from a Poisson distribution with mean equal to Λ_i . We extract n_i values of masses and binary properties from $p(\theta_{\text{IMF},i})$ and $p(\theta_{\text{Bin},i})$, respectively. We also extract n_i ages, metallicities, extinction values in each band, and distances from the immutable prior distributions of Section 4.4.

We then compute the corresponding magnitudes given mass, age, binary fraction, metallicity, extinction, and distance, using the stellar models and the artificial star experiments. Some of the simulated objects will not be observable due to either stellar evolution or incompleteness. We thus end up with $m_i \leq n_i$ stars in the CMD.

We compute the kernel distance between the m_i simulated points and the actual observations (see Section 4.2), and accept/reject the proposed $(\Lambda_i, \theta_{\text{IMF},i}, \theta_{\text{Bin},i})$ tuple according to the ABC method described in Section 4.1. We finally iterate using the ABC-MCMC scheme described in Section 4.1 as well.

5. Results

We run our ABC-MCMC experiments adopting uniform priors on the IMF model parameters. It is interesting to note, however, that for the single power-law case, there is a strong correlation between the intensity and the power-law slope, especially at shallower slopes. Triangle plots showing the pairwise relations between the parameters are presented in Appendix B in Figures 8–13 (these plots make use of the `corner.py` python code, Foreman-Mackey 2016).

The correlation between slope and intensity arises because for a higher total number of stars but a shallower (less negative) slope, the number of observed stars remains almost constant. Given the narrow mass range probed, the shape of the luminosity function (and the CMD) also shows only very minor changes. It is well established, however, that galaxies with fewer stars are more probable than larger galaxies. Thus we explored the idea of assuming a prior on the total number that is flat in logarithm, rather than in actual number. We do not run separate MCMC experiments, but we rather weight the results assigning to each draw a weight that is inversely proportional to the value of the intensity for that draw. This is equivalent to assuming a logarithmic prior in the intensity. For

completeness we report results obtained both ways. We note that in some cases, such as for Boo I or Com Ber, the changes are minor, thus the prior seems to be of little importance here. However, for other galaxies such as CVn II, Hercules, and Leo IV, the difference in the average slope is on the order of 0.2, which is not negligible. In general, the prior, by downweighting the highest intensities, correspondingly downweights the shallowest slopes, thus in all cases, a steeper slope is reported when using the more realistic logarithmic prior.

The results are summarized in Table 2, which reports the best-fit values for the IMF model parameters. A visualization of the results is also given in Figure 2. A more exhaustive summary of the results is given in Appendix C, where Tables 4 and 5 contain not only the best-fit values, but also the 68%, 95%, and 99% credible intervals (CIs) for the parameters. Figure 3 instead shows a comparison of the observed m_{814} luminosity functions and of the luminosity function that can be obtained by using the best-fit parameters to simulate a CMD. This figure also shows the histogram of the masses of the simulated stars, both for those that land in the CMD (the present-day mass function), and for all the extracted ones, including those that are too faint to be observed and those that have evolved off the main sequence (the IMF).

In reference to the mass histograms of Figure 3, we stress that even though the observed mass range for single stars is limited to about 0.5–0.8 $M_\odot$, the IMF model for stellar systems needs to be specified on a wider mass range, 0.35–8.0 $M_\odot$, see Section 4.3. This extension ensures that the secondary stars from more massive systems as well as low-mass single stars with large photometric errors that are up-scattered in the CMD are correctly accounted for. We also enforce that the model range specification is the same for all UFDs. The implication of our assumptions obviously is that if the IMF functional form is very different outside the specified range, then the model parameters will be systematically affected. On the other end, a partial model specification, limited to a very narrow mass range, would equally bias the results by neglecting binary secondaries and low-mass objects.

5.1. Single Power-law Model

Even if we consider the logarithmic prior case, which produces steeper slopes, for all six galaxies, the slope values are less negative (flatter IMF) than the value for the Milky Way disk IMF, -2.35 (Salpeter 1955). By looking at Figure 2 and by reading the values out of Table 4, it is clear that for all six UFDs the Salpeter value falls outside the 68% CI. For Boo I and UMa I, the Salpeter slope is still within the 95% CI; in the case of Com Ber, it falls outside the 95% CI, but barely within the 99%; while for CVn II, Hercules, and Leo IV, the Salpeter slope is outside the 99% CI, with the latter two UFDs having 99% of their IMF slope posterior probability below -2.0 .

We note, however, that a single power-law model for the Milky Way IMF is not a valid description down to very low stellar masses. Kroupa (2001) shows that below 0.5 $M_\odot$, the IMF for the Galactic Field population I stars flattens, changing slope from -2.3 to -1.3 ; he thus prefers a broken power-law model to describe the Milky Way IMF. Equivalently, Chabrier (2003) parametrizes this flattening as a gradual change, described by a log-normal function, below $\sim 1 M_\odot$, while he adopts the Salpeter value above this mass. It could be argued that the low-mass limit for our measurements is near the Kroupa (2001) inflection point (or within the Chabrier 2003

Table 2
Best-fit Parameters for Both the Adopted IMF Models

Galaxy	Single Power Law						Log-normal							
	Intensity		Slope		Bin. Frac		Intensity		$m_c[M_\odot]$		σ		Bin. Frac	
	unif	log	unif	log	unif	log	unif	log	unif	log	unif	log	unif	log
Boo I	8.39e+3	8.28e+3	-1.84	-1.87	0.28	0.28	7.75e+3	7.46e+3	0.33	0.34	0.59	0.56	0.27	0.28
CVn II	1.58e+4	1.36e+4	-1.00	-1.17	0.05	0.06	9.91e+3	9.51e+3	0.56	0.54	0.70	0.67	0.06	0.07
Com Ber	3.71e+3	3.67e+3	-1.66	-1.68	0.61	0.61	3.52e+3	3.45e+3	0.30	0.30	0.67	0.65	0.60	0.60
Hercules	2.75e+4	2.32e+4	-0.93	-1.11	0.12	0.13	1.64e+4	1.60e+4	0.60	0.59	0.69	0.67	0.11	0.11
Leo IV	1.20e+4	1.01e+4	-0.82	-1.01	0.13	0.14	6.40e+3	6.14e+3	0.62	0.61	0.67	0.64	0.14	0.14
UMa I	1.05e+4	9.56e+3	-1.46	-1.58	0.30	0.31	8.19e+3	7.92e+3	0.53	0.51	0.64	0.61	0.30	0.31

Note. Unif and log column labels stand for parameter estimations with a uniform prior in intensity or a prior that is uniform in $\log(\text{Intensity})$, respectively. Full versions of the results, also including the credible intervals, are given in Appendix C, Tables 4 and 5 for the single power law and log-normal IMF models, respectively.

flattening region), thus our results for a single power-law model could be artificially “flatter” if the underlying mass distribution for the UFDs does indeed have a similar break point, or is in truth a log-normal, as cautioned, for example, by El-Badry et al. (2017). On the other hand, studies of Milky Way star-forming regions and young cluster (see, e.g., Da Rio et al. 2012; Andersen et al. 2017) show that for recent star formation in the Milky Way, the distribution of stellar masses can be described by a single power law with a slope of ~ -2.3 down to $0.2\text{--}0.3 M_\odot$ (it should be noted that Da Rio et al. (2012) also show that the uncertainties on pre-main sequence evolutionary models imply very large uncertainties on the derived stellar masses, and thus IMFs, in young star-forming regions). It is still useful to compare our slopes with the traditional -2.3 value for the Milky Way, with the caveat that this single power-law parametrization may not truly apply to the Galaxy throughout the full stellar mass range, with several authors slightly disagreeing about where exactly the IMF changes its behavior.

Regardless of the limitations of a simplified single power-law description, it must be noted that in several cases (CVn II, Hercules, and Leo IV) our best-fit slopes are flatter than even the Kroupa slope for $M < 0.5 M_\odot$, -1.3 . Thus if we take this as the lower limit that could be expected by erroneously adopting a single power-law model for a true underlying broken power-law distribution, our results appear to be even more intriguing. Globally, these results suggest that when the IMF is described as a power law, the UFDs have a less negative IMF slope than the Milky Way in the probed mass range (see also Section 6.3). An extension of our results beyond this range may be prone to systematic errors, however.

5.2. Log-normal Model

A parallel result is that obtained for the log-normal IMF model. For all six UFDs, the mean characteristic system mass is found to be higher than the Milky Way disk value, $m_c = 0.22$ to $0.25 M_\odot$ (Chabrier 2003; Bochanski et al. 2010), with mean values as high as $m_c \sim 0.5 M_\odot$ for CVn II, Hercules, and Leo IV. The shallower slope for the single power-law model translates into a higher characteristic mass, but the result implied by the data is the same: the UFDs have produced relatively fewer low-mass stars with respect to the Milky Way. The widths of the UFDs best-fit log-normal models are in the $0.59\text{--}0.70$ range, similar but also slightly larger than the Milky Way one, $\sigma = 0.57$.

The significance of the deviations of the UFDs values with respect to the Milky Way ones is weaker for the log-normal model than it is for the single power-law one, however. In more detail, the m_c 68% CIs for Boo I, Com Ber, and Uma I contain the Chabrier (2003) value of $0.22 M_\odot$. For CVn II, Hercules, and Leo IV, the same characteristic mass is outside the 68% CIs, but still within the 95% CIs.

Given the limited range of masses we actually observe, we do not interpret this fact as the log-normal being “better” at giving a more universal representation of the IMF. We rather interpret it as this model being more “flexible” in reproducing the data under study. Multiple combinations of (m_c, σ) can be used to obtain similar “slopes” for the log-normal distribution within the observed mass range. The flexibility added by one extra parameter makes it such that it is harder to constrain them individually, again, due to the limited mass range (see El-Badry et al. 2017).

6. Discussion

6.1. Dynamical Effects

The implicit assumption in our analysis is that the distribution of stellar masses in the UFDs did not change with time. Dynamical evolution could potentially affect the IMF measurements and make our assumption invalid. For example, in globular clusters it is observed that two-body relaxation, which causes mass segregation, alters the present-day mass function. Lower-mass stars have higher velocities and are found at larger radii, thus escape the clusters more easily than higher-mass ones. Leigh et al. (2012) claim that when dynamical modeling is taken into account, no significant differences are measured between the IMFs of Milky Way globular clusters, and thus the stellar masses in globular clusters are consistent with being drawn from the same universal IMF. The recent study by Webb et al. (2017), however, argues that at least for the NGC 5466 and NGC 6101 globular clusters, a non-universal IMF may be required to explain the observations.

Two-body relaxation times in globular clusters are on the order of 1 Gyr, thus shorter than the age of the clusters themselves, on the order of 10 Gyr, therefore these systems are completely collisional and strongly affected by internal dynamics. On the other hand, UFDs are systems with total masses, and thus typical internal velocities, similar to globular clusters, while stellar encounters, which are responsible for the relaxation, are much more infrequent, given the smaller number of stars. This is a consequence of the UFDs being dark-matter dominated. The relaxation time, which can be defined as the

Figure 2. Gaussian kernel density estimates of the marginal probability distributions for the fitted parameter for all the UFDs. The kernel width, w , is estimated using Scott’s rule (Scott 2015): $w = \sigma n^{-1/(d+4)}$, with σ being the data variance, n the number of points, d the number of dimensions (always 1 in this case). The orange lines correspond to the best-fit values, i.e., the average of the marginals. The shaded areas indicate credibility intervals, estimated as the smallest intervals containing a (0.6827, 0.9545, and 0.9973) fraction of the probability. The vertical gray dashed lines correspond to the standard Milky Way disk values, slope = -2.3 (Salpeter 1955), for the single power-law model, and $(m_c, \sigma) = (0.22, 0.57)$ for the log-normal model (values from the system mass function of Chabrier 2003). For this figure we use the posterior probability obtained using a flat prior on the logarithm of the intensity.

time it takes for two-body encounters to change a star’s velocity by an amount equal to itself, is thus much longer for UFDs than for globular clusters.

Using Equations (2)–(62) of Spitzer (1987), the relaxation time can be written as

$$t_{\text{rel}} [\text{Gyr}] \sim \frac{3.4}{\ln\left(\frac{b_{\text{max}}}{b_{\text{min}}}\right)} \left(\frac{v}{\text{km s}^{-1}}\right)^3 \left(\frac{m}{M_{\odot}}\right)^{-2} \left(\frac{n}{\text{pc}^{-3}}\right)^{-1}. \quad (4)$$

Twice the system size $2R$ can be used as maximum impact parameter, b_{max} . The minimum impact parameter can be fixed as the distance for which the potential energy in an encounter equals the kinetic energy, $b_{\text{min}} = 2Gm/v^2$, i.e., the radius for which the weak encounters approximation, used to derive Equation (4), fails. Typical numbers for the UFDs are $N \sim 10^4$, $v \sim 4 \text{ km s}^{-1}$, $m \sim 0.6 M_{\odot}$, and $R \sim 100 \text{ pc}$. The Coulomb logarithm evaluates to ~ 13 . Substituting $n = \frac{3N}{4\pi R^3}$, we have

$$t_{\text{rel}} [\text{Gyr}] \sim \frac{1.1}{N} \left(\frac{v}{\text{km s}^{-1}}\right)^3 \left(\frac{m}{M_{\odot}}\right)^{-2} \left(\frac{R}{\text{pc}}\right)^3. \quad (5)$$

With the above typical UFDs numbers, we thus obtain $t_{\text{rel}} \sim 2 \times 10^4 \text{ Gyr}$, much longer than the age of the Universe. It follows that as opposed to globular clusters, dynamics is not expected to affect the measurement of the IMF in UFDs.

6.2. Comparison with Previous Results

Geha et al. (2013) studied the IMF in Hercules and Leo IV using the same data as we used here. Our new results confirm what was found previously. For the single power-law model, the slope values agree very well within the reported 1σ uncertainties (we use the 68% credibility interval as the equivalent measure), as can be seen in Table 3. The table also shows how each paper’s results for the IMF slope compare to the classical values of -2.35 for $M > 0.5 M_{\odot}$.

For Hercules, our results are nearly identical to those of Geha et al. (2013). Comparing our credibility interval to their “sigma” levels in a quantitative way is not straightforward. However, we note that our “steep” limit for the 99% credibility interval is -2.0 , while Geha et al. (2013) report that they rule out Salpeter at the 5.8σ level. For both papers it is extremely unlikely that the masses of the stars in Hercules and in the Milky Way disk were

Figure 3. Simulations obtained using the best-fit values for each UFD for the two IMF models: single power law (left), and log-normal (right). The left columns for each IMF model show the simulated F814W luminosity function (red), compared to the observed one (blue). The right columns within each IMF model case show the full ensemble of system mass values that are drawn to obtain the luminosity function (the IMF, red), as well as the masses that survive into the CMD (the present-day mass function, PDMF, blue). The dotted lines indicate the ratios between the two, i.e., the completeness as a function of system mass; the y-range of completeness values is not reported for clarity, but it extends from 0 (bottom of the y-axis) to 1 (top of the y-axis). Completeness never reaches values of 1 as a function of system mass because there are always combinations of system mass and binary mass ratio that make it such that some systems are unobservable. The drop of completeness above $\sim 0.8 M_{\odot}$ corresponds to the turnoff mass. There are systems, however, that are observable even at system masses higher than the turnoff mass. These are systems where the total mass is higher than the turnoff one, but the individual components are on the main sequence and observable, thus the system is “complete.”

Table 3

Comparison of the Results for Hercules and LeoIV between this Work (TW) and Geha et al. (2013) (G13)

Galaxy	Power Law Slope		Comparison with	
	TW	G13	TW	G13
Hercules	$-1.11^{+0.50}_{-0.46}$	$-1.2^{+0.5}_{-0.4}$	Outside 99% CI	5.8σ
Leo IV	$-1.01^{+0.61}_{-0.44}$	$-1.3^{+0.8}_{-0.8}$	Outside 99% CI	1.9σ

drawn from the same distribution, independent of the adopted parametrization (similar results, although less significant, are obtained for the log-normal model as well).

For LeoIV, the reported best-fit slopes are again very similar and well within the reported (1σ , or 68% CI) uncertainties. However, our current analysis allows a more precise estimate, with uncertainties almost halved, and thus constitutes a significant improvement in terms of strength with which a Galaxy-like IMF slope can be ruled out for a single power-law model. While Geha et al. (2013) reported that a slope of -2.3 could be only ruled out at 1.9σ level, we obtain a 99% credibility interval steeper limit of -2.0 (the same as for Hercules).

Our analysis, even limited to just these two objects, would already strengthen the conclusions of Geha et al. (2013) that UFD galaxies have a different IMF than the Milky Way.

Figure 4. Left: comparison between galaxies of the IMF slope distributions for the single power-law model. Right: comparison between galaxies of the IMF characteristic mass for the log-normal model. Top: complete results. Bottom: results when only MCMC draws with a binary fraction $0.15 < \text{B.F.} < 0.25$ are considered; this excludes Com Ber from the comparison, whose MCMC solution never reaches this low binary fraction range. In all panels, the black dashed line represents the distribution over the whole UFDs ensemble, obtained by combining the individual UFDs MCMC draws. The gray shaded areas represent the credibility intervals, estimated as the smallest intervals containing a (0.6827, 0.9545, and 0.9973) fraction of the probability.

6.3. Ensemble Properties

All the UFDs in our sample show a flatter slope, or alternatively, a higher characteristic mass, when compared to the typical values found for the Milky Way disk and young cluster. The significance of the discrepancy varies with UFD. To evaluate the overall significance for the ensemble of UFDs, we consider the union of their MCMC draws, and look at the resulting global distributions of slope and m_c , shown in Figure 4.

The figure shows that for both slope and m_c , the resulting distribution over the ensemble is not much narrower than the individual distributions. For the single power-law model, the value -2.3 is at the limit of the 95% CI, while the m_c value of 0.22 falls within the 68% one. It is quite clear that the variations between galaxies are comparable to the individual uncertainties, thus combining the draws does not reduce the global uncertainty by the expected square root of six.

In reference to the same figure, it is very interesting to note that most of the scatter in the sample is caused by the differences between Boo I and Com Ber with respect to CVn II, Hercules, and Leo IV, with UMa I somewhat in between the two groups. The common denominator between Boo I and Com Ber is that they are the closest UFDs in our sample with apparent distance moduli of $(m - M)_V = 19.11$ mag and 18.08 mag, respectively, followed by UMa I at 20.10 mag. CVn II, Hercules, and Leo IV have distance moduli of $(m - M)_V = 21.04, 20.92, 20.10$ mag instead.

These differences may suggest that some physical process, whose strength depends on the distance of these satellites from the Milky Way, may be responsible for the observed differences in IMF. However interesting such an interpretation may be, an alternative is that the differences reside in some systematics within the data. Boo I, Com Ber, and UMa I, being closer, subtend a larger angle in the sky and required more *HST*/*ACS* pointings to be observed (5, 12, and 9, respectively). CVn II, Hercules, and Leo IV required only 1, 2, and 1 pointings. The wider observations translate into a larger number of fore- and background contaminants. Even though

we take great care in the selection process (Section 2), our catalogs might be suffering from a different level of impurity for the different galaxies. It is reasonable to expect that most of the undetected contaminants lie toward the faint end of our data set. In this hypothesis, the effect of such interlopers would be to increase the star counts at low masses, and make the IMF of the closest UFDs steeper.

6.4. IMF Slope and Galaxy Properties

We investigate possible dependencies of the derived IMF slopes on several galactic properties: mean metallicity, mean age, mean individual stellar mass in the observed range, velocity dispersion, total dynamical mass, ellipticity, half-light radius, and absolute magnitude. For each of these quantities, we report the Pearson correlation coefficient, ρ , as well as the p -value. The latter roughly indicates the probability of an uncorrelated system producing data sets that have a Pearson correlation at least as high as the one computed from these data sets. We used the python `scipy.stats.pearsonr` module for such calculations. The p -value cannot be considered completely reliable for our small sample of six galaxies, but we still choose to report it for completeness. The p -value is used to estimate the extent to which the within-sample covariance is consistent with no correlation. It does not quantify the amount of correlation that might be induced by the uncertainties on the individual points, however. Given that these are themselves large, our correlation coefficients cannot be considered extremely robust, but a rather qualitative assessment.

Mean metallicity. The mean metallicities plotted in Figure 5 are derived from the MDFs described in Section 2.2 of Brown et al. (2014). The latter are based on spectra obtained with *Keck*/DEIMOS, also described in the aforementioned paper. The error bars on the $[\text{Fe}/\text{H}]$ values represent the 16% and 84% quantiles of the MDFs, *not* the errors on the mean metallicities themselves. The anticorrelation between IMF slope and mean metallicity is the strongest we observe, and also the one with the lowest p -value, i.e., the lowest chance of being a false

Figure 5. IMF slope trends with several galactic properties. ρ is the Pearson correlation coefficient, and the p -values represents the probability of obtaining an equivalent or higher value of ρ from a sample with the same number of data but no correlation. For the metallicity (top left) the range bars indicate the range of observed metallicities, not errors in the metallicities or the mean metallicity.

positive. The observed correlation suggests that metallicity regulates the star formation process (e.g., by affecting accretion, or by massive star feedback).

Mean age. The mean ages for the stellar populations within each UFD are derived from Table 2 of Brown et al. (2014). These mean values are the weighted average of the ages of a two-burst star formation history model, fitted to the same data as we used here. For all UFDs but UMa I, one of the age components is truly dominant, thus the mean age reported here is close to one of the peaks of the bimodal age distribution model. For UMa I, the mean age is 12.7 Gyr, in between the two best-fit burst of 14.1 and 11.6 Gyr. No correlation is observed between mean age and IMF slope for our sample of UFDs.

Mean stellar mass. These mean mass values are obtained by using the best-fit parameters for the single power-law model IMF and the logarithmic prior to simulate CMDs with a large number of stars. Then, only the systems that are not binaries are considered; they are used to determine the average observed mass, as well as the extrema of the observable mass range. Obviously, the mean mass value anticorrelates with the IMF slope because by construction, steeper IMF slopes correspond to lower average masses, provided that the mass intervals considered are similar. Even though the correlation with mean mass is not truly meaningful, we still plot the slope versus mass, because it is worth noting that the IMF slopes do not correlate with the probed mass interval (shown as the horizontal error bar). For example, the probed intervals for Leo IV and UMa I as well as for Boo I and CVn II are basically the same, but the best-fit slopes differ substantially. This sanity check ensures that at least one type of systematic effect can be ruled out, i.e., the effect of an artificially changing slope that

would be obtained by trying to fit a truly log-normal distribution using power-law models in different intervals.

Velocity dispersion. Velocity dispersion values are obtained from the same spectra for which we obtained the MDFs. We use an approach similar to that of Li et al. (2017). We assume an underlying single-Gaussian distribution for the ensemble of measured line-of-sight velocities and MCMC methods to derive the posterior distribution of both the mean velocity and velocity dispersion, σ_v . We use the median value as the best-fit estimate, and use the 16% and 84% quantiles for the errors. There is some correlation of IMF slope and σ_v , but not as strong as that with metallicity. Moreover, its p -value is quite high. The σ_v values derived here are used to derive the dynamical masses (see below). We caution that further analysis is ongoing and that a more complete and thorough study of the individual velocities as well as of the galaxies' velocity dispersions are the objective of a future paper.

Dynamical mass. The dynamical mass values we use are derived from our velocity dispersion values, using Equation (3) of Simon & Geha (2007):

$$M_{\text{tot}} = 167\beta r_c \sigma_v^2; \quad (6)$$

with $\beta = 8$, with r_c being the King core radius (see Illingworth 1976, for a derivation of the equation). As suggested in Simon & Geha (2007), the King core radius can be related to the measurable Plummer radius by $r_c = 0.64r_{\text{Plummer}}$. We use the tabulated list of Plummer radii in Table 6 of Simon & Geha (2007). The correlation strength of IMF and dynamical mass is similar and has a similar p -value as that with σ_v . As mentioned above, further analysis is ongoing to improve on the velocity

Figure 6. Left: correlation between IMF slope and binary fraction. Right: correlation between m_c , the characteristic mass for a log-normal model and the binary fraction. ρ is the Pearson correlation coefficient, and the p -values represent the probability of obtaining an equivalent or higher value of ρ from a sample with the same number of data, but no correlation.

dispersion estimates, and the simple relation of Equation (6) may not be best suited for all the six galaxies in our sample, making our estimates of the dynamical mass quite uncertain. In the future, more carefully derived estimates will give the ability to confirm or dispute the currently observed (albeit weak) correlation of IMF slope with mass.

Ellipticity, half-light radii, and absolute magnitudes. These values are taken from Martin et al. (2008). No correlation between the IMF slope and any of these parameters is observed.

6.5. Binary Fraction, Slope, and m_c

There appears to be a correlation between the best-fit IMF slope values and binary fraction (see Figure 6). A similar correlation is observed between m_c and the binary fraction for the log-normal model. The fact that both slope and m_c show this correlation with the binary fraction is a consequence of the fact that the slopes and m_c are themselves correlated. Steeper slopes correspond to more bottom-heavy distributions, with lower mean mass, thus lower m_c .

This correlation with binary fraction is striking, and we investigated whether part of it could be spurious in origin. The individual triangle plots of Appendix B show a weak correlation for the (slope, binary fraction) or (m_c , binary fraction) parameter pairs of the individual MCMC runs. Computing the Pearson correlation coefficient for the MCMC draws of the individual galaxies, we find $\rho = (-0.09, -0.25, 0.05, -0.20, -0.12, \text{ and } -0.23)$ for the (slope, binary fraction) pairs and $\rho = (-0.10, 0.01, 0.10, 0.02, -0.03, \text{ and } -0.11)$ for the (m_c , binary fraction) pairs, where the numbers correspond in order to Boo I, CVn II, Com Ber, Hercules, Leo IV, and UMa I. We note that the correlation between value pairs across the UFDs sample (Figure 6) is much stronger than each of the individual correlations in the MCMC solution.

Given their weakness, these intrinsic correlations in the individual posterior distributions may only be partly responsible for the behavior observed in Figure 6. To corroborate this, we show in the bottom panel of Figure 4 the slope and m_c marginal distributions when considering only MCMC draws around a fixed binary fraction value: $0.15 < \text{Binary Fraction} < 0.25$. This range is appropriate for comparison with the system mass function of Chabrier (2003), which assumes a binary fraction of 0.20 for the Milky Way. Figure 4 shows that even when limited to this binary fraction range, our results for individual galaxies do not change significantly, and neither do the results over the UFDs sample. The fact that the results do not change when fixing the binary fraction further demonstrates that the observed correlation of slope and m_c with binary fraction over the UFDs cannot be

solely explained by the intrinsic small correlation of the two parameter pairs within the individual UFDs MCMC draws.

Still, it is hard to understand why some of the UFDs would “prefer” a shallower slope and smaller binary fraction while others prefer the opposite. We do not try to overinterpret the results for the binary fractions themselves. The binary fraction is clearly poorly constrained due to the small number statistics and all uncertainties related to, e.g., the assumed mass ratio distribution. Leaving the binary fraction free to vary allows us to marginalize the pdf for the slope (or equivalently m_c) over this unknown parameter. Thus, rather than fixing the binary fraction to, e.g., the typical Milky Way value, we include it in our calculation to account for the uncertainty induced by its unknown value.

7. Summary and Conclusions

We have derived the parameters for a single power-law and a log-normal model of the IMF of six UFD satellites of the Milky Way. Our results show that in the probed mass range of $0.5\text{--}0.8 M_\odot$, the stellar populations of these objects have a mass distribution that differs from that of the Milky Way disk and young clusters stars, i.e., population I stars with solar metallicity. When using a single power-law model description for the IMF, the UFDs show a flatter slope than the -2.3 for the Galactic disk and open clusters (Salpeter 1955; Kroupa 2001; Da Rio et al. 2012; Andersen et al. 2017). When parametrizing the IMF as a log-normal, the UFDs show a higher characteristic mass than the galactic value of $0.22\text{--}0.25 M_\odot$ (Chabrier 2003; Bochanski et al. 2010). In either case, the data show us that in the $0.5\text{--}0.8 M_\odot$ interval, there are more low-mass stars in the Milky Way relative to the UFDs. The significance of the discrepancy with the Milky Way IMF is stronger for the single power-law model. The two-parameter log-normal functional form is more flexible, thus constraints on the individual parameters, characteristic mass, and width are necessarily less stringent, and we cannot rule out the galactic values at more than 1σ for all of the UFDs. Given the limited mass range probed, we do not attempt a quantitative model comparison between these two adopted parametrizations.

We investigate the possible environmental effects that may be driving this behavior and find that the average galactic metallicity moderately correlates with the IMF slope, while a lower, less significant correlation is observed between IMF slope and the velocity dispersion and the total galaxy mass.

Even though all UFDs have more bottom-light IMFs than the Milky Way disk, we observe variations within the UFDs sample. The farthest UFDs, CVn II, Hercules, and Leo IV, have more bottom-light IMFs than the Milky Way with respect

to the UFDs closest to it, Boo I, Com Ber, and UMa I. This could indicate some physical process that depends on the distance from the Milky Way as the responsible factor for the observed differences. However, we caution that a higher fraction of undetected residual fore- and background contamination may be affecting the closest UFDs, which have greater extent on the sky, and required more *HST*/ACS tiles to be observed. If that is the case, the less contaminated catalogs for the farthest UFDs would be showing a more authentic picture of these small satellites, with a much increased significance in IMF variations with respect to the Galaxy.

Further work is needed to try and increase the sample of studied systems, e.g., by using the *James Webb Space Telescope* and the *Wide-Field Infrared Space Telescope*, thus improving the reliability of the results suggested by our small sample.

We would like to thank Professor Don Vandenberg (University of Victoria, British Columbia, Canada) for providing the oxygen-enhanced models used in this work, as well as the *Fortran* routines used for their interpolation.

Support for program GO-12549 was provided by NASA through a grant from the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-26555.

This work was supported by a NASA Keck PI Data Award, administered by the NASA Exoplanet Science Institute under RSA number 1474359. Data presented herein were obtained at the W.M. Keck Observatory from telescope time allocated to NASA through the agency's scientific partnership with the California Institute of Technology and the University of California. The Observatory was made possible by the generous financial support of the W.M. Keck Foundation. The authors wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Maunakea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain.

E.N.K. acknowledges support from NSF grant AST-1614081.

Appendix A Validating the Fitting Technique

In order to validate our method, we use simulated catalogs on which we run the algorithm described in Section 4. For the simulations, we adopt the star formation history and MDF of

Hercules. We simulated populations with both an underlying single power-law IMF and a log-normal one. All of the simulations have an intensity such that the number of observed stars is about 2000, similar to the numbers for our galaxies. All simulations assume a binary fraction of 30%. For the single power-law case, the adopted slope is -1.5 , while for the log-normal case, we adopted the Chabrier (2003) values for the system IMF, $(m_c, \sigma) = (0.22, 0.57)$.

Figure 7 shows the results for some of these simulations. The left and right panels show results for simulated data of depths comparable to the observations of Hercules. The central panel shows results from an MCMC run on a simulated catalog where observations have been artificially made 3 mag deeper, reaching about $0.15 M_\odot$. The results in this case are obviously more precise than those for shallower observations.

We note that the estimates for the IMF parameters are always unbiased and the true values are easily recovered within the 68% CIs for all cases and all variables. We are confident that provided that the IMF model specification (single power law or log-normal) is the correct description of the underlying IMF for the UFD populations, our method gives robust estimates for the model parameters.

In the single power-law (shallow) case, the estimate of the best-fit value of the intensity falls outside the 68% CI. This is related to the fact that the marginal distributions are very skewed, with a long tail at high intensities. The mean, which we use to estimate the best-fit value, falls far from the mode. The credibility intervals instead typically straddle the mode, because for these pdfs, which are generally unimodal, the smallest interval containing a certain fraction of the probability (our definition of CI) does contain the mode region. We note that for the distributions obtained using the actual observations, this problem does not present itself, even though sometimes the mean is close to one edge of the 68% CI. This apparent problem arises from the fact that describing a complex multidimensional pdf in terms of a single central tendency value per dimension and some estimate of its uncertainty is a less than ideal process, because it can hardly capture the full available information. We have chosen to use the mean as central tendency indicator, with the understanding that the uncertainty in our estimates can only be captured by using the full pdfs, which we always display.

We also note that the typical uncertainty from these tests is similar to the uncertainties on the fitted parameters for our real galaxies, thus our estimated uncertainties are reasonable.

Figure 7. Results of a subset of the tests performed to validate the method. The curves represent kernel density estimates of the marginal distribution of the MCMC draws. The kernel width, w , is estimated using Scott's rule (Scott 2015): $w = \sigma n^{-1/(d+4)}$, with σ being the data variance, n the number of points, and d the number of dimensions (always 1 in this case). The true values are indicated by the dashed magenta vertical lines, and the best-fit values by the solid orange ones. The shaded areas indicate credibility intervals, estimated as the smallest intervals containing a (0.6827, 0.9545, and 0.9973) fraction of the probability. Left: single power-law IMF case for simulations of depths similar to our data. Center: single power-law case, but for observations that are 3 mag deeper; note the shrinking of the uncertainty intervals with respect to the left panel. Right: log-normal IMF case, again for simulations of depths similar to our data. Note that the m_c value for the logarithmic prior on the intensity is recovered so precisely that the magenta and orange lines overlap. Also note that even though the peak value of $0.22 M_\odot$ is outside the observed range (our simulations, like the observation, reach about $0.5 M_\odot$), it can be easily recovered within the uncertainties.

Appendix B

Detailed MCMC Result Plots

Figure 8. Samples from the posterior probability distribution function for the IMF model parameters, as given by our MCMC fitting technique. The case of the Boo I galaxy.

Figure 9. Samples from the posterior probability distribution function for the IMF model parameters, as given by our MCMC fitting technique. The case of the CVn II galaxy.

Figure 10. Samples from the posterior probability distribution function for the IMF model parameters, as given by our MCMC fitting technique. The case of the Com Ber galaxy.

Figure 11. Samples from the posterior probability distribution function for the IMF model parameters, as given by our MCMC fitting technique. The case of the Hercules galaxy.

Figure 12. Samples from the posterior probability distribution function for the IMF model parameters, as given by our MCMC fitting technique. The case of the Leo IV galaxy.

Figure 13. Samples from the posterior probability distribution function for the IMF model parameters, as given by our MCMC fitting technique. The case of the UMa I galaxy.

Appendix C

Full Tables of Results, Including Credible Intervals

Table 4
Results for the Single Power-law IMF Model

Galaxy	Interval	Intensity				Slope				Binary Fraction			
		uniform		logarithmic		uniform		logarithmic		uniform		logarithmic	
Boo I	Avg.	8.39e+3		8.28e+3		−1.84		−1.87		0.28		0.28	
	68%	7.15e+3	8.81e+3	7.12e+3	8.68e+3	−2.19	−1.50	−2.20	−1.53	0.14	0.40	0.14	0.40
	95%	6.77e+3	1.05e+4	6.81e+3	1.03e+4	−2.44	−1.20	−2.44	−1.21	0.02	0.52	0.02	0.52
	99%	6.52e+3	1.32e+4	6.30e+3	1.22e+4	−2.59	−0.83	−2.71	−0.97	0.01	0.65	0.01	0.65
CVn II	Avg.	1.58e+4		1.36e+4		−1.00		−1.17		0.05		0.06	
	68%	7.63e+3	1.82e+4	6.93e+3	1.49e+4	−1.31	−0.43	−1.41	−0.45	0.00	0.07	0.00	0.08
	95%	6.93e+3	2.89e+4	6.93e+3	2.53e+4	−2.02	−0.32	−2.02	−0.36	0.00	0.14	0.00	0.15
	99%	6.93e+3	3.46e+4	6.93e+3	3.46e+4	−2.22	−0.16	−2.22	−0.16	0.00	0.18	0.00	0.18
Com Ber	Avg.	3.71e+3		3.67e+3		−1.66		−1.68		0.61		0.61	
	68%	3.24e+3	4.01e+3	3.19e+3	3.94e+3	−1.89	−1.40	−1.94	−1.45	0.47	0.75	0.47	0.75
	95%	2.93e+3	4.52e+3	2.92e+3	4.45e+3	−2.14	−1.18	−2.15	−1.19	0.36	0.88	0.37	0.88
	99%	2.88e+3	5.25e+3	2.88e+3	5.25e+3	−2.37	−1.06	−2.37	−1.06	0.28	0.98	0.28	0.99
Hercules	Avg.	2.75e+4		2.32e+4		−0.93		−1.11		0.12		0.13	
	68%	1.37e+4	3.05e+4	1.37e+4	2.50e+4	−1.36	−0.33	−1.57	−0.61	0.00	0.15	0.00	0.16
	95%	1.26e+4	5.43e+4	1.26e+4	4.44e+4	−1.80	−0.02	−2.00	−0.32	0.00	0.27	0.00	0.29
	99%	1.31e+4	7.56e+4	1.26e+4	7.56e+4	−2.00	−0.02	−2.00	−0.02	0.00	0.34	0.00	0.34
Leo IV	Avg.	1.20e+4		1.01e+4		−0.82		−1.01		0.13		0.14	
	68%	5.28e+3	1.38e+4	5.22e+3	1.12e+4	−1.15	−0.17	−1.45	−0.40	0.00	0.16	0.00	0.18
	95%	5.28e+3	2.29e+4	5.18e+3	1.98e+4	−1.83	−0.13	−1.82	−0.20	0.00	0.31	0.01	0.34
	99%	5.18e+3	2.95e+4	5.18e+3	2.85e+4	−1.86	−0.00	−2.00	−0.01	0.00	0.38	0.00	0.38
UMa I	Avg.	1.05e+4		9.56e+3		−1.46		−1.58		0.30		0.31	
	68%	7.17e+3	1.07e+4	7.16e+3	9.72e+3	−2.15	−0.82	−2.21	−1.00	0.09	0.43	0.11	0.44
	95%	6.75e+3	1.86e+4	6.79e+3	1.50e+4	−2.71	−0.22	−2.75	−0.44	0.00	0.61	0.00	0.62
	99%	6.82e+3	3.17e+4	6.63e+3	2.99e+4	−3.02	−0.01	−3.03	−0.06	0.00	0.74	0.00	0.74

Note. Uniform and logarithmic column labels stand for parameter estimations with a uniform prior in intensity or a prior that is uniform in $\log(\text{Intensity})$, respectively. The credible intervals are defined as containing (0.6827, 0.9545, and 0.9973) of the total number of draws from the posterior (weighted by intensity^{−1} in the case of the logarithmic prior).

Table 5
Results for the Log-normal IMF Model

Galaxy	Interval	Intensity				m_c				σ				Binary Fraction			
		uniform		logarithmic		uniform		logarithmic		uniform		logarithmic		uniform		logarithmic	
Boo I	Avg.	7.75e+3		7.46e+3		0.33		0.34		0.59		0.56		0.27		0.28	
	68%	6.40e+3	9.39e+3	5.78e+3	8.71e+3	0.13	0.51	0.12	0.52	0.29	0.80	0.29	0.80	0.13	0.38	0.13	0.38
	95%	4.89e+3	1.07e+4	4.69e+3	1.05e+4	0.05	0.66	0.04	0.65	0.21	0.99	0.17	0.97	0.03	0.50	0.04	0.52
	99%	4.13e+3	1.14e+4	4.13e+3	1.14e+4	0.01	0.77	0.01	0.76	0.10	0.99	0.10	0.99	0.01	0.65	0.01	0.65
CVn II	Avg.	9.91e+3		9.51e+3		0.56		0.54		0.70		0.67		0.06		0.07	
	68%	8.32e+3	1.21e+4	7.34e+3	1.14e+4	0.42	0.98	0.25	0.82	0.58	0.96	0.58	1.00	0.00	0.08	0.00	0.08
	95%	6.04e+3	1.33e+4	5.55e+3	1.29e+4	0.11	0.99	0.10	0.99	0.32	1.00	0.28	0.99	0.00	0.19	0.00	0.19
	99%	5.04e+3	1.42e+4	5.02e+3	1.40e+4	0.03	1.00	0.02	1.00	0.22	1.00	0.20	1.00	0.00	0.24	0.00	0.24
Com Ber	Avg.	3.52e+3		3.45e+3		0.30		0.30		0.67		0.65		0.60		0.60	
	68%	2.99e+3	3.93e+3	2.92e+3	3.84e+3	0.04	0.41	0.03	0.39	0.55	0.98	0.49	0.94	0.44	0.71	0.45	0.72
	95%	2.63e+3	4.52e+3	2.56e+3	4.38e+3	0.01	0.66	0.02	0.65	0.31	1.00	0.29	1.00	0.35	0.88	0.35	0.88
	99%	2.51e+3	4.95e+3	2.43e+3	4.95e+3	0.01	0.95	0.01	0.95	0.18	1.00	0.18	1.00	0.30	0.97	0.30	0.97
Hercules	Avg.	1.64e+4		1.60e+4		0.60		0.59		0.69		0.67		0.11		0.11	
	68%	1.43e+4	1.91e+4	1.41e+4	1.91e+4	0.46	0.94	0.46	0.95	0.58	1.00	0.49	0.94	0.00	0.13	0.01	0.14
	95%	1.16e+4	2.12e+4	1.04e+4	2.03e+4	0.16	0.99	0.15	0.98	0.36	1.00	0.34	1.00	0.00	0.25	0.00	0.25
	99%	9.90e+3	2.23e+4	9.78e+3	2.23e+4	0.05	1.00	0.05	1.00	0.24	1.00	0.22	1.00	0.00	0.36	0.00	0.36
Leo IV	Avg.	6.40e+3		6.14e+3		0.62		0.61		0.67		0.64		0.14		0.14	
	68%	5.72e+3	7.84e+3	5.18e+3	7.59e+3	0.48	0.99	0.47	0.98	0.59	1.00	0.51	0.98	0.02	0.18	0.02	0.19
	95%	4.38e+3	8.36e+3	3.77e+3	8.39e+3	0.17	1.00	0.15	0.99	0.27	1.00	0.22	1.00	0.00	0.34	0.00	0.35
	99%	1.92e+3	8.51e+3	1.92e+3	8.48e+3	0.05	1.00	0.05	1.00	0.09	1.00	0.09	1.00	0.00	0.39	0.00	0.39
UMa I	Avg.	8.19e+3		7.92e+3		0.53		0.51		0.64		0.61		0.30		0.31	
	68%	6.88e+3	9.77e+3	6.19e+3	9.14e+3	0.20	0.77	0.19	0.74	0.49	0.94	0.40	0.88	0.08	0.40	0.12	0.45
	95%	5.68e+3	1.11e+4	5.12e+3	1.10e+4	0.12	0.99	0.10	0.98	0.27	1.00	0.23	1.00	0.02	0.59	0.01	0.59
	99%	4.14e+3	1.15e+4	4.02e+3	1.15e+4	0.03	1.00	0.02	1.00	0.13	1.00	0.12	1.00	0.00	0.67	0.00	0.67

Note. Uniform and logarithmic column labels stand for parameter estimations with a uniform prior in intensity or a prior that is uniform in $\log(\text{Intensity})$, respectively. The credible intervals are defined as containing (0.6827, 0.9545, and 0.9973) of the total number of draws from the posterior (weighted by intensity⁻¹ in the case of the logarithmic prior).

ORCID iDs

Mario Gennaro <https://orcid.org/0000-0002-5581-2896>
 Thomas M. Brown <https://orcid.org/0000-0002-1793-9968>
 Puragra Guhathakurta <https://orcid.org/0000-0001-8867-4234>
 Evan N. Kirby <https://orcid.org/0000-0001-6196-5162>
 Jason Tumlinson <https://orcid.org/0000-0002-7982-412X>
 Luis C. Vargas <https://orcid.org/0000-0001-7155-0040>

References

- Andersen, M., Gennaro, M., Brandner, W., et al. 2017, *A&A*, **602**, A22
 Aparicio, A., & Hidalgo, S. L. 2009, *AJ*, **138**, 558
 Bastian, N., Covey, K. R., & Meyer, M. R. 2010, *ARA&A*, **48**, 339
 Berntson, E., Jacob, P. E., Gerber, M., & Robert, C. P. 2017, arXiv:1701.05146
 Bochanski, J. J., Hawley, S. L., Covey, K. R., et al. 2010, *AJ*, **139**, 2679
 Brown, T. M., Tumlinson, J., Geha, M., et al. 2014, *ApJ*, **796**, 91
 Cameron, E., & Pettitt, A. N. 2012, *MNRAS*, **425**, 44
 Cappellari, M., McDermid, R. M., Alatalo, K., et al. 2012, *Natur*, **484**, 485
 Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, *ApJ*, **345**, 245
 Chabrier, G. 2003, *PASP*, **115**, 763
 Cignoni, M., Degl'Innocenti, S., Prada Moroni, P. G., & Shore, S. N. 2006, *A&A*, **459**, 783
 Conroy, C., & van Dokkum, P. G. 2012, *ApJ*, **760**, 71
 Da Rio, N., Robberto, M., Hillenbrand, L. A., Henning, T., & Stassun, K. G. 2012, *ApJ*, **748**, 14
 Dolphin, A. E. 2002, *MNRAS*, **332**, 91
 El-Badry, K., Weisz, D. R., & Quataert, E. 2017, *MNRAS*, **468**, 319
 Fitzpatrick, E. L. 1999, *PASP*, **111**, 63
 Foreman-Mackey, D. 2016, *JOSS*, **24**
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
 Geha, M., Brown, T. M., Tumlinson, J., et al. 2013, *ApJ*, **771**, 29
 Gennaro, M., Tchernyshyov, K., Brown, T. M., & Gordon, K. D. 2015, *ApJ*, **808**, 45
 Goodman, J., & Weare, J. 2010, *Communications in Applied Mathematics and Computational Science*, **5**, 65
 Harris, J., & Zaritsky, D. 2001, *ApJS*, **136**, 25
 Illingworth, G. 1976, *ApJ*, **204**, 73
 Kalirai, J. S., Anderson, J., Dotter, A., et al. 2013, *ApJ*, **763**, 110
 Kroupa, P. 2001, *MNRAS*, **322**, 231
 Kroupa, P. 2002, *Sci*, **295**, 82
 Leigh, N., Umbreit, S., Sills, A., et al. 2012, *MNRAS*, **422**, 1592
 Li, T. S., Simon, J. D., Drlica-Wagner, A., et al. 2017, *ApJ*, **838**, 8
 Marjoram, P., Molitor, J., Plagnol, V., & Tavaré, S. 2003, *PNAS*, **100**, 15324
 Martin, N. F., de Jong, J. T. A., & Rix, H.-W. 2008, *ApJ*, **684**, 1075
 Newman, A. B., Smith, R. J., Conroy, C., Villaume, A., & van Dokkum, P. 2017, *ApJ*, **845**, 157
 Ng, Y. K., Brogt, E., Chiosi, C., & Bertelli, G. 2002, *A&A*, **392**, 1129
 Offner, S. S. R., Clark, P. C., Hennebelle, P., et al. 2014, in *Protostars and Planets VI*, ed. H. Beuther et al. (Tucson, AZ: Univ. of Arizona Press), **53**
 Phillips, J. M., & Venkatasubramanian, S. 2011, CoRR, arXiv:1103.1625
 Pritchard, J. K., Seielstad, M. T., Perez-Lezaun, A., & Feldman, M. W. 1999, *Mol. Biol. Evol.*, **16**, 1791
 Robin, A. C., Reylé, C., Derrière, S., & Picaud, S. 2003, *A&A*, **409**, 523
 Rubin, D. B. 1984, *AnSta*, **12**, 1151
 Salpeter, E. E. 1955, *ApJ*, **121**, 161
 Scott, D. W. 2015, *Multivariate Density Estimation: Theory, Practice, and Visualization* (Hoboken, NJ: John Wiley and Sons)
 Simon, J. D., & Geha, M. 2007, *ApJ*, **670**, 313
 Smith, R. J. 2014, *MNRAS*, **443**, L69
 Spitzer, L. 1987, *Dynamical Evolution of Globular Clusters* (Princeton, NJ: Princeton Univ. Press)
 VandenBerg, D. A., Bergbusch, P. A., Ferguson, J. W., & Edvardsson, B. 2014, *ApJ*, **794**, 72
 van Dokkum, P. G., & Conroy, C. 2010, *Natur*, **468**, 940
 van Dokkum, P. G., & Conroy, C. 2011, *ApJL*, **735**, L13
 van Dokkum, P. G., & Conroy, C. 2012, *ApJ*, **760**, 70
 Vergely, J.-L., Köppen, J., Egret, D., & Bienaymé, O. 2002, *A&A*, **390**, 917
 Webb, J. J., Vesperini, E., Dalessandro, E., et al. 2017, *MNRAS*, **471**, 3845
 Weyant, A., Schafer, C., & Wood-Vasey, W. M. 2013, *ApJ*, **764**, 116