

A CLEAR AGE–VELOCITY DISPERSION CORRELATION IN ANDROMEDA’S STELLAR DISK

CLAIRE E. DORMAN¹, PURAGRA GUHATHAKURTA¹, ANIL C. SETH², DANIEL R. WEISZ^{3,7}, ERIC F. BELL⁴, JULIANNE J. DALCANTON³,
 KAROLINE M. GILBERT⁵, KATHERINE M. HAMREN¹, ALEXIA R. LEWIS³, EVAN D. SKILLMAN⁶, ELISA TOLOBA¹, AND
 BENJAMIN F. WILLIAMS³

¹UCO/Lick Observatory, University of California at Santa Cruz, 1156 High Street, Santa Cruz, CA 95064, USA; cdorman@ucolick.org, raja@ucolick.org,
khamren@ucolick.org, toloba@ucolick.org

²Department of Physics & Astronomy, University of Utah, Salt Lake City, UT 84112, USA; aseth@astro.utah.edu

³Department of Astronomy, University of Washington, Box 351580, Seattle, WA 98195, USA; dweisz@astro.washington.edu, jd@astro.washington.edu,
arlewis@astro.washington.edu, ben@astro.washington.edu

⁴Department of Astronomy, University of Michigan, 500 Church Street, Ann Arbor, MI 48109, USA; ericbell@umich.edu

⁵Space Telescope Science Institute, Baltimore, MD 21218, USA; kgilbert@stsci.edu

⁶Minnesota Institute for Astrophysics, University of Minnesota, Minneapolis, MN 55455, USA; skillman@astro.umn.edu

Received 2014 October 9; accepted 2015 February 11; published 2015 April 9

ABSTRACT

The stellar kinematics of galactic disks are key to constraining disk formation and evolution processes. In this paper, for the first time, we measure the stellar age–velocity dispersion correlation in the inner 20 kpc (~ 3.5 disk scale lengths) of M31 and show that it is dramatically different from that in the Milky Way (MW). We use optical *Hubble Space Telescope*/Advanced Camera for Surveys photometry of 5800 individual stars from the Panchromatic Hubble Andromeda Treasury survey and Keck/DEIMOS radial velocity measurements of the same stars from the Spectroscopic and Photometric Landscape of Andromeda’s Stellar Halo survey. We show that the average line-of-sight dispersion is a steadily increasing function of age exterior to $R = 10$ kpc, increasing from 30 km s^{-1} for the main-sequence stars to 90 km s^{-1} for the red giant branch stars. This monotonic increase implies that a continuous or recurring process contributed to the evolution of the disk. Both the slope and normalization of the dispersion versus age relation are significantly larger than in the MW, allowing for the possibility that the disk of M31 has had a more violent history than the disk of the MW, more in line with Λ cold dark matter predictions. We also find evidence for an inhomogeneous distribution of stars from a second kinematical component in addition to the dominant disk component. One of the largest and hottest high-dispersion patches is present in all age bins and may be the signature of the end of the long bar.

Key words: galaxies: individual (M31) – galaxies: kinematics and dynamics – galaxies: spiral – galaxies: stellar content

1. INTRODUCTION

Most late-type galaxies have multiple disk populations with distinct scale heights. The vertical surface brightness profiles of edge-on spirals are well fit by double exponential profiles (Yoachim & Dalcanton 2006), while dynamically hot “thick disk” populations have been found in both the Milky Way (MW; e.g., Chiba & Beers 2000) and Andromeda (Collins et al. 2011). While it is unclear whether galactic disks typically contain two distinct components or a continuum of progressively thicker populations as argued in Bovy et al. (2012), the structure of the stellar disk is an important key to understanding a galaxy’s formation history.

Thick disks are thought to have formed through some combination of the following three processes. First, stars can be formed in a cold, thin disk and later be dynamically heated by satellite impacts (Quinn et al. 1993; Velazquez & White 1999; Purcell et al. 2010; Tissera et al. 2013) or by internal perturbers such as spiral arms, bars, or scattering by giant molecular clouds (GMCs; Ida et al. 1993). Second, they can be formed “in situ,” from a thick, clumpy gas disk at high redshift whose remnants collapse further over time to form progressively younger, thinner, more metal-rich stellar disks (Bournaud et al. 2009; Forbes et al. 2012). Third, the dynamically hot population can be accreted from satellite galaxies through tidal interactions (Abadi et al. 2003). Each scenario should produce

a different relationship between age, metallicity, and degree of heating. For example, accretion of metal-poor satellites onto a thin disk would create a binary disk structure with thin and thick components, each with a distinct vertical scale height, age distribution, and metallicity distribution. In contrast, a continuous process such as collapse of a clumpy gas disk or heating from frequent low-mass satellite impacts would produce disk layers whose thickness (degree of heating) increases with age.

In nearby, low-mass, edge-on spiral galaxies, stellar populations’ vertical scale heights increase with age over three age bins, suggesting that a continuous process such as disk heating plays a role in the evolution of those galaxies (Seth et al. 2005). An alternative heating diagnostic to scale height is velocity dispersion (σ_v). Measuring σ_v as a function of age can yield even stronger physical constraints on possible disk evolution mechanisms. Dispersion measurements are also less sensitive to dust than are scale height measurements and are possible in galaxies that are not perfectly edge-on.

However, meaningful kinematical measurements are difficult to make. In distant galaxies, kinematics derived from integrated-light spectra cannot differentiate between the contributions from old, intermediate-age, and young stellar populations. In particular, near-infrared light from old red giant branch (RGB) populations is contaminated by flux from younger asymptotic giant branch (AGB) stars. In the MW itself, we can more easily separate stars by age. Velocity

⁷ Hubble Fellow.

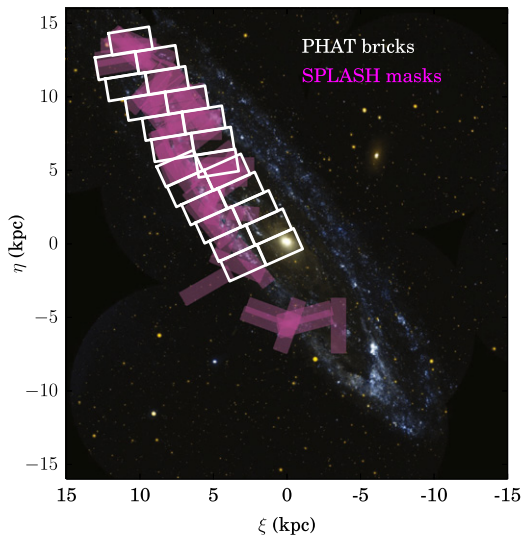


Figure 1. Spatial coverage of spectroscopic (SPLASH) and photometric (PHAT) surveys from which our data are drawn, overlaid on a *GALEX* UV image of M31 for reference. Magenta regions demarcate the 47 Keck/DEIMOS spectroscopic slitmasks used in the SPLASH survey, whereas white rectangles outline the 23 PHAT “bricks” (clusters of *HST* pointings). In this paper, we use only stars in the intersection of these two surveys: those with both PHAT optical photometry and reliable SPLASH-derived radial velocities.

dispersion appears to increase monotonically with age in the solar neighborhood (Nördstrom et al. 2004), but the rest of the disk is obscured to the extent that it is impossible to tell if the solar neighborhood is representative of the entire disk, much less to trace large-scale kinematical structure across the Galaxy.

Additionally, a study of only one disk galaxy (for example, the MW) does not allow us to draw general conclusions about the structure of disk galaxies in general. In the case of the MW, this is an important concern. Λ cold dark matter (CDM) cosmology predicts that galaxies are built up via accretion of smaller satellites. While most collisions between host and satellite occur in the halo of the host, a disk the mass of the MW’s should have experienced at least one encounter with a massive ($\sim 3M_{\text{disk}}$) satellite (Stewart et al. 2008). Such an encounter should significantly thicken and heat the disk (Purcell et al. 2009), but the disk of the MW does not exhibit any such signs of a cosmologically common heating event. Studying the detailed structure of a second galaxy disk—like that of Andromeda (M31)—can provide an important constraint on whether Λ CDM needs to be revisited or whether the MW is simply an outlier in the collision frequency distribution.

M31 is an ideal candidate for mapping a disk’s velocity dispersion as a function of age. It is close enough (785 kpc; McConnachie et al. 2005) that we can isolate and map the velocity dispersions of the RGB, AGB, and young upper main-sequence (MS) populations separately, but distant enough that we can see the entire disk. We take advantage of data from two surveys of the disk-dominated region of M31. The Spectroscopic and Photometric Landscape of Andromeda’s Stellar Halo (SPLASH) survey has used the Keck/DEIMOS multi-object spectrograph to measure radial velocities of thousands of individual bright stars in the inner 20 kpc (~ 3.5 disk scale lengths) of M31 (Gilbert et al. 2009; Dorman et al. 2012, 2013; Howley et al. 2013). Meanwhile, the recently completed Panchromatic Hubble Andromeda Treasury (PHAT) survey, a *Hubble Space Telescope* MultiCycle Treasury (*HST*/MCT) program, has obtained six-filter photometry of 117 million

individual stars in the same portion of the galaxy (Dalcanton et al. 2012; Williams et al. 2014), allowing clean color/magnitude-based separation of RGB, AGB, and MS stars.

Previously, we analyzed the kinematics of *only* the RGB stars in the intersection of the SPLASH and PHAT surveys. We found that 20% of the RGB stars in our survey belonged to a population with spheroid-like kinematics: with a velocity dispersion of 150 km s^{-1} and $v_{\text{rot}}/\sigma \sim 1/3$ (the “inner spheroid”; Dorman et al. 2012). Later, we found that the inner spheroid population has a disk-like luminosity function despite its spheroid-like kinematics (Dorman et al. 2013). In the current paper, we expand our survey to include three younger PHAT photometry-defined age bins in addition to the RGB population, and map the line-of-sight velocity dispersion (LOS) of stars in each age bin. In this paper, we aim to characterize the bulk properties of the disk in each age bin and so do not distinguish between “disk” and “spheroid” members. However, in Section 5, we discuss the evidence for and possible origins of subsets of the RGB and AGB populations with atypical LOS velocities.

This paper is organized as follows. First, in Section 2, we present the data set composed of Keck/DEIMOS radial velocity measurements and *HST* optical photometry of over 8200 individual stars. In Section 3, we separate the stars into four age bins based on their position in the optical color–magnitude diagram (CMD) and estimate RGB photometric metallicities. We also define our smoothed velocity dispersion statistic. We discuss trends and possible biases in Section 4. In Section 5, we discuss the constraints our results place on disk evolution scenarios. Finally, we summarize in Section 6.

2. DATA ACQUISITION AND REDUCTION

Our data set is a subset of two larger surveys. Figure 1 shows the spatial coverage of both overlaid on a *GALEX* image of M31. We start with radial velocity measurements and optical *HST* photometry of 8265 stars in the inner 20 kpc of M31, the region dominated by the visible disk. The radial velocities were measured using the Keck/DEIMOS multi-object spectrograph between 2007 and 2012 as part of the SPLASH survey (Gilbert et al. 2007; Dorman et al. 2012, 2013). The photometry is from *HST*/ACS/WFC via the PHAT survey (Dalcanton et al. 2012). The filters used were ACS F475W and F814W, which, for context, are roughly equivalent to Sloan *g* and Cousins *I* bands. Here we briefly describe the spectroscopic sample; for information on the PHAT photometry, see Dalcanton et al. (2012) and Williams et al. (2014).

This paper combines spectroscopic data from several smaller projects (Gilbert et al. 2007; Dorman et al. 2012, 2013). Therefore, the spectroscopic target selection function is not homogeneous. 44% of the targets, those observed between 2007 and 2010, were selected based on their apparent degree of isolation in a single-filter *i'* CFHT/MegaCam mosaic image. These targets primarily trace the dominant population in M31’s stellar disk: old, metal-rich RGB stars. Details on target selection techniques for these masks can be found in Dorman et al. (2012).

Because these objects are from masks designed prior to acquiring PHAT photometry at that location, they had to be later matched to their corresponding sources in the six-filter PHAT photometric catalogs. First, astrometric offsets between each DEIMOS mask and the PHAT coordinate system were obtained and applied to the spectroscopic sample. Then, for

each object in the spectroscopic sample, the brightest (in m_{F814W}) PHAT star within a search radius of $0''.5$ was chosen as the match.

The remaining 56% of the spectroscopic targets used in this work, observed in 2011 and 2012, were chosen based on existing PHAT photometry, eliminating the need for post-spectroscopy cross-catalog matching. Since we had color information at the target selection stage, we were able to prioritize under-represented populations over the dominant metal-rich red giants. For the five slitmasks, targeting about 1000 sources, observed in 2011, we chose red giants across the broad range of photometric metallicities $-2.0 < [M/H] < 0.2$. We restricted our sample to stars that were in the magnitude range $20 < m_{F814W} < 22$, but otherwise chose stars randomly in position and magnitude space. (The quoted magnitudes, like all magnitudes in this paper, are in the Vega system.) See Dorman et al. (2013) for more information on the *HST*-aided spectroscopic target selection in 2011.

In 2012, we targeted about 4500 more stars across a broad range of ages, including young massive MS stars, intermediate-age AGB stars, and RGB stars. We also targeted a few young clusters identified by the PHAT team (Johnson et al. 2012), although those are not used in this work. All of our targets were brighter than either $m_{F814W} = 22$ or $m_{F475W} = 24$. Faint RGB stars (with $m_{F814W} > 21.5$) and faint MS stars (with $m_{F814W} > 21$ or $m_{F475W} > 23$) were given very low priority and were only used on the rare occasion that there was unused space to fill on the masks, but otherwise targets were chosen randomly within each evolutionary stage. These data are presented here for the first time; see Appendix for a full data table.

After observing, each raw 2D spectrum is collapsed in the spatial direction and cross-correlated against a suite of template rest-frame spectra to measure its radial velocity. Figure 2 shows some representative spectra and zoomed-in views of the absorption lines that dominate the radial velocity measurement. Though the entire spectrum is used in the cross-correlation, certain absorption lines are most important in determining the radial velocity. The Ca II triplet near 8500 \AA is strong in RGB and many AGB stars. Temperature-sensitive TiO bands across the red side of the spectrum, including a strong triplet near 7050 \AA and another band near 8850 \AA , determine the velocity for many of the redder AGB and RGB stars, including those without strong CaT. For some of the blue MS stars, $H\alpha$ and $H\beta$ are the only reliable lines present, but the Paschen series is also present for young supergiants. After the automatic cross-correlation, each spectrum is inspected by eye, and only robust velocities (those based on at least two strong spectral features) are used for kinematical analysis. About 1/3 of the MS stars do not pass the quality cut—since hot stars have so few spectral features in the optical, it was harder to recover a robust velocity.

The velocity precision varies based on the instrument settings that were used. Through 2011, we used the 1200 line/mm grating on DEIMOS, resulting in an approximate wavelength range of $6500\text{--}9100 \text{ \AA}$ and a spectral resolution of $R = 6000$. In 2012, we used the coarser 600 line/mm grating to gain spectral coverage as blue as 4500 \AA and better characterize the younger massive upper MS stars, yielding a resolution of about $R = 2000$. In both cases, the width of each slitlet was $0''.8$. The radial velocity uncertainties derived from the cross-

correlation are on the order of a few km s^{-1} for the higher-resolution spectra, and $\sim 10 \text{ km s}^{-1}$ for the spectra from 2012.

The stars in our full spectroscopic sample fall in the region of the galaxy dominated by the disk. Based on surface brightness profile decompositions (Courteau et al. 2011; Dorman et al. 2012), the bulge contributes only 2% of the *I*-band surface brightness at the innermost portion of the survey and essentially zero exterior to about 8 kpc. Therefore, we do not need to remove bulge stars from the sample.

3. METHODS: VELOCITY DISPERSION AS A FUNCTION OF AGE AND OF METALLICITY

Our goal is to measure the LOS velocity dispersion of the stellar disk as a function of age and of metallicity. In this section, we describe our analysis procedure. First, we define regions in the optical CMD corresponding to very young MS, younger AGB, older AGB, and old RGB populations using a simulated optical CMD. Next we split our spectroscopic sample into those four bins using the observed optical PHAT CMD. Then, we map the LOS dispersion of each component. Finally, to look at the old population in more detail, we further split the RGB population into two metallicity bins and construct a dispersion map for the stars in each bin.

3.1. Definition of Age Bins

We first define four age bins in the optical F814W/F475W CMD using two criteria: first, we use both photometric and spectral discriminants to identify stars that are unlikely to be MW foreground (MWFG) stars. Second, we use a simple simulated CMD to identify regions containing stars of similar ages. We then roughly estimate the average ages of our four age bins using the simulated CMD and point out that the age estimates of the older bins depend significantly on the assumed star formation history.

3.1.1. Foreground Contamination

MWFG dwarfs can lie in the same magnitude window as our M31 spectroscopic targets. We took steps, both pre- and post-spectroscopy, to eliminate them from the catalog.

For PHAT-selected targets, we avoided likely MW members in the target selection stage using UV-IR color-color cuts. Here the UV-color is ACS F336W–F475W and the IR color is F110W–F160W. A comparison of the MWFG as simulated by the TRILEGAL galaxy model (Vanhollebeke et al. 2009) and a toy model of M31 shows that the foreground dwarfs are exclusively bluer in the IR color and redder in the UV-color than the M31 giant sequence for stars with $F475W\text{--}F814W > 2$ —that is, all of the RGB and AGB stars in our M31 sample. We use the TRILEGAL MWFG simulation to define a box in color-color space to exclude from our spectroscopic target selection.

For the CFHT-selected targets, color information from PHAT was not available as a tool for excluding MWFG stars at the target selection stage. Instead, we identified red foreground dwarfs based on the presence of the Na I doublet at 8190 \AA , which is a useful giant/dwarf discriminant due to its sensitivity on surface gravity and temperature (Schiavon et al. 1997). We visually inspected each spectrum taken between 2007 and 2011, and about half of the spectra taken in 2012, for the presence of the Na I doublet. We also calculated the equivalent width (EW) and the uncertainty on the EW

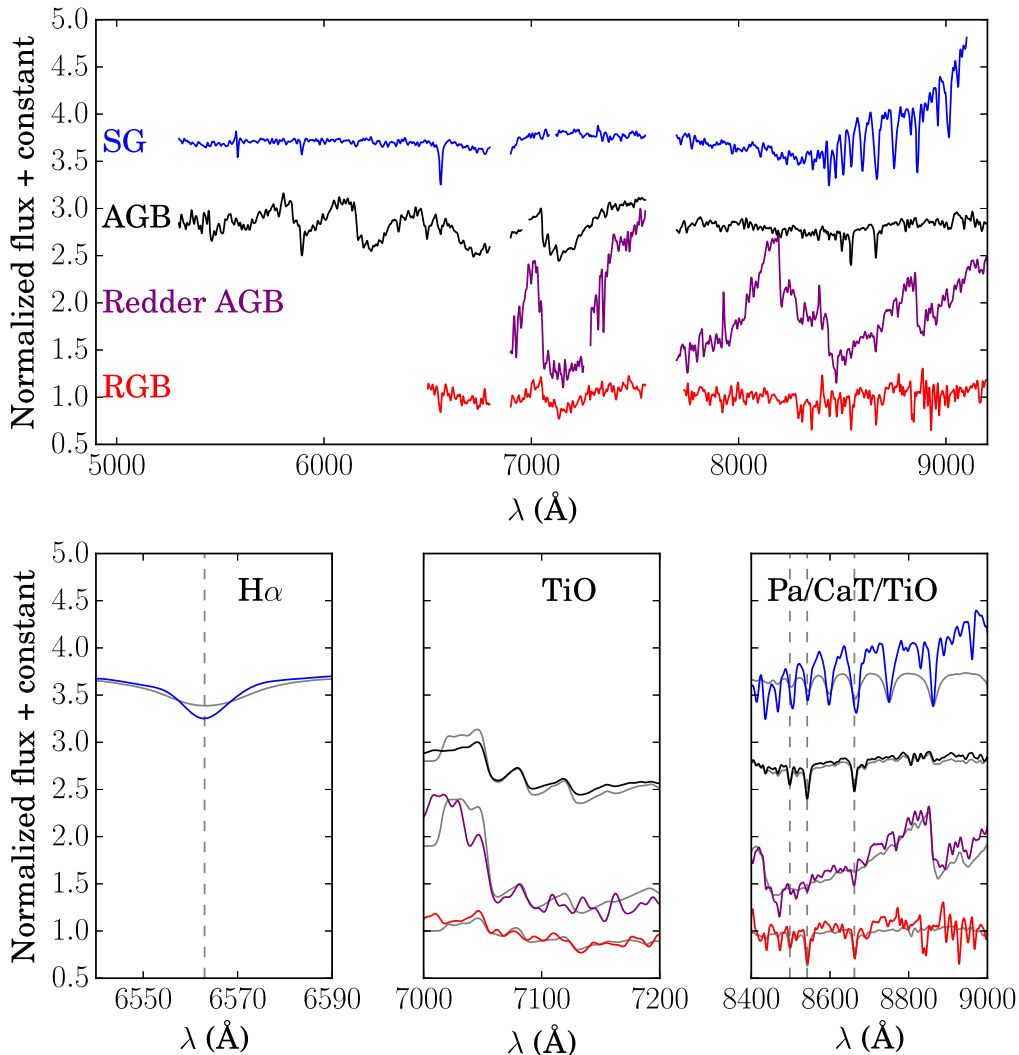


Figure 2. Example spectra from a variety of stellar types. Spectra have been normalized and shifted to rest frame. Top: representative spectra from a young supergiant, young-intermediate age AGB, older and redder AGB, and RGB stars. The two bottom spectra are from very red stars, and so their flux on the blue half of the spectrum is very noisy and not shown. Gaps in the spectrum correspond to either the DEIMOS chip gap or the atmospheric A-band, whose positions vary from spectrum to spectrum in the rest frame of the star. Bottom: zoom-in views of the portions of the spectrum useful for obtaining velocities. The science spectra are color-coded as in the top panel, while template rest-frame spectra used to measure radial velocities via cross-correlation are shown in gray. The wavelengths of $H\alpha$ and the Ca II triplet are shown in gray vertical lines to aid the eye. (Note that the young (blue) star in the rightmost panel displays the Paschen series, not the calcium triplet, although some of the lines fall at similar wavelengths.)

measurement (σ_{EW}) across the doublet bandpass (8179–8200 Å) relative to the adjacent continuum (8130–8175 Å and 8210–8220 Å) as in Gilbert et al. (2006). The set of stars with both EW measurements and visual flagging formed a training set from which we found a diagnostic that can be used to automatically identify definite foreground dwarfs. Figure 3 shows the training set (left) and all stars with spectra (right). Stars with $EW > 3.2$ and significance $EW/\sigma_{EW} > 8$ (those inside the red box) are foreground dwarfs and were eliminated from the spectroscopic sample. This discriminant was applied to all stars in the sample, including PHAT-selected targets that had already survived the pre-spectroscopy color-color cut.

The steps described above do not necessarily eliminate every foreground star. To estimate the number of contaminants that remain in the sample as a function of optical color and magnitude, we employ the TRILEGAL simulation of the MW in the direction of M31. In each color/magnitude bin (CM), we

calculate the number of contaminants $N_{MW,expected,CM}$ expected in the SPLASH survey:

$$N_{MW,expected,CM} = N_{SPLASH,CM} \frac{N_{Trilegal,CM}}{N_{PHAT,CM}}. \quad (1)$$

We then subtract the number of contaminants already identified in and removed from that color/magnitude bin based on the Na I doublet discriminant:

$$N_{contaminants,CM} = N_{MW,expected,CM} - N_{removed,CM}. \quad (2)$$

Figure 4 shows the resulting distribution of $N_{contaminants,CM}$. The number is only significant (>10% of the sample) brighter than $m_{F814W} = 21$ and between $1 < m_{F475W} - m_{F814W} < 2$. Stars in this portion of the CMD are excluded from the analysis. After exclusion, fewer than 0.1% of the stars in the sample are expected to be MWFG stars.

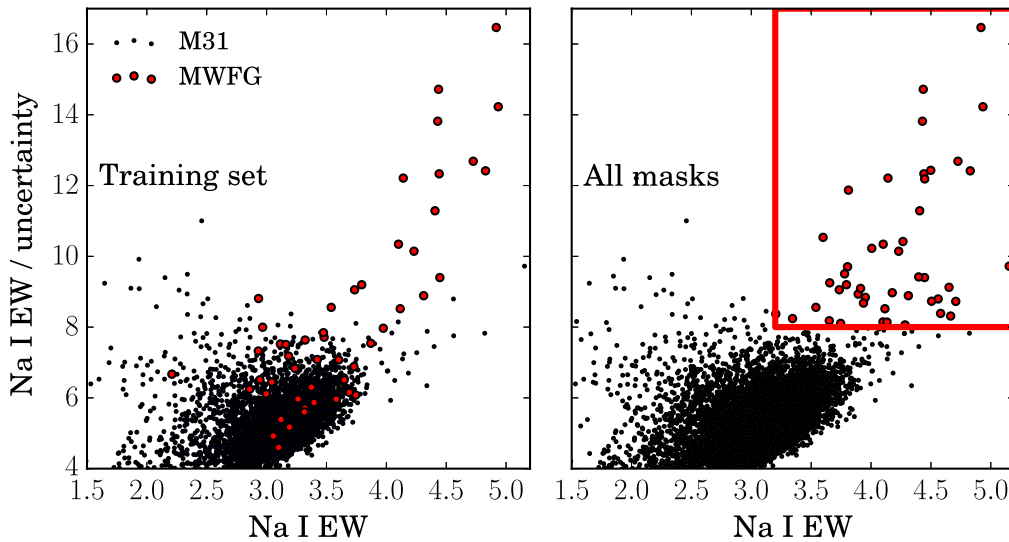


Figure 3. Identifying and removing likely MW foreground dwarfs based on the presence of the surface gravity-sensitive Na I doublet. Left: training set based on about 8000 stars that were visually checked for the presence of the doublet. Stars with equivalent widths (EW) greater than 3.2 and significance (EW / uncertainty on EW) greater than 8 were almost universally flagged as dwarfs. Right: same discriminant applied to the full set of spectra (not all of which had been manually inspected). Stars within the red box are almost certainly MW dwarfs and were excluded from the analysis.

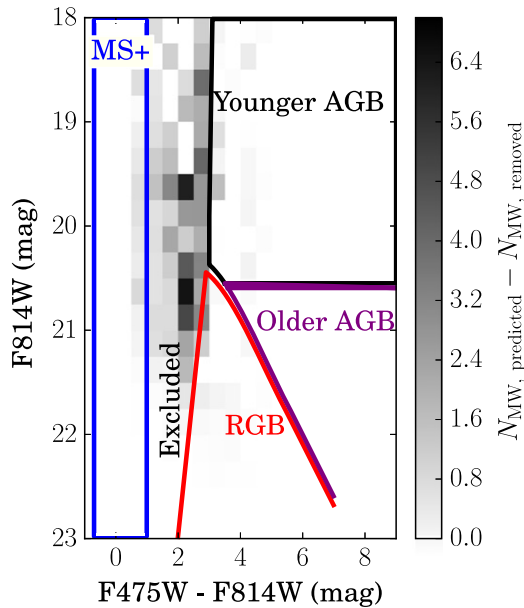


Figure 4. Optical Hess diagram showing the number of foreground MW stars expected in our spectroscopic sample after cutting out stars with strong surface gravity-sensitive Na I doublets, which are very likely to be foreground MW dwarfs. The blue, black, violet, and red lines outline the MS, younger AGB, older AGB, and RGB regions, respectively, described and defined in Section 3.1.2 below. The most contaminated portion of the CMD—brighter than $m_{F814W} = 21$ and between $1 < m_{F475W} - m_{F814W} < 2$ —is excluded from the analysis.

3.1.2. Estimating Ages Using a Simulated CMD

To estimate the ages of stars as a function of CMD position, we simulate a simple stellar population in the optical CMD.

Simulating a CMD requires choosing an age–metallicity relation (AMR) and star formation rate (SFR). A few measurements from the outer regions of M31’s disk ($R \sim 20$ – 30 kpc) using deep *HST* photometry are available (Brown et al. 2006; Bernard et al. 2012). Both show a clear inverse relationship between stellar age and metallicity. For our CMD, we are more interested in separating the CMD into

regions with distinct average ages than pinning down those ages precisely. We assume a constant SFR of $1M_{\odot} \text{ yr}^{-1}$. This assumption is overly simplistic—for example, Bernard et al. (2012) measure a burst of star formation in the outer disk at around 2 Gyr ago—but is acceptable as we only aim to separate our CMD into age bins, rather than measure the precise ages of stars in those bins. We estimate the AMR empirically from the PHAT RGB data in the following way: we use 10 Gyr old PARSEC 1.1 (Bressan et al. 2012) isochrones in the metallicity range $-2.1 \leq [M/H] \leq 0.3$ to estimate the metallicity distribution function (MDF) of all of the bright ($m_{F814W} < 23$) RGB stars in the PHAT survey. The AMR is then constructed so as to replicate the MDF, assuming a constant star formation rate and a metallicity that never decreases with time. For example, since there are seven times as many stars at $[M/H] = -0.5$ as at $[M/H] = -1$, the simulation is allowed to spend seven times as long producing stars with $[M/H] = -0.5$ as with $[M/H] = -1$. We use Girardi et al. (2010) isochrones, a Kroupa initial mass function (IMF), assume that 35% of stars are in binaries, and apply a constant foreground reddening of $A_v = 0.2$.

While this technique ignores the age–metallicity degeneracy on the RGB, it generates a reasonable CMD. The CMD generated using this AMR is very similar to one produced using an AMR adapted from the empirical one presented in Brown et al. (2006), which used deep *HST* photometry of a small field in M31’s disk about 25 kpc from the galactic center. Our CMD assumes a constant SFH and thus is not an SFH fit. Due to the age–metallicity–extinction degeneracy along the RGB, the assumed SFH may differ dramatically from the true one. However, it enables us to select four CMD regions with different average ages and measure approximate values for those average ages.

We use the simulated CMD to define four regions containing stars of increasing average age, while avoiding the highly contaminated region described in the previous section. For simplicity, we refer to these bins as “MS+,” “younger AGB,” “older AGB,” and “RGB,” although stars in a given bin do not necessarily all belong to the exact same evolutionary stage. The

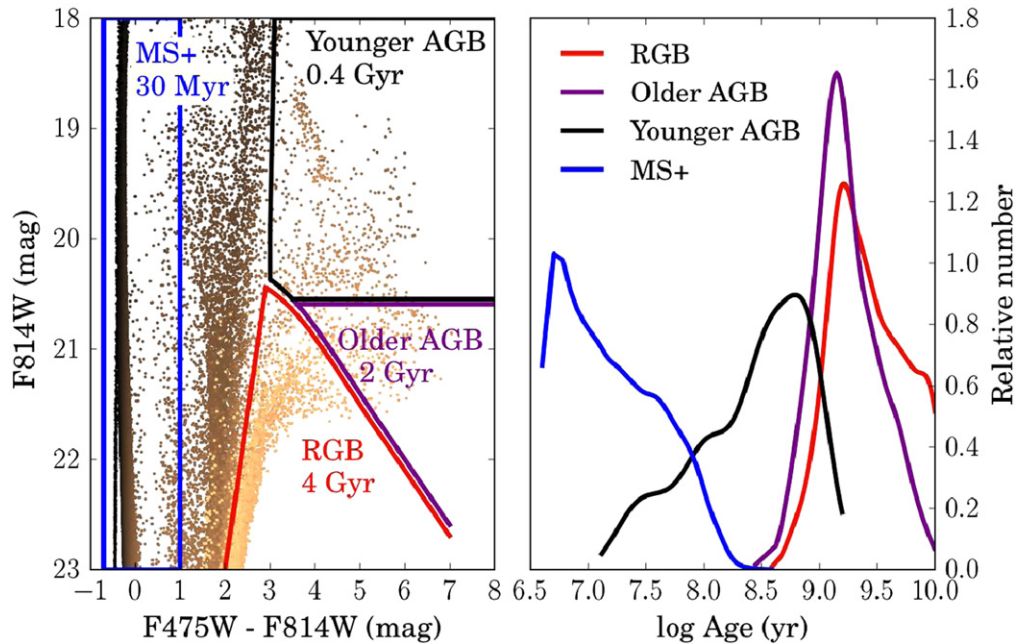


Figure 5. Left: simulated CMD assuming a constant star formation rate. The blue, black, violet, and red outlines define the “MS,” “younger AGB,” “older AGB,” and “RGB” age bins, respectively. Stars are color coded by $\log(\text{age})$, with brighter (yellow) colors corresponding to older ages. Numbers correspond to the average age of stars within each bin. Right: age distributions for the same four regions. Age bin colors are the same as the left panel and throughout this paper.

constant-SFR CMD and age bin outlines are shown in the left panel of Figure 5; in the rest of this section, we explain the choice of bins.

All stars bluerward of $m_{F475W} - m_{F814W} = 1$ are classified as “MS+” stars. These bright blue stars are primarily massive upper MS members younger than 100 Myr, although some may be blue supergiant stars, which have similar ages. The RGB region includes stars redward of the line that passes through $(m_{F475W} - m_{F814W}, m_{F814W}) = (2, 23)$ and $(2.7, 20.4)$ and fainter than the tip of the red giant branch (TRGB). The blue limit is chosen by eye so as to minimize contamination with young MS stars, which overlap the oldest RGB stars in high-metallicity systems such as M31. We measure the TRGB using the Bressan et al. (2012) isochrones described earlier. The brightest RGB stars on these isochrones trace the TRGB as a function of $F475W - F814W$ color.

We classify most of the rest of the red side of the CMD as AGB. To avoid contaminating the AGB bins with older RGB stars, we do not use stars less than 0.1 mag brighter than the TRGB (those within photometric scatter of the TRGB).

Stars of a large range of ages, from a few hundred Myr to several Gyr, can lie on the AGB. At a given metallicity, younger AGB isochrones are brighter and bluer than older AGB tracks. In the simulated CMD, age roughly tracks luminosity, with younger stars toward the top of the CMD. We use this age–luminosity dependence to split the AGB in half (along the line $m_{F814W} = 20.5$) into two age bins: “younger AGB” and “older AGB.”

The final classification scheme into four age bins is shown in the left panel of Figure 5. The MS+ bin is outlined in blue on the blue side of the optical CMD, RGB stars in red below the TRGB, older AGB stars in purple brightward of the TRGB, and younger AGB stars in black. The right panel shows the age distributions of stars in these bins. The age distributions are broad and overlap, but have increasing mean ages: the MS+ bin has a mean age of 30 Myr, while the younger AGB, older

AGB, and RGB stars have average ages of 0.4, 2, and 4 Gyr, respectively. The RGB bin has a low average age for two reasons: first, because we imposed an AMR, the older, bluer metal-poor RGB stars actually overlap in CMD space with young red helium-burning stars, while younger, metal-rich RGB stars do not suffer from this ambiguity. Therefore, we use only the younger (redder) portion of the RGB in this work. Second, and more importantly, the RGB age distribution in any magnitude-limited sample with a constant SFR is biased toward younger ages since the rate of stars moving off the MS onto the RGB is higher for younger stars.

For comparison, we also generate a CMD with an exponentially decreasing SFR with timescale $\tau = 4$ Gyr. This SFR is much steeper (skewed toward older stars) than seen in the outer disk of M31 (Bernard et al. 2012), but gives an interesting boundary case. The boundaries of the most reasonable four age bins in CMD space are the same as for the constant-SFR simulation, but the average ages of the older age bins increase: the older AGB has an average age of 3.5 Gyr, and the RGB has an average age of 5.5 Gyr. We do not show this CMD here since it does not influence our choice of age bin boundaries, but we discuss both sets of age estimates later in the paper.

3.2. Separation of Data into Age Bins

We divide the stars in the kinematical sample into the four age bins defined in the previous section, using optical PHAT photometry.

The left panel of Figure 6 shows an optical Hess diagram of all stars in the PHAT survey in a representative region of the galaxy. Spectroscopic targets fall into the color/magnitude range outlined in green. The right panel of Figure 6 shows a CMD of the spectroscopic sample only, divided into age bins as in Figure 5.

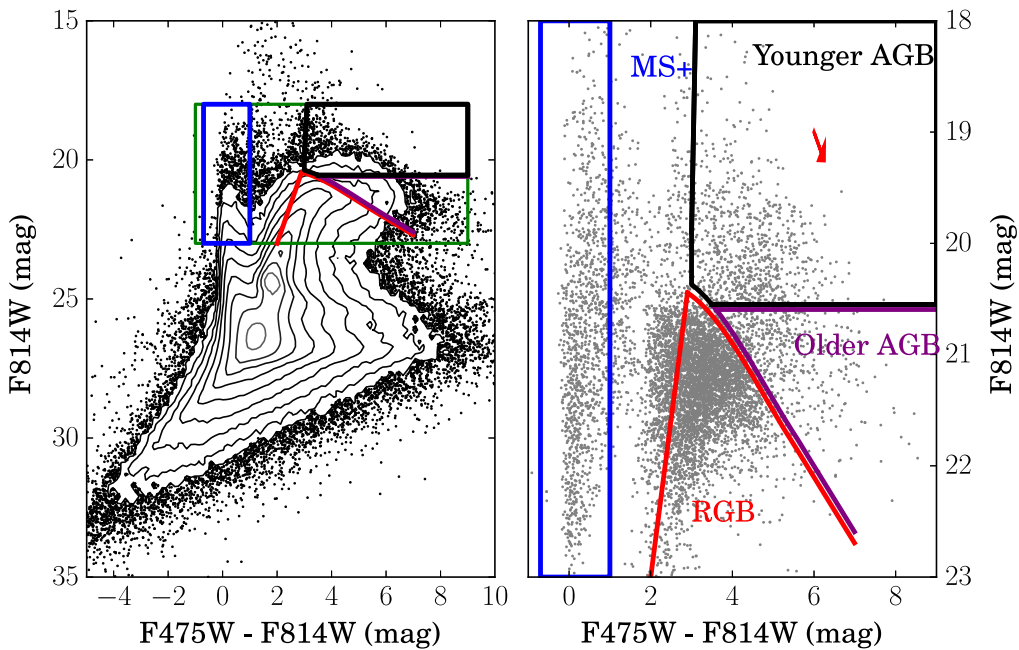


Figure 6. Left: optical PHAT CMD from Brick 12, a representative brick near the middle of the PHAT survey area. The green box marks the portion of the CMD sampled by the spectroscopic survey. The age bin region outlines are the same as in Figure 5. Right: optical PHAT CMD of spectroscopic targets in the PHAT survey region that have optical PHAT photometry and reliable velocities and are unlikely to be foreground MW stars based on the strength of their Na I doublet lines. The age bin region outlines are the same as in the left panel and in Figure 5. Stars that fall outside the age bin boundaries, have ambiguous ages, and/or may be foreground stars are not used in this work. The red arrow shows the median reddening vector due to dust in M31’s disk, as measured by Dalcanton et al. (2015), assuming that half of the spectroscopic targets lie behind M31’s dust layer. Photometric errors are typically less than 0.01 mag in each filter (Dalcanton et al. 2012).

While our simulated CMD does include reddening from the MWFG, it does not account for differential extinction by dust within the disk, and we do not attempt to account for this shortcoming in classifying the data into age bins. This means that some of the CMD regions may be contaminated by stars from bluer age bins. However, the direction of the reddening vector is such that only two bins are likely to be contaminated: a few younger AGB stars may be reddened into the older AGB region. The RGB region is largely immune from contamination due to reddening, since the shape of the TRGB is such that old AGB stars will never be reddened into the RGB region, and the other two bins are far enough away in CMD space that their members will not be reddened onto the RGB either. Additionally, the red helium-burning stars, which are not included in any of our bins, can be reddened into the RGB bin; however, these are much smaller in number than the RGB stars and thus largely insignificant even with reddening. However, the region in the CMD occupied by RGB stars contains stars of a range of ages, as discussed earlier and pictured in Figure 5.

We estimate the average reddening vector in our sample using the M31 dust map presented in Dalcanton et al. (2015). This map, constructed by comparing the infrared colors of RGB stars in the PHAT survey to the unreddened RGB, gives the average A_V as a function of location across the disk. Only stars that lie behind the dust layer are affected. The average reddening vector at the locations of our spectroscopic targets, assuming half the targets lie behind the dust layer, is shown in Figure 6.

3.3. Velocity Dispersion Maps

We now map the LOS velocity dispersion of stars in each of the four age bins using a smoothing technique and display the result in two ways: as a 2D sky map (Figure 7) and as a 1D

dispersion distribution (Figure 9). We do not fit separately for the radial, azimuthal, and vertical components of the dispersion. Such a decomposition requires assumptions on the rotation curve and the geometry of the disk, and we choose to keep our analysis purely empirical. A future paper will examine the shape of the velocity ellipsoid. For reference, because of M31’s nearly edge-on inclination, the vertical component of the velocity dispersion has negligible contribution to the LOS dispersion anywhere on the disk. In general, the LOS dispersion is a combination of the radial and azimuthal dispersion components, though the azimuthal component dominates near the major axis where most of the survey field lies.

For each target, we measure the weighted second moment of the velocity distribution of all neighbors that belong to the same age bin and also fall within some radius of that target using the maximum likelihood method described in Pryor & Meylan (1993). The weights are the inverse square of the velocity measurement uncertainty. This technique takes into account random scatter from individual velocity measurement uncertainties and also allows a straightforward computation of the uncertainty on the velocity dispersion estimate as long as the number of points is at least ~ 15 –20.

To create 2D dispersion maps, displayed in Figure 7 and Figure 8, the smoothing radius is fixed to $200''$ for the MS and RGB bins and $275''$ for the less densely populated AGB bins. Stars with fewer than 15 neighbors are dropped from the sample, as their dispersions and associated uncertainties are unreliable. These stars are still available to serve as “neighbors” for nearby targets, but the dispersion measurements centered on them are not used. Using smaller smoothing lengths increases spatial resolution in densely packed regions of the survey area, but results in many points being cut because they do not have enough neighbors to yield reliable dispersions and

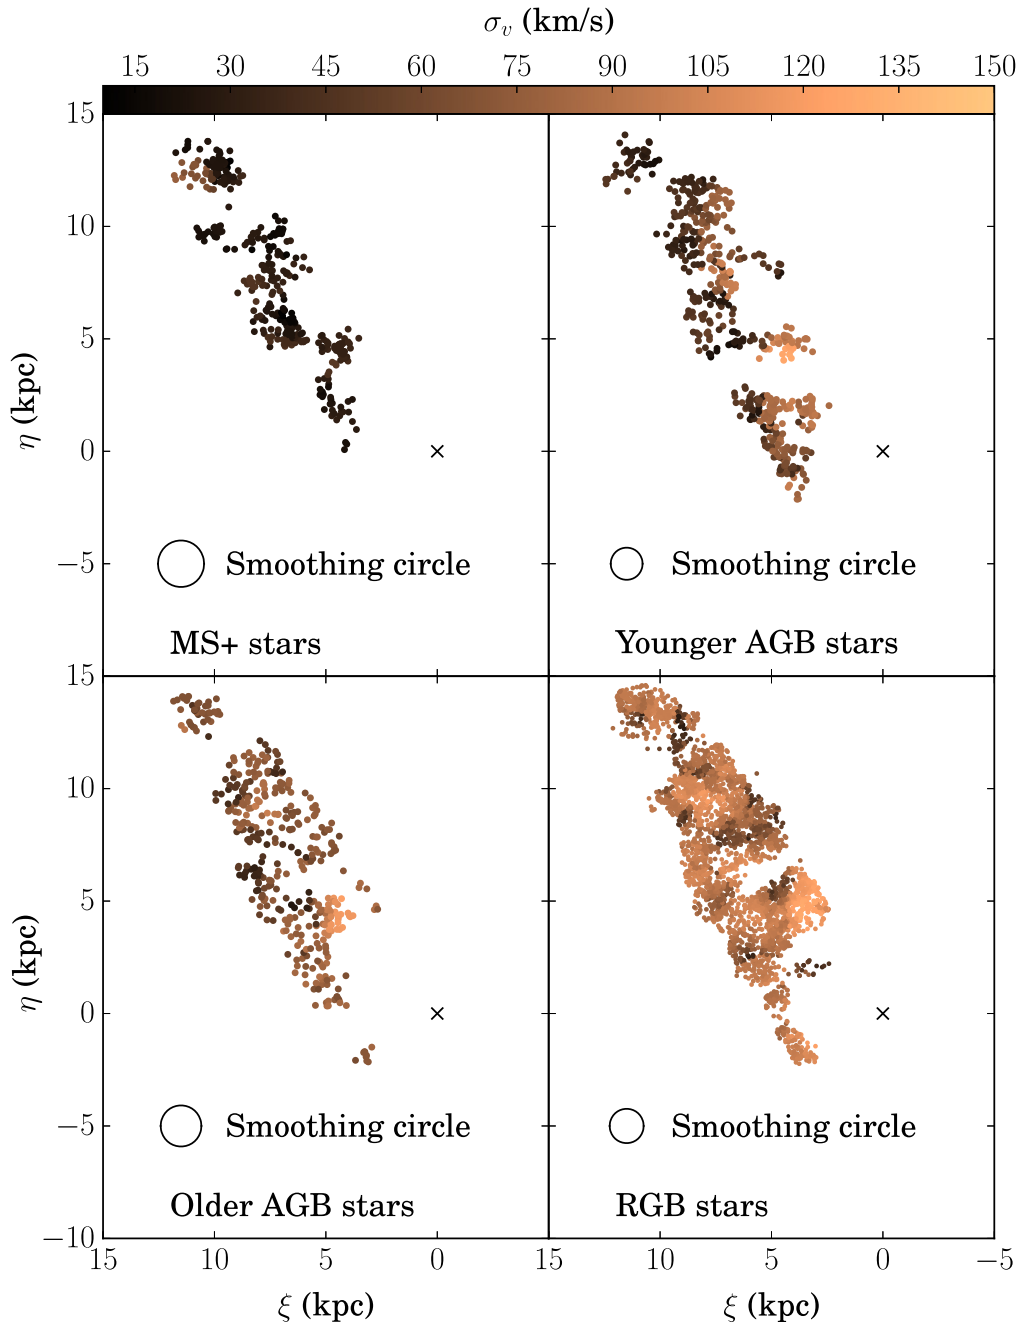


Figure 7. Smoothed, local, LOS velocity dispersion of stars of different evolutionary stages: young MS (upper left); younger AGB (upper right); older AGB (lower left); and RGB (lower right). The circles show the sizes of the smoothing circles in which the dispersion is calculated; smaller circles can be used for populations with higher number density. The typical dispersion increases with average age. The dispersion varies across the face of the disk within each age bin in a way that can be explained by our finite sampling density, as described in the Discussion section.

uncertainties. Using larger smoothing lengths further reduces spatial resolution in the maps, reducing the contrast between adjacent high-and low-dispersion patches.

To create the 1D dispersion distributions in Figure 9, the smoothing radius is chosen independently for each target such that the uncertainty on the velocity dispersion is constant within an age bin. A constant uncertainty makes it easier to understand the spread in the dispersion distribution due to measurement uncertainty. We arbitrarily choose uncertainties of 5 km s^{-1} for the MS+ bin, 7 km s^{-1} for the RGB stars, and 10 km s^{-1} for the AGB stars. Requiring a constant uncertainty means that stars in regions of lower target density or with larger individual

measurement uncertainties have larger smoothing circles. Again, dispersions measured using fewer than 15 neighbors are dropped. These uncertainty choices result in an average smoothing circle size of about $200''$ for all age bins. If we allow the dispersion uncertainty to be larger, the dispersion distributions are smeared out too much to see individual features. If we cap the uncertainty at a small value, too many points are removed due to small numbers of neighbors.

It is clear from both the 2D dispersion map in Figure 7 and the 1D dispersion distributions in Figure 9 that the typical dispersion increases with average age—that is, older populations are dynamically hotter than younger populations. In

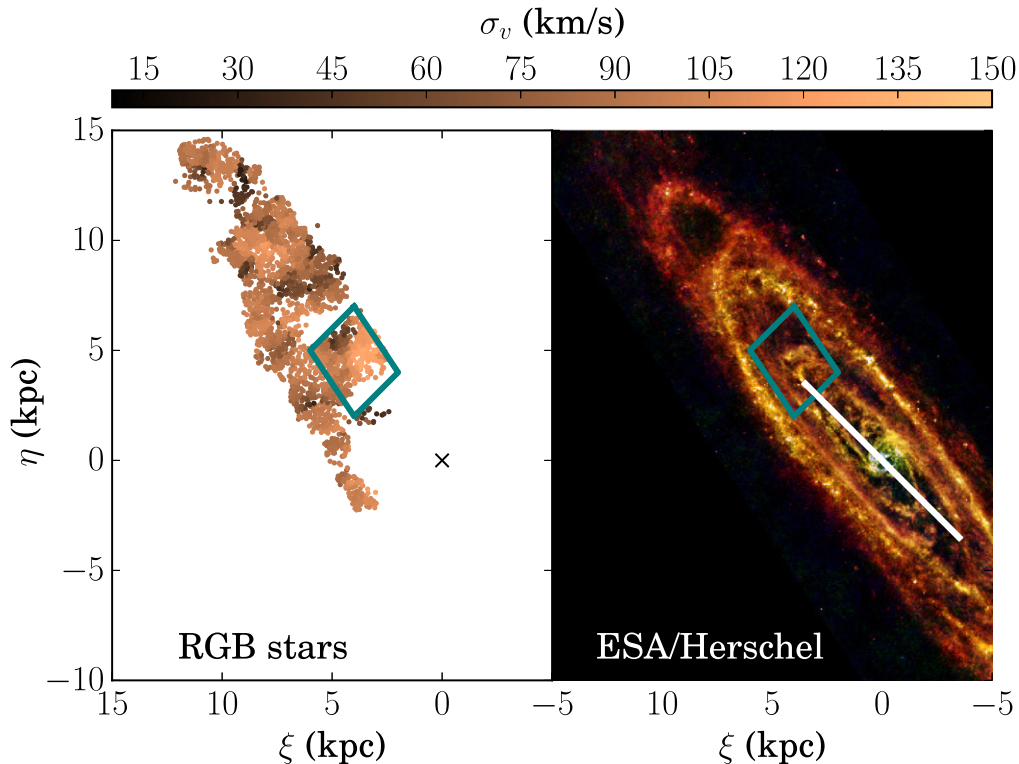


Figure 8. RGB dispersion map from Figure 7 next to a Herschel image of M31 for reference. A high-dispersion region that we term the “Brick 9 Region” is outlined in teal in each panel. The Brick 9 region aligns with the end of the inner ring, also roughly consistent with the end of the long bar according to the simulations by Athanassoula & Beaton (2006). The estimated size and orientation of the bar from Athanassoula & Beaton (2006) are marked with a white line.

addition, the 2D maps appear to be patchy, but we will show in Section 4.3 that much of the small-scale spatial variation in dispersion is an effect of finite sampling and is not physical.

3.4. RGB Velocity Dispersion Maps as a Function of Metallicity

We also examine how the RGB dispersion map varies with metallicity.

We estimate the metallicities of the RGB stars by interpolating on the grid of 10 Gyr old Bressan et al. (2012) isochrones described in Section 3.1. The resulting metallicity distribution ranges from $-2.3 < [M/H] < 0.3$ and is strongly skewed toward high metallicities, with a peak at $[M/H] = -0.2$. We then split the RGB sample approximately in half, into a high-metallicity ($[M/H] > -0.25$) and a lower-metallicity bin. The stars in our RGB bin are not all 10 Gyr old; the age–metallicity degeneracy on the RGB means that the metal-poor bin has a slightly higher average age than the metal-rich bin, although there is significant overlap in the age distributions. When we divide the RGB stars in the simulated CMD into the same two bins by position in the CMD, the low-metallicity bin has a median age of 2.7 Gyr and a mean age of 4.9 Gyr, with a broad, relatively flat age distribution. Meanwhile, the high-metallicity bin is slightly younger with a median age of 2.4 Gyr and a mean age of 3.5 Gyr. For each bin, we construct a smoothed velocity dispersion map and a dispersion histogram as described in Section 3.3. The maps are displayed in Figure 10 and the 1D distributions in Figure 11. As before, for ease of interpretation, the smoothing circle radius is held constant at $200''$ for the maps and is allowed to

vary in order to reach a target dispersion uncertainty of 7 km s^{-1} for the 1D distributions.

4. RESULTS

4.1. Age–Dispersion Relation

We can now measure, for the first time, the age–velocity dispersion relation in an external galaxy. Figure 12 shows the average velocity dispersion of stars in our four age bins versus their average ages, as estimated from the simulated CMDs. The solid black line shows the best fit to the points assuming a constant SFR, and the dashed black line assuming a decreasing SFR. For comparison, we also show the age–dispersion relation for F and G dwarfs in the solar neighborhood of the MW from Nördstrom et al. (2004). We show both the σ_R and σ_ϕ components of the MW’s velocity ellipsoid. σ_ϕ is the best comparison to our major-axis-dominated M31 data set, but in general the LOS dispersion at any location is some combination of the two components (with a negligible contribution from the vertical component σ_z). The difference in both slope and normalization is striking. The dispersion of M31 stars increases with age more than three times faster than the dispersion of MW stars in the case of the exponentially decreasing SFR, and five times faster in the case of the constant SFR. Additionally, the average dispersion of the RGB bin (90 km s^{-1}) is nearly three times as high as the oldest, hottest population probed in the MW thin disk by Nördstrom et al. (2004): 10 Gyr stars with $\sigma_\phi \sim 32 \text{ km s}^{-1}$. The MW’s metal-poor thick disk is slightly more disturbed, but still 50% cooler than M31’s disk, at $\sigma_R = 60 \text{ km s}^{-1}$ (Budenbender et al. 2014).

This age–dispersion correlation is robust against contamination of our age bins. Two bins—the RGB and older AGB—are

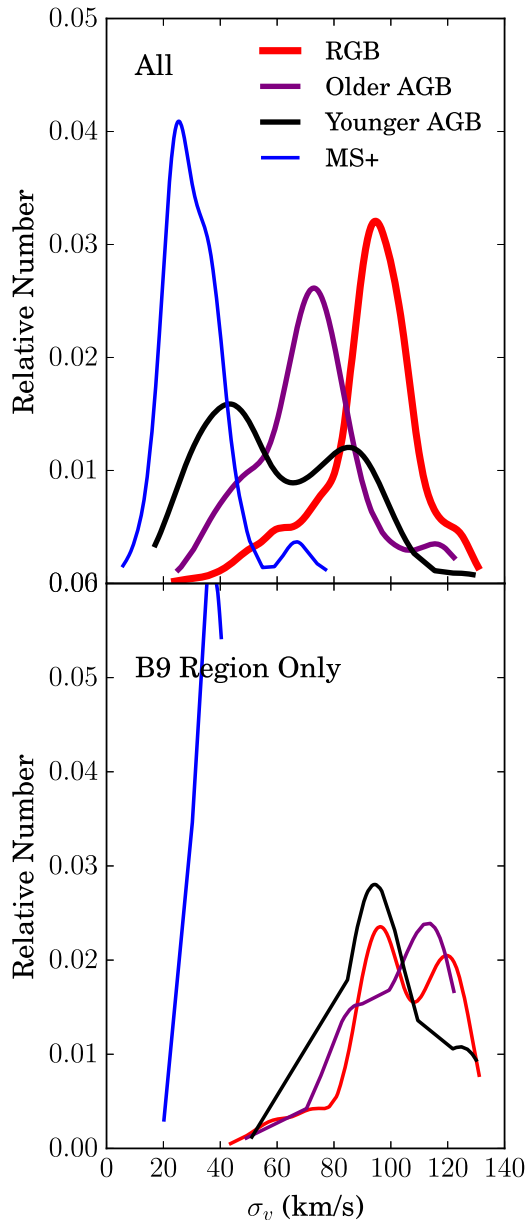


Figure 9. Local LOS velocity dispersion distributions of stars in the four age bins for the entire survey area (top) and only stars in the Brick 9 region outlined in Figure 8 (bottom). For this plot, smoothing circle sizes were chosen independently for each point such that the uncertainty in velocity dispersion was 7, 10, and 5 km s^{-1} for the RGB, AGB, and MS+ age bins, respectively. Dispersion distributions are constructed from the data using a kernel density estimator, with optimum bandwidth chosen using Silverman’s method. In the galaxy as a whole (top panel), average dispersion increases with average age. The B9 region has significantly hotter kinematics than the rest of the galaxy.

probably contaminated by a few younger stars. The RGB age bin consists primarily of old stars, but also includes some higher-mass, intermediate-age stars. This contamination biases the RGB dispersion distribution toward that of the faint AGB; in other words, an exclusively old population would be at least as dynamically hot as our RGB bin. The other affected bin is the older (fainter) AGB group, contaminated by younger AGB stars that have been reddened by dust in the disk of M31. Again, this contamination biases the faint AGB dispersion distribution toward smaller values, so that the difference in typical dispersions between uncontaminated young-

intermediate and older-intermediate age populations is at least as big as the one we report in Figure 9.

The high dispersion of M31’s stellar disk is also robust against “smearing” of the LOS component of the rotation velocity v_{DiskLOS} within the finite-sized smoothing circles. To confirm this, we conduct a test in which we map the dispersion in $(v - v_{\text{DiskLOS}})$ instead of just v . Here, v_{DiskLOS} is computed assuming that the inclination $i = 77^\circ$ and major axis P.A. = 38° over the entire disk, and that the rotation velocity of the disk v_{DiskRot} is a constant within each age bin. v_{Rot} is calculated by fitting to the deprojected velocity field of stars in each age bin, and comes out to (260, 250, 220) km s^{-1} for the (MS+, AGB, RGB) stars, respectively. The resulting dispersion in $(v - v_{\text{DiskLOS}})$ corresponds to a median correction of just 0.5% for the RGB dispersion, 1% for the AGB dispersion, and 6% for the MS+ dispersion.

We will discuss the evolutionary implications of the age–dispersion relation in more detail in Section 5.

4.2. Age–Metallicity Relation among RGBs

Figures 10 and 11 show that the metal-poor RGB population is dynamically hotter than the metal-rich population, by a factor of almost 50%. The metal-rich component has kinematics only about 15% hotter than the older AGB population. It is dynamically hot ($\sigma_v \sim 90 \text{ km s}^{-1}$) in two spots on the disk. It should be noted that the metal-rich bin likely contains some reddened metal-poor stars; this may explain the broad dispersion distribution of the metal-rich bin.

4.3. Structure in Dispersion Map

The spatial dispersion maps in Figure 7 show that the dispersion of all four age bins varies across the face of the disk.

One of the most obvious features in all four age bins is a kinematically hot patch about 6 kpc from the center of the galaxy on the major axis. This patch is marked by a teal line and outlined in both panels of Figure 8. We dub the region enclosed by the black box the “Brick 9 region,” since it is centered on the same area as the set of *HST* pointings known as “Brick 9” in the PHAT survey tiling pattern. The two panels in Figure 9 contrast the dispersion distributions for the full survey and for the brick 9 region. For every age group, the dispersion of stars in the Brick 9 region alone is much higher than for the overall survey.

In Section 5, we will explore the possible relationship between this hot patch and M31’s bar.

Outside of the Brick 9 region, the dispersion varies with position (is “kinematically clumpy”), although the dispersions of different populations do not necessarily follow the same spatial pattern. Much of this clumpiness is simply a result of low spectroscopic target density, as we now show using a modified version of the disk toy model described earlier. We construct and analyze the toy model for the RGB disk as an example; similar experiments could be run for the younger stellar populations as well.

We construct a toy model of a disk by scattering N stars uniformly in a volume 0.8 kpc thick whose base is a circle of radius 20 kpc. The velocity vectors of the stars are drawn from a three-dimensional random normal distribution corresponding to a rotation velocity of 220 km s^{-1} and a velocity dispersion of 90 km s^{-1} (the average observed RGB dispersion) in each direction. (While it is unlikely that the velocity ellipsoid of

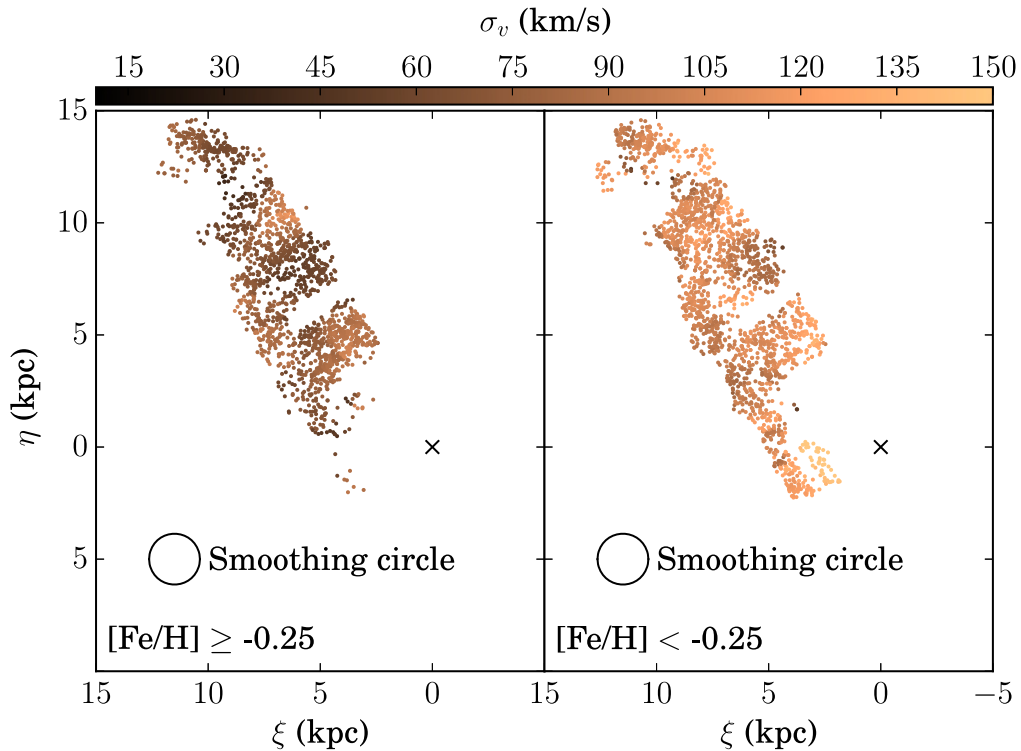


Figure 10. Smoothed velocity dispersions of RGB stars in high- (left) and lower-metallicity (right) portions of the RGB sample. Metallicities are measured using 10 Gyr old Bressan et al. (2012) isochrones. On average, the metal-rich population has a lower velocity dispersion. Some of the hot patches in the high-metallicity bin may be due to contamination from reddened low-metallicity stars.

M31’s disk is actually isotropic, this choice does not affect the qualitative results here.) We choose N such that it most accurately reproduces the sampling density of the SPLASH survey.

We then sample from our toy disk exactly as we do in our spectroscopic sample: at the location of each SPLASH target, we compute the LOS velocity dispersion of all neighbors within a $200''$ radius. The resulting dispersion map, shown in Figure 13, is kinematically clumpy. In Figure 14, we quantify the clumpiness and compare the simulation’s 1D dispersion distribution with M31’s RGB dispersion distribution. The two distributions have similar FWHM. To test the effect of sampling density, we also run a simulation with a number density five times higher. The resulting dispersion distribution is only half as wide as the original simulation, corresponding to a less patchy dispersion map.

Why does finite sampling density cause patchiness in the dispersion map? With few samples, velocity outliers (those on the tails of the velocity distribution) will be scarce and will not fall into every smoothing circle. Velocity dispersion estimates from different smoothing circles may be inflated or deflated based on how many velocity outliers they happen to include. The average dispersion measured accurately reflects the input (intrinsic) velocity distribution (90 km s^{-1} in this case), but the spread in dispersions from star to star is an artifact of low sampling density.

Figure 14 shows that M31’s dispersion distribution has large high- and low-dispersion tails in excess of the simulation—that is, there are dynamically hot and cold patches that cannot be explained by finite sampling density. This excess clumpiness is not a result of patchy dust: the spatial distributions of dust (Dalcanton et al. 2015) and velocity dispersion are

uncorrelated. In Section 5 we will argue that the dynamically hot patches within each age bin are regions where a few stars from a second kinematical component (for example, a bar, halo, or tidal debris) are overlaid on the uniform underlying disk.

5. DISCUSSION

5.1. High-dispersion Population: Superposition of Two Kinematical Components

Thus far, we have shown that the second moment of the local velocity distribution (LOSVD)—the local LOS velocity dispersion—varies substantially between age bins and also across the disk of the galaxy, and that much of the variation across the disk can be explained by finite target density in the spectroscopic survey. Here, we show that the shape of the LOSVD within each age bin changes with dispersion in a way that cannot be explained by sampling density alone, but can be explained by the presence of a second kinematical component superimposed on a uniform disk.

If some regions do have intrinsically higher velocity dispersions (not caused by sampling effects), it is not immediately clear whether the high-dispersion patches contain a single dynamically hotter disk component, or whether they contain a higher fraction of stars from a second kinematical component (such as a spheroid, bar, or tidal stream debris). The two scenarios can be distinguished by the shape of the velocity distribution: if we assume that the local LOSVD of the disk is Gaussian (which is a reasonable approximation in a galaxy with the rotation speed of M31), the former scenario would produce a wide symmetric Gaussian distribution, whereas the




Figure 11. Same as Figure 9, for the RGB stars in the two metallicity bins from Figure 10. The top panel includes all stars, whereas the bottom panel only includes stars in the Brick 9 region. The metal-rich population is dynamically colder on average, and its dispersion is higher on the B9 region than off.

latter would produce a superposition of two distributions: a narrow central distribution with a wide tail to one side.

Here we test the symmetry of the velocity distribution of low- and high-dispersion patches in the RGB age bin. We repeat the same process for the other three bins.

We first split the RGB velocity field into nonoverlapping square pixels, 1 kpc on a side. Then, for each pixel, we zero and normalize the velocity distribution by subtracting the mean and scaling by the standard deviation of velocities within that pixel: $x = (v - \bar{v})/\text{std}(v)$. We now stack all of the scaled velocity distributions x . If the local LOSVD were always Gaussian, the stacked distribution would be a unit normal symmetric about $x = 0$.

Figure 15 compares the scaled, stacked velocity distributions and a unit normal for RGB stars in pixels with lower dispersions (less than 100 km s^{-1} ; top panel) and higher dispersions (bottom panel). The velocity distributions are

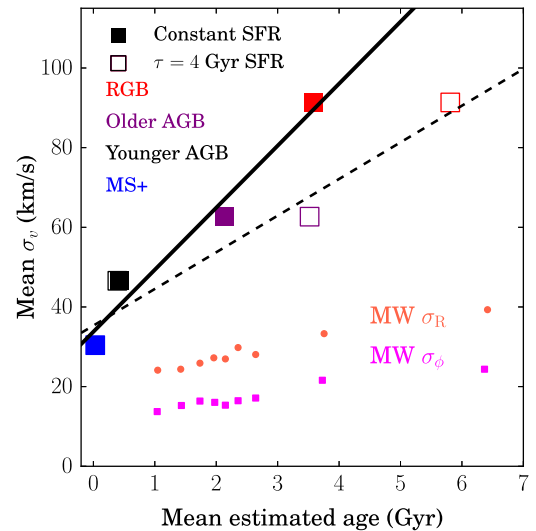


Figure 12. Comparison of dispersion vs. age relations for M31 and the MW solar neighborhood. Large squares: age–dispersion relation for four M31 age bins. Ages are the mean ages from the simulated CMD; dispersions are the mean values from Figure 9. Large solid squares and large open squares correspond to ages measured from the constant SFR and decreasing SFR simulated CMDs, respectively. The solid and dashed lines are the best fit to the constant SFR and decreasing SFR data points, respectively. Though this is a very rough comparison, the dispersion of a population appears to be a monotonic function of its average age. The orange circles and pink squares trace, respectively, the radial and azimuthal dispersion–age profiles for F and G dwarfs from the Geneva–Copenhagen survey (Nördstrom et al. 2004). Both the slope and normalization of the dispersion vs. age relation are much higher in M31 than in the MW, regardless of the assumed SFR.




Figure 13. Velocity dispersion map for a toy model of a disk with a Gaussian velocity distribution with $v_\phi = 220 \text{ km s}^{-1}$ and $\sigma_R = \sigma_\phi = 90 \text{ km s}^{-1}$.

skewed to negative velocities, which is toward M31’s systemic velocity on this redshifted side of M31’s disk. The stacked distribution composed of the low-dispersion pixels is nearly symmetric about zero. However, the high-dispersion velocity distribution is skewed: while its mean is at $x = 0$ by

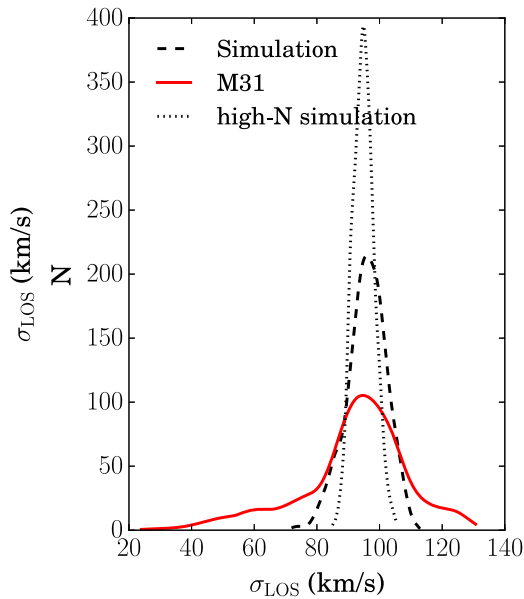


Figure 14. Velocity dispersion distribution for RGB stars from the toy model of a disk with target density comparable to our spectroscopic survey (black dashed line), from a toy model with target density five times higher (black dotted line), and from the M31 data (red line). Finite sample density can account for most of the width of the observed dispersion distribution.

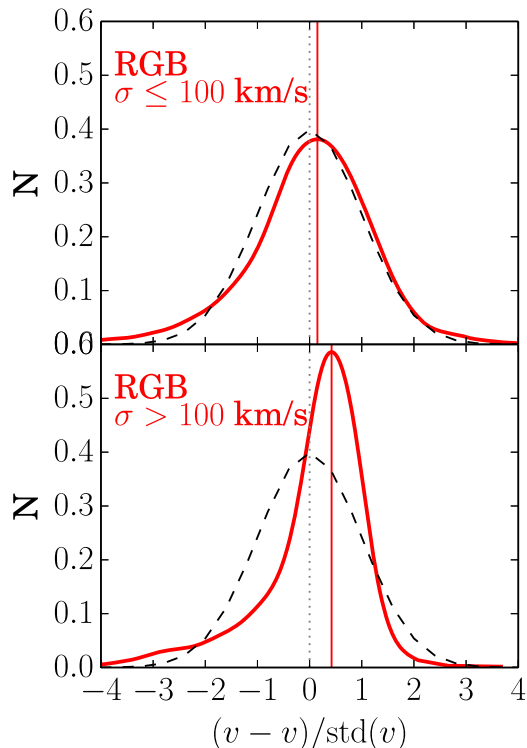


Figure 15. Comparison of composite velocity distribution shapes for RGB stars (solid red curves) to a normal distribution (dashed black curves). The top panel includes pixels with dispersion less than 100 km s^{-1} , and the bottom panel includes the rest of the pixels. By construction, all distributions have the same mean of 0 and variance of 1. In each panel, the modes of the LOSVD and normal distributions are marked by the solid and dotted vertical lines, respectively. The velocity distribution of the high dispersion pixels (bottom panel) has a significant low-velocity tail that is barely present in the low-dispersion distribution. The inflated dispersion of the dynamically hot patches in the RGB dispersion map, then, is due to a contribution from a few low-velocity stars in a second kinematical component (spheroid or tidal debris), rather than to an intrinsically hotter underlying disk. Though not shown, the same trend is seen in the AGB bins, but not the MS bin or in the simulated disk.

construction, its mode is at a positive velocity and it has a wide tail toward negative velocities. The same trend is seen, to a smaller degree, in both the AGB bins, but not in the MS+ bin.

The skew in the high-dispersion pixels is notably absent when the same analysis is run on the simulated toy disk from Section 4.1. In that case, both the high- and low-dispersion pixels have symmetric (that is, one-component) LOSVDs.

In the data, the lower-dispersion pixels always have approximately symmetric velocity distributions, suggesting that they are also dominated by stars from a single kinematical component and that the median dispersions are reliable indicators of the width of the LOSVD of that component. For these lower-dispersion pixels, the correlation between average dispersion and age does indeed reflect a real trend in the degree of heating of the dominant disk. However, within an age bin, the dispersion in dynamically hotter patches is inflated due to a contribution from a few low-velocity stars, rather than to a single, dynamically hot disk component. The spatial variation in dispersion simply reflects inhomogeneity in the number density of this second population.

The spatially inhomogeneous dynamically hot component has been identified before: first, we identified and mapped the population of stars whose high velocities relative to the local disk LOS velocity rendered them likely “inner spheroid” members (Dorman et al. 2012). Then, we found that there are more kinematically identified inner spheroid members than can be accounted for by the disk-dominated surface brightness profile, and that the spheroid members have a luminosity function nearly identical to that of the disk (Dorman et al. 2013). The high-dispersion regions identified in this paper are simply regions that include some of these “inner spheroid” stars. The nonuniform spatial distribution of this population could easily be produced by localized effects such as satellite accretion or disk heating via interactions with the bar or with satellites (as proposed in Dorman et al. 2013) among other things. As in Dorman et al. (2013), the hot component is only distinguishable from the rest of the disk via kinematics, not photometry: RGB stars that form the low-velocity tail have the same distribution in color–magnitude space as the rest of the RGB population.

The hot kinematical component inflates the smoothed velocity dispersion in certain pixels above that of the underlying disk and therefore complicates the comparison to the MW data in Figure 12. The MW data from the Geneva–Copenhagen survey sample only stars within $\sim 40 \text{ pc}$ of the Sun (Nördstrom et al. 2004); since the scale height of old stars in the Galaxy’s thin disk is several hundred parsecs (e.g., Bovy et al. 2012), the GCS data set is strongly dominated by disk stars and includes only minimal contribution from a spatially extended, dynamically hot halo. If the hot population in M31 is much more vertically extended than the underlying colder component, then part of the steep slope of the M31 age–dispersion relation may be due to “contamination” from a halo-like component that was not even probed in the MW survey. However, even in this case, the qualitative point that M31’s disk is hotter than the MW’s still stands: if we eliminate the likely “hot component” stars (those on the low-velocity tail of the high-dispersion patches in Figure 15), the remaining RGB population still has an average dispersion of 80 km s^{-1} , twice as high as in the oldest MW stars in the Nördstrom et al. (2004) sample.

5.2. Disk Evolution Scenarios

In this section, we discuss possible disk formation and evolution scenarios’ ability to reproduce the observed features of M31’s stellar disk: a positive correlation between age and velocity dispersion, a negative correlation between RGB metallicity and velocity dispersion, a high overall dispersion, and the skewed LOSVD in the high-dispersion patches. Dynamically hot populations spatially coincident with a galaxy disk can arise from three broad categories of processes: dynamical heating of an initially thin disk; in situ formation of a thick disk at high redshift; or accretion of satellite debris onto the disk. In this section we walk through the three scenarios. We do not attempt to choose the “correct” scenario; that will require more detailed modeling. However, we show that we can explain all of the observed features with a combination of some mechanism to create an age–dispersion relation (either thin disk heating or thick disk collapse) plus some mechanism to heat the disk nonuniformly (such as accretion of satellites or kicking up of existing disk stars via satellite impacts).

5.2.1. Heated Thin Disk

The canonical explanation of thick disk formation involves heating an initially thin disk via perturbations from internal structures such as a bar, spiral arms, or GMCs (e.g., Sellwood 2014, and references therein); N -body simulations suggest that impacts from satellite galaxies can also significantly heat an existing disk (Purcell et al. 2009; McCarthy et al. 2012; Tissera et al. 2013). If these heating processes happen over an extended period of time, then the heating a stellar population has undergone should correlate with its age, consistent with the trend we see in Figures 7, 9, and 12. It is unclear whether internal heating from GMC scattering or spiral arm perturbations could inflate dispersions by 60 km s^{-1} over a period of a several Gyr, but stars kicked out of the disk as a result of satellite impacts can reach dispersions of over 100 km s^{-1} (Purcell et al. 2010; McCarthy et al. 2012; Tissera et al. 2013).

The young stars in our sample have a velocity dispersion $>50\%$ higher than the young stars in the MW (Nördstrom et al. 2004). If disk evolution is entirely due to heating of an initially thin disk, then M31’s disk must have already been born with a higher dispersion.

Heating from several discrete satellite impacts can explain the skewed LOSVD in patches in the dispersion map.

5.2.2. Progressive Collapse of a Thick Gas Disk

An alternative to the heated thin disk scenario described above involves the collapse of an initially thick gas disk. At early times, a thick, clumpy disk of gas can form a thick disk of stars. The remaining gas later collapses into progressively thinner, dynamically colder layers, forming stars along the way (Bournaud et al. 2009; Forbes et al. 2012). In this way, younger (more recently formed) stars lie in a colder, more metal-rich disk than older stars. Such a scenario produces a continuous age–dispersion correlation. Combined with a small amount of heating from satellite encounters and radial migration, it can approximately reproduce both the magnitude and slope of the age–dispersion relationship in the disk of the MW (Bird et al. 2013; Bird 2015).

While this disk settling formation mechanism can explain the continuous nature of the age–dispersion correlation in M31 as seen in Figure 12, it cannot reproduce the steepness of the age–

dispersion trend ($16 \text{ km s}^{-1} \text{ Gyr}^{-1}$ for the constant SFR or 9 for the declining $\tau = 4 \text{ Gyr}$ SFR) and the very high velocity dispersion. The fiducial model in Forbes et al. (2012) shows that, at 3 radial scale lengths ($R \sim 15 \text{ kpc}$ in M31), the dispersion in an MW-mass galaxy with a decreasing gas accretion rate and a standard radial migration prescription increases only $7 \text{ km s}^{-1} \text{ Gyr}^{-1}$ from 0 to 2 Gyr—half that of M31 assuming a constant SFR, and still 20% less than in M31 in the case of the heavily old star-weighted declining SFR.

Similarly, additional heating is need to explain the extremely high velocity dispersion, which, for the RGB stars, is 40% of the circular velocity of 250 km s^{-1} (Chemin et al. 2009; Corbelli et al. 2010). In contrast, the highest stellar dispersion in models presented in Forbes et al. (2012) is 40 km s^{-1} , or $<20\%$ of the circular velocity.

This scenario also explains the inverse correlation between metallicity and velocity dispersion for the RGB stars, assuming that the more metal-rich RGB population is the younger one.

5.2.3. Satellite Accretion

Cosmological simulations predict that stellar halos are built up via hierarchical merging with smaller galaxies, and the prominent substructure in the outskirts of the MW (e.g., Mathewson et al. 1974; Ibata et al. 1994; Majewski et al. 2003; Newberg et al. 2003; Rocha-Pinto et al. 2003; Yanny et al. 2003) and M31 (Ferguson et al. 2002; Ibata et al. 2004, 2007; Kalirai et al. 2006; Gilbert et al. 2009, 2012; McConnachie et al. 2009; Tanaka et al. 2010) lends support to this idea. Because the total area of the disk is smaller than that of the extended halo, satellite–disk interactions should be rarer than satellite–halo interactions, but not nonexistent. In Section 5.2.1, we discussed how these interactions can heat the existing disk. But stars from the satellite galaxies can also accrete onto the disk, changing the age and metallicity distributions of the disk in addition to dynamically heating it. Debris from at least one tidal disruption event has been found across the face of M31. The Giant Southern Stream, Northeast Shelf, and Western Shelf all show up as kinematically cold substructure (Ibata et al. 2001; Kalirai et al. 2006; Gilbert et al. 2009; Dorman et al. 2012). Accretion from satellites can naturally explain the asymmetric velocity distributions in our RGB population. Since the velocity of the satellite would likely be quite different than the rotation velocity of the disk, regions with accreted stars would end up with inflated LOS velocity dispersions, producing localized structure. However, accretion alone cannot account for the continuous increase in dispersion with age. It also cannot account for the high dispersion of the young MS+ stars relative to the young stars in the MW unless M31’s young stars were born with a higher dispersion.

5.3. Brick 9 Feature: Associated with the Bar?

Many of the high-dispersion measurements lie in the “Brick 9 Region” centered at $(\xi, \eta) = (4.5, 4.5) \text{ kpc}$, marked by the ends of the teal line in Figure 8. This high-dispersion region was identified in RGB stars in Dorman et al. (2012) and Dorman et al. (2013) as having the highest fraction of “inner spheroid” or “kicked-up disk” stars. Here, we see that it is not limited to the RGB sample. The Brick 9 feature is present in all age bins and in both RGB metallicity bins, suggesting that it is not an accreted population or formed by an event that ended more than 1 Gyr ago. It is located on the major axis of the

galaxy, so it is not subject to the geometric dispersion inflation that happens near the minor axis. It does happen to be located on the inner ring (visible in the Herschel image in the lower right panel of Figure 7) and may be coincident with the end of a long bar.

Athanassoula & Beaton (2006) compared the infrared isodensity contours and radial luminosity profiles to N -body simulations of barred galaxies to show that M31 likely has a bar that may extend out to $1300''$ (5 kpc) at a position angle of 45° —placing its NE end just interior to the Brick 9 feature. The effect of bars on stellar kinematics is not well understood, but Athanassoula & Beaton (2006) show that the presence of a bar, when misaligned with the major axis of its host disk, can result in a broad velocity distribution near the end of the bar. Typically, bar members have different orbits from disk members, and so smoothing circles containing both bar and disk members would contain a large spread in LOS component of stellar velocities.

5.4. Comparison to MW

The work presented in this paper shows that M31’s stellar disk is dynamically hotter than the MW’s: the average dispersion of the RGB bin (90 km s^{-1}) is nearly three times as high as the oldest, hottest population probed in the MW by N rdstrom et al. (2004), and the average dispersion of our youngest age bin is more than 50% higher than that in the MW. In addition, the slope of the dispersion versus age relation is more than three times higher in M31 than our own Galaxy. Both suggest that M31 has had a more violent accretion history than the MW in the recent past.

This result is encouraging for Λ CDM cosmology, which predicts that 70% of disks the size of those in the MW and M31 should have interacted with at least one satellite of mass $\sim 3M_{\text{disk}}$ in the past 10 Gyr (Stewart et al. 2008). The MW’s disk is far too cold for it to have undergone such an encounter (Purcell et al. 2010), but the fact that M31’s disk is dynamically warmer leaves open the possibility that Λ CDM predictions are correct and that the MW is simply an outlier with an unusually quiescent history.

Our results support a growing body of evidence that M31 has experienced a more violent merger history than our own Galaxy. Hammer et al. (2007) found that, relative to other local disk galaxies with similar rotation velocities, the MW is deficient in stellar mass, angular momentum, disk radius and metallicity, whereas M31 looks more typical. M31’s halo is littered with substructure that can be identified both in photometry (Ibata et al. 2014) and kinematics (Chapman et al. 2008; Gilbert et al. 2012), and fields at similar radii can host clearly distinct stellar populations, indicating that not all stars share a common origin (Ferguson et al. 2002; Richardson et al. 2008). The number of giant streams discovered in M31 outnumbers that of the MW (Ibata et al. 2007). Furthermore, the slope of the surface brightness profile of M31’s outer halo is shallow, with projected power-law slope $\alpha \sim -2$ (Gilbert et al. 2012; Dorman et al. 2013; Ibata et al. 2014). In comparison, the surface brightness of the MW halo exterior to 50 kpc as traced by BHB stars is much steeper, falling as r^{-6} , suggesting a more quiescent recent accretion history (Deason et al. 2014). M31’s gas disk is also strongly warped in the outer regions: the major-axis position angle of the gas disk changes by 10° , and the inclination by 7° , between $R = 20$ and 40 kpc (Chemin et al. 2009; Corbelli et al. 2010).

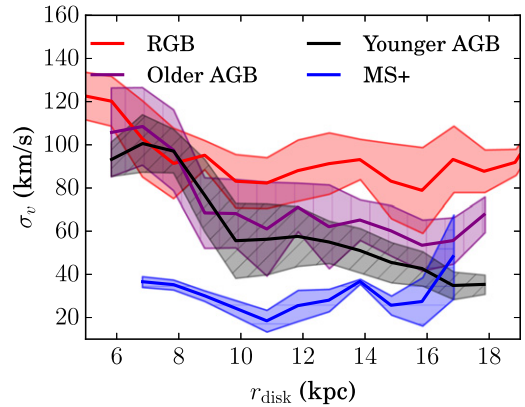


Figure 16. Dispersion as a function of radius in the plane of the disk r_{disk} for each age bin. Age bins are color coded as in previous plots. The solid line and shaded region correspond to the median and standard deviation of the velocity dispersion in each 1 kpc wide radial bin. r_{disk} is computed assuming a constant major-axis position angle of 38° and disk inclination of 77° . Only stars within $\pm 20^\circ$ of the major axis are used, since near the minor axis distance measurements become more sensitive to precision in inclination, and a single smoothing circle may cover a wide range of radii. Velocity dispersion decreases with radius in all bins out to about 10 kpc.

A quantitative comparison between the velocity dispersion of M31’s old stellar disk and cosmological predictions requires the measurement of the shape of M31’s velocity ellipsoid. We save this measurement for a future paper (C. E. Dorman et al. 2015, in preparation) because it requires the choice of a particular disk rotation curve model.

5.5. Radial Trends

In Figure 16, we show velocity dispersion as a function of radius: the median and standard deviation of the velocity dispersion in each 1 kpc wide radial bin in the plane of the disk for each age bin. We only include stars within $\pm 20^\circ$ of the major axis, because near the minor axis a single smoothing circle over which dispersion is calculated may cover a wide range of radii.

The figure reveals two distinct radial regions in the galaxy. Inside 10 kpc, the velocity dispersions of all age groups are inflated by a factor of 1.5–2 from their values outside, and the dispersions of the three older age groups are similar to each other, especially interior to 8 kpc. The dispersions of all age groups decline steeply with radius in this inner region. Outside the B9 region, the velocity dispersions of the different age groups separate from one another, revealing the clean age versus dispersion relationship, and decline less steeply with radius. The transition radius may correspond to a transition in the dominant heating mechanism: in the outer regions, the clean age–dispersion relation implies a continuous heating process over the past several gigayears. Inside, a single perturbation has modified the orbits of stars older than 100 Myr. This could be a single event such as a very recent satellite impact, or a large fraction of the older stars in this region could belong to a single, kinematically distinct component such as the bar or a thick disk with short scale length.

The 10 kpc break corresponds to a bit less than two radial scale lengths (Ibata et al. 2005; Worthey et al. 2005; Seigar et al. 2008; Dorman et al. 2013). This is consistent with the “mass follows light” radial trend seen in stellar kinematics in the DiskMass survey of nearly face-on, late-type spiral galaxies

Table 1
Keck/DEIMOS Multiobject Slitmask Exposures from Fall 2011

Mask Name	Observation Date (UT)	α [J2000] (h m s)	δ [J2000] ($^{\circ}$ ' ")	P.A. ($^{\circ}$)	t_{exp} (s)	Seeing FWHM	No. of Slits	No. of Usable Target Velocities (Success Rate)	No. of Usable Velocities of Serendipitously Detected Stars
mct6C	2012 Sept 18	00 44 47.17	41 22 00.0	-140.0	3×1020	0".6	225	184 (82%)	31
mct6D	2012 Sept 20	00 44 34.60	41 29 44.6	170.0	3×1200	0".8	208	167 (80%)	38
mct6E	2012 Sept 18	00 44 13.52	41 19 05.5	-20.0	2×1080	0".6	228	107 (47%)	16
mct6F	2012 Sept 19	00 45 54.72	41 41 58.6	-164.0	3×1020	0".6	221	192 (87%)	31
mct6G	2012 Sept 20	00 45 38.34	41 43 37.4	-155.0	3×1200	0".9	231	178 (77%)	24
mct6H	2012 Sept 18	00 45 26.92	41 44 04.3	+15.0	3×1020	0".6	244	185 (76%)	24
mct6I	2012 Sept 19	00 44 10.25	41 25 16.2	-95.0	3×1080	0".7	209	138 (66%)	52
mct6K	2012 Sept 18	00 44 38.26	41 37 22.5	-95.0	3×1020	0".7	207	155 (75%)	45
mct6L	2012 Sept 19	00 46 05.81	42 02 28.1	-20.0	3×1020	0".8	246	212 (86%)	8
mct6M	2012 Sept 20	00 44 36.98	41 32 39.5	-20.0	3×1080	0".75	227	153 (67%)	40
mct6O	2012 Sept 20	00 45 08.17	41 52 34.0	-80.0	$2 \times 1080 + 1 \times 1140$	0".93	240	169 (70%)	24
mct6P	2012 Sept 19	00 45 30.50	41 55 37.6	-80.0	$2 \times 1019 + 1 \times 855$	0".6	229	161 (70%)	10
mct6Q	2012 Sept 18	00 45 35.98	42 00 17.9	-80.0	3×1080	0".8	225	155 (69%)	21
mct6R	2012 Sept 19	00 47 02.28	42 12 07.3	-80.0	3×1080	0".8	227	175 (77%)	9
mct6S	2012 Sept 20	00 47 02.27	42 09 25.2	-80.0	3×1200	0".8	224	163 (73%)	18
mct6T	2012 Sept 18	00 45 57.60	42 01 12.0	-40.0	3×1020	0".6	245	210 (86%)	10
mct6U	2012 Sept 20	00 46 13.24	42 14 35.6	+50.0	3×1200	0".8	254	185 (73%)	5
mct6V	2012 Sept 18	00 46 40.80	42 13 48.0	-130.0	$2 \times 1080 + 1 \times 1020$	0".6	226	126 (56%)	13
mct6W	2012 Sept 19	00 46 51.73	42 12 46.0	-130.0	3×1020	0".6	225	175 (78%)	8
mct6X	2012 Sept 19	00 46 22.94	42 01 58.0	140.0	3×1020	0".7	225	173 (77%)	15
Total:							4566	3368 (74%)	442

(Martinsson et al. 2013): the LOS stellar velocity dispersion typically decreases with radius out to two scale lengths, after which the data become noisy enough that the trend becomes consistent with flat. It is worth noting that because the DiskMass survey galaxies are much closer to face-on than Andromeda is ($<30^{\circ}$ versus 77°), the DiskMass LOS dispersion profiles are sensitive to the vertical component of the velocity dispersion σ_z while our LOS velocity dispersion has a negligible contribution from σ_z .

6. SUMMARY

We have split a spectroscopic sample of 5800 stars from the disk of M31 into four age groups based on optical *HST* photometry and mapped the LOS velocity dispersion of each age group. We have found:

1. Stellar velocity dispersion and age are directly correlated in the disk of M31. The slope of the best-fit line assuming a constant SFR is approximately $16 \text{ km s}^{-1} \text{ Gyr}^{-1}$, more than five times higher than that of the MW.
2. Among RGB stars in our sample, the metal-poor half is 50% kinematically hotter than the metal-rich half.
3. The stellar disk is kinematically clumpy on scales smaller than about $200''$ (760 pc). Most of the structure can be attributed to sampling effects. However, the LOS velocity distribution of the RGB and AGB stars in the highest-dispersion regions is quite asymmetric, with a low-velocity tail. This asymmetry indicates the presence of an additional kinematic component that is not smoothly distributed across the galaxy.
4. There is a patch on the major axis about 6 kpc from the galactic center where the dispersions of all four age

groups are inflated by a factor of 1.5–2. This patch may correspond to the end of the long bar.

P.G. and C.D. acknowledge NSF grants AST-1010039 and AST-1412648 and NASA grant *HST*-GO-12055. C.D. was supported by an NSF Graduate Research Fellowship and a Eugene Cota-Robles Graduate Fellowship. D.R.W. is supported by NASA through Hubble Fellowship grant *HST*-HF-51331.01 awarded by the Space Telescope Science Institute. We thank Emily Cunningham, John Forbes, Leo Girardi, Dylan Gregersen, Cliff Johnson, Connie Rockosi, and Graeme Smith for helpful conversations. Finally, we are grateful to the anonymous referee for constructive suggestions that greatly improved the clarity of the paper. We acknowledge the very significant cultural role and reverence that the summit of Mauna Kea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain.

APPENDIX

This paper presents, for the first time, resolved stellar kinematics derived from Keck/DEIMOS observations from Fall 2012. Table 1 lists the slitmasks observed and the number of quality radial velocities recovered from each.

The target selection was done in a non-uniform way in order to maximize the number of younger (MS, AGB, blue and red supergiant) stars relative to the dominant RGB population. As such, the spectroscopic selection function varies with color, magnitude, and position on the sky, though it can be easily recovered by comparing the distribution of targets to the full PHAT catalog.

The raw data were flat-fielded, locally sky subtracted, collapsed into 1D spectra, and cross-correlated against template spectra to measure radial velocities as described in Dorman et al. (2012, 2013).

REFERENCES

- Abadi, M. G., Navarro, J. F., Steinmetz, M. S., & Eke, V. R. 2003, *ApJ*, 597, 21
- Athanassoula, E., & Beaton, R. L. 2006, *MNRAS*, 370, 1499
- Bernard, E. J., Ferguson, A. M. N., Barker, S. L. H., et al. 2012, *MNRAS*, 420, 2625
- Bird, J. C. 2015, *AAS*, 225, 227.03
- Bird, J. C., Kazantzidis, S., Weinberg, D. H., et al. 2013, *ApJ*, 773, 43
- Bovy, J., Rix, H.-W., & Hogg, D. W. 2012, *ApJ*, 751, 131
- Bressan, A., Marigo, P., Girardi, L., et al. 2012, *MNRAS*, 427, 127
- Brown, T. M., Smith, E., Ferguson, H. C., et al. 2006, *ApJ*, 652, 323
- Bournaud, F., Elmegreen, B. G., & Martig, M. 2009, *ApJL*, 707, L1
- Budembender, A., van de Ven, G., & Watkins, L. L. 2014, *MNRAS*, submitted (arXiv:1407.4808)
- Chapman, S. C., Ibata, R., Irwin, M., et al. 2008, *MNRAS*, 390, 1437
- Chiba, M., & Beers, T. C. 2000, *AJ*, 119, 2843
- Chemin, L., Carignan, C., & Foster, T. 2009, *ApJ*, 705, 1395
- Collins, M. L. M., Chapman, S. C., Ibata, R. A., et al. 2011, *MNRAS*, 413, 1548
- Corbelli, E., Lorenzoni, S., Walterbos, R., Braun, R., & Thilker, D. 2010, *A&A*, 511, A89
- Courteau, S., Widrow, L. M., McDonald, M., et al. 2011, *ApJ*, 739, 20
- Dalcanton, J. J., Fouesneau, M., Hogg, D. W., et al. 2015, *ApJ*, submitted
- Dalcanton, J. J., Williams, B. F., Lang, D., et al. 2012, *ApJS*, 200, 18
- Deason, A. J., Belokurov, V., Koposov, S. E., & Rockosi, C. M. 2014, *ApJ*, 787, 30
- Dorman, C. E., Guhathakurta, P., Fardal, M., et al. 2012, *ApJ*, 752, 147
- Dorman, C. E., Guhathakurta, P., Widrow, L., et al. 2013, *ApJ*, 779, 103
- Ferguson, A. M. N., Irwin, M. J., Ibata, R. A., Lewis, G. F., & Tanvir, N. R. 2002, *AJ*, 124, 1452
- Forbes, J., Krumholz, M., & Bukert, A. 2012, *ApJ*, 754, 48
- Gilbert, K. M., Fardal, M., Kalirai, J. S., et al. 2007, *ApJ*, 668, 245
- Gilbert, K. M., Font, A. S., Johnston, K. V., & Guhathakurta, P. 2009, *ApJ*, 701, 776
- Gilbert, K. M., Guhathakurta, P., Beaton, R. L., et al. 2012, *ApJ*, 760, 76
- Gilbert, K. M., Guhathakurta, P., Kalirai, J. S., et al. 2006, *ApJ*, 652, 1188
- Girardi, L., Williams, B. F., Gilbert, K. M., et al. 2010, *ApJ*, 724, 1030
- Hammer, F., Puech, M., Chemin, L., Flores, H., & Lehnert, M. D. 2007, *ApJ*, 667, 322
- Howley, K. M., Guhathakurta, P., van der Marel, R., et al. 2013, *ApJ*, 765, 65
- Ida, S., Kokuba, E., & Makino, J. 1993, *MNRAS*, 263, 875
- Ibata, R., Chapman, S., Ferguson, A. M. N., et al. 2004, *MNRAS*, 351, 117
- Ibata, R., Chapman, S., Ferguson, A. M. N., et al. 2005, *ApJ*, 634, 287
- Ibata, R., Irwin, M., Lewis, G., Ferguson, A. M. N., & Tanvir, N. 2001, *Natur*, 412, 49
- Ibata, R., Martin, N. F., Irwin, M., et al. 2007, *ApJ*, 671, 1591
- Ibata, R. A., Gilmore, G., & Irwin, M. J. 1994, *Natur*, 370, 194
- Ibata, R. A., Lewis, G. F., McConnachie, A. W., et al. 2014, *ApJ*, 780, 128
- Johnson, L. C., Seth, A. C., Dalcanton, J. J., et al. 2012, *ApJ*, 752, 95
- Kalirai, J. S., Guhathakurta, P., Gilbert, K. M., et al. 2006, *ApJ*, 641, 268
- Majewski, S. R., Skrutskie, M. F., Weinberg, M. D., & Ostheimer, J. C. 2003, *ApJ*, 599, 1082
- Martinsson, T. P. K., Verheijen, M. A. W., Westfall, K. B., et al. 2013, *A&A*, 557, A130
- Mathewson, D. S., Cleary, M. N., & Murray, J. D. 1974, *ApJ*, 190, 291
- McCarthy, I. G., Font, A. S., Crain, R. A., et al. 2012, *MNRAS*, 420, 2245
- McConnachie, A. W., Irwin, M. J., Ferguson, A. M. N., et al. 2005, *MNRAS*, 356, 979
- McConnachie, A. W., Irwin, M. J., Ibata, R. A., et al. 2009, *Natur*, 461, 66
- Newberg, H. J., Yanny, B., Grebel, E. K., et al. 2003, *ApJL*, 596, L191
- Nordstrom, B., Mayor, B., Andersen, J., et al. 2004, *A&A*, 418, 989
- Pryor, C., & Meylan, G. 1993, *ASPC*, 50, 375
- Purcell, C. W., Bullock, J. S., & Kazantzidis, S. 2010, *MNRAS*, 404, 1711
- Purcell, C. W., Kazantzidis, S., & Bullock, J. S. 2009, *ApJL*, 694, L98
- Quinn, P. J., Hernquist, L., & Fullager, D. P. 1993, *ApJ*, 403, 74
- Richardson, J. C., Ferguson, A. M. N., Johnson, R. A., et al. 2008, *AJ*, 135, 1998
- Rocha-Pinto, H. J., Majewski, S. R., Skrutskie, M. F., & Crane, J. D. 2003, *ApJL*, 594, L115
- Schiavon, R. P., Barbuy, B., Rossi, S. C. F., & Milone, A. 1997, *ApJ*, 479, 902
- Seigar, M. S., Barth, A. J., & Bullock, J. S. 2008, *MNRAS*, 389, 1911
- Sellwood, J. A. 2014, *RvMP*, 86, 1
- Seth, A. C., Dalcanton, J. J., & de Jong, R. S. 2005, *ApJ*, 130, 1574
- Stewart, K. R., Bullock, J. S., Wechsler, R. H., Maller, A. H., & Zentner, A. R. 2008, *ApJ*, 683, 597
- Tanaka, M., Chiba, M., Komiyama, Y., et al. 2010, *ApJ*, 708, 1168
- Tissera, P. B., Scannapieco, C., Beers, T. C., & Carollo, D. 2013, *MNRAS*, 432, 3391
- Vanhollebeke, E., Groenewegen, M. A. T., & Girardi, L. 2009, *A&A*, 498, 95
- Velazquez, H., & White, S. D. M. 1999, *MNRAS*, 304, 254
- Williams, B. F., Lang, D., Dalcanton, J. J., et al. 2014, *ApJ*, 215, 9
- Worthey, G., Espana, A., MacArthur, L. A., & Courteau, S. 2005, *ApJ*, 631, 820
- Yanny, B., Newberg, H. J., Grebel, E. K., et al. 2003, *ApJ*, 588, 824
- Yoachim, P., & Dalcanton, J. J. 2006, *AJ*, 131, 226