

UC Santa Cruz

UC Santa Cruz Electronic Theses and Dissertations

Title

Structure and Dynamics of Andromeda's Stellar Disk

Permalink

<https://escholarship.org/uc/item/67h0z1b6>

Author

Dorman, Claire E.

Publication Date

2015

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives License, available at <https://creativecommons.org/licenses/by-nc-nd/4.0/>

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

**STRUCTURE AND DYNAMICS OF
ANDROMEDA'S STELLAR DISK**

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

ASTRONOMY & ASTROPHYSICS

by

Claire Elise Dorman

June 2015

The Dissertation of Claire Elise Dorman
is approved:

Professor Puragra GuhaThakurta, Chair

Professor Constance Rockosi

Professor Anil Seth

Dean Tyrus Miller
Vice Provost and Dean of Graduate Studies

Copyright © by

Claire Elise Dorman

2015

Table of Contents

List of Figures	vi
List of Tables	viii
Abstract	ix
Acknowledgments	xi
1 Introduction	1
2 Spectroscopic Survey of M31's Stellar Disk	6
3 Age-dispersion relation	12
3.1 INTRODUCTION	12
3.2 Data acquisition and reduction	16
3.3 Methods: Velocity dispersion as a function of age and of metallicity	19
3.3.1 Definition of age bins	22
3.3.2 Separation of data into age bins	30
3.3.3 Velocity dispersion maps	32
3.3.4 RGB velocity dispersion maps as a function of metallicity	38
3.4 Results	41
3.4.1 Age-dispersion relation	41
3.4.2 Age-metallicity relation among RGBs	44
3.4.3 Structure in dispersion map	44
3.5 Discussion	46
3.5.1 High-dispersion population: superposition of two kinematical components	46
3.5.2 Disk evolution scenarios	51
3.5.3 Brick 9 Feature: Associated with the Bar?	56
3.5.4 Comparison to Milky Way	57
3.5.5 Radial trends	58
3.6 Summary	59

4	M31's Inner Spheroid	62
4.1	INTRODUCTION	62
4.2	OBSERVATIONS	68
4.2.1	Data Sets	70
4.2.2	Source Catalogs	71
4.2.3	Isolated Target Selection	71
4.2.4	Observations	77
4.2.5	Data Reduction	79
4.2.6	Cross-Correlation Analysis	79
4.2.7	Quality Assessment	80
4.2.8	Serendipitous Sources	80
4.3	DATA ANALYSIS	81
4.3.1	Choice of Subregions and Expected Disk LOS Velocity Pattern	82
4.3.2	Fitting the Velocity Distribution Model	84
4.3.3	Accounting for Tidal Debris Associated with the Giant Southern Stream	86
4.4	DISCUSSION	89
4.4.1	Kinematical Parameters of the Inner Spheroid	90
4.4.2	Dispersion and Velocity of the Cold Population	92
4.4.3	Effect of Tidal Debris Associated with the GSS on Spheroid Kinematics	94
4.4.4	Spheroid/Disk Membership Probability and Extreme Velocity Stars	95
4.4.5	Anisotropy	96
4.4.6	Effect of Target Selection Criteria on Velocity Distribution	98
4.5	SUMMARY & CONCLUSIONS	110
4.6	MCMC results for regions NE2, NE3, SE, and SSW	111
4.6.1	NE2 Region	111
4.6.2	NE3 Region	112
4.6.3	SE Region	112
4.6.4	SSW Region	113
5	Structural Decomposition	120
5.1	INTRODUCTION	120
5.2	Overview of Analysis Procedure	125
5.3	Observational Constraints	126
5.3.1	Surface Brightness	127
5.3.2	Velocity distribution	130
5.3.3	Luminosity Function	136
5.4	Analysis	139
5.4.1	Model	140
5.4.2	Model-Data Comparison: Likelihood Function	148
5.4.3	MCMC Sampler	149
5.5	Results	150
5.5.1	Confidence Intervals & Correlations	150
5.5.2	Quality of Profile Fits	151

5.5.3	Comparison to previous measurements	153
5.5.4	Conversion between integrated-light and star-count disk fractions	156
5.6	Discussion	157
5.6.1	Kicked-up disk	157
5.6.2	Relationship between bulge and halo	162
5.6.3	Radial Gradient in the Disk LF	165
5.7	Summary	167
5.8	MCMC Sampler	187
5.9	Parameter Distributions	189
6	Rotation curves, velocity dispersion profiles, and the asymmetric drift	192
6.1	INTRODUCTION	192
6.2	Data	194
6.2.1	Stars: Photometry	194
6.2.2	Stars: Spectroscopy	195
6.2.3	Atomic Hydrogen Gas	196
6.3	Age Groups	196
6.3.1	Using PHAT photometry: Summary of previous work	196
6.3.2	Using CFHT-based photometry	197
6.4	Kinematics	198
6.4.1	Projected velocity and dispersion maps	198
6.4.2	Deprojection	201
6.5	Rotation curve and dispersion profiles	206
6.6	Asymmetric drift	208
6.7	Summary	214
7	Conclusions and Future Directions	217

List of Figures

3.1	Spatial coverage of SPLASH and PHAT surveys after 2012 run	20
3.2	Example SG, AGB, and RGB Keck/DEIMOS spectra	21
3.3	Na I discriminant for identifying likely MW dwarfs	25
3.4	Expected density of MW foreground contaminants in SPLASH survey .	26
3.5	Simulated CMD assuming a constant SFR	30
3.6	PHAT and SPLASH optical CMDs	32
3.7	Line-of-sight velocity dispersion maps for stars in 4 age groups	35
3.8	RGB dispersion map compared to Herschel M31 image	36
3.9	1D velocity dispersion distributions for stars in 4 age bins	37
3.10	Velocity dispersion maps for low-and high-metallicity RGB stars	39
3.11	1D velocity dispersion distributions for low- and high-metallicity RGB stars	40
3.12	Age-dispersion relation for M31 and the Milky Way	43
3.13	Velocity dispersion map for a toy model of a disk	47
3.14	1D velocity dispersion distribution for toy model of disk	48
3.15	Skewed RGB velocity distribution	52
3.16	Velocity dispersion profiles for 4 age bins	60
4.1	Keck/DEIMOS multiobject slitmasks observed in 2010	69
4.2	<code>chi/sharp</code> criteria for identifying isolated targets	77
4.3	RMS criterion for identifying isolated targets	78
4.4	Division of 2010 data set into spatial regions and subregions	99
4.5	Radial velocities for stars in 2010 data set.	100
4.6	Maximum likelihood disk/spheroid fits to the LOS velocity distribution in 5 regions.	101
4.7	Radial velocity vs. projected radius: evidence for stream debris	102
4.8	Disk/spheroid kinematical decompositions in the NE1 region	103
4.9	Inner spheroid velocity and dispersion profiles	104
4.10	Inner spheroid velocity and dispersion profiles, excluding tidal debris members	105
4.11	Disk/spheroid membership probability map	106
4.12	Anisotropy diagram for inner spheroid and other systems	107

4.13	Disk/spheroid kinematical decompositions for the NE2 region	114
4.14	Disk/spheroid kinematical decompositions for the NE3 region	115
4.15	Disk/spheroid kinematical decompositions for the SE region	115
4.16	Disk/spheroid kinematical decompositions for the SSW region	116
4.17	Velocity and dispersion profiles for the cold disk	117
5.1	Flow chart illustrating the structural decomposition analysis procedure .	126
5.2	Locations of fields used for surface brightness measurements	169
5.3	Zoom-in on SB field map of the inner 15 kpc of the galaxy	170
5.4	Major- and minor-axis projections of the SB profiles	171
5.5	Map of target locations from the PHAT and SPLASH surveys	172
5.6	Disk/spheroid velocity decomposition in the NE2 region	173
5.7	Probability distribution of the kinematically derived disk fraction in sub- region NE2 ₄	174
5.8	PHAT and SPLASH luminosity functions and SPLASH selection function	175
5.9	Correlated pairs of structural parameters	176
5.10	Minor axis bulge/disk/halo surface brightness decomposition	177
5.11	Best-fit luminosity function decomposition	178
5.12	Best fit bulge, disk, and halo luminosity functions	179
5.13	Tension between kinematically-derived cold fraction and best-fit model disk fraction	180
5.14	Kicked-up fraction in each subregion	181
5.15	PDF of total kicked-up fraction	182
5.16	Map of bulge/spheroid SB fraction	183
5.17	Comparison between best-fit halo SB and BHB profiles	184
5.18	Best-fit bulge, disk, and halo luminosity functions, assuming a constant disk LF shape	185
5.19	1D posterior probability distribution of each parameter	190
5.20	Figure 5.19, continued.	191
6.1	Division into 3 age bins	199
6.2	Distinguishing between AGB and RGB stars in the sample without multi- band PHAT photometry.	200
6.3	Velocity, mean velocity, and dispersion fields for stars in three age bins .	202
6.4	Same as Figure 6.3, for neutral hydrogen data from Chemin et al. (2009).	203
6.5	Effect on σ_ϕ of anisotropy parameter assumption	205
6.6	Rotation curves and dispersion profiles for stars and HI gas	207
6.7	Stellar velocity dispersion as a function of lag for the RGB and AGB stars.	211
6.8	Comparison between predicted and observed lag-dispersion relations. . .	215

List of Tables

2.1	Spectroscopic Data	8
2.1	Spectroscopic Data	9
2.1	Spectroscopic Data	10
2.1	Spectroscopic Data	11
4.1	Spheroid Kinematical Parameters	108
4.2	Spheroid Kinematical Parameters Corrected for Tidal Debris	109
4.3	Disk Kinematical Parameters Corrected for Tidal Debris	118
4.3	Disk Kinematical Parameters Corrected for Tidal Debris	119
5.1	Skew-Normal Parameters of Disk Fraction Probability Distribution $p_s(f_d)$	136
5.2	Model Parameters	186
5.2	Model Parameters	187
6.1	RGB Stellar Disk	211
6.2	AGB Stellar Disk	212
6.3	MS+ Stellar Disk	213
6.4	HI Disk	213

Abstract

Structure and Dynamics of Andromeda's Stellar Disk

by

Claire Elise Dorman

Lambda cold dark matter (Λ CDM) cosmology predicts that the disks of Milky Way-mass galaxies should have undergone at least one merger with a large (mass ratio 1:10) satellite in the last several Gyr. However, the stellar disk in the solar neighborhood of the Milky Way is too thin and dynamically cold to have experienced such an impact. The dynamics of the nearby Andromeda galaxy can serve as a second data point, and help us understand whether the Milky Way may simply have had an unusually quiescent merger history, or whether LCDM theory needs to be revisited. Over the last few years, we have carried out a detailed study of the resolved stellar populations in the disk of the Andromeda galaxy using data from two surveys: six-filter Hubble Space Telescope photometry from the recently-completed Panchromatic Hubble Andromeda Treasury (PHAT) survey, and radial velocities derived from Keck/DEIMOS optical spectra obtained as part of the Spectroscopic and Photometric Landscape of Andromeda's Stellar Halo (SPLASH) program. These detailed, multidimensional data sets allow us to decouple the structural subcomponents and characterize them individually. We find that an old, dynamically hot (velocity dispersion ~ 150 km/s) RGB population extends out to 20 kpc (the edge of the visible disk) but has a disk-like surface brightness profile and

luminosity function. This population may have originated in the disk but been kicked out subsequently in impacts with satellite galaxies. We also study the kinematics of the disk as a function of the age of stellar tracers, and find a direct correlation between age and velocity dispersion, indicating that Andromeda has undergone a continuous heating or disk settling process throughout its lifetime. Overall, both the velocity dispersion of Andromeda's disk and the slope of the velocity dispersion vs. stellar age curve are several times those of the Milky Way's, suggesting a more active merger history more in line with Λ CDM cosmological predictions.

Acknowledgments

I owe my PhD to everyone who was part of my life over the last six years, including the following people who have made an especially big impact.

I am first grateful to my family. My parents' commitment to and passion for education gave me the skills necessary to effectively execute a six-year-long, self-paced project. In addition to teaching me to love learning and not to give up, my parents instilled in me some bits of wisdom that have served me well in writing this thesis. I know that I should not begin my dissertation with the line "The thrust of this project is newness," or include in it a description of "a hill o'beans." Thank you to my siblings, Eric, Leah, and Sarah, for making me laugh whenever I need it, for achieving lots of achievements that I can brag about, and for staying in California so that we can still celebrate birthdays and holidays together.

Thank you to my advisors: First, to my PhD advisor, Raja GuhaThakurta, for teaching me nearly everything I know about astronomy and being infinitely patient with my unpredictable work schedule. Knowing that I'd always be welcomed back made my illness infinitely more tolerable. Thank you for providing me with so many teaching opportunities: teaching at UCSC and at Castilleja School and mentoring high school and undergraduate students enriched my graduate experience. Thank you also to my UCLA REU advisors, Sylvana Yelda and Andrea Ghez, for introducing me to astronomical research and teaching me how to give accessible presentations.

The graduate students at UCSC have made grad school worth it. Thank you to my classmates, Rosalie McGurk, Angie Wolfgang, Nathan Goldbaum, and Jerome

Fang, for making an effort to build a community right when we got here. We all made it through! Thank you to my middle academic sister, Katie Hamren, for making me feel useful by asking me questions to which I knew the answers. Thank you to my other sister, Emily Cunningham, for your friendship and hugs and for consuming innumerable cocktails and truffle fries at the 515 with me. Thank you to my two-time co-czar-czar and four-year officemate, Caroline Morley, for striking the the right balance of pushing me to action and preventing me from doing dumb things too hastily. Anyone I work with in the future will have a lot to live up to: you've set that collaboration bar pretty high. Thank you to Anne Medling for being the first person to say hi to me when I arrived at UCSC and for being my friend ever since. Some day we'll make another quilt together, even if we have to send pieces across the Pacific. Thank you to Zach Jennings for sending me dumb jokes even after you were released from beer czar duties. Thank you to Rosalie McGurk for the incredible feat of celebrating everyone's birthday with delicious baked goods. Thank you to Justin Brown for being my most entertaining officemate and providing me with endless amusing stories and putting up with my incessant teasing. Thank you to Kat Feng, Tiffany Hsyu, Rosalie McGurk, Anne Medling, and Camille Leibler for letting me hijack your kitchens to make cupcakes and other tasty treats. Thank you to all my officemates over the years: Edmond Cheung, Eric Lopez, Rosalie McGurk, Nathan Goldbaum, Jerome Fang, Angie Wolfgang, Anna Rosen, Jennifer Burt, Caroline Morley, James Guillochon, Marie Lau, Justin Brown, and Neil Miller. Going into the office and seeing you for even five minutes always gives me the energy to keep trying for another day.

Dancing has helped keep me sane over the last several years. I would like to thank Lisa Norris for introducing me to jazz and Rena Cochlin for constantly pushing me to use ballet technique as a foundation for dance rather than a set of limits, and also overcome my fear of choreography. I also thank my friends Masha Bluestein, Megan Boysen, Kaylie Caires, Kristen Kovaletz, Alex Law, Amelia Nommensen, Juliet Paramor, Helen Porter, Alyssa Soderberg, Caroline Trent, and David Vo for including me in UCSC's beautiful performances.

I am proud to be an Ohlone Community College alumna. All the best teachers I've had and all the best classes I've taken have been at Ohlone. It was there that I discovered my love for science and teaching. It was there that I learned how to study, to manage my time, to critically analyze text, to write papers, to work with people with different academic backgrounds, and to teach.

I am thankful to the National Science Foundation and the University of California for funding all six years of my graduate program.

The text of this dissertation includes reprints of the previously published material. Chapter 3 was published as Dorman et al., *ApJ* 2015, 803, 24, in collaboration with Puragra Guhathakurta, Anil Seth, Dan Weisz, Eric Bell, Julianne Dalcanton, Karoline Gilbert, Katherine Hamren, Alexia Lewis, Evan Skillman, Elisa Toloba, and Benjamin Williams. Chapter 4 was published as Dorman et al., *ApJ* 2012, 752, 147, in collaboration with Puragra Guhathakurta, Mark Fardal, Dustin Lang, Marla Geha, Kirsten Howley, Jason Kalirai, James Bullock, Jean-Charles Cuillandre, Julianne, Dalcanton, Karoline Gilbert, Anil Seth, Erik Tollerud, Benjamin Williams, and Basilio Yniguez.

Chapter 5 was published as Dorman et al., ApJ 2013, 779, 103, in collaboration with Lawrence Widrow, Puragra Guhathakurta, Anil Seth, Daniel Foreman-Mackey, Eric Bell, Julianne Dalcanton, Karoline Gilbert, Evan Skillman, and Benjamin Williams. Chapter 6 is in preparation; Puragra Guhathakurta, Laurent Chemin, and Katherine Hamren have contributed. I performed the observations and data reduction for all of the data except those centered around the galaxy M32 (which were performed by Kirsten Howley) and on the M31 minor axis (which were performed by Karoline Gilbert). I led the analysis and paper writing for all of the work described in this thesis. I thank all of my collaborators for their support and permission to include this work in this dissertation.

Chapter 1

Introduction

The inner regions of a late-type galaxy, including our own Milky Way, are composed of several cospatial, but kinematically and chemically distinct, structural subcomponents: a bulge, a halo, and multiple disks of various scale heights (e.g., Chiba & Beers, 2000; Yoachim & Dalcanton, 2006). This picture is overly simplistic, of course: not all galaxies have both a thin and a thick disk, and there is evidence that the observed disk bimodality is artificial and instead galaxies, including the Milky Way, include more complex disks where scale height increases and scale length decreases smoothly with stellar age (Bovy et al., 2012). Moreover, tidal debris from recently stripped satellites may be littered across any part of the galaxy. In any galaxy, the structure of and relationship among its constituent subcomponents is influenced by the evolution processes that the galaxy has undergone throughout its history. Therefore, we can constrain possible evolution mechanisms via detailed observations of stellar disks.

There are many possible factors that can influence the structure and dynamics

of the disk-dominated region of a galaxy. First, it is not clear how disks form in the first place: oldest stars can form either in an initially thin disk or in a thick, clumpy gas disk that collapses over time to produce progressively thinner, dynamically cooler layers of younger stars (Bournaud et al., 2009). Any time after the initial formation of the old disk, the existing disk can be dynamically heated either by internal perturbers such as the bar, spiral arms, or giant molecular clouds (Ida et al., 1993), or by external perturbers such as satellite galaxies on radial orbits (Purcell et al., 2010; McCarthy et al., 2012; Tissera et al., 2013; Pillepich et al., 2015). Additional stars can also be accreted onto the disk via tidal disruption of satellites; these accreted populations appear as tidal streams or clumps and relax and dynamically mix over time with the rest of the *in situ* stellar population.

However, the observations necessary to understand the relative contributions of these mechanisms do not yet exist. The Milky Way is the only galaxy in which it is possible to measure both photometry and kinematics of resolved stellar populations, and measurements in our home galaxy have been the basis of comparison to many sophisticated galaxy evolution simulations. However, Milky Way studies suffer from two limitations: First, the Milky Way is only one system, and even the most complete description of it cannot help us understand the *diversity* of galactic histories. Second, though, our internal perspective of our own galaxy means that we are limited to *local* measurements of the properties of the solar neighborhood; we have no way of describing the global properties of the disk or variations across the disk. These are both serious limitations that have cosmological limitations: the solar neighborhood of the Milky

Way is kinematically much less disturbed than simulations predict for a "typical" MW-mass galaxy with a cosmologically motivated merger history. Observations of additional systems help us understand whether the Milky Way is simply unusually quiescent or whether Λ CDM systematically overpredicts merger frequency and needs to be revisited.

In this thesis, I address both of the above limitations by exploring the detailed structure and kinematics of the inner 20 kpc (~ 4 disk scale lengths) of a second large spiral galaxy, one in which we are not embedded: Andromeda (M31). We address several different aspects of the structure and dynamics of our galactic neighbor, with the ultimate goal of understanding the factors that contributed to the dynamical heating and/or cooling (settling) of the stellar disk.

In Chapter 2, I briefly introduce our new survey, in which I used the Keck II 10-meter telescope with the DEIMOS multi-object spectrograph to measure radial velocities of over ten thousand individual stars in Andromeda. The data from this survey form the core of this thesis. The kinematical survey was complemented with data from the Hubble Space Telescope/Advanced Camera for Surveys and Wide Field Camera 3 data from the Panchromatic Hubble Andromeda Treasury survey and additional existing ground-based data sets to perform the most detailed kinematical survey to date of an external galaxy disk. The multiple data sets allow us to disentangle subcomponents based simultaneously on their kinematics, surface brightness profiles, and stellar populations. In Chapter 3, I use ages derived from PHAT optical photometry along with the stellar kinematics presented in Chapter 2 to measure the age-velocity dispersion relation in the stellar disk. I find a direct correlation between stellar age and velocity

dispersion between ages of 30 Myr and 6 Gyr, indicating that a continuous heating or cooling mechanism has been occurring in M31 for at least the last six billion years. This measurement provides two new challenges for modelers: to produce a continuous age-dispersion relation (which is already done well by most chemodynamical simulations) and to produce a very high “heating rate” (slope of the dispersion vs. age relation, which is much higher than anything seen in either the Milky Way or in simulations.) I also find something that could not have been discovered simply by exploring the Milky Way solar neighborhood in more detail: that the velocity dispersion of stars in a given age bin varies over the face of the disk because of a spatially varying second kinematical component with a very high, spheroid-like, velocity dispersion. In Chapters 4 and 5, I explore this hot component in more detail. First, in Chapter 4, I statistically isolate the hot component members and characterizing their kinematics. Then, in Chapter 5, I explore the relationship between the hot component and the bulge, disk, and halo structural subcomponents via a detailed structural decomposition using multiple data sets. I find that the hot component is kinematically associated with the halo but that its luminosity function is the same as that of the disk: in other words, these stars may have been born in the disk but been dynamically kicked out as the result of a satellite impact. Lastly, in Chapter 6, I measure one more indicator of heating: the asymmetric drift of the stars relative to the circular velocity of the galaxy as traced by neutral hydrogen gas. I find that the velocity lag and velocity dispersion are correlated but not as tightly as in the Milky Way. While these measurements alone are insufficient to uniquely reconstruct the history of M31 — let alone those of other Milky Way-mass

galaxies — they set the stage for more detailed simulations. Importantly, our results support suggestions from previous studies that M31 has had a more typical merger history (in the Λ CDM context) than the Milky Way, so, in the context of understanding common galaxy evolution mechanisms, it will be useful for simulations to be able to match the observed properties of M31 in addition to those of the Milky Way.

Chapter 2

Spectroscopic Survey of M31's

Stellar Disk

Table 2.1 summarizes our spectroscopic observations of M31's stellar disk between 2007 and 2012. All observations were carried out with the Keck/DEIMOS optical multiobject spectrograph on Mauna Kea; however, because the data were obtained with separate projects in mind, the target selection strategy varied widely from year to year. Through 2010, we blindly targeted stars that were isolated and bright in the CFHT/MegaCam i' filter, and therefore mostly obtained spectra of the dominant bright population in M31's disk, metal-rich red giants. In 2011 and 2012, we used photometry from the PHAT survey to more accurately identify isolated spectroscopic targets and also to prioritize targets from rarer populations. In 2011 we focused on red giant branch stars of a wide range of metallicities, whereas in 2012 we targeted young upper main sequence stars and intermediate-age asymptotic giant branch stars in addition to

red giants. In every observing run, we also observed a few ancillary non-stellar targets such as optical X-ray counterparts and young star clusters from the PHAT survey, but analysis of these targets is not included in this thesis. All of these data, except for those to the south of the galaxy (mask name M32*) and some to the east of the galaxy (mask name SE) now have associated six-filter PHAT photometry. The remainder of the data only have the CFHT/MegaCam i' photometry that was used for their target selection.

Through 2011, we used the 1200 l/mm grating on DEIMOS, which yields wavelength coverage that includes $H\alpha$ through the calcium triplet. In 2012, in an effort to obtain useful spectra of younger, hotter stars, we switched to the lower-resolution 600 l/mm grating, obtaining spectra as blue as 4500Å while still covering the calcium triplet in the near-infrared.

Each of the chapters in this thesis describes work done with a different subset of these data, based primarily on what data were available at the time of writing and on whether PHAT photometry was required for the science. In all cases, only stellar targets with reliable velocity measurements were used. Chapter 3 uses all of the stars with PHAT photometry. Chapter 4 uses all of the stars observed through 2010, including targets without PHAT photometry. Chapter 5 uses stars with PHAT photometry that were observed through 2011. Chapter 6 uses all stars. The Data section of each chapter includes a summary of the target selection, observations, and data reduction specific to the analysis in that chapter.

Table 2.1: Spectroscopic Data

Mask Name	Observation Date (UT)	α [J2000] (h m s)	δ [J2000] ($^{\circ}$ ' ")	P.A. ($^{\circ}$)	t_{exp} (sec)	Seeing (FWHM)	# of Slits	# of Usable Target Velocities	# of Usable Serendip Velocities
M32_1	2007 Nov 14	00 42 38.28	+40 51 34.0	+160.0	2×20	0".5	199	188 (94%)	72
M32_2	2008 Aug 03	00 43 03.82	+40 55 07.7	+70.0	3×20	0".6	189	166 (88%)	27
M32_3	2008 Aug 03	00 43 11.60	+40 52 34.7	-110.0	3×20	0".7	203	132 (65%) ^b	10
M32_4	2008 Aug 04	00 42 13.87	+40 54 44.2	+105.0	3×20	0".6	165	137 (83%)	119
M32_5	2008 Aug 04	00 42 13.88	+40 52 02.6	-75.0	3×20	0".6	177	152 (86%)	72
∞ M32_6	2008 Aug 31	00 41 20.41	+41 51 32.2	0.0	$3 \times 20 + 1 \times 10$	0".7	169	152 (90%)	128
SE7	2008 Sept 01	00 43 38.74	+41 10 17.4	+39.0	2×10	0".8	170	148 (87%)	50
SE8	2008 Sept 30	00 44 00.82	+41 09 27.1	-113.0	3×15	0".5	197	178 (90%)	27
SE9	2008 Oct 01	00 44 49.26	+41 03 27.6	-60.0	$2 \times 12.5 + 1 \times 15$	0".4	204	185 (91%)	11
mctA5	2010 Oct 07	00 44 18.33	+41 39 28.1	+270.0	3×16	0".6	212	197 (93%)	82
mctB4	2010 Oct 08	00 44 29.04	+41 35 10.0	+270.0	3×16	0".7	177	172 (97%)	81
mctC3	2010 Oct 07	00 45 11.89	+41 53 37.7	+90.0	3×16	0".5	198	170 (86%)	69
mctD3	2010 Oct 07	00 45 09.46	+41 49 08.8	+90.0	3×18	0".6	209	185 (84%)	69
mctE3	2010 Oct 08	00 46 53.23	+42 14 59.3	+90.0	3×17	0".7	221	202 (91%)	18

Table 2.1: Spectroscopic Data

Mask Name	Observation Date (UT)	α [J2000] (h m s)	δ [J2000] ($^{\circ}$ ' ")	P.A. ($^{\circ}$)	t_{exp} (sec)	Seeing (FWHM)	# of Slits	# of Usable Target Velocities	# of Usable Serendip Velocities
mct04p	2010 Oct 07	00 44 51.81	+41 25 19.2	-142.3	3×16	0".6	254	223 (88%)	24
mct05p	2010 Oct 08	00 44 19.70	+41 32 53.5	-52.3	3×16	0".5	254	188 (74%)	100
mct06p	2010 Oct 07	00 44 33.71	+41 36 17.6	-52.3	3×16	0".5	251	211 (84%)	43
mct07p	2010 Oct 08	00 44 41.77	+41 40 03.0	-52.3	3×16	0".6	264	210 (80%)	71
mct09p	2010 Oct 08	00 45 39.23	+41 38 39.1	-142.3	3×16	0".6	252	213 (85%)	22
mct10p	2010 Oct 07	00 45 08.24	+41 46 19.2	-52.3	3×18	0".9	255	207 (82%)	34
mct12p	2010 Oct 07	00 45 28.34	+41 53 23.3	-52.3	3×18	0".9	265	212 (80%)	12
mct13p	2010 Oct 08	00 45 42.02	+41 56 42.4	-52.3	3×18	0".7	259	217 (84%)	23
mct15p	2010 Oct 08	00 45 54.36	+41 59 43.1	-52.3	3×17	0".9	261	206 (79%)	10
mct16p	2010 Oct 08	00 46 08.44	+41 02 58.6	-52.3	3×18	0".7	258	221 (86%)	5
mctF	2011 Nov 23	00 44 24.00	+41 36 00.0	-30.0	2900	0".6	246	179 (73%)	47
mctG	2011 Nov 23	00 45 53.03	+41 42 05.1	+25.0	2900	0".5	259	207 (80%)	24
mctJ	2011 Nov 24	00 45 10.80	+41 55 48.0	+35.0	3600	0".6	253	182 (72%)	18
mctK	2011 Nov 24	00 46 46.85	+42 13 35.3	+45.0	3400	0".8	270	208 (78%)	10
mctL	2011 Nov 23	00 46 19.97	+42 14 05.2	-65.0	3680	1".0	257	182 (71%)	7

Table 2.1: Spectroscopic Data

Mask Name	Observation Date (UT)	α [J2000] (h m s)	δ [J2000] ($^{\circ}$ ' ")	P.A. ($^{\circ}$)	t_{exp} (sec)	Seeing (FWHM)	# of Slits	# of Usable Target Velocities	# of Usable Serendip Velocities
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

Table 2.1: Spectroscopic Data

Mask Name	Observation Date (UT)	α [J2000] (h m s)	δ [J2000] ($^{\circ}$ ' ")	P.A. ($^{\circ}$)	t_{exp} (sec)	Seeing (FWHM)	# of Slits	# of Usable Target Velocities	# of Usable Serendip Velocities
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
Totals							10860	8793 (80%)	1727

Chapter 3

Age-dispersion relation

3.1 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 (e.g., Chiba & Beers, 2000) and Andromeda (Collins et al., 2011) galaxies. 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 (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 & Navarro, 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 & Dalcanton, 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 Milky Way itself, we can more easily separate stars by age. Velocity dispersion appears to increase monotonically with age in the solar neighborhood (Nordstrom 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 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 Milky Way 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 multiobject 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 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 (LOS) velocity dispersion 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 § 5.6, 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 §6.2, we present the dataset composed of Keck/DEIMOS radial velocity measurements and HST optical photometry of over 8200 individual stars. In §3.3, 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 §5.5. In §5.6, we discuss the constraints our results place on disk evolution scenarios. Finally, we summarize in §5.7.

3.2 Data acquisition and reduction

Our dataset is a subset of two larger surveys. Figure 3.1 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 multiobject 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 selec-

tion 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 A 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 3.2 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 8500Å is strong in RGB and many AGB stars. Temperature-sensitive TiO bands across the red side of the spectrum, including a strong triplet near 7050Å and another band near 8850Å, 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 α and H β 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 – 9100Å and a spectral resolution of $R = 6000$. In 2012, we used the coarser 600 line/mm grating to gain spectral coverage as blue as 4500Å 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.

3.3 Methods: Velocity dispersion as a function of age and of metallicity

Our goal is to measure the line-of-sight (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.




Figure 3.1: 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.




Figure 3.2: 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 Calcium 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.)

Next we split our spectroscopic sample into those four bins using the observed optical PHAT CMD. Then, we map the line-of-sight 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.

3.3.1 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.

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 MW foreground 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 $F475W - F814W > 2$ — 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 Å, which is an 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 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 3.3 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 Milky Way

in the direction of M31. In each color/magnitude bin (CM), we calculate the number of contaminants $N_{\text{MW,expected,CM}}$ expected in the SPLASH survey:

$$N_{\text{MW,expected,CM}} = N_{\text{SPLASH,CM}} \frac{N_{\text{Trilegal,CM}}}{N_{\text{PHAT,CM}}} \quad (3.1)$$

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

$$N_{\text{contaminants,CM}} = N_{\text{MW,expected,CM}} - N_{\text{removed,CM}} \quad (3.2)$$

Figure 3.4 shows the resulting distribution of $N_{\text{contaminants,CM}}$. The number is only significant ($> 10\%$ of the sample) brighter than $m_{\text{F814W}} = 21$ and between $1 < m_{\text{F475W}} - m_{\text{F814W}} < 2$. 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.

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




Figure 3.3: 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 3.4: Optical Hess diagram showing the number of foreground Milky Way stars expected in our spectroscopic sample after cutting out stars with strong surface gravity-sensitive NaI 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 described and defined in § 3.3.1 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.

($R \sim 20-30$ 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 $1 M_{\odot} \text{ yr}^{-1}$. 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 7 times as many stars at $[M/H] = -0.5$ as at $[M/H] = -1$, the simulation is allowed to spend 7 times as long producing stars with $[M/H] = -0.5$ as with $[M/H] = -1$. We use Girardi et al. (2010) isochrones, a Kroupa 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 a 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 constant-SFR CMD and age bin outlines are shown in the left hand panel of Figure 3.5; in the rest of this section, we explain the choice of bins.

All stars bluewards 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 (BSG) stars, which have similar ages.

The RGB region includes stars redwards 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 magnitude 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 towards 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-hand panel of Figure 3.5. 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 hand 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 age-metallicity relation, 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 towards younger ages since the rate of stars moving off the main sequence onto the red giant branch 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 towards 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

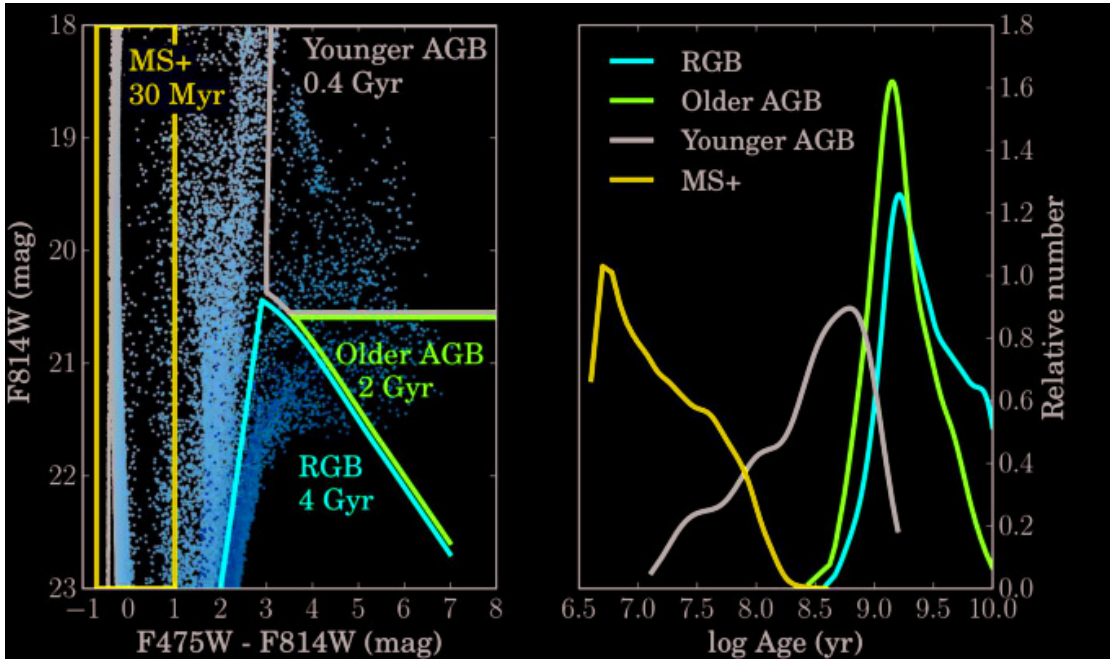


Figure 3.5: *Left*: Simulated CMD assuming a constant star formation rate. The blue, black, violet, and red outlines define the “MS,” “younger AGB,” “older AGB,” and “RGB” age bins, respectively. Stars are color coded by $\log(\text{age})$, with brighter (yellow) colors corresponding to older ages. Numbers 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.

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.

3.3.2 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 hand panel of Figure 6.1 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 hand panel of Figure 6.1 shows a CMD of the spectroscopic sample only, divided into age bins as in Figure 3.5.

While our simulated CMD does include reddening from the MW foreground, 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 3.5.

We estimate the average reddening vector in our sample using the M31 dust map presented in Dalcanton et al. (2014). 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

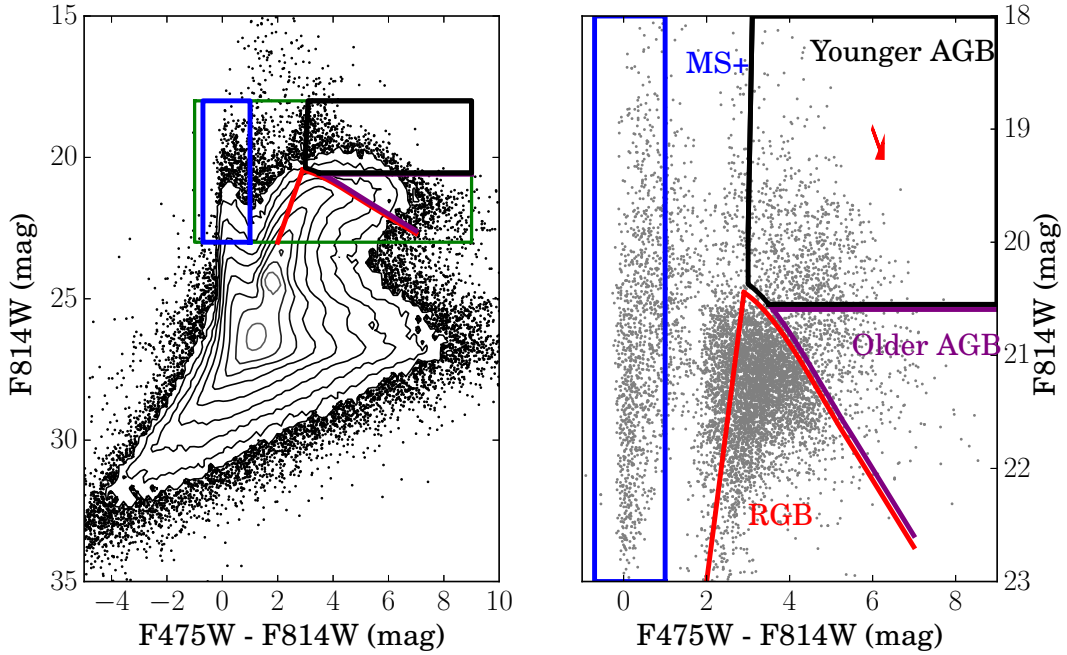


Figure 3.6: *Left:* Optical PHAT CMD from Brick 12, a representative brick near the middle of the PHAT survey area. The green box marks the portion of the CMD sampled by the spectroscopic survey. The age bin region outlines are the same as in Figure 3.5. *Right:* Optical PHAT CMD of spectroscopic targets 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 hand panel and in Figure 3.5. 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. (2014), 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).

targets, assuming half the targets lie behind the dust layer, is shown in Figure 6.1.

3.3.3 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 2-D sky map (Figure 3.7) and as a 1-D dispersion distribution (Figure 3.9). 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 $\sim 15 - 20$.

To create 2D dispersion maps, displayed in Figure 3.7 and Figure 3.8, 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 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 3.9, 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 3.7 and the 1D dispersion distributions in Figure 3.9 that the typical dispersion increases with average age – that is, older populations are dynamically hotter than younger populations. In addition, the 2D maps appear to be patchy, but we will show in § 3.4.3 that much of the small-scale spatial variation in dispersion is an effect of finite sampling and is not physical.

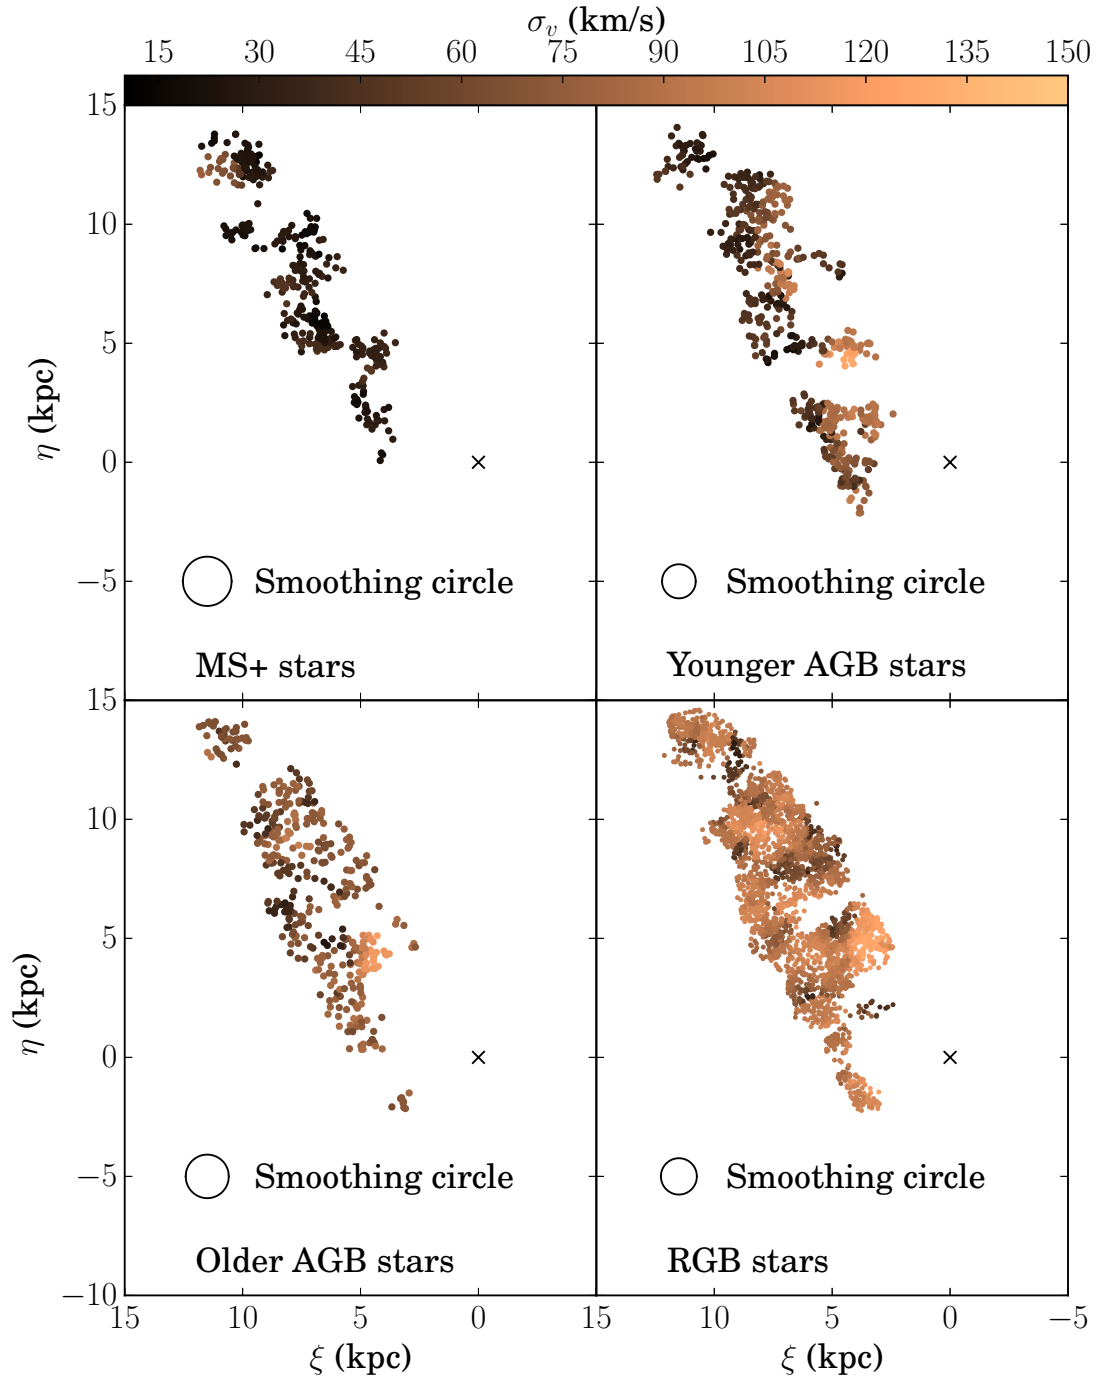


Figure 3.7: Smoothed, local, line-of-sight velocity dispersion of stars of different evolutionary stages: young main-sequence (*Upper Left*); younger AGB (*Upper Right*); older AGB (*Lower Left*); and RGB (*Lower Right*). The circles show the sizes of the smoothing circles in which the dispersion is 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.




Figure 3.8: RGB dispersion map from Figure 3.7 next to a Herschel image of M31 for reference. A high-dispersion region which 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) is marked with a white line.




Figure 3.9: Local line-of-sight 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 3.8 *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.

3.3.4 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 towards 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 an 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 §3.3.3. The maps are displayed in Figure 3.10 and the 1D distributions in Figure 3.11. 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.

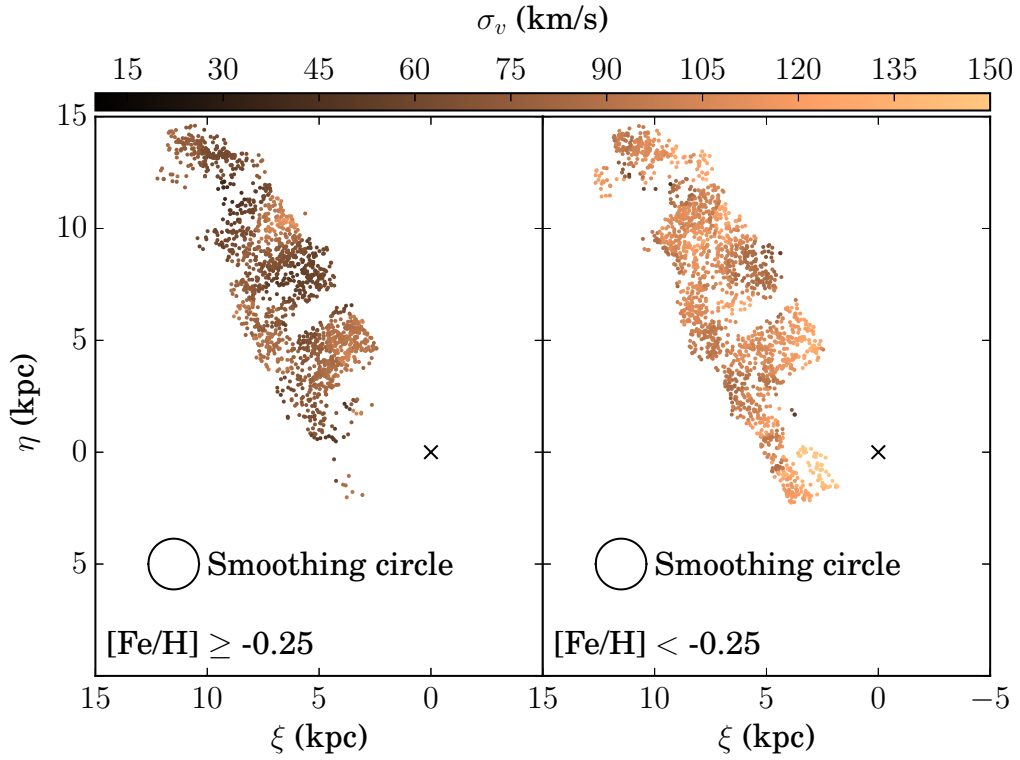


Figure 3.10: Smoothed velocity dispersions of RGB stars in high- (*Left*) and lower-metallicity (*Right*) portions of the RGB sample. Metallicities are measured using 10 Gyr old Bressan et al. (2012) isochrones. On average, the metal-rich population has a lower velocity dispersion. Some of the hot patches in the high-metallicity bin may be due to contamination from reddened low-metallicity stars.




Figure 3.11: Same as Figure 3.9, for the RGB stars in the two metallicity bins from Figure 3.10. 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.

3.4 Results

3.4.1 Age-dispersion relation

We can now measure, for the first time, the age-velocity dispersion relation in an external galaxy. Figure 3.12 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 Milky Way from Nordstrom et al. (2004). We show both the σ_R and σ_ϕ components of the Milky Way’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 3 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 Nordstrom 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 $\sigma_R = 60 \text{ km s}^{-1}$ (Büdenbender et al., 2014).

This age-dispersion correlation is robust against contamination of our age bins. Two bins — the RGB and older AGB — are 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 towards 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 towards 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 3.9.

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 § 5.6.




Figure 3.12: Comparison of dispersion versus 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 3.9. 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 (Nordstrom et al., 2004). Both the slope and normalization of the dispersion versus age relation are much higher in M31 than in the MW, regardless of the assumed SFR.

3.4.2 Age-metallicity relation among RGBs

Figures 3.10 and 3.11 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.

3.4.3 Structure in dispersion map

The spatial dispersion maps in Figure 3.7 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 3.8. 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 3.9 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 §5.6, 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 “kine-

matically 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 3-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 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 3.13, is kinematically clumpy. In Figure 3.14, we quantify the clumpiness and compare the simulation’s 1D dispersion distribution with M31’s RGB dispersion distribution. The two distributions have similar full widths at half maximum. 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 3.14 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, in prep.) and velocity dispersion are uncorrelated. In §5.6 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.

3.5 Discussion

3.5.1 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 line-of-sight velocity dispersion — varies substantially between age bins and also across the disk of the galaxy, and that much of the variation




Figure 3.13: Velocity dispersion map for a toy model of a disk with a Gaussian velocity distribution with $v_\phi = 220 \text{ km s}^{-1}$ and $\sigma_R = \sigma_\phi = 90 \text{ km s}^{-1}$.

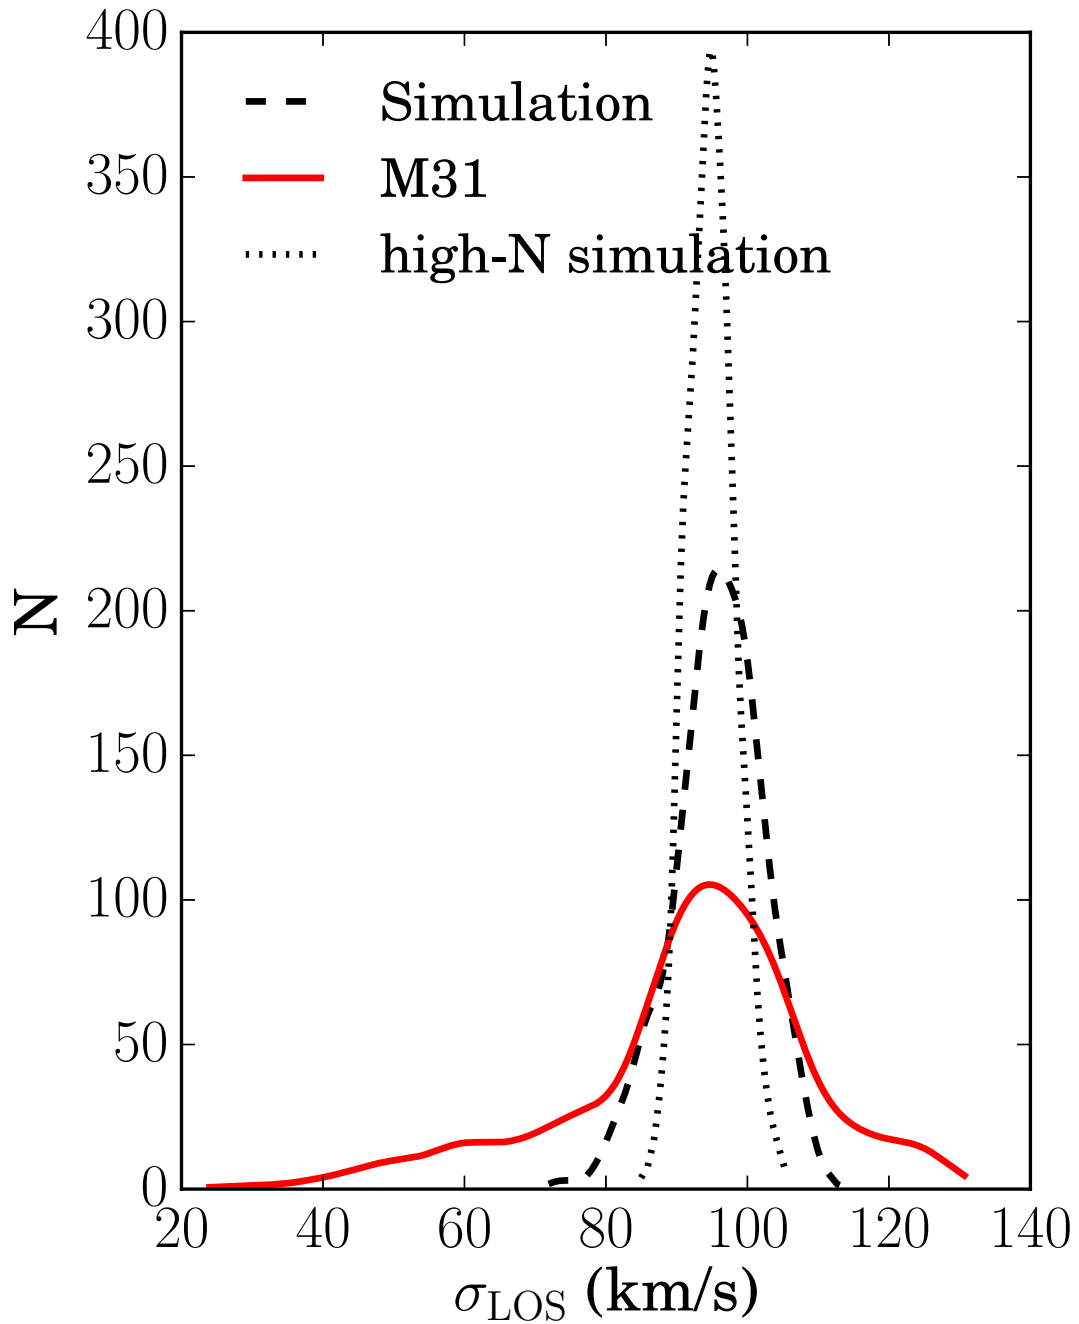


Figure 3.14: Velocity dispersion distribution for RGB stars from the toy model of a disk with target density comparable to our spectroscopic survey (black dashed line), from a toy model with target density 5 times higher (black dotted line), and from the M31 data (red line). Finite sample density can account for most of the width of the observed dispersion distribution.

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 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, 0.5 kpc on a side (1.5 kpc on a side for the less-dense AGB and MS+ populations). 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 3.15 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 skewed to negative velocities, which is towards 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 construction, its mode is at a positive velocity and it has a wide tail towards 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 §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 are 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.

3.5.2 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

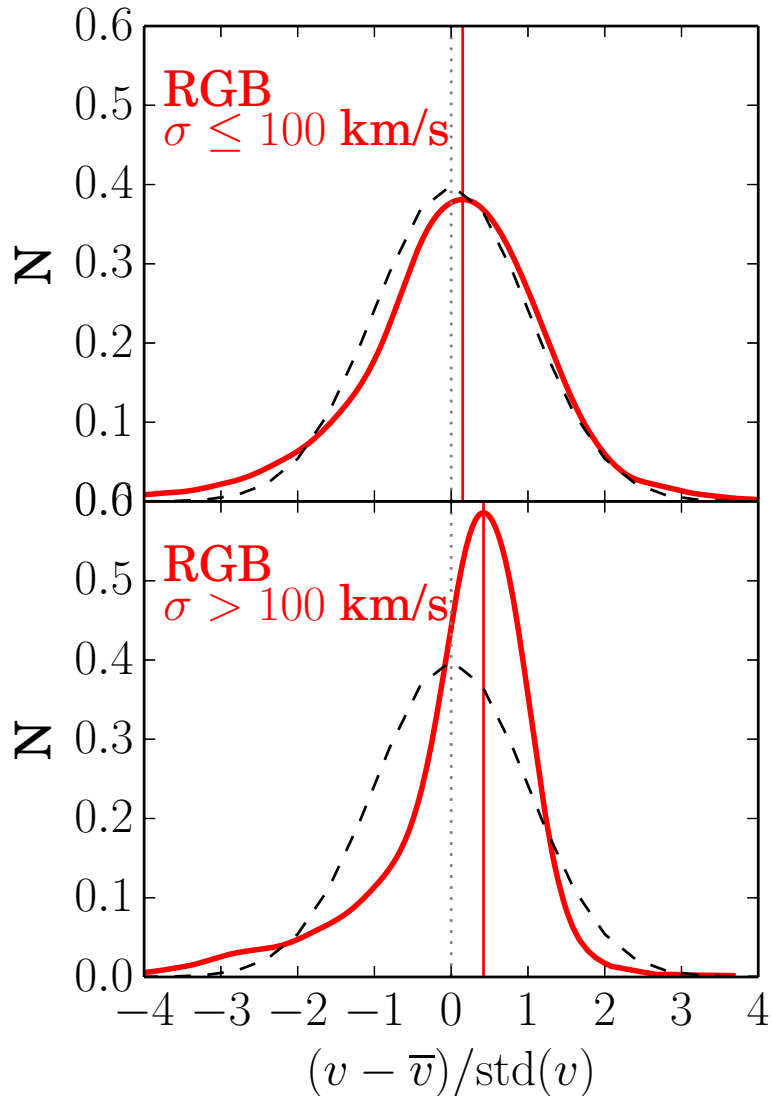


Figure 3.15: Comparison of composite velocity distribution shapes for RGB stars (solid red curves) to a normal distribution (dashed black curves). The top panel includes pixels with dispersion less than 100 km s^{-1} , and the bottom panel includes the rest of the pixels. By construction, all distributions 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.

collapse) plus some mechanism to heat the disk nonuniformly (such as accretion of satellites or kicking up of existing disk stars via satellite impacts).

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 giant molecular clouds (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 3.7, 3.9, and 3.12. 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 Milky Way (Nordstrom 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.

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 Milky Way (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 3.12, 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 a 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 needed 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.

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 Milky Way (e.g., Mathewson et al., 1974; Yanny et al., 2003; Rocha-Pinto et al., 2003; Ibata et al., 1994; Majewski et al., 2003; Newberg et al., 2003) and M31 (Ibata et al., 2004, 2007; Ferguson et al., 2002; Kalirai et al., 2006a; Gilbert et al., 2009; McConnachie et al., 2009; Tanaka et al., 2010; Gilbert et al., 2012) lend 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 § 3.5.2, 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., 2006a; 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 line-of-sight

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.

3.5.3 Brick 9 Feature: Associated with the Bar?

Many of the high-dispersion measurements lie in the “Brick 9 Region” centered at $(\xi, \eta) = (4.5, 4.5)$ kpc, marked by the ends of the teal line in Figure 3.8. 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 a Gyr ago. It is located on the major axis of the galaxy, so 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 3.7) 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 line-of-sight component of stellar velocities.

3.5.4 Comparison to Milky Way

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 Nordstrom et al. (2004), and the average dispersion of our youngest age bin is more than 50% higher than that in the Milky Way. In addition, the slope of the dispersion versus age relation is more than 3 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 last 10 Gyr (Stewart et al., 2008). The Milky Way'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. 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 powerlaw 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).

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 (Dorman et al. 2014b, in preparation) because it requires the choice of a particular disk rotation curve model.

3.5.5 Radial trends

In Figure 3.16, 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 shows that 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. The dispersions of all age groups decline steeply out to 10 kpc and then decline less steeply. Outside the B9 region, the velocity dispersions of the different age groups separate from one another, revealing the clean age vs. dispersion relationship.

The 10 kpc break corresponds to a bit less than two radial scale lengths (Worthey et al., 2005; Seigar et al., 2008; Ibata et al., 2005; 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 (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 ($< 30^\circ$ vs. 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 .

3.6 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 line-of-sight velocity dispersion of each age group. We have found:

- 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 5 times higher than that of the Milky Way.




Figure 3.16: 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.

- Among RGB stars in our sample, the metal-poor half is 50% kinematically hotter than the metal-rich half.
- 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 line of sight 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.
- 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.

Chapter 4

M31's Inner Spheroid

4.1 INTRODUCTION

The elegant progression of the Hubble sequence from ellipticals to spirals demonstrates that galaxy morphology can be described in large part based on the relative importance of spheroid and disk subcomponents. While it is now clear that the simple evolutionary path from elliptical “early-type” to disk-dominated “late-type” galaxies that Hubble originally proposed is incorrect, the physical origin of the Hubble sequence and the formation of and relationship between the different structural subcomponents remain subjects of vigorous research.

The central spheroids of spiral galaxies fall into two categories, which can be explained by distinct formation mechanisms as reviewed by Kormendy & Kennicutt (2004). Classical bulges, which are typically described as elliptical galaxy analogs with random stellar velocity distributions, large velocity dispersions, and $r^{1/4}$ de Vaucouleurs surface brightness profiles, are likely formed through violent merger/accretion

events. Pseudobulges, which are more flattened, have more ordered kinematics, and have roughly exponential brightness cutoffs (or, more generally, Sérsic profiles with low $n_{\text{Sérsic}}$ values) are likely formed through secular heating of the disk. More detailed observations, yielding constraints on the structure and dynamics of bulges, will lead to a clearer understanding of possible formation scenarios.

Any study of the inner regions of a galaxy is complicated by the presence of several spatially overlapping structural subcomponents, such as the disk, spheroid, and halo. Deconvolving these subcomponents to reveal the behavior of a single one is difficult. Traditionally, codes such as GALFIT (Peng et al., 2002) or GIM2D (Simard et al., 2002) are employed to fit galactic integrated light profiles with the sum of a Sérsic bulge and exponential disk (e.g. Courteau, 1996; Courteau et al., 2011). This technique is the only possible method for characterizing the structure of distant galaxies, but it suffers from strong assumptions about the characteristic light profiles of bulges and disks. In addition, degeneracy in the best-fit derived parameters can cloud interpretation of the results.

Resolved stellar kinematics offer a complementary approach to structural deconvolution of the nearest galaxies. Instead of assuming specific surface brightness profiles of disks and spheroids, one must only make the geometrical argument that a stable disk – a thin, flat structure – is kinematically colder (has a higher v_{rot}/σ_v) than a stable spheroid. Separate components can then be identified and characterized by their distinct stellar velocity distributions. The proximity of Andromeda (M31) at about 785 kpc (e.g., McConnachie et al., 2005) renders it the only large spiral galaxy other

than the Milky Way (MW) where detailed photometric *and* kinematical observations are possible with current observing facilities.

We use resolved stellar kinematics to study M31’s kinematically hot “inner spheroid” at projected radii of 2–20 kpc. Any description of this region – or that in the intermediate region of any large galaxy – is necessarily complex; the literature is full of vocabulary such as “bulge,” “spheroid,” “inner spheroid,” “outer spheroid,” “disk,” “thin disk,” “thick disk,” “extended disk,” and so on. There is not yet a consensus on the best combination of these nouns to represent M31. For the purposes of this paper, we use the word “spheroid” to describe a kinematically hot component: some combination of bulge, halo, and/or any other spheroidal component. Likewise, we refer to the kinematically colder population as the “disk,” where this term includes any distinct disk components that may be present, such as the thin, thick or extended disks.

Despite the possibility of multiple components, the disk is likely to be locally kinematically cold. Collins et al. (2011) claim that M31’s stellar disk at $r_{\text{proj}} \sim 10\text{--}40$ kpc consists of a cold thin disk and a warm thick disk, as is the case for the MW. Given that most of our fields are closer to M31’s center than the innermost field of Collins et al. (2011), and given their finding that the thin disk has twice the density of the thick disk and a shorter radial scale length, we expect the cold thin disk to dominate the stellar disk population in our fields. Similarly, Ibata et al. (2005) suggest that no more than about 10% of the total disk luminosity may be due to an extended disk component which also lags the cold disk. Though we do not know a priori the relative contributions of the thin, thick and extended disks, in §4.4.2 we show that our assumption of a dominant

thin component is justified for the purposes of measuring the kinematical parameters of the inner spheroid.

Unlike the case of its stellar disk, M31's inner spheroid has no analog in the MW and is therefore of great interest. The spheroidal system at these radii in the MW is relatively metal-poor ($\langle Fe/H \rangle \sim -1.6$; Carollo et al., 2007), is composed entirely of old stars, and has a power-law spatial density profile of the form r_{deproj}^{-3} , which corresponds to an r^{-2} power-law surface brightness profile. Models such as those proposed by Bullock & Johnston (2005) and Zolotov et al. (2010) suggest that the MW halo represents a population of accreted dwarf satellite galaxies. In contrast, the inner spheroid in M31 more closely resembles a bulge than a halo. It is more metal-enhanced than the MW halo, with $[Fe/H] \sim -0.7$ (Kalirai et al., 2006b), and has a Sérsic surface brightness profile with $n_{\text{Sérsic}} \sim 2-4$ (Pritchet, C.J., van den Bergh, 1994; Guhathakurta et al., 1911; Courteau et al., 2011). In addition, the stellar population of M31's spheroid is younger than that of the MW inner halo on average, with 40% of the stars younger than 10 Gyr (Brown et al., 2006), and the stellar density is also significantly higher than that at an equivalent location in the MW (Reitzel et al., 1998).

The inner spheroid straddles territory between two well-studied components of the spheroid: the classical and boxy bulges interior to ~ 1 kpc (Athanassoula & Beaton, 2006; Beaton et al., 2007; Courteau et al., 2011), and the outer halo which dominates past $R_{\text{proj}} \sim 30$ kpc (e.g. Guhathakurta et al., 1911; Irwin et al., 2005; Ibata et al., 2007). In the central kpc, where the density is too high for resolved stellar population spectroscopy, Saglia et al. (2010) analyzed integrated-light kinematics to reveal a bulge

rotation speed of 70 km s^{-1} and a velocity dispersion of 140 km s^{-1} at $R_{\text{proj}} = 1.1 \text{ kpc}$ on the major axis. However, they cautioned that this measurement is contaminated by the kinematically cold disk which may contribute nearly a third of the light at this radius.

Farther out in M31's halo, kinematical surveys of the resolved stellar population using the Keck/DEIMOS multiobject spectrograph have mapped the cold substructure, as well as the underlying smooth virialized population, out to $R_{\text{proj}} > 150 \text{ kpc}$ (Guhathakurta et al., 1911; Chapman et al., 2006; Kalirai et al., 2006b; Gilbert et al., 2007; Gilbert et al., 2009). Chapman et al. (2006) compiled kinematics of ~ 1200 red giant branch (RGB) halo stars in scattered fields between $R_{\text{proj}} = 8$ and 70 kpc . Using a windowing technique to eliminate stars whose velocities were consistent with that of the disk, they found that the velocity dispersion of the remaining population decreased radially outwards: $\sigma_v(R_{\text{proj}}) = (152 - 0.9 R_{\text{proj}} \text{ kpc}^{-1}) \text{ km s}^{-1}$. Subsequently, as part of the Spectroscopic and Panchromatic Landscape of Andromeda's Stellar Halo (SPLASH) survey, Gilbert et al. (2007) fit a double Gaussian profile to the velocity distribution of RGB stars in a large contiguous region along the southeastern minor axis of the galaxy and measured a constant velocity dispersion of 128.9 km s^{-1} between $R_{\text{proj}} = 10$ and 30 kpc .

In recent years, the focus of SPLASH has migrated inwards, first to target the dwarf galaxies Andromeda I and Andromeda X (Tollerud et al., 2012), NGC 205 (Geha et al., 2006; Howley et al., 2008) and M32 (Howley et al., 2013), and now towards the disk- and bulge-dominated inner regions of M31. The majority of the data analyzed in

the present paper come from the most crowded area targeted to date: a large contiguous disk-dominated area on the NE major axis with $R_{\text{proj}} = 2\text{--}19$ kpc. This area was selected to overlap the coverage of the Panchromatic Hubble Andromeda Treasury (PHAT) survey, a five-year Hubble Space Telescope (HST) MultiCycle Treasury (MCT) program that began in 2010 (Dalcanton et al., 2012).

The disk and spheroid share the inner regions of the galaxy with remnants of tidally disrupted galaxies. The dominant features in star-count maps of the 2–20 kpc region are the Giant Southern Stream (GSS; the remnant of a tidally stripped satellite galaxy) and the shelves (sharp edges in stellar density) created by it (Ibata et al., 2001; Fardal et al., 2007). Both the GSS and a “secondary stream,” which is cospatial with the GSS but separated by 100 km s^{-1} in velocity, have been kinematically detected in multiple fields south of M31 (Kalirai et al., 2006a; Gilbert et al., 2009). Fardal et al. (2007) identified the northern extension of the stream in the Chapman et al. (2006) and Ibata et al. (2005) sample of RGBs and in the planetary nebulae of Merrett et al. (2006).

There are two principal challenges to a resolved stellar population kinematical study of the crowded inner spheroid of M31. First, we must select intended stellar targets whose spectra are least likely to be contaminated by close (in projection) stellar neighbors. Second, we must disentangle the stellar disk from the spheroid population we wish to characterize. This is especially important in disk-dominated fields, where fewer than 20% of the stars may belong to the spheroid. Note that though a large part of this paper will focus on accounting for the disk contribution, our analysis method is

designed to elucidate the nature of the inner spheroid rather than the disk. We do not attempt to make a statement here about the rotation curve of the stellar disk or the presence of a thick disk. We plan to analyze these components in a future paper.

This paper is organized as follows. In §4.2 we explain our target selection techniques, spectroscopic observations and radial velocity extraction. In §5.4 we describe our method for isolating and characterizing the spheroid velocity distribution in each of five spatial bins. In §5.5 we discuss the implications of our results; finally, we summarize our findings in §5.7.

4.2 OBSERVATIONS

Our data set for this project is a compilation of three sets of RGB spectra, two of which are presented here for the first time. A detailed technical description of the spectroscopic slitmask design and data reduction is given in Howley et al. (2013). In this section, we describe the target selection criteria for the different data sets and give an overview of the data acquisition and reduction methods common to all the observations.

In §4.2.1, we outline the three data sets used in this paper. In §4.2.2, we describe the source catalogs from which we select our spectroscopic targets. In §4.2.3, we explain our target selection criteria. In §4.2.4, we provide the observing details. In §4.2.5, we give a rundown of the data reduction process. In §4.2.6 and §4.2.7, we measure velocities of individual stars and determine the quality of those measurements, respectively. Finally, in §4.2.8, we discuss the detection and velocity measurement of serendipitously detected stars.

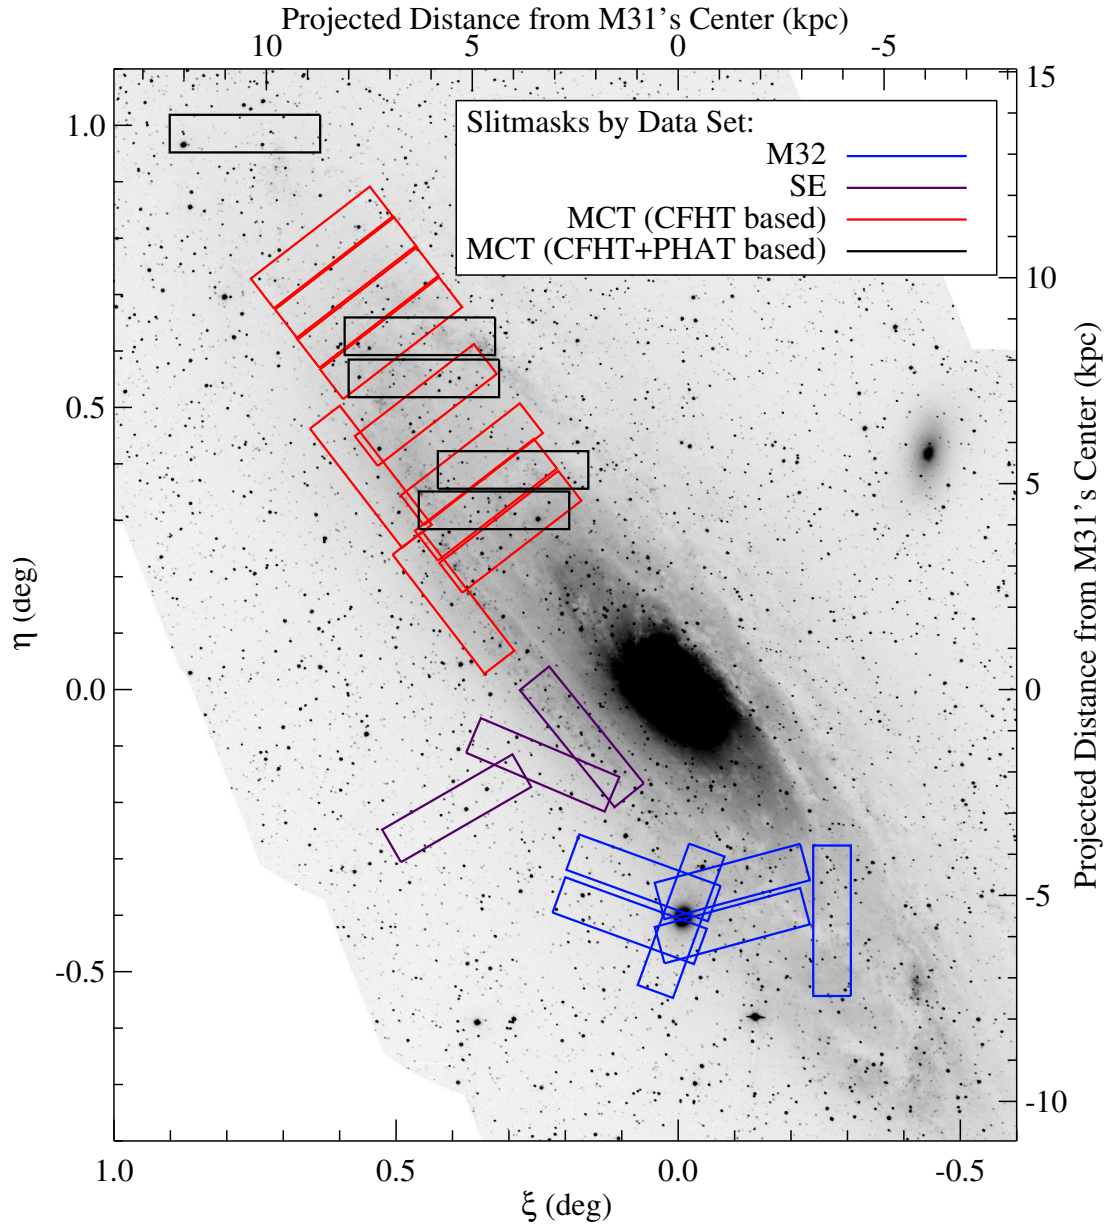


Figure 4.1: Twenty-four Keck/DEIMOS multiobject slitmasks overlaid on the Choi et al. (2002) KPNO Burrell Schmidt B -band mosaic image of M31. Colored rectangles outline the slitmasks from the different data sets: SE (violet), M32 (blue), and MCT (red and black, respectively, for slitmasks for which target selection was based on only the ground-based CFHT/MegaCam photometry/astrometry catalog versus slitmasks for which target selection was based on a combination of the PHAT and CFHT/MegaCam photometry/astrometry catalogs).

4.2.1 Data Sets

The spatial coverage of our three data sets is shown in Figure 4.1. Each rectangular outline represents a Keck/DEIMOS slitmask, which covers approximately $16' \times 4'$ and yields 200–270 useful spectra. (The actual footprint of a DEIMOS slitmask is not perfectly rectangular.) The SE data sets consists of three slitmasks oriented along the eastern minor axis of M31 (violet in Figure 4.1). The M32 data set includes the five slitmasks covering the compact elliptical galaxy M32 directly south of M31 and one slitmask on the SW major axis (blue in Figure 4.1). The SE and M32 data sets were observed during the 2007 and 2008 fall seasons.

Our newest data set, from October 2010, covers a portion of M31’s northeastern major axis spanning the projected radial range 0.33° – 1.38° or 2–19 kpc from the center of M31 in the plane of the disk (red and black in Figure 4.1). Five of these 15 slitmasks (black) were chosen to maximally overlap the regions of existing photometry from the first year of the PHAT program.

The pre-imaging and reduction processes are identical for all three data sets. The primary difference in data acquisition is in the target selection: isolated sources are hand-selected for the SE and M32 slitmasks from a single monochromatic ground-based source catalog (Howley et al., 2013), while we use an automated series of statistical techniques as well as limited color information from the PHAT survey to choose targets for the MCT slitmasks.

4.2.2 Source Catalogs

All targets are chosen from an i' -band $2^\circ \times 2^\circ$ CFHT/MegaCam mosaic centered on M31, obtained in November 2004. We run the software package DAOPHOT (Stetson, 1994) on the image to identify sources, fit PSFs, and produce a PSF-subtracted residual image. The final catalog consists of nearly 2 million unique sources.

For our 2010 Keck/DEIMOS observing run, we also had access to data from the first round of observations of the PHAT program. The data are organized into $12'0 \times 6'5$ bricks; three half-bricks were available at the time of our slitmask design. From these data, we created lists of different stellar populations: metal-poor ($[\text{Fe}/\text{H}] \lesssim -1.3$), metal-intermediate ($-1.3 \lesssim [\text{Fe}/\text{H}] \lesssim -0.7$), and metal-rich ($[\text{Fe}/\text{H}] \gtrsim -0.7$) RGB stars, and hot, massive main sequence stars selected on the basis of SED fitting to six-filter HST photometry. Though we do not treat stars differently based on subcategory membership in the present paper, the kinematics of these stellar populations and their relationship to the kinematics of the RGB populations will be presented in a future work.

4.2.3 Isolated Target Selection

Not all the sources in our catalogs are equally good candidates for multi-object spectroscopy; we prefer isolated targets, those whose spectra are least likely to be contaminated by light from neighboring objects. We classify this contamination as either *crowding* or *blending*. We define a crowded source as one that has at least one neighbor detected by DAOPHOT that is bright and close enough to potentially interfere

with the spectrum of the source. In contrast, we define a blended source as one that is identified by DAOPHOT as a single source but for which visual inspection of the PSF-subtracted image indicates that more than one object may be present.

Furthermore, we only target stars in the apparent magnitude range $20 < i' < 22$ for M32 and SE slitmasks (modified to $20 < i' < 21.5$ for MCT slitmasks). We select a bright limit of $i' = 20$ because the tip of the red giant branch is at $i' = 20.5$ in M31 and so the MW contamination fraction increases significantly in brighter stars. The faint-end limit is chosen because in very crowded areas, it is difficult to recover high-quality spectra of stars fainter than about $i' = 22$. The surface density of sources in the area covered by the MCT slitmasks is so high that we can efficiently pack targets on our slitmasks even with a conservative faint-end limit of $i' = 21.5$.

To choose targets for all slitmasks, we first sort possible targets into three lists (1, 2, 3) in decreasing order of isolation, with list 3 reserved for rejects (§ 4.2.3 – § 4.2.3). Within each list, we prioritize possible targets by magnitude, giving highest priority to intermediate-brightness stars with $20.5 < i' < 21.0$. In the final target selection process, we exhaust each list before moving on to the next. § 4.2.3 describes additional selection criteria specific to the five PHAT-based slitmasks.

Crowding in the SE & M32 Data Sets

In the M32 and SE data sets, we use a neighbor-rejection test to eliminate crowded sources. We reject any star with at least one neighbor with a sufficient combination of proximity and relative brightness, i.e., satisfies the following empirical criterion determined from visual inspection of the CFHT/MegaCam image:

$$I_{\text{nbr}} < I_{\text{tgt}} - \left(\frac{d}{0''.8} \right)^2 + 3.0 \quad (4.1)$$

Here, I_{tgt} and I_{nbr} are the i' -band magnitude of the target source and the neighbor, respectively, and d is the distance in arcseconds between the target and the neighbor. This cut eliminates about 90% of the stars in the M32 and SE data sets (Howley et al., 2013).

Blending in the SE & M32 Data Sets

We identify likely blends in the SE and M32 data sets by visually inspecting the high-pass filtered and PSF-subtracted versions of the i' -band CFHT/MegaCam image at the locations of the stars that survive the crowding test. Each target is flagged as unblended, marginally blended, or badly blended depending on the degree to which its image resembles the PSF on the high-pass filtered image and the strength of systematic residuals at its location on the residual image.

Crowding in the MCT Data Set

Based on our experience with the M32 and SE spectroscopic data sets, we decide to use a slightly relaxed crowding criterion for the MCT data set, rejecting (assigning to list 3) any catalog entry with at least one neighbor which satisfies the following:

$$I_{\text{nbr}} < I_{\text{tgt}} - \left(\frac{d}{0''.8} \right)^{3/2} + 3.0 \quad (4.2)$$

This change accounts for the fact that the seeing during our spectroscopic observations tends to be slightly better than the $0''.8$ CFHT seeing upon which the original criterion was based. This cut eliminates 80%–90% of the stars in the inner slitmasks, 50%–70% at intermediate radii, and only 40% in the furthest slitmask, mctE3.

Blending in the MCT Data Set

In order to avoid a tedious visual inspection of the large area covered by the MCT data set, we design two empirically-based statistical tests to detect possible blends. We visually inspect and flag as “blended” or “non-blended” $\lesssim 100$ objects in each of three small representative image sections at different distances from the center of M31. We then design quantitative tests that approximately reproduce our visual classifications.

The first test is based on the DAOPHOT-generated goodness-of-fit parameter `chi` and shape parameter `sharp`. Objects that appear isolated based on visual inspection fall into a well-defined locus in `chi/sharp` space, as shown in Figure 4.2. Based on this relationship, we retain only objects with `sharp` < 0.2 . We assign a radially dependent linear cut in `chi`, accepting all stars with `chi` < 0.3 in the crowded areas and progressing to the more stringent criterion `chi` < 0.2 in the least crowded outermost slitmask. This cut eliminates 20% of the remaining candidate targets. Because the dividing line

between isolated sources and likely blends is less well defined in `chi` than `sharp`, we allow a buffer zone 0.5 units wide in `chi`. The 1% of stars in the buffer zone are relegated to list 2, but not rejected outright.

In the second test for possible blends, we compare the apparent quality of subtraction to the normalized RMS flux of a 5×5 pixel square of the PSF-subtracted image centered on the source. This value tends to increase with apparent degree of blending. We determine that the best cut is a linear function of magnitude, where blended sources have $\frac{\text{RMS}}{\langle \text{flux} \rangle} > 0.3$ at $i' = 20$ and $\frac{\text{RMS}}{\langle \text{flux} \rangle} > 1.2$ at $i' = 21.5$ (Figure 4.3). The 36% of stars that fail this test are flagged as “possibly isolated” and pushed to list 2. Because the correlation between apparent PSF subtraction quality and RMS is not as tight as those in the `chi/sharp` test, we do not use the RMS cut to reject (assign to list 3) targets that are “isolated” according to both the neighbor-rejection and `chi/sharp` tests.

Target Selection for PHAT-based Masks

We also design five slitmasks based jointly on CFHT data and the PHAT survey-based stellar population lists described in §4.2.2. We require that a star from the PHAT catalog be in the CFHT catalog and have passed through the filters described above to be considered for selection. To ensure final selection of the most isolated (list 1) PHAT objects, we push all non-PHAT objects down one list. For these five slitmasks, then, the priority scheme is as follows: list 1 consists of only isolated, PHAT-selected sources; list 2 includes the isolated CFHT-only sources; and list 3 includes all of the possibly-isolated sources. Because the shape of the PHAT survey bricks is different

from the shape of DEIMOS slitmasks, large areas of the DEIMOS slitmasks that overlap PHAT survey bricks do not target PHAT-based sources. In these areas, the mask design software automatically proceeds to list 2 to select CFHT-based objects.

Summary of Target Selection

To summarize, each source passes through three tests for isolation: neighbor-rejection, `chi/sharp`, and RMS. Each source receives a score for each test: 0 if isolated, 0.6 if marginally isolated, and 2.0 if not isolated. The sum of the three scores determines which list the source belongs in. A total score of 0 maps to list 1; a total of less than 2.0 maps to list 2; and a star with a total score of 2.0 or greater is assigned to list 3. Hence, any star that fails either the neighbor-rejection or `chi/sharp` test cannot be selected, and any star that passes all three filters is given highest priority.

None of the methods described here can fully eliminate the possibility of placing one slit over several objects. Especially in the crowded inner areas, it is very common to obtain multiple spectra in one slit. We discuss our handling of these serendipitous detections in §4.2.8. In addition, a small percentage of the target spectra may still be contaminated with light from nearby objects; objects with unusable spectra are identified by eye and removed from the sample at the end of the data reduction process as described in §4.2.7. The target selection process outlined above serves simply to make educated guesses about the objects best suited for spectroscopy.

In §4.4.6 we show that the spectroscopic target selection criteria (PHAT CMD vs. magnitude only) and actual degree of crowding (isolated vs. sharing a slit with another bright object) have minimal, if any, effect on the measured velocity distribution.

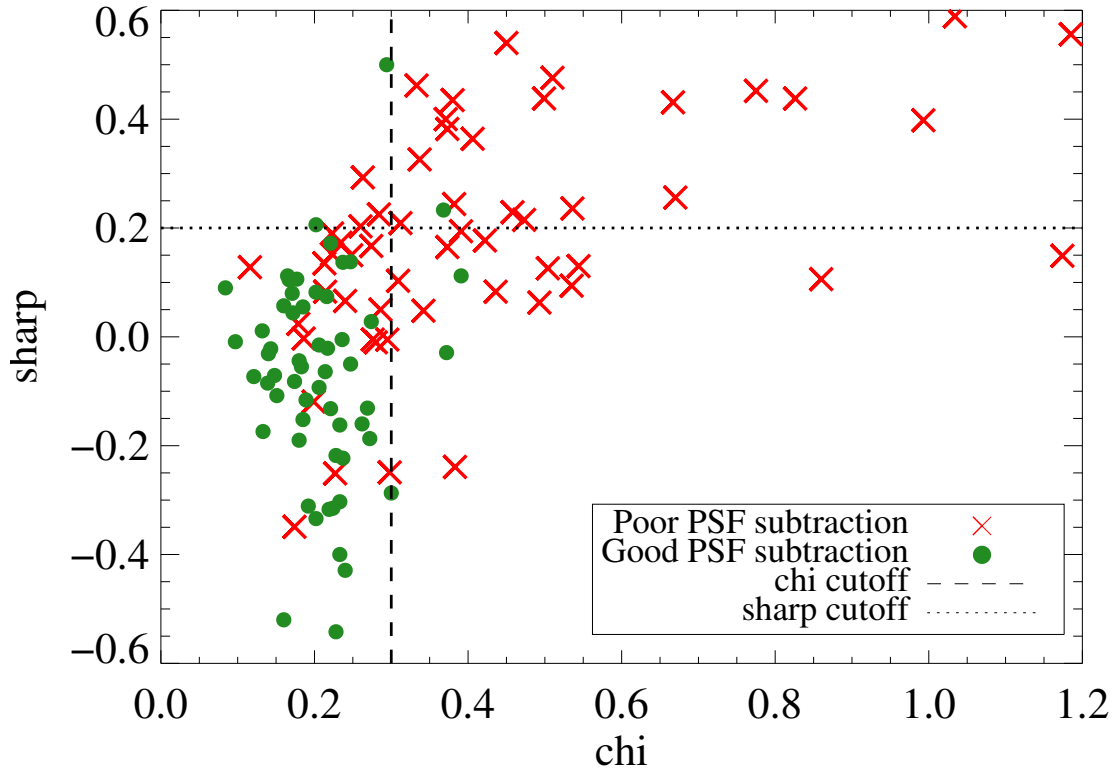


Figure 4.2: χ / sharp isolation criteria. Visually identified blends (red crosses) and non-blends (green circles) are shown from a small representative area close to the center of M31. We reject all sources with DAOPHOT goodness-of-fit parameters $\chi > 0.3$ (to the right of the dashed line) or shape parameter $\text{sharp} > 0.2$ (above the dotted line). The χ cutoff value is lowered in less dense fields farther from the galactic center.

4.2.4 Observations

All slitmasks are observed using Keck/DEIMOS with the 1200 line mm^{-1} grating. This configuration yields a spatial scale of $0''.12 \text{ pixel}^{-1}$ and a spectral dispersion of $0.33 \text{ \AA pixel}^{-1}$. We set the central wavelength to 7800 \AA , corresponding to a wavelength range of $\sim 6450 - 9150 \text{ \AA}$. The exact wavelength range for each slit varies as a result of location on the slitmask and/or truncation due to vignetting. The wavelength region is chosen to target the Ca II triplet absorption feature present in RGB stars. The anamorphic distortion factor for this grating and central wavelength is 0.606. There-

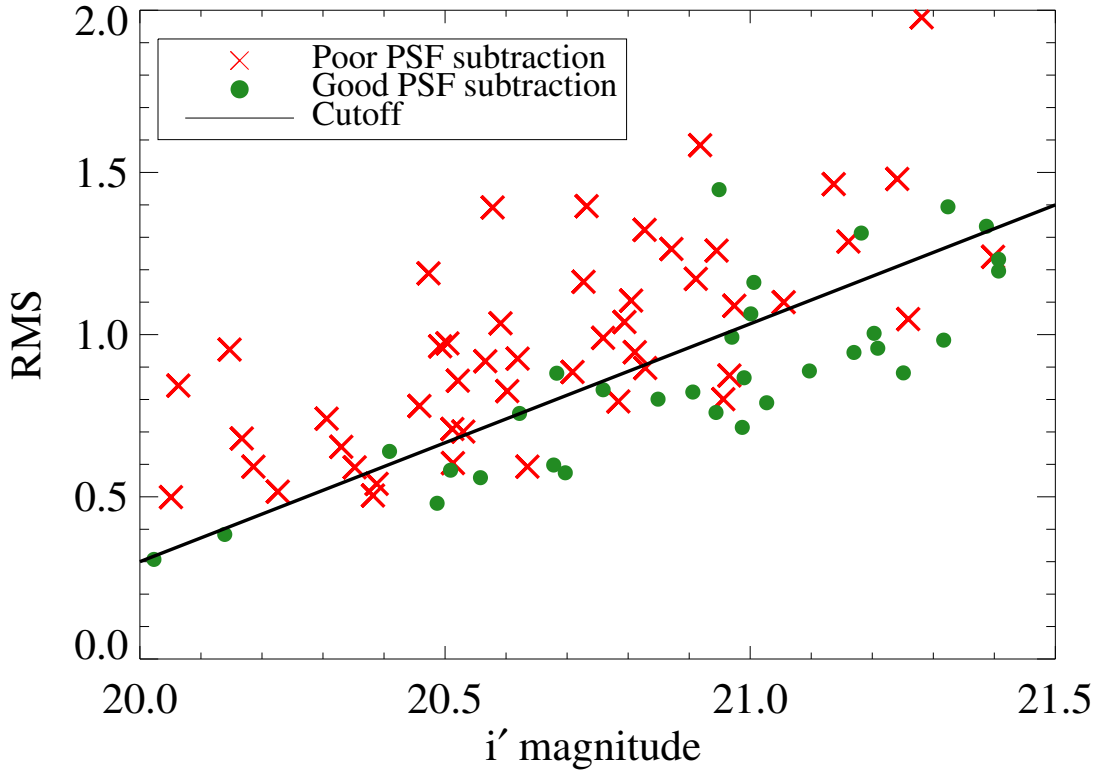


Figure 4.3: RMS isolation criterion. Visually identified blends (red crosses) and non-blends (green circles) are shown from a small representative area about 8 kpc from the center of M31. The RMS deviation of the pixel values of the residual image in a 5×5 pixel square centered on each source is plotted against the i' magnitude of the source. We reject all sources above the solid black line.

fore, each $0''.8$ wide slitlet subtends 4.1 pixels. Better still, excellent seeing conditions ($\sim 0''.6$) during observations can provide somewhat better spectral resolution yielding an average resolution of 3.1 pixels = 1.0 \AA .

Reliable spectra (those that yield secure velocities, as described in § 4.2.7), are obtained from 4465 of the 5263 slitlets. Approximately 30% of the slitlets from slitmask M32_3 did not produce useful spectra due to a warp in the slitmask.

See Table 1 for information on the positions of the slitmasks and the number of useful spectra recovered from each one. This chapter uses all the data taken through

the end of 2010.

4.2.5 Data Reduction

The Keck/DEIMOS multiobject slitmasks are processed using the *spec2d* and *spec1d* software (version 1.1.4) developed by the DEEP Galaxy Redshift Survey team at the University of California, Berkeley (Davis et al., 2002)¹. Briefly, the reduction pipeline rectifies, flat-field and fringe corrects, wavelength calibrates, sky subtracts, and cosmic ray cleans the two-dimensional spectra, and extracts the one-dimensional spectra. For more details, see Howley et al. (2013).

4.2.6 Cross-Correlation Analysis

Line-of-sight (LOS) velocities for resolved targets are measured from the one-dimensional spectra using a Geha et al. (2010) modified version of the visual inspection software *zspec*, developed by D. Madgwick for the DEEP Galaxy Redshift Survey at the University of California, Berkeley. The software determines the best-fit LOS velocity for a target by cross-correlating its one-dimensional science spectrum with high signal-to-noise stellar templates in pixel space and locating the best fit in reduced- χ^2 space.

A-band telluric corrections and heliocentric corrections are calculated and applied to the measured LOS velocities. The A-band telluric corrections, which account for velocity errors associated with the slight mis-centering of a star in a slit, are determined using the method discussed in Sohn et al. (2007) and Simon & Geha (2007).

LOS velocity errors are determined for each star by adding in quadrature

¹<http://astro.berkeley.edu/~cooper/deep/spec2d/>

the cross-correlation based velocity error and a systematic error estimated by repeat observations, as described in Howley et al. (2013). The typical LOS velocity error in our sample is 4–5 km s⁻¹.

4.2.7 Quality Assessment

Each two-dimensional spectrum, one-dimensional spectrum, and corresponding Doppler shifted template match are visually inspected in *zspec* and assigned a quality code based on the reliability of the fit. This process allows the user to evaluate the quality of a spectrum and reject instrumental failures and poor quality spectra. Velocity measurements based on two or more strong spectral features are labeled “secure.” Velocity measurements based on one strong feature plus additional marginal features are labeled “marginal.” Spectra that contain no strong features, low S/N and/or instrumental failures are considered unreliable, and so are not included in our analysis. Additional details on quality code assignment can be found in Guhathakurta et al. (2006). During this process, we also identify and flag 43 likely MW M dwarfs based on their strong surface-gravity sensitive Na I 8190Å doublet. These are excluded from the radial velocity analysis.

4.2.8 Serendipitous Sources

Upon visual inspection of the one-dimensional and two-dimensional spectra during the quality assessment phase outlined in § 4.2.7, some fraction of the slits clearly show that the slitlet intersects more than one star: the target star and one or more serendipitously detected stars, or *serendips*. Serendips are detected via one of two

methods: through continuum detections that are offset from the primary target in the spatial direction, or by the detection of spectral features that are offset from the primary target in the spectral direction. Serendip detections occur frequently in the inner parts of M31 and close to M32 due to the severe crowding and blending in the CFHT/MegaCam data. Serendips are also assigned quality codes and those with secure or marginal velocities are included in the radial velocity analysis. More information on serendipitous detections can be found in Howley et al. (2013).

4.3 DATA ANALYSIS

We perform our analysis in each of the five spatial regions labeled in Figure 4.4. These regions are the SE minor axis (not expected to yield constraints on the spheroid rotation velocity); the SSW quadrant (expected to yield a constraint on the rotation velocity via a negative velocity offset from the systemic velocity of M31); and three regions along the northeast major axis, which we name NE1, NE2 and NE3 in order of increasing projected radial distance (expected to yield three independent estimates of the rotation velocity via positive velocity offsets from the systemic velocity of M31). Note that these regions are defined by lines of constant position angle and projected radius, and are not quite the same as the three data sets shown in Figure 4.1. The positions of and number of stars within each region are given in the first six columns of Table 4.1.

The kinematically cold peak at around -100 km s^{-1} in Figure 4.5 suggests that our measured velocity distribution has a significant contribution from the disk. Hence,

it is imperative that we realistically account for disk contributions. Instead of adopting a specific model for its velocity field, we only assume that the stellar disk is locally cold with a symmetric velocity distribution. As explained in § 4.3.1, we apply this assumption to divide each region into several subregions. To each subregion we fit two Gaussian distributions, corresponding to a kinematically cold and a hot component, where the hot component is required to have the same mean velocity and velocity dispersion across all subregions in a region (described in § 4.3.2). Lastly, in § 4.3.3 we describe how we modify our analysis to account for the possibility of contamination by the GSS and associated tidal debris.

4.3.1 Choice of Subregions and Expected Disk LOS Velocity Pattern

Because we assume that the disk is only locally cold, we fit for the disk in each of many small subregions. The spatial boundaries of individual subregions are dictated by two competing desirable factors, namely a small spread in disk mean velocity and high number statistics. Our assumption of a perfectly cold disk is only strictly true in the limit of infinitely small subregions. On the other hand, a multi-Gaussian fit to a velocity distribution requires a somewhat large number of points. We arbitrarily decide that 100 is the minimum satisfactory number of points.

To estimate the spread of the mean disk velocity as a function of position, we employ a simple geometrical model for the rotation pattern of an inclined disk with perfectly circular motion (Guhathakurta et al., 1988):

$$v_{\text{obs}}(\xi, \eta) = v_{\text{sys}} \pm \frac{v_{\text{rot}}(r_{\text{deproj}}) \sin(i)}{\sqrt{1 + \tan^2(\Delta\text{PA}) / \cos^2(i)}} \quad (4.3)$$

where ξ, η are tangent-plane coordinates with origin at the center of M31, $i = 77^\circ$ is the inclination of the disk of M31, $v_{\text{sys}} = -300 \text{ km s}^{-1}$ is the systemic heliocentric velocity of M31, ΔPA is the position angle projected onto the plane of the sky measured relative to the major axis of the disk of M31, r_{deproj} is the radial position measured in the plane of the disk, v_{rot} is the disk rotation speed, and the $+$ and $-$ signs apply to the NE and SW halves of the disk, respectively. An azimuthally averaged estimate for $v_{\text{rot}}(r_{\text{deproj}})$, based on HI kinematics, ranges from $\sim 250 \text{ km s}^{-1}$ at $r_{\text{deproj}} = 15 \text{ kpc}$ to $\sim 175 \text{ km s}^{-1}$ at 5 kpc (Corbelli et al., 2010). The expected velocity spread calculated in this way will not be exact, because in addition to ΔPA , the true spread in mean stellar velocity over a subregion is influenced by Δr_{deproj} of the subregion, any departure from perfectly circular motion, and the intrinsic local velocity distribution (due to a multiple-component disk, for example).

Using Equation 4.3 with $v_{\text{rot}} = 250 \text{ km s}^{-1}$, we bin our data in each region into subregions based on position angle. The angle subtended by a single subregion is approximately the greater of two ΔPA criteria: 1) the ΔPA such that the change in v_{obs} over a subregion due to ΔPA is 10 km s^{-1} , or 2) the ΔPA that includes 100 data points. Our final subregions are shown in Figure 4.4. In the NE1–NE3 and SSW regions, we identify these subregions with subscripts that increase with distance from the nearest major axis: NE1₁, NE1₂, ..., NE1₇; SSW₁, ..., SSW₅, etc. In the SE region, the outer, inner south, and inner north subregions are named SE₁, SE₂, SE₃, respectively. Note

that we use this rotation pattern only to estimate appropriate bin sizes, not to determine disk rotation speeds. The final positions of and number of stars in each subregion are presented in the first six columns of Table 4.3 in the Appendix.

4.3.2 Fitting the Velocity Distribution Model

We use a Markov chain Monte Carlo (MCMC) sampler to find the velocity distributions of the disk and spheroid in each region. The spheroid is modeled as a single Gaussian distribution for each region, while the disk is modeled as a Gaussian distribution for each subregion. The model likelihood at each point in parameter space is then:

$$L = \prod_{j=1}^N \left[f_{s(j)} \mathcal{N}(v_j | \bar{v}_{\text{sph},r(j)}, \sigma_{\text{sph},r(j)}) + (1 - f_{r(j)}) \mathcal{N}(v_j | \bar{v}_{\text{disk},s(j)}, \sigma_{\text{disk},s(j)}) \right] \quad (4.4)$$

where N is the number of stars in our data set, and star j with measured velocity v_j is found in region $r(j)$ and subregion $s(j)$. The notation $\mathcal{N}(x | \mu, \sigma)$ indicates the Gaussian distribution with mean μ and standard deviation σ evaluated at x . The scalar f_s , subject to $0 \leq f_s \leq 1$, is the fraction of the stars in subregion s that belong to the spheroid. The spheroid in region r has mean velocity $\bar{v}_{\text{sph},r}$ and velocity dispersion $\sigma_{\text{sph},r}$, while the disk in subregion s is characterized by velocity $\bar{v}_{\text{disk},s}$ and dispersion $\sigma_{\text{disk},s}$.

Subregion SSW₄ includes stars from the galaxy M32, so for that subregion only we replace the single-Gaussian disk model $\mathcal{N}(v_j | \bar{v}_{\text{disk},s(j)}, \sigma_{\text{disk},s(j)})$ with a double-

Gaussian model.

The likelihood in Equation 4.4 can be sampled independently in each of the five regions. We use the code `emcee` (Foreman-Mackey et al., 2015), which implements the affine-invariant ensemble sampler of Goodman & Weare (2010), to perform the MCMC algorithm. In addition to the likelihood above, we must specify prior probability distributions for the parameters. For f_s we use a uniform distribution on the interval $[0, 1]$; for the mean velocities $\bar{v}_{\text{sph},r}$ and $\bar{v}_{\text{disk},s}$ we use flat priors; and for the dispersions $\sigma_{\text{sph},r}$ and $\sigma_{\text{disk},s}$ we demand positive values. If we sample each region independently, then the parameter space for region r includes $\bar{v}_{\text{sph},r}$, $\sigma_{\text{sph},r}$, plus the set of f_s , $\bar{v}_{\text{disk},s}$, $\sigma_{\text{disk},s}$ for each subregion s within r . We initialize the MCMC at $\bar{v}_{\text{sph},r} = -300 \text{ km s}^{-1}$, $\sigma_{\text{sph},r} = 150 \text{ km s}^{-1}$, each $f_s = \frac{1}{2}$, and $\bar{v}_{\text{disk},s}$ and $\sigma_{\text{disk},s}$ = the sample mean and standard deviation of the velocities in subregion s .

Briefly, the MCMC ensemble sampler works as follows. It explores the parameter space by maintaining a set of *walkers*. Each walker represents a point in the parameter space. At each iteration of the MCMC, each walker takes a step in the parameter space by choosing another walker and stepping along the line in parameter space connecting itself to the other walker. The step size is chosen stochastically and allows interpolation as well as extrapolation. In effect, the walkers choose their steps based on the covariance of the set of walkers. After each step is taken, the posterior probability distribution at the new point in parameter space is evaluated. Steps that increase the probability are always accepted, while steps that decrease the probability are *sometimes* accepted. After a large number of steps, the ensemble of walkers will sample the pa-

parameter space with frequency proportional to the posterior probability distribution; we can draw fair samples from the distribution by selecting points from the histories of the walkers. We estimate the mean and variance of each parameter (assuming unimodal distributions) based on sample statistics of the histories of the walkers. In particular, we allow each of 32 walkers to take 10,000 steps. We then compute the mean spheroid velocity as the mean value of $\bar{v}_{\text{sph},r}$ over the last 2000 steps of all the walkers (i.e., the mean of 64,000 points), and the 68% confidence interval as the standard deviation of that quantity over the last 2000 steps of all of the walkers. We use the same method to estimate the values and uncertainties of the spheroid dispersion and the disk parameters. We report these values in columns 7 and 8 of Table 4.1 and illustrate them in Figure 4.9.

4.3.3 Accounting for Tidal Debris Associated with the Giant Southern Stream

The analysis described thus far separates the spheroid from the disk (and from M32 in subregion SSW₄). The mean velocity and dispersion we extract describe the average properties of the spheroid, regardless of its underlying structure.

Several lines of evidence, however, suggest that our measurements do not well represent a rotation curve of a smooth spheroid. First, the mean spheroid velocities in the five regions do not follow a physical rotation pattern: the mean velocity of the NE2 region is consistent with zero, while the surrounding regions (NE1 and NE3) are rotating at about 50 km s⁻¹ in the same direction as the disk. Second, the sum of hot and cold Gaussians does not well represent the velocity distributions. A chi-squared

analysis applied to the data binned by 20 km s^{-1} reveals that the probability of the model representing the velocity distribution is relatively low (see the final column of Table 4.1).

Velocity histograms, such as those in Figure 4.6, suggest that cold substructure, possibly tidal debris from the GSS, could be skewing our measurements. A close-up view of the negative-velocity tail of the NE2 histogram reveals a cold spike of about 10 stars in excess of the Gaussian tail at -580 km s^{-1} (Figure 4.6, top right). A slight overdensity of points at this velocity can be seen in the NE1 and NE3 histograms as well. Though nothing is immediately visible in the SE or SSW regions, a few extra stars around -580 km s^{-1} would be partially concealed by the bulk of the velocity distribution.

Figure 4.7 shows another projection of these data: the velocity of each star in the NE1, NE2, and NE3 regions plotted against R_{proj} . Also shown (in turquoise triangles) are data from two GSS fields at 17 and 21 kpc (Gilbert et al., 2009). The concentrations of stars at approximately -500 and -400 km s^{-1} in these data represent the GSS and the secondary stream, respectively. The magnitudes of the central stream velocities increase with decreasing R_{proj} as the streams fall into the potential well of M31 from the south. The black crosses in Figure 4.7 show the expected stream velocity as a function of radius closer to the center of M31 (Fardal et al., 2013). The NE region data from this work show clear concentrations of objects near the predicted stream velocity, continuing the trend seen south of the galaxy. Hence, it seems unlikely that the peak in the Figure 4.6 histograms is simply a binning artifact, and probable that it comes from the northern extension of a cold tidal stream.

To account for the presence of the GSS and its associated tidal debris, we repeat the MCMC fits in regions NE1, NE2, NE3, and SSW after removing all stars within $\sigma = \pm 30 \text{ km s}^{-1}$ of the predicted velocities of the two streams. (The measured velocity dispersion of the stream from Gilbert et al. (2009) is 20 km s^{-1} ; however, we use the larger value to ensure that we exclude all stream stars despite the slightly uncertain mean stream velocity.) In the SE minor axis region, we exclude all stars with velocity $v < -600 \text{ km s}^{-1}$ or $v > 0 \text{ km s}^{-1}$. The former is to account for GSS debris, and the latter for SE Shelf stars (GSS debris from third pericentric passage) (Gilbert et al., 2007).

We account for the fact that we have removed stars within a range of velocities by renormalizing the Gaussian distributions. That is, given a Gaussian velocity distribution, we compute the fraction of the probability mass that falls within the excised velocity range and scale up the remainder of the distribution so that the integral over the remaining (unexcised) velocities is unity. The resulting kinematical spheroid parameters are reported in columns 7 and 8 of Table 4.2.

The resulting fit for region NE1 is plotted in Figure 4.8. The first seven panels show the seven subregions in NE1. In each panel, each walker at the end of 10,000 steps is represented by a set of colored lines: a red Gaussian with parameters f_s , $\bar{v}_{\text{sph},r}$, $\sigma_{\text{sph},r}$ and a blue or green Gaussian with parameters $(1 - f_s)$, $\bar{v}_{\text{disk},s}$, $\sigma_{\text{disk},s}$. The sum of these two is shown in violet. These distributions are overlaid on the velocity histogram of stars in the subregion. The velocity ranges excised for stream debris are shown in two ways: the velocity range is shaded in light gray, and the stars in this range are colored

red.

The bottom middle panel shows the cumulative best-fit distributions: the red line is the Gaussian corresponding to the mean $(\bar{v}_{\text{sph},r}, \sigma_{\text{sph},r})$. The blue curve is the cumulative disk distribution: the sum of the mean disk Gaussians from the seven subregions. The violet is the sum of the other two curves. The excised velocity ranges are represented in the same way as in the subregion panels.

The bottom right panel shows the positions of the walkers in parameter space after 10,000 steps. There are eight concentrations of points corresponding to the eight (\bar{v}, σ_v) pairs (one spheroid and seven disk). Each pair is summarized by an ellipse displaying the mean and dispersion of that distribution.

4.4 DISCUSSION

This section is organized as follows. We first present the rotation curve and velocity dispersion profile for the inner spheroid in §4.4.1. In §4.4.2 we compare the velocity dispersion of the cold component to previous measurements of the stellar disk as a sanity check on our analysis method. In §4.4.3 we show that the exclusion of velocity ranges corresponding to tidal debris from the GSS significantly impacts the measured kinematical parameters of the spheroid, but has minimal effect on the cold component. In §4.4.4 we compute a spheroid membership probability for each star in our sample. In §4.4.5 we explain that the spheroid is likely supported by velocity anisotropy in addition to rotation. Finally, in §4.4.6 we show that neither spectroscopic target selection criteria nor degree of crowding introduces a significant bias towards the

kinematically cold or hot population.

4.4.1 Kinematical Parameters of the Inner Spheroid

The spheroid distributions corresponding to the best-fit \bar{v}_{sph} and σ_{sph} (hereafter σ_v) are plotted in Figure 4.6 for each of the five regions. (This figure is simply a compilation of the summary panels of Figures 4.8 and 13–16.) Kinematical and goodness-of-fit parameters, accounting for the GSS and associated tidal debris, are reported in Table 4.2. The χ^2 probabilities have improved significantly from the uncorrected values in Table 4.1. Four of the five regions now have probabilities greater than 90%, and that of the fifth (SSW) has increased by a factor of 15.

The dispersion profile is shown in the right panel of Figure 6.6. The profile appears to decrease smoothly with radius, though it is consistent with flat to 2σ . The best-fit line to this profile is

$$\sigma_v(a_{\text{eff}}) = (159.5 \pm 10.8) - (2.9 \pm 1.3) \frac{a_{\text{eff}}}{1 \text{ kpc}} \text{ km s}^{-1}. \quad (4.5)$$

where a_{eff} is the effective spheroid major axis coordinate, assuming a 5:3 axis ratio (Pritchet, C.J., van den Bergh, 1994). Our dispersion profile is consistent with that measured by Gilbert et al. (2007) but is slightly offset from other existing measurements, including the Saglia et al. (2010) integrated-light measurement at 1.1 kpc on the major axis of the bulge (black square in Figure 6.6) and the linear dispersion profile of Chapman et al. (2006). These differences should be taken lightly, though, because the measurements are not directly comparable. The Saglia et al. (2010) point may be

slightly deflated by contributions from the low-dispersion disk stars, which those authors estimate contribute about 30% of the light in the slit. The Chapman et al. (2006) profile, meanwhile, is the innermost limit of measurements primarily made farther out in the halo, so we do not necessarily expect agreement at the radii covered by our study.

We also compare our results to the dispersion profile produced by a model with spherical, isotropic, non-rotating bulge and halo stellar components as well as a stellar disk and halo ($M/L = 2.5$) and bulge ($M/L = 5.6$) (dotted orange line in Figure 6.6; Fardal et al., 2013). The model falls below our data at the 2σ level in the NE2 and NE3 regions and at the 3σ level in the SE and NE1 regions. The mismatch suggests that, for example, the mass profile of the model galaxy is too shallow, or the gradient of the density profile of the tracer population is too large. A more detailed analysis is beyond the scope of this paper.

Extracting the intrinsic rotation curve of the spheroid is nontrivial. The relationship between the mean LOS velocity and the intrinsic rotation velocity depends on the orbital dynamics of the spheroid, and in general is difficult to determine without detailed 2D or 3D kinematical mapping. Therefore, in the left panel of Figure 6.6 we simply present the mean LOS component of the velocity versus the effective major axis coordinate (based on a 5:3 axis ratio). Though two points are consistent with zero to 2σ , we detect significant rotation in the SSW, NE1 and NE3 regions. The average value of $|v - v_{M31}|$, excluding the SE minor axis point, is $52.6 \pm 6.8 \text{ km s}^{-1}$ (solid line in Figure 6.6). All four off-minor axis points are consistent with this mean velocity to better than 1σ . This is the first measurement of significant rotation in the inner spheroid.

The velocity dispersion of M31’s inner spheroid is similar to that the halo of the MW at similar radii, but its mean velocity is significantly larger. The velocity ellipsoid of the MW’s inner halo is $(\sigma_{V_R}, \sigma_{V_\phi}, \sigma_{V_Z}) = (150 \pm 2, 95 \pm 2, 85 \pm 1)$ (Carollo et al., 2010), on the same order as the LOS component of the M31 spheroid dispersion. However, the MW’s inner halo has a mean rotation velocity consistent with zero (Carollo et al., 2010).

4.4.2 Dispersion and Velocity of the Cold Population

The cold component (\bar{v}, σ_v) for each subregion are plotted against r_{deproj} in Figure 4.17. As a sanity check on our analysis method of fitting multiple Gaussians to subregions, we compare these values to previously measured kinematical parameters of the stellar disk. The cold component in a single subregion has an average dispersion of 58 km s^{-1} , reasonably consistent with the Collins et al. (2011) dispersion measurements of 51 km s^{-1} and 36 km s^{-1} for the thick and thin stellar disks, respectively. We expect our dispersion to be larger than the local value, primarily because the finite spatial extent of our subregions necessarily smears out the velocity distribution. For example, if the spread in mean velocity due to ΔPA is $\sim 20 \text{ km s}^{-1}$ in a subregion, and the true velocity dispersion is $\sim 40 \text{ km s}^{-1}$, then we would expect to measure an effective dispersion of $\sqrt{20^2 + 40^2} = 45 \text{ km s}^{-1}$. This effect is accentuated if there is an additional spread in mean velocity with radius. Therefore, the fact that our cold component looks like the thick disk of Collins et al. (2011) does not imply that the thick component dominates the stellar disk.

A possible concern is that contributions from multiple stellar disk components

(a combination of thin, thick, or extended) may invalidate our assumption of a locally cold disk. If this were the case, we would expect to see non-Gaussianity in the velocity signature of the cold component. However, Table 4.3 shows that that the velocity distributions of the cold components in most of the subregions are well fit by a single Gaussian. While this observation does not say anything about the kinematical structure of the disk, it does suggest that our simple assumption is adequate for our purpose of describing the spheroid.

It is also possible that a kinematically warm thick disk component would be incorporated into our Gaussian representation of the spheroid; however, the high goodness-of-fit statistic again suggests that this is not a significant effect. In any case, the bias on our kinematical spheroid parameters induced by thick disk contamination confirms, rather than invalidates, our qualitative results, as discussed in §4.4.5.

A closer look at the best-fit cold and hot components by subregion in Figures 4.8 and 13–16 show that the trends in the cold component with position angle match those expected for an inclined rotating disk. For example, in Figure 4.8, the mean velocity of the cold component (blue and green curves) transitions from -100 km s^{-1} in subregion NE1₁, along the major axis, to -200 km s^{-1} in subregion NE1₇, farthest from the major axis. In other words, the absolute value of the offset of the mean velocity of the cold component from the systemic velocity of M31 moves closer to zero as we march away from the major axis. This progression is clear in Figures 4.13 and 4.14 for regions NE2 and NE3 as well, although it is less pronounced because these regions subtend a smaller range in PA.

4.4.3 Effect of Tidal Debris Associated with the GSS on Spheroid Kinematics

Accounting for the GSS and its associated debris has a significant effect on the measured kinematical parameters of the underlying smooth inner spheroid, supporting the observation of Fardal et al. (2007) that the debris was visible in the samples of Ibata et al. (2005), Chapman et al. (2006), and Merrett et al. (2006).

It is unclear whether the secondary stream identified by Kalirai et al. (2006a) and confirmed by Gilbert et al. (2009) is present in the NE1–NE3 regions. Figure 4.7 reveals a second clump of objects in the NE2 region, offset from the primary stream by 75–100 km s⁻¹, the same separation as that between the two streams south of the galaxy. However, such clumps are barely, if at all, visible in the NE1 and NE3 regions, and their exclusion does not significantly affect the kinematical parameters of the spheroid. Further observations are necessary to determine if there is a second substructure and, if so, whether it has a physical connection to the GSS.

Is the GSS the only source of nonvirialized substructure biasing our measurements of the mean velocity and dispersion of the inner spheroid? While we cannot prove that all stars except those in the excised velocity ranges belong to either a disk or a perfectly smooth spheroid, we can show that the effect of other nonvirialized substructure on our kinematical characterization of the underlying smooth spheroid is negligible. In star-count maps of the inner regions of M31, such as those in Ibata et al. (2001), the GSS is by far the most prominent substructure. Even so, GSS stars account for only a small fraction of our spectroscopic sample (see pink shaded regions in Figure 4.6), and

have a relatively small effect on our results. Other, less-prominent substructures would (1) be nearly impossible to detect and account for, and (2) have a negligible effect on the measured kinematical parameters of the inner spheroid.

4.4.4 Spheroid/Disk Membership Probability and Extreme Velocity Stars

We can apply our subregion fits to quantify the spheroid membership probability for any star. At the location and velocity of each star in our sample, we calculate the ratio of the values of the best-fit hot component in that region and the disk and M32 components in that subregion. The disk and M32 membership likelihoods are calculated in a similar fashion. The results are sorted into three categories as follows: stars that are at least three times as likely to be disk members as anything else (yellow in Figure 4.11); stars at least three times as likely to be spheroid members as anything else (magenta); and other objects, including likely M32 members (green). As expected, objects on the major axis are much more likely to be disk members than are those on the minor axis. Most important, we see likely spheroid members at all radii covered by our sample.

Recently, Caldwell et al. (2010) reported discovery of an “extreme velocity” star at a projected radius of 4 kpc (0.3°) along the SW major axis of M31. The star has a velocity of -780 km s^{-1} , essentially excluding it from membership in the thin, cold stellar disk. Those authors attribute the star’s highly negative velocity to possible membership in the GSS, even though it would have to be a 6σ outlier of the stream velocity distribution. Our probability map demonstrates that this object can more

easily be interpreted as a member of the spheroid, even at these large radii: it may be a 4σ outlier of the spheroid distribution.

It is tempting to interpret this distribution of probabilities as a map of bulge-to-disk fraction. However, this statistic can be more reliably constrained using photometric light-profile fitting in conjunction with kinematical decomposition. This study is the subject of Dorman et al. (2013).

4.4.5 Anisotropy

We investigate the degree to which the flattening of the spheroid may be due to rotation by comparing the spheroid ellipticity $\epsilon = 1 - b/a$ to the ratio v/σ_v . The value of ϵ is uncertain at these radii. Pritchett, C.J., van den Bergh (1994) measure $\epsilon = 0.4$ with limited data at $R_{\text{proj}} = 10$ kpc. Courteau et al. (2011) perform a fit to a more extended data set to obtain values between 0.21 and 0.37, depending on their bulge/disk decomposition and modeling technique, for a relatively small bulge with scale length ~ 1.0 kpc.

Despite this range of possible ellipticities, \bar{v}/σ_v of the spheroid in every region is lower than that of a rotationally flattened oblate isotropic rotator (black line in Figure 4.12). We measure $\bar{v}/\sigma_v = 0.23\text{--}0.52$, but v_{rot}/σ_v of 0.87 or 0.54 would be required for rotation to flatten the spheroid to an ellipticity of 0.4 or 0.21, respectively. Hence, rotation alone probably does not account for all of the flattening of the spheroid.

It is possible that anisotropy in the velocity ellipsoid can provide the remainder of the flattening. Anisotropy can be parameterized by

$$\delta \equiv 1 - \frac{\Pi_{zz}}{\Pi_{xx}} \quad (4.6)$$

where z points along the axis of symmetry of the spheroid, x is any orthogonal direction (say, r), and Π_{kk} represents the pressure from velocity dispersion along direction k (e.g. Cappellari et al., 2007). In a system with an axisymmetric velocity ellipsoid, $\delta = 0$ corresponds to a system whose flattening is unaffected by anisotropy, while $\delta \sim 1$ corresponds to a system whose flattening is almost entirely due to anisotropy. Using the iterative method described in Cappellari et al. (2007), we find that an anisotropy of $\delta = 0.05$ – 0.27 is required to explain a spheroid ellipticity of 0.21 – 0.4 , given the mean value of \bar{v}/σ_v in regions NE1–NE3 and SSW.

We compare the inner spheroid of M31 to other spheroidal systems. The bulges of spiral galaxies generally fall on or above the line for rotationally-flattened systems (Kormendy & Illingworth, 1982; Kormendy, 1982). In contrast, so-called “fast-rotating” ellipticals tend to lie between this line and the line $\delta \sim 0.7\epsilon_{\text{intr}}$ (magenta in Figure 4.12) (Cappellari et al., 2007). However, these comparison systems are observed at around one effective radius (r_{eff}), whereas our kinematical measurements range from 1.5 – $14 r_{\text{eff}}$. The inner spheroid of M31, then, more closely resembles the inner r_{eff} of a fast-rotating elliptical than the central bulge of a spiral galaxy.

As suggested earlier in §4.4.2, it is possible that the spheroid velocity distribution may be contaminated by thick disk stars. If so, the true spheroid mean velocity is lower than we report, and the true spheroid dispersion higher. However, note that the \bar{v}/σ_v of the spheroid is already small enough to look more like an elliptical galaxy

than the bulge of a spiral galaxy. The possible bias induced by a thick disk component would merely increase this effect, confirming our conclusion that M31's inner spheroid rotates unusually slowly.

4.4.6 Effect of Target Selection Criteria on Velocity Distribution

The majority of our spectroscopic targets were chosen on the basis of i' magnitude only, but a small fraction of the targets in the MCT region were identified based on position in the PHAT CMD. This latter category contains proportionally more of the rarer metal-poor and metal-intermediate RGB populations. If metal-poor and metal-intermediate RGBs preferentially trace the kinematically hot population, then inclusion of the PHAT-selected objects could bias our measurements, especially our spheroid membership probabilities. To test for bias, we performed a Komolgorov-Smirnov (K-S) test comparing the velocity distributions of the PHAT-selected targets and the magnitude-selected targets in each subregion. The test confirmed that the two distributions were indistinguishable.

Similarly, we used a K-S test to confirm that inclusion of crowded objects does not bias us towards one structural subcomponent. We created two lists of objects: "crowded" (those sharing a slit with at least one serendipitously detected neighbor) and "isolated" (those without any such neighbors). Again, there was no significant difference between the velocity distributions of the two categories in any subregion.




Figure 4.4: Division into regions and subregions. Each point corresponds to a single velocity measurement. Solid lines delineate the five regions for which we obtain independent kinematical parameters of the spheroid: inner northeastern (NE1), intermediate northeastern (NE2), outer northeastern (NE3), southeastern (SE), and south-southwestern (SSW). Colors show the subregions within each region that are used to determine the disk contribution to the velocity distribution. With the exception of the SE region, our subregion naming convention is based on distance from the major axis. For example, in the NE3 region, NE3₁ (light gray) straddles the major axis, NE3₂ (black) is slightly farther out, and NE3₃ (dark gray) is farthest from the major axis. In the SE region, we use projected distance from the center of M31 in addition to lines of constant position angle to define three subregions: SE₁ (outer, yellow), SE₁₂ (inner south, dark brown) and SE₁₃ (inner north, light brown).




Figure 4.5: Distribution of heliocentric radial velocities for all 3697 stars with reliable velocity measurements from the MCT slitmasks. The peak around -100 km s^{-1} corresponds to the disk, the LOS projection of whose rotation speed causes stars to be offset by as much as $+200 \text{ km s}^{-1}$ with respect to the -300 km s^{-1} systemic velocity of M31 (dotted line).

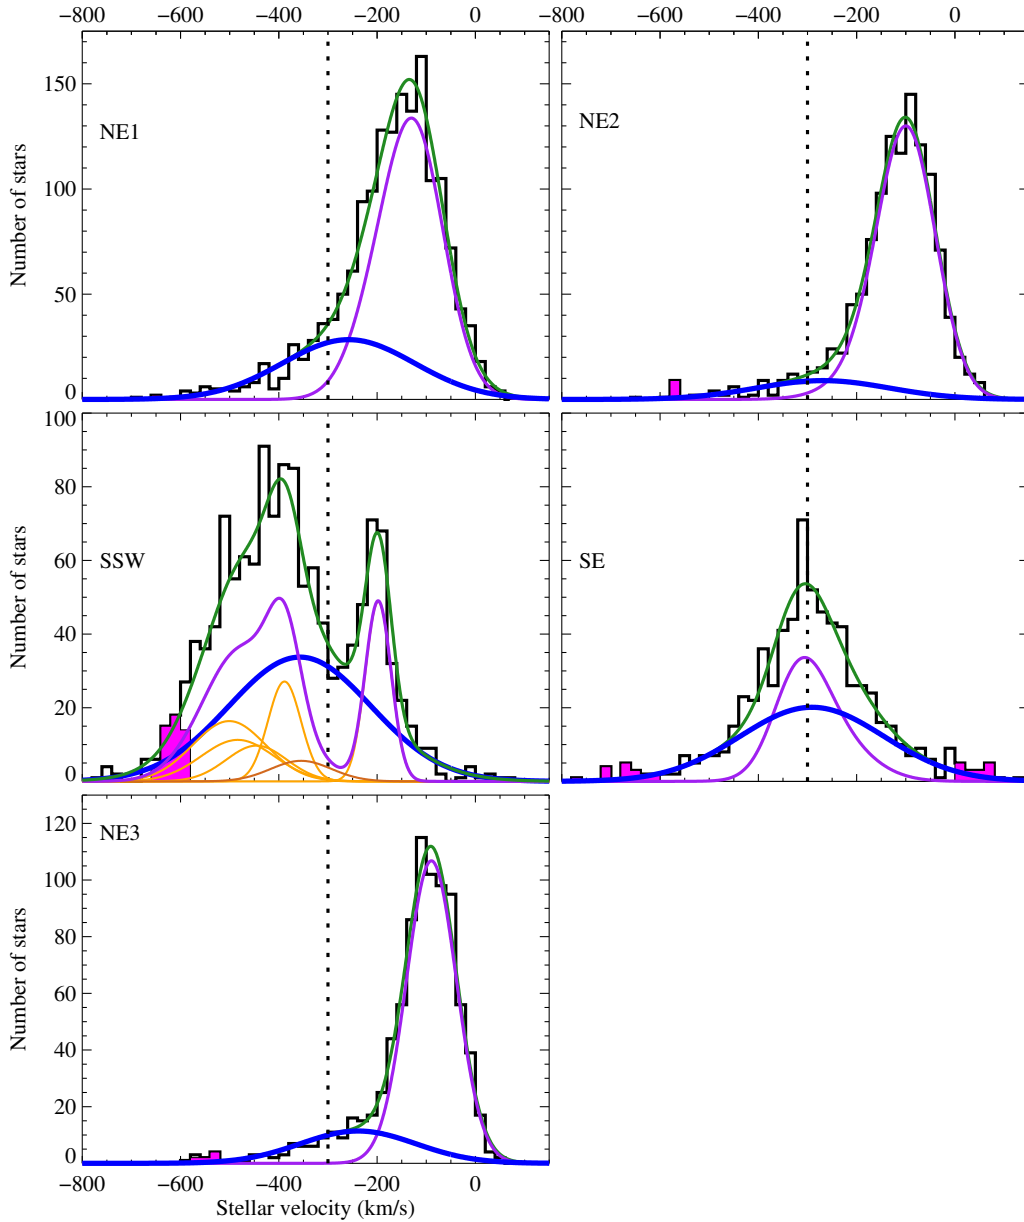


Figure 4.6: Maximum likelihood fits of a kinematically hot spheroid to each of the five regions in our sample, after excluding the velocity range encompassing the GSS and its associated tidal debris (shaded pink). Violet lines show the cumulative region cold component; blue show the best-fit spheroid Gaussian; green show the sum of these two components. The dashed lines show the systemic velocity of M31 relative to the MW. Individual subregion cold components are shown in orange in the SSW region panel, but left out of the other panels for clarity. We remind the reader that our fitting procedure makes use of the distinct velocity profiles of the individual subregions as shown in Figures 4.8 and 13-16, so that the fits are much better constrained than may be apparent from the combined distributions here.




Figure 4.7: Radial velocity vs. projected radius of the most negative-velocity stars in the NE1, NE2, and NE3 regions (red, blue, and magenta, respectively). Overplotted are data from two fields centered on the GSS south of M31 (turquoise triangles) (Gilbert et al., 2009). The clusters of these turquoise triangles around $v \sim -500 \text{ km s}^{-1}$ and $v \sim -390 \text{ km s}^{-1}$ are the GSS and the secondary stream, respectively. The black crosses show six points for the predicted velocity of the NE Shelf (Fardal et al., 2013). The GSS appears to be present in the NE fields as slight concentrations around the black crosses.

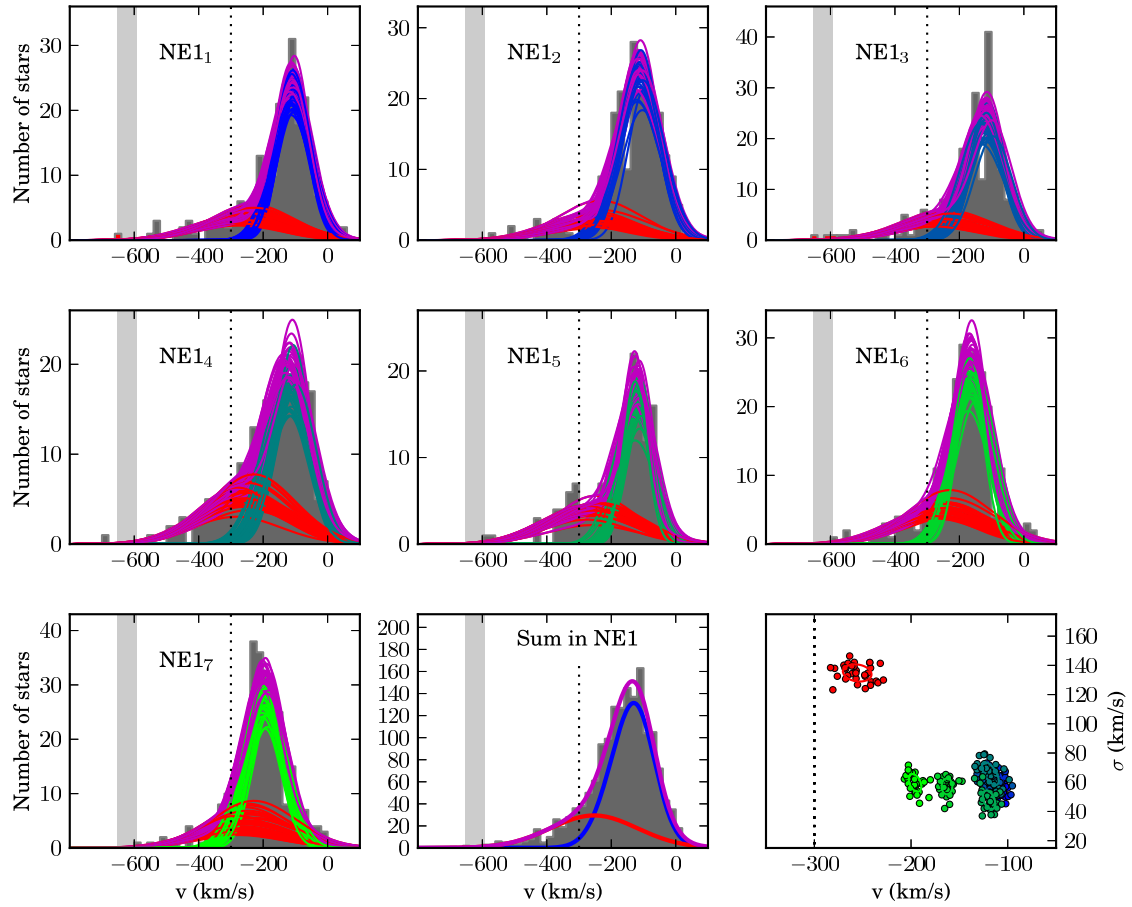


Figure 4.8: MCMC fits of kinematically hot (red) and cold (blue to green) components to each of 7 subregions in the NE1 (inner northeast major axis) region plotted over velocity histograms of stars in each subregion. Subregions progress from NE1₁, straddling the NE major axis, to NE1₇, farthest from the major axis. One curve is drawn for each component for each of the 32 walkers in the MCMC. Each violet distribution is the sum of the hot and cold distributions corresponding to a single walker. Velocity ranges excluded due to possible contamination by tidal debris from the GSS are shown in two ways: by the gray shaded regions, and by the stars shaded red on the histogram. The bottom middle panel shows the best-fit hot component (red) that, when added to the cumulative region cold distribution (blue), best fits the observed velocity distribution. The bottom right panel shows the final position of the walkers in parameter space. Ellipses show the mean and uncertainty of each of the parameters (\bar{v}, σ_v) for each of the kinematical subcomponents. The dotted line in each panel marks the systemic velocity of M31 relative to the MW.

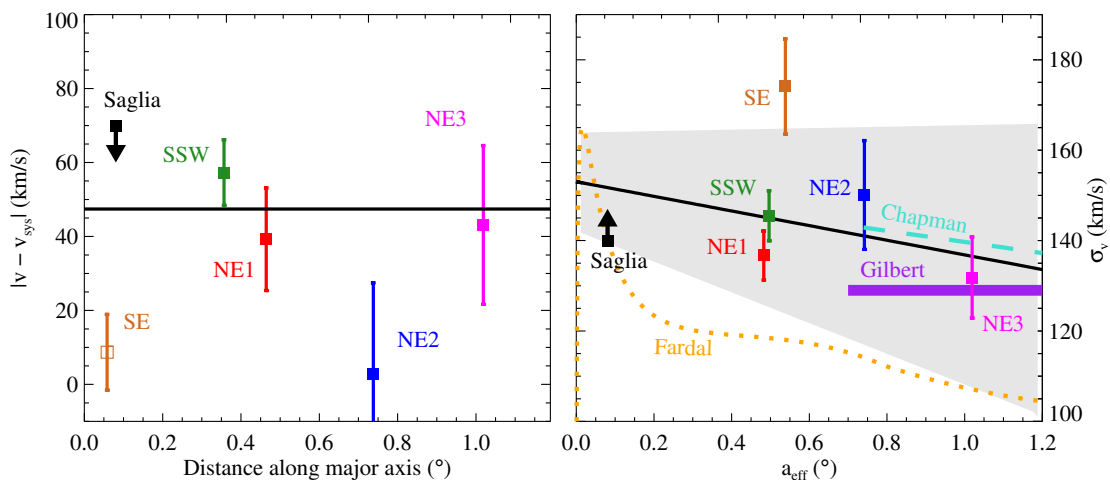


Figure 4.9: Velocity (left) and dispersion (right) profiles for M31’s inner spheroid without accounting for the GSS or associated tidal debris. Left: Mean velocity is plotted against the projection along the major axis coordinate. The SE minor axis (brown open square) is close to the minor axis, so we do not expect to measure a significant $|\bar{v} - v_{\text{sys}}|$ due to rotation. The black line in the left-hand plot is the mean value of the four off-minor-axis velocity measurements. Right: Dispersion is plotted versus a_{eff} , the effective major axis coordinate for a spheroid with a 5:3 axis ratio. The best fit line to our dispersion measurements is shown by the solid black line, with the region within $\pm 1\sigma$ of the best-fit line shaded in gray. Also plotted are limits on the bulge mean rotation velocity and velocity dispersion from Saglia et al. (2010, black points with arrows), and velocity dispersion measurements further out in the spheroid by Chapman et al. (2006) and Gilbert et al. (2007, horizontal purple line). The orange dotted line represents a model dispersion profile that includes isotropic bulge and halo stellar components, as well as a stellar disk and dark matter halo (Fardal et al., 2013).




Figure 4.10: Same as Figure 4.9, excluding stars in the velocity ranges corresponding to the GSS and associated tidal debris.




Figure 4.11: Probable subcomponent membership of each star based on location and velocity. Stars at least 3 times more likely to be a disk member than anything else are shown in yellow; those at least 3 times more likely than anything else to be a spheroid member are shown in pink; all other objects, including likely M32 members, are shown in green. Region divisions are shown as solid lines.

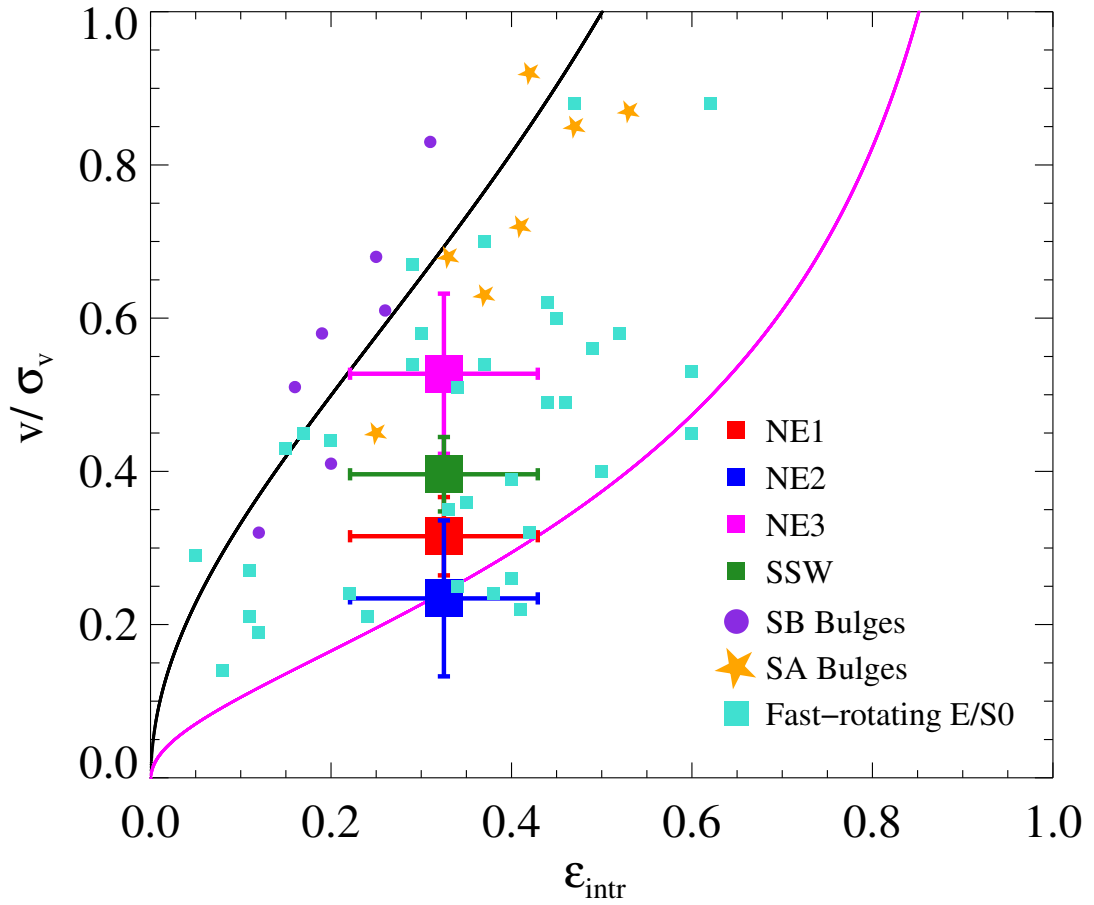


Figure 4.12: Anisotropy diagram for spheroidal systems. The large filled squares show the ratio of the observed mean velocity and the velocity dispersion against the intrinsic ellipticity of M31’s spheroid. Error bars represent the range of reported ellipticities converted to edge-on intrinsic ellipticities as explained in Cappellari et al. (2007). The black curve shows the expected value of the ratio \bar{v}/σ_v as a function of ellipticity for an isotropic, rotating galaxy. The magenta line approximates the behavior of an oblate anisotropic galaxy with anisotropy parameter $\delta \sim 0.7\epsilon$ (Cappellari et al., 2007). Also shown are fast-rotating E/S0 galaxies from SAURON (turquoise squares) and bulges of barred [violet circles, (Kormendy, 1982)] and unbarred [orange stars, (Kormendy & Illingworth, 1982)] spirals. The inner spheroid of M31 appears to be flattened by velocity anisotropy in addition to rotation, and therefore more closely resembles an elliptical galaxy than a spiral galaxy bulge.

Table 4.1: Spheroid Kinematical Parameters

Region	ξ	η	R_{disk}	a_{eff}	N_{region}	$\bar{v} - \bar{v}_{M31}$ (km/s)	σ_v (km/s)	$N_{\text{spheroid}}/N_{\text{total}}$	χ^2 prob
NE1	0.40	0.31	6.45	6.81	1615	$+39.3 \pm 13.9$	136.7 ± 5.4	0.288	0.731
NE2	0.48	0.55	10.11	10.31	1180	$+2.8 \pm 24.6$	150.1 ± 12.0	0.120	0.136
NE3	0.64	0.79	13.93	14.00	859	$+48.0 \pm 24.0$	131.4 ± 9.9	0.203	0.930
SE	0.26	-0.14	4.30	6.91	684	$+8.7 \pm 10.3$	174.1 ± 10.5	0.522	0.121
SSW	-0.09	-0.40	5.69	6.73	1313	-57.3 ± 8.9	145.5 ± 5.5	0.465	1.5×10^{-5}

Table 4.2: Spheroid Kinematical Parameters Corrected for
Tidal Debris

Region	ξ	η	R_{disk}	a_{eff}	N_{region}	$\bar{v} - \bar{v}_{M31}$ (km/s)	σ_v (km/s)	$N_{\text{spheroid}}/N_{\text{total}}$	χ^2 prob
NE1	0.40	0.31	6.45	6.81	1615	$+41.8 \pm 13.3$	134.4 ± 5.5	0.297	0.945
NE2	0.48	0.55	10.11	10.31	1180	$+31.3 \pm 24.5$	135.3 ± 12.5	0.130	1.000
NE3	0.64	0.79	13.93	14.00	859	$+61.2 \pm 16.44$	117.5 ± 8.8	0.198	0.998
SE	0.26	-0.14	4.30	6.91	684	$+9.0 \pm 12.2$	144.5 ± 5.5	0.571	0.819
SSW	-0.09	-0.40	5.69	6.73	1313	-57.1 ± 8.2	145.4 ± 6.5	0.484	2.1×10^{-4}

4.5 SUMMARY & CONCLUSIONS

We have measured reliable radial velocities of over five thousand red giant branch stars in the inner 20 kpc of M31 using the Keck/DEIMOS multiobject spectrograph. Targets were selected using a series of statistical tests designed to identify isolated M31 members bright enough to yield quality spectra.

By fitting a locally cold disk and a kinematically hot spheroid to the velocity distribution with an MCMC algorithm, we have measured the most probable kinematical parameters \bar{v} and σ_v of the red giant stellar population of the inner spheroid of M31 in each of five spatial bins. We find that, though the raw values are inconsistent with a physical rotation pattern, accounting for the presence of tidal debris due to the GSS allows us to detect a significant spheroid rotation velocity.

We find that probable spheroid members are present at all radii in our sample. When used in conjunction with integrated-light measurements, these membership probabilities will be a powerful tool for rigorous bulge/disk decomposition in the inner parts of M31.

We also compare the \bar{v}/σ_v of the inner spheroid to those of other spheroidal structures. We find that rotation is insufficient to explain the flattening of the inner spheroid; it more closely resembles an anisotropic fast-rotating elliptical than the bulge of a spiral or barred spiral galaxy as measured at $\sim 1 r_{\text{eff}}$.

Our magnitude-limited survey biases us towards the bright, old stellar population. As more PHAT data becomes available over the next few years, we will be able to select spectroscopic targets that represent a range of stellar populations within our

magnitude range. We plan to present an analysis of the kinematics of different stellar populations in an upcoming paper.

4.6 MCMC results for regions NE2, NE3, SE, and SSW

Figures 13–16 are the analogs of Figure 4.8 for regions NE2, NE3, SE, and SSW, respectively. In this appendix, we briefly describe the results in each of these regions.

4.6.1 NE2 Region

Eleven stars are excluded from the MCMC fit in the NE2 region due to possible tidal debris membership; as seen in Figure 4.13, all are from the four subregions closest to the major axis. Exclusion drastically improves the chi-squared probability that the spheroid plus disk Gaussians are a good representation of the velocity distribution, raising it from 0.136 to 1.000. Exclusion also increases the mean velocity and decreases the velocity dispersion of the spheroid, bringing them much more into agreement with those measured in the other regions.

Several trends characterize the subregion panels. First, the fraction of stars in the spheroid (parameterized by ratio of the areas of the red spheroid and blue/green disk curves) increases with increasing distance from the major axis. Second, the mean velocity of the cold component decreases with increasing distance from the major axis. Reading off the velocity axis of the parameter-space view in the bottom right panel, we see that the mean velocity of the cold component moves from about -90 to -170 km s⁻¹.

The minimum velocity is less negative than that in region NE1 because region NE2 subtends a smaller range in PA.

4.6.2 NE3 Region

Nine stars are excluded from region NE3 as possible tidal debris contaminants. Exclusion slightly improves the reduced chi-squared of the fit. It also slightly increases the mean velocity of the hot component and decreases the velocity dispersion, bringing them more into agreement with those measured in the other regions.

Because this region is centered farthest from the center of M31, it subtends the smallest range in PA of the three NE regions, and so the variation in cold component mean velocity between subregions is small. However, the parameter-space view in Figure 4.14 shows that the cold component kinematical parameters are very well defined.

4.6.3 SE Region

We exclude 36 stars from the tails of the velocity distribution in the SE minor axis region. While this does not skew the results of the MCMC fits in one direction or another, it does serve to increase the uncertainty in the final kinematical parameters; this effect can be seen in the large scatter in the walker positions in the parameter-space view in Figure 4.15.

The minor axis is a saddle point in the velocity field of M31. Here, the observed rotation velocity changes from less negative than -300 km s^{-1} on the north side to more negative than -300 km s^{-1} on the south side. The mean velocities of the cold component

reflect this transition. The parameter-space view in the final panel of Figure 4.15 shows that in subregion SE_1 (blue), which lies parallel to and just north of the minor axis, we measure a mean velocity slightly less negative than the systemic velocity of M31. In subregions SE_2 (turquoise) and SE_3 (green), situated south and north of the minor axis, respectively, we measure mean velocities more negative and less negative than -300 km s^{-1} . Finally, we notice that the fraction of spheroid members is smallest in subregion SE_1 , which is centered farthest from the center of the galaxy than the other two subregions.

4.6.4 SSW Region

Eleven stars fall into the velocity range excluded due to stream contamination in the SSW region. The exclusion does not have a significant effect on the mean velocity or dispersion of the hot component. The chi-squared probability of a good fit is lower in the SSW region than in the others, but exclusion of the possible stream stars improves the probability by a factor of about 15.

Subregion SSW_4 completely contains the galaxy M32, which is treated as a second cold component. (In a previous trial run, we allowed for the possibility of a second cold component in each of the five subregions, but the best-fit fraction of M32 stars in subregions $SSW_{1-3,5}$ was zero.) The parameter-space view in Figure 4.16 shows that the mean velocity and dispersion of M32 are very well constrained around -200 and 27 km s^{-1} , respectively.

As in the NE regions, the fraction of stars in the spheroid is higher in the subregions farther from the major axis.

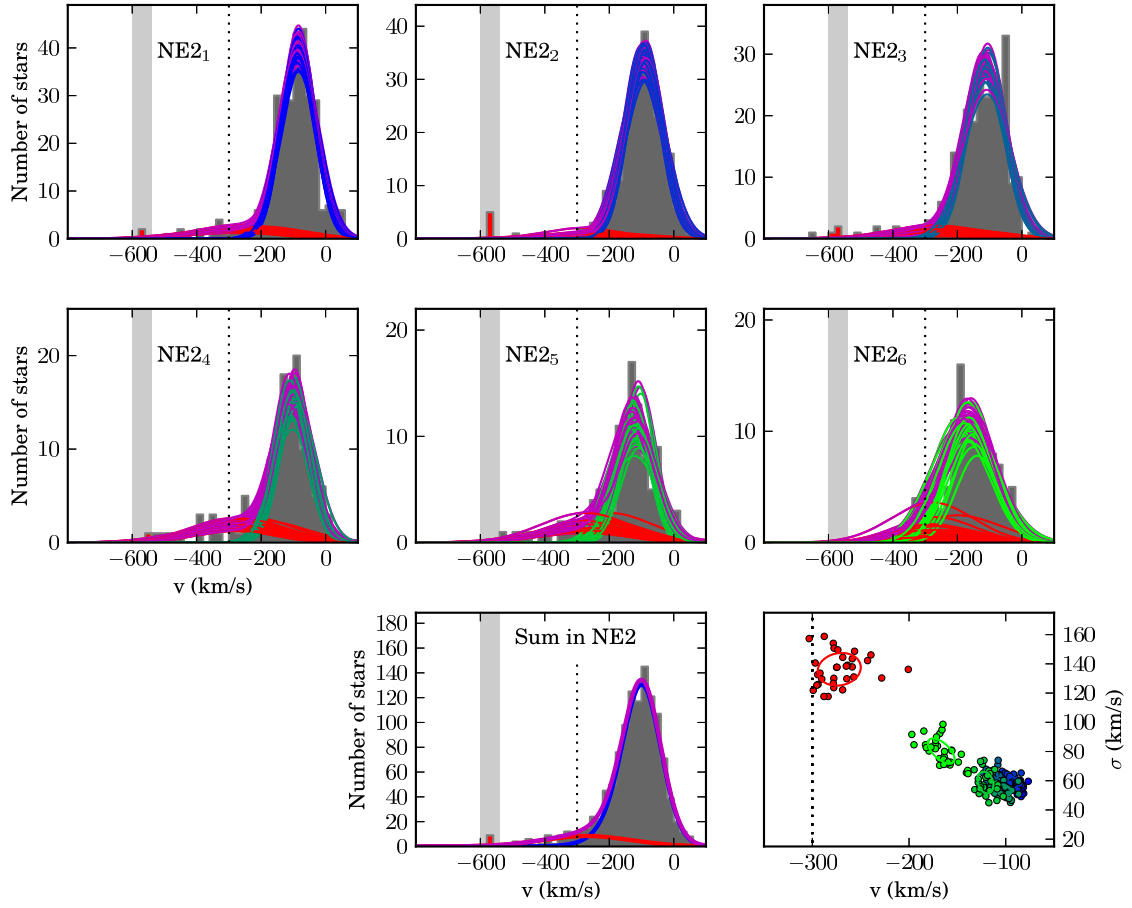


Figure 4.13: Same as Figure 4.8, but for the NE2 region.




Figure 4.14: Same as Figure 4.8, but for the NE3 region.

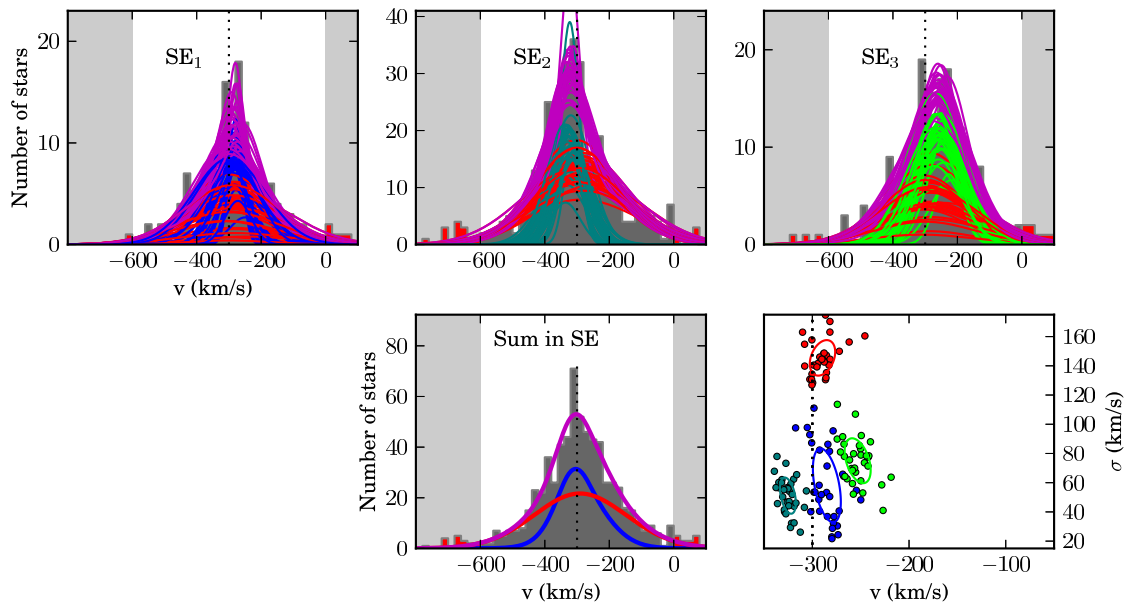


Figure 4.15: Same as Figure 4.8, but for the SE region. SE₁, SE₂ and SE₃ are the outer, inner south, and inner north subregions, respectively.




Figure 4.16: Same as Figure 4.8, but for the SSW region.




Figure 4.17: Kinematical parameters of the cold component in each subregion, with 1σ error bars, before exclusion of velocity ranges corresponding to GSS debris (black points) and after exclusion (colored points). Note that exclusion of tidal debris has a negligible effect on the cold component. These cold components correspond to M31’s stellar disk. The left panel shows the rotation velocity of each cold component (computed using Equation 4.3) vs. the mean r_{deproj} of each subregion. The three SE subregions have been excluded because they do not yield meaningful constraints on the rotation velocity. The right panel shows the velocity dispersion. The velocity and dispersion follow the trends expected for M31’s stellar disk, supporting the idea that that our minimal assumption of a locally cold disk adequately accounts for disk contamination.

Table 4.3: Disk Kinematical Parameters Corrected for Tidal

Debris

Subregion	$\xi(^{\circ})$	$\eta(^{\circ})$	r_{proj} (kpc)	r_{disk} (kpc)	$N_{\text{subregion}}$	$\bar{v} - \bar{v}_{M31}$ (km/s)	σ_v (km/s)	$N_{\text{cold}}/N_{\text{total}}$	χ^2 prob
NE1 ₁	0.284	0.363	6.311	6.35	220	+192.4 \pm 5.8	56.4 \pm 5.4	0.725	0.706
NE1 ₂	0.287	0.364	6.36	6.61	227	+186.8 \pm 6.1	64.2 \pm 4.4	0.806	0.753
NE1 ₃	0.286	0.359	6.33	7.05	236	+179.4 \pm 6.0	58.6 \pm 5.3	0.729	0.191
NE1 ₄	0.321	0.349	6.56	8.06	234	+184.5 \pm 8.9	63.8 \pm 6.7	0.602	0.220
NE1 ₅	0.362	0.328	6.74	9.22	164	+181.8 \pm 6.4	47.0 \pm 5.5	0.599	0.933
NE1 ₆	0.394	0.281	6.67	11.01	252	+137.0 \pm 6.0	60.2 \pm 6.5	0.700	0.862
NE1 ₇	0.422	0.180	6.30	14.64	282	+103.6 \pm 5.2	60.0 \pm 6.8	0.733	0.520
NE2 ₁	0.463	0.597	10.36	10.42	306	+213.9 \pm 3.7	56.1 \pm 2.8	0.892	0.767
NE2 ₂	0.459	0.594	10.30	10.71	264	+206.9 \pm 4.2	52.3 \pm 3.3	0.943	1.000
NE2 ₃	0.449	0.583	10.15	11.30	247	+188.4 \pm 4.7	62.5 \pm 4.0	0.893	0.005
NE2 ₄	0.481	0.562	10.25	12.48	135	+198.5 \pm 6.3	53.5 \pm 5.6	0.749	0.979
NE2 ₅	0.551	0.434	9.63	14.17	108	+179.0 \pm 8.7	58.3 \pm 6.6	0.743	0.987
NE2 ₆	0.568	0.366	9.26	16.17	120	+133.5 \pm 12.1	80.1 \pm 9.2	0.861	1.000
NE3 ₁	0.653	0.837	14.56	14.65	339	+206.6 \pm 8.2	47.7 \pm 6.0	0.770	1.000
NE3 ₂	0.639	0.805	14.12	14.68	280	+218.7 \pm 9.4	51.1 \pm 7.1	0.827	0.104
NE3 ₃	0.609	0.704	12.83	14.51	240	+206.2 \pm 7.9	54.6 \pm 9.2	0.820	0.999
SE ₁	0.426	-0.241	6.68	29.38	130	+16.0 \pm 14.0	52.8 \pm 26.1	0.448	0.424

Table 4.3: Disk Kinematical Parameters Corrected for Tidal

Debris

Subregion	$\xi(^{\circ})$	$\eta(^{\circ})$	r_{proj} (kpc)	r_{disk} (kpc)	$N_{\text{subregion}}$	$\bar{v} - \bar{v}_{M31}$ (km/s)	σ_v (km/s)	$N_{\text{cold}}/N_{\text{total}}$	χ^2 prob
SE ₂	0.193	-0.170	3.62	15.61	355	-25.5 ± 8.9	49.8 ± 13.2	0.345	0.340
SE ₃	0.281	-0.036	3.95	15.00	199	$+46.3 \pm 14.4$	77.9 ± 16.2	0.567	0.912
SSW ₁	-0.262	-0.358	6.07	6.38	254	-200.9 ± 10.1	74.2 ± 8.2	0.647	0.489
SSW ₂	-0.219	-0.425	6.54	8.34	149	-183.7 ± 9.8	74.5 ± 8.0	0.779	0.998
SSW ₃	-0.111	-0.371	5.31	9.84	175	-145.3 ± 12.5	61.6 ± 10.8	0.449	0.433
SSW ₄	-0.011	-0.404	5.54	15.18	608	-88.5 ± 5.7	33.6 ± 6.4	0.191	0.002
SSW ₅	-0.098	-0.366	5.19	18.58	127	-54.7 ± 21.6	61.8 ± 21.5	0.340	5.34×10^{-9}
SSW ₄ (M32) ^a	-0.011	-0.404	5.54	15.18	608	$+102.2 \pm 2.7$	26.4 ± 3.2	0.270	0.002

Chapter 5

Structural Decomposition

5.1 INTRODUCTION

Spiral galaxies like the Milky Way (MW) and Andromeda (M31) are composed of several structural subcomponents, each with its own formation history. Tracing the evolution of such a galaxy, then, depends on characterization of the individual subcomponents. This decomposition is especially difficult in the inner regions, where the cospatial bulge, disk, and halo complicate characterization of any single subcomponent.

The most commonly used way to characterize the stellar component of a large spiral galaxy is via surface brightness (SB) decomposition. SB profiles of disks tend to fall off exponentially, whereas bulges follow more general Sérsic profiles (Sersic, 1968). Widely used SB decomposition codes such as GIM2D (Simard et al., 2002) and GALFIT (Peng et al., 2002) fit a sum of a Sérsic spheroid and an exponential disk to the SB distribution of a galaxy to estimate the relative contributions of the components. Unfortunately, such fitting is plagued by degeneracies that arise because the differ-

ent subcomponents are cospatial and because the procedure generally relies on ad hoc fitting formulae that do not necessarily separate the galaxy into dynamically distinct subcomponents (e.g., Abadi & Navarro, 2003).

In this paper, we present a new technique for decomposing the stellar component of M31 into distinct bulge, disk, and halo subcomponents. This technique can be applied to any galaxy close enough to measure radial velocities and a luminosity function of stars down to about 1.5 magnitudes fainter than the tip of the red giant branch (TRGB). In addition to fitting a toy model (exponential disk, powerlaw halo, and Sérsic bulge) to the SB distribution, we attempt to break the aforementioned degeneracies by including an additional constraint: the fraction of stars that belong to the disk (“disk fractions”), as measured from stellar kinematics of individual red giant branch (RGB) stars. We perform the decomposition using Bayesian techniques so that we can identify covariances between model parameters, quantifying lingering degeneracies.

Two large-scale, ongoing resolved stellar population surveys of M31 make it an ideal galaxy in which to develop and test our decomposition technique. The Spectroscopic and Panchromatic Landscape of Andromeda’s Stellar Halo (SPLASH) survey has used the Keck/DEIMOS multiobject spectrograph to measure radial velocities of over 15,000 stars in M31 (Guhathakurta et al., 1911, 2006), including over 10,000 in the crowded inner 20 kpc (Dorman et al., 2012; Gilbert et al., 2012; Howley et al., 2013). Meanwhile, the Panchromatic Hubble Andromeda Treasury (PHAT) survey, a five-year Hubble Space Telescope/Multi-Cycle Treasury (HST/MCT) program, has so far imaged over 10^8 stars in six filters in the UV, optical, and near-IR in the same quadrant of the

disk most densely sampled by SPLASH (Dalcanton et al., 2012).

The decomposition technique presented here builds directly on the work of Courteau et al. (2011) and Dorman et al. (2012). The former performed SB-only decompositions of M31 using two different techniques: Markov chain Monte Carlo (MCMC) sampling and a nonlinear least-squares (NLLS) fitting method. In each case, the authors fit a sum of Sérsic bulge, exponential disk, and powerlaw halo to I-band SB profiles of the galaxy. The profiles were a composite of major- and minor- axis cuts of an *I*-band image of the central ~ 20 kpc of M31 (Choi et al., 2002) plus minor-axis profiles measured from RGB star counts (Pritchett, C.J., van den Bergh, 1994; Irwin et al., 2005; Gilbert et al., 2009). They found that the best-fit bulge, disk, and halo structural parameters depended significantly on the decomposition method, data points used (major vs. minor axis), and even on the binning of the data. Despite the uncertain results, their model was still relatively restrictive, not allowing for a flattened halo or for variations from the canonical values in the ellipticities or position angles of the bulge and disk.

In Dorman et al. (2012), we performed a *kinematical* decomposition of M31's dynamically cold disk (without distinguishing between thin and thick components) and dynamically hot spheroid (without distinguishing between bulge or halo components) using radial velocity measurements of ~ 6000 bright ($20 < I < 22$) RGB stars in the inner parts of the galaxy ($5 < R_{\text{proj}} < 20$ kpc; the region dominated by the disk in optical and UV images). In each of 24 small spatial subregions, we measured the fraction of stars that belonged to the hot component. This fraction is nonzero ($> 10\%$) everywhere in the survey region. The origin of this dynamically hot population is unclear a priori:

is it an outward extension of the central bulge or the innermost reaches of the extended halo? On one hand, the bulge appears too small in SB decompositions to contribute stars past a few kpc (Courteau et al., 2011); on the other hand, CMD analyses of fields between 11 and 45 kpc on the minor axis reveal populations with young/intermediate ages and intermediate metallicities (Brown et al., 2003; Richardson et al., 2008) — quite unlike the old, metal-poor halo field stars in the MW (e.g., Carollo et al., 2007; Kalirai, 2012).

In this paper, we use kinematically derived disk fractions measured as in Dorman et al. (2012) as constraints in a SB decomposition of M31’s *I*-band SB profile. Simultaneously fitting to the total SB profile of the galaxy and the fraction contributed by the disk may help to reduce degeneracies in the best fit parameters — and possibly understand the structural association and origin of the dynamically hot population identified in Dorman et al. (2012). With more constraints, we can relax some of the assumptions made in Courteau et al. (2011), fitting for the ellipticities and position angle of the bulge, disk, and halo.

A primary challenge in our work is that our disk fraction measurements come from star count data, whereas we model the contributions to the integrated surface brightness. Thus, we must convert a disk fraction in star counts to a disk fraction in SB units. In general, this conversion factor is such that the kinematical survey undersamples the disk: the SPLASH target selection function happens to peak at magnitudes where the spheroid and disk both contribute stars, and to fall off at brighter magnitudes where only the disk contributes light. In addition, while it is easily recoverable, the selection

function is somewhat arbitrary, varying across the galaxy. In order to quantify how fairly it samples the three subcomponents, we model and fit for the disk, bulge, and halo luminosity functions (LFs).

The model presented here incorporates stellar population, kinematics, and SB data for the first time. Including complexities like the boxy bulge (Beaton et al., 2007) or separate thin, thick, and extended disks (Collins et al., 2011; Ibata et al., 2005) are beyond the scope of the paper, which nonetheless our analysis represents a significant improvement over using only SB profile fits to characterize a galaxy's physical components.

This paper is organized as follows. First, in § 5.2, we outline the analysis procedures that will be described in detail in § 6.2-5.4. In § 6.2 we summarize the kinematical, integrated-light, and resolved stellar photometric data sets used as constraints. In § 5.4 we present the mathematical formalism used to find the probability distribution functions (PDFs) of each of the model parameters. In § 5.5 we discuss the probability distributions of our parameters, comparing them to previous measurements of M31's structural parameters. In § 5.6, we show that the dynamically hot population in our kinematical sample is more closely associated with the extended halo than the central bulge, and propose that some fraction of that population originated in the disk. Finally, we summarize in § 5.7.

5.2 Overview of Analysis Procedure

Simultaneously fitting even a simple toy model to the observed SB map, disk fractions, and LF is relatively complicated. We must process three different data sets, develop a multiparameter model, combine the model and data into a likelihood function, and sample that likelihood to obtain a probability distribution function (PDF) for each parameter. Each of these steps will be described in detail in § 6.2-5.4, but first we will present a brief road map to those sections.

The flow chart in Figure 5.1 illustrates the path from data processing to model parameter PDFs. We start with three sets of observational constraints: SB profiles, the kinematically-derived disk fraction probability distribution in each of the spatial subregions, and the $F814W$ (I) LF in each subregion. The processing of these three data sets is described in § 6.2.1-5.3.3.

As the left-hand side of Figure 5.1 shows, each data set has an accompanying toy model. We model the SB as the sum of three profiles: a Sérsic bulge, an exponential disk, and a cored power-law halo (§ 5.4.1). Similarly, we model the luminosity function as a sum of three broken powerlaws, one each for the disk, bulge, and halo, in the magnitude range $m_{F814W} = [20, 22]$ (§ 5.4.1). The model for the disk fraction comes from a combination of the previous two: the disk fraction in flux units at a given location is the ratio of the disk to total SB at that location; however, we must convert this integrated-light disk fraction to star counts using the disk LF (§ 5.4.1).

Given a certain vector w in parameter space, we compute the probability $P_{\text{SB}}(w)$ that the SB model represents the observed SB map; similarly, we define and

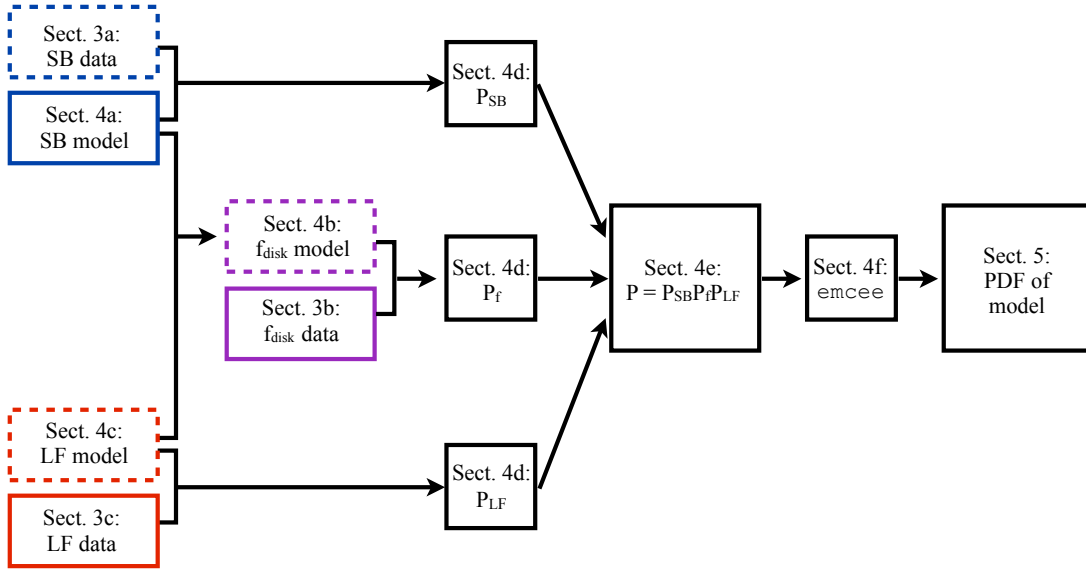


Figure 5.1: Flow chart illustrating the analysis procedure. Each box is labeled with the section of the text in which that step is described. See text of § 5.2 for description.

compute $P_f(w)$ and $P_{LF}(w)$ as the probabilities that the model matches the kinematical and LF data sets, respectively. The total probability of that vector w is then $P(w) = P_{SB}P_fP_{LF}$. § 5.4.2 gives the equation for P .

We then sample this probability distribution P using the Markov chain Monte Carlo (MCMC) sampler `emcee` (Foreman-Mackey et al., 2015), as described in § 5.4.3. The sampler yields a PDF of each model parameter. In § 5.5, we discuss the median values and confidence intervals of the PDFs in the context of previous measurements.

5.3 Observational Constraints

The integrated-light, resolved stellar photometric, and spectroscopic data sets used in this work are described in more detail in Courteau et al. (2011), Gilbert et al. (2012), Dalcanton et al. (2012), and Dorman et al. (2012). In this section we briefly

recap the salient information.

5.3.1 Surface Brightness

The SB map is derived from two classes of data: major- and minor-axis wedge cuts from a mosaic image, and RGB star counts. All of the profiles are in the I band to minimize the effects of inhomogeneities such as dust lanes or the UV-bright star-forming rings. The SB data used in this work are identical to those used in Courteau et al. (2011), with the addition of new halo fields out to a galactocentric projected radius of $R = 225$ kpc first presented in Gilbert et al. (2012).

Integrated Light Profiles

We use the major- and minor-axis SB profiles of M31 that Courteau et al. (2011) constructed from an I -band CCD mosaic image of M31 taken with the Kitt Peak National Observatory Burrell Schmidt telescope and presented in Choi et al. (2002). Rather than relying on azimuthally averaged isophotal fitting, which smears together components with very different ellipticities such as bulge and disk, they cut wedges from the major and minor axes of the galaxy. They measured the median SB in each of 186 bins on the major axis and 152 on the minor axis. The sizes of the bins increased with galactocentric radius. Bright stars were removed via an iterative sigma-clipping process. The quoted error bar in each bin is the RMS deviation about the median value of the pixels in that bin. The locations of these bins on the sky are marked as teal circles in Figure 5.2 and in Figure 5.3, the latter zooming in on the the portion of the galaxy dominated by the bright disk and bulge.

Because the CCD pixel size is $2''$ and the telescope was not well focused during some of the observations, some of the images that went into the final mosaic image are out of focus at the $2'' - 5''$ level. (See §2 of Choi et al. (2002).) To avoid the effects of finite spatial resolution, we exclude from our analysis the two bins within projected radius $R = 10''$ (40 pc) of the center of M31.

Star Counts

In the outer regions of the galaxy, where crowding is less severe, the SB profile can be measured, up to a normalization factor, from RGB star counts. Like Courteau et al. (2011), we have combined the extended star counts of the M31 stellar halo by Pritchet, C.J., van den Bergh (1994), Irwin et al. (2005), Gilbert et al. (2009), and Gilbert et al. (2012). These profiles cover the range $20 \text{ kpc} < R < 225 \text{ kpc}$.

The Pritchet, C.J., van den Bergh (1994) data set — which combines digital star counts to measure the SB in three fields along the minor axis of the galaxy — is identical to that used in Courteau et al. (2011). These fields are marked in yellow squares in Figures 5.2 and 5.3.

The Irwin et al. (2005) data set is also identical to that used in Courteau et al. (2011). Irwin et al. (2005) combined the data from Pritchet, C.J., van den Bergh (1994) with faint RGB star counts to trace the minor-axis stellar distribution out to a projected radius of 55 kpc. They used star counts from the Isaac Newton Telescope, exposing typically 800 – 1000 s per field in the Gunn i band. These fields are shown as blue triangles in Figures 5.2 and 5.3.

The final SB data set, presented in Gilbert et al. (2009) and Gilbert et al.

(2012), is based on RGB star counts from a Keck/DEIMOS spectroscopic survey of 38 fields between 9 – 225 kpc in projected radius [rather than just the 12 fields used by (Courteau et al., 2011)]. Members of M31’s smooth halo were identified and distinguished from foreground MW dwarf star contaminants and substructure in the M31 halo using a combination of photometric and spectroscopic diagnostics (Guhathakurta et al., 2006; Gilbert et al., 2006). To estimate M31’s stellar surface density as a function of radius, the observed ratio of M31 red giants to MW dwarf stars was multiplied by the surface density of MW dwarf stars predicted by the Besançon Galactic star-count model (e.g., Robin et al., 2003, 2004). The surface density was converted to SB units by scaling to match the minor axis profile from Choi et al. (2002). These fields are shown as magenta stars in Figures 5.2 and 5.3.

In summary, we have 637 SB measurements μ_{obs} from 4 data sets. Though the SB data are given in magnitudes, we perform the SB fits using flux units; the conversion is given through the usual formula:

$$\Sigma_{\text{obs}}(R, \Delta PA) = 10^{-(\mu_{\text{obs}} + zp)/2.5} \quad (5.1)$$

where Σ_{obs} is in flux units, μ_{obs} is in magnitudes, and zp is the zeropoint. To ensure consistency between datasets, we anchor the zeropoints of the Pritchett, C.J., van den Bergh (1994); Gilbert et al. (2012), and Irwin et al. (2005) datasets to that of the Choi et al. (2002) image by first fitting to the composite SB map *only* (i.e., not including constraints on the disk fraction), leaving the zeropoint as a free parameter with a flat

prior between ± 2 magnitude of the Choi et al. (2002) value. For each dataset, the median (best fit) zeropoint of the resulting posterior distribution is adopted for the rest of the analysis.

The composite, zeropoint-adjusted major- and minor-axis SB profiles are shown in the two panels of Figure 5.4.

5.3.2 Velocity distribution

The constraints on the fraction of stars dynamically associated with the disk are obtained via a kinematical decomposition of the stellar velocity distribution into dynamically cold (disk) and hot (spheroid, or combined bulge and halo) components, described in detail in Dorman et al. (2012).

Kinematical data

We obtain spectra of 5257 stars in the inner 20 kpc of M31 using the DEIMOS multi-object spectrograph on the Keck II telescope as part of the SPLASH survey (Guhathakurta et al., 1911, 2006; Gilbert et al., 2009; Dorman et al., 2012; Howley et al., 2013). The positions of the 29 DEIMOS slitmasks are shown as red rectangles overlaid on a GALEX image of M31 in Figure 5.3. The targets on 19 of these masks contain targets selected from an i' -band CFHT/MegaCam image, while the targets on 5 masks were selected exclusively from the PHAT catalog and the targets on the remaining 5 were selected using a combination of PHAT and CFHT data. In each case, the targets are selected based on a complex set of factors: apparent degree of isolation in the source catalog, magnitude bright enough to yield a high-S/N spectrum but not so

bright as to be a likely foreground MW dwarf, and position on the sky so as to optimize the use of the slitmask. Hence, the luminosity function of DEIMOS-observed stars is different in each slitmask, but in general is nonzero between $20 < i' < 22$ and peaks somewhere around $i' \sim 21$.

Each spectrum is collapsed in the spatial direction and cross-correlated against a suite of template rest-frame spectra to measure its radial velocity. The velocity measurement is then quality-checked by eye, and only robust velocities (those based on at least two strong spectral features) are included in the analysis.

We perform a number of cuts on this preliminary kinematical dataset. First, we retain only the 3247 (63% of) targets that are also present in the existing PHAT catalog. (The other targets lie in regions on the sky not sampled by the PHAT survey at the time of writing.) Second, we cut out the 25% of the remaining objects that are located in regions with $A_v > 1.0$ mag. The low extinction regions were identified by modeling the NIR CMD as a sum of a foreground unreddened RGB and a background RGB, reddened by a log-normal distribution of dust reddening, in 10 arcsecond bins. This modeling produced a map of the median A_V across the disk, independent of the fraction of reddened stars along a given line of sight. Details can be found in Dalcanton et al. (2013, in prep), and an overview of the technique can be found in Dalcanton et al. (2012). All of the 2443 remaining stars are redder than $m_{F475W} - m_{F814W} = 1.5$, so are likely to be RGB stars rather than bright main sequence (MS) objects. In the following text, we describe how we use these 2443 stars to measure disk fractions in each of 14 spatial subregions.

Review of Disk Fraction Measurements

In Dorman et al. (2012), the SPLASH and PHAT survey region was divided into four spatial “regions”: three straddling the NE major axis and one along the SE minor axis. The boundaries of the regions (lines of constant R_{proj} or position angle P.A.) were chosen such that each region contained enough SPLASH targets to constrain the velocity distribution of the subdominant spheroid component. Each region was then subdivided along lines of constant P.A. into multiple “subregions,” each large enough to contain > 100 stars but small enough that the rotation velocity of the disk component does not vary substantially. We measured the mean velocity v , velocity dispersion σ_v , and fraction f of the disk and spheroid components in each of the resulting 24 subregions by fitting a sum of two Gaussians (representing a dynamically cold disk and dynamically warmer spheroid) to the velocity distribution in each subregion.

Because in the current paper we use only the subset of spectroscopic targets also identified in the PHAT survey, we must combine some subregions from Dorman et al. (2012) to have sufficient number statistics to reliably separate the disk and spheroid. The 14 subregions we use are shown in the right-hand panel of Figure 5.5. The subregions are named using the same formalism as in Dorman et al. (2012). The survey area is still divided into four regions: SE (along the SE minor axis), and NE1, NE2, and NE3 (along the NE major axis, in order of increasing projected radius). The subregions within each region are identified with subscripts that increase with distance from the major axis: NE1₁, NE1₂, ..., NE1₅; NE2₁, ..., NE2₄, and NE3₁, NE3₂, NE3₃. In the SE region, the inner south, and inner north subregions are named SE₂ and SE₃,

respectively. (Subregion SE_1 is not used because it does not overlap with the PHAT survey region.)

We fit a sum of two Gaussians (representing a dynamically cold disk and dynamically warmer spheroid) to the velocity distribution in each subregion with the MCMC sampler `emcee` (Foreman-Mackey et al., 2015, see the Appendix for more information). We allow the disk velocity distribution to change from subregion to subregion, but require that the mean velocity and dispersion of the spheroid remain constant between subregions within a region. The decomposition makes almost no assumptions about the nature of the disk, except that its velocity distribution is symmetric and locally colder than that of the spheroid.

The latter is perfectly reasonable; however, one could imagine that asymmetric drift in a stellar disk could skew the distribution of circular velocities towards lower speeds and upweighting the spheroid in the decomposition. We found that asymmetric drift does *not* in fact have a significant effect on the kinematically derived disk fractions. We tested the effects of asymmetric drift using the velocity distribution presented in Schönrich & Binney (2012) to measure disk fractions in the major axis subregions, where the effect of asymmetric drift should be most significant. In near-major-axis subregions ($NE1_{1-3}$, $NE2_{1-2}$, $NE3_{1-3}$), the disk fractions computed by assuming symmetric and asymmetric velocity distributions are consistent to within 1σ . In subregions that are closer to the minor axis, it is difficult to fit the asymmetric function to the velocity distribution because the line-of-sight velocity measurements do not well constrain the rotation velocity. However, for the same reason, we expect the effects of asymmetric

drift on the measured disk fraction to be decrease towards the minor axis. The line-of-sight velocity of a star on the minor axis is independent of the circular velocity of that star, and so a slight change in the circular velocity distribution should not affect the line-of-sight distribution. Therefore, we can safely ignore the effect of asymmetric drift and use a Gaussian disk velocity distribution without affecting the disk fraction measurements.

Figure 5.6 illustrates the velocity decomposition in each subregion within a representative region, NE2, which is centered along the northeast major axis at a projected radius of about 10 kpc. Note that in each subregion within the NE2 region, some nonzero fraction of stars belong to the dynamically hot component. Nearly half of these spheroid stars are *counterrotating* relative to the disk (that is, they have radial velocities less than $v_{M31} = -300$ km/s), and thus must be spheroid stars regardless of the detailed shape of the disk velocity distribution. In § 5.6, we will show that these stars are more closely associated with the extended halo than the central bulge, and that some of them likely originated in the disk.

The cold fraction in each subregion s is the ratio of the areas of the cold Gaussian and the total velocity distribution. At the end of the 10,000-step MCMC chain, the disk fractions measured from the 32 walkers over the final 2000 steps of the chain are compiled and binned by 0.5% into a normalized PDF $p_s(f_d)$.

Because it will be more useful to have an analytic form for $p_s(f_d)$, we fit a skew-normal function ϕ_s to $p_s(f_d)$ in each subregion s using a Levenberg-Marquardt minimization. The form of ϕ_s is

$$\phi_s(f) = A_s e^{-(z_s/\sqrt{2})^2} \left(1 + \operatorname{erf} \left(\frac{\tau_s z_s}{\sqrt{2}} \right) \right) \quad (5.2)$$

where

$$z_s = (f_s - \mu_s)/\sigma_s. \quad (5.3)$$

The best-fit skew-normal parameters μ_s, σ_s, τ_s , along with the medians and standard deviations of the p_s distributions, are given in Table 5.1. Figure 5.7 shows an example fit to $p_s(f_d)$ in representative subregion NE2₄. The PDFs are very well approximated by a skew-normal function: the χ^2 p -statistic is 1.000 in every subregion. The uncertainties in the disk fraction measurements in each subregion are parameterized by the width of the skew-normal function; the entire function (not just an error bar) will be used as a prior on the disk fraction in the decomposition presented. Note that because the subregions are our smallest resolution elements in our kinematical analysis, we cannot map the variation in kinematically-derived disk fraction *within* a subregion. However, in subregions where the true disk fraction varies significantly, the PDF is broad — in other words, systematic uncertainties in the measured disk fractions derived from finite spatial binning are incorporated into the error bars.

Table 5.1: Skew-Normal Parameters of Disk Fraction Probability Distribution $p_s(f_d)$

Subregion	μ	σ	τ	A	Median	Std. dev.
NE1 ₁	0.75	0.11	-1.64	1.83	0.68	0.08
NE1 ₂	0.82	0.11	-1.73	1.86	0.75	0.08
NE1 ₃	0.56	0.11	-0.78	1.88	0.51	0.09
NE1 ₄	0.83	0.22	-1.65	0.95	0.70	0.15
NE1 ₅	0.77	0.10	-1.11	1.97	0.72	0.08
NE2 ₁	0.96	0.04	-2.43	4.61	0.93	0.03
NE2 ₂	0.86	0.08	-2.04	2.4	0.81	0.06
NE2 ₃	0.87	0.08	-2.00	2.57	0.82	0.05
NE2 ₄	0.84	0.15	-2.53	1.38	0.75	0.10
NE3 ₁	0.83	-0.04	1.47	5.12	0.81	0.03
NE3 ₂	0.86	0.05	-1.35	4.04	0.84	0.04
NE3 ₃	0.84	0.06	-1.34	3.50	0.81	0.04
SE ₂	0.30	0.25	-0.07	4.53	0.33	0.21
SE ₃	-2.68	3.16	-1.26	19.30	0.37	0.29

5.3.3 Luminosity Function

The bright end of M31's LF is the crucial link between the kinematical and SB data. The kinematics measure the fraction of stars, as sampled by the SPLASH survey, that belong to the disk in each small subregion, while a SB decomposition yields the fraction of integrated light contributed by the disk. To convert between these units, we must know how fairly the SPLASH survey samples the bulge, disk, and halo LFs. We

therefore fit three model LFs to the observed PHAT LF in each subregion. This section describes the observed LF; the model LFs are discussed in § 5.4.1.

We measure the PHAT LF in the magnitude range sampled by SPLASH ($20 < m_{F814W} < 22$). While the PHAT survey is crowding-limited and thus incomplete at faint magnitudes, within our radial range, $6 < R < 20$ kpc in deprojected radius, it is 100% complete down to $m_{F814W} = 22$ for colors $m_{F475W} - m_{F814W} < 4.5$. (We have tried a decomposition excluding the few percent of stars redder than this cutoff, but the qualitative results are not affected.) We clean the sample in the same way as for the SPLASH data set, using only stars redder than $m_{F475W} - m_{F814W} = 1.5$ and in pixels with $A_v < 1.0$ mag. We bin these stars by 0.1 mag in the range $20 < m_{F814W} < 22$. Figure 5.5 shows the PHAT stars in this magnitude range before (left) and after (right) excluding stars in extincted regions.

In each subregion s , we measure two normalized luminosity functions in the magnitude range sampled by the SPLASH spectroscopic survey ($I = [20, 22]$). $\mathcal{L}_{\text{PHAT},s}$ contains stars from the PHAT catalog that fall into subregion s . $\mathcal{L}_{\text{SPLASH},s}$ contains stars found in both the SPLASH and PHAT surveys in subregion s . Each has units of number/arcsec²/mag. The uncertainty on each observed LF is simply the Poisson counting uncertainty on each bin \sqrt{N} , where N is the number of stars in the subregion that fall into that magnitude bin.

We display the PHAT and SPLASH luminosity functions from two representative subregions in Figure 5.8. The shape of the PHAT LF is similar in all subregions: the slope is shallow at the faint end, gradually steepens brightward of $m_{F814W} \sim 21$,

and finally flattens again brightward of $m_{F814W} \sim 20.5$. These changes in slope can be explained by the position of the TRGB in M31 and the presence of intermediate-age stars brighter than the TRGB. The magnitude of the TRGB at the distance of M31 is $m_I = 20.35$ near $[M/H] = -1$, changing by only about 0.1 dex between $-2.5 < [M/H] < -0.57$ (Salaris & Cassisi, 1997). However, the metallicity distribution of M31’s disk and bulge extend to supersolar values (Sarajedini, 2005; Brown et al., 2006). At such high metallicities, line blanketing in the red can push the TRGB even fainter than $m_I = 21$. The wide shoulder seen in the PHAT LF between the TRGB and $m_{F814W} \sim 21$ is a signature of a broad metallicity distribution with a broad range of TRGB magnitudes. The shallow slope brightward of the metal-poor TRGB at $m_{F814W} = 20.35$ suggests a high fraction of young or intermediate-age populations with bright AGB stars, as described by Mendez et al. (2002).

The SPLASH luminosity function varies from subregion to subregion, but usually peaks in the range $20.5 < m_{F814W} < 21.5$ due to the design of the spectroscopic survey. (Targets in this magnitude range were bright enough to yield high-quality spectra, but faint enough to likely be M31 members rather than foreground MW dwarfs, so they were given highest priority in the target selection process.) Because there are so many fewer stars in the SPLASH sample than in PHAT, for display purposes we have scaled the SPLASH LF by a factor of 100, and scaled the Poisson errors accordingly.

We can measure the spectroscopic selection function $S_s(M)$ in each subregion s :

$$S_s(M) = \mathcal{L}_{\text{SPLASH},s} / \mathcal{L}_{\text{PHAT},s} \quad (5.4)$$

$S_s(M)$ is a purely empirical measure of the complicated target selection criteria of the spectroscopic survey. Examples of $S_s(M)$ in two representative subregions are shown in the bottom panels of Figure 5.8. The variation in the selection function with subregion and with magnitude means that the SPLASH survey may over- or under-sample the spheroid relative to the disk. We correct for this effect, as described later in the following section.

5.4 Analysis

Our goal is to find the most probable combination of structural and LF parameters for a Sérsic bulge, exponential disk, and power-law halo given the three sets of constraints described above: surface brightness, PHAT LF in each of 14 subregions, and kinematically-derived disk fraction in each subregion. As shown in Figure 5.1, we build a toy model to represent the SB, disk fraction and LF across the galaxy. In § 6.5 we describe our model: a 2D surface brightness profile and LF for each structural subcomponent. In § 5.4.2 we show the likelihood function to be sampled. Finally, in § 5.4.3, we describe the Markov chain Monte Carlo sampler `emcee` (Foreman-Mackey et al., 2015) that we use to sample the parameter space.

5.4.1 Model

The model parameters are listed in the first column of Table 5.2. All parameters have flat priors within the range specified in the fourth column of the table. The only fixed parameter is the bright-end slope ($\log N/\text{mag} = 500/\text{mag}$) of the halo LF, chosen for reasons described in §5.4.1.

Surface Brightness Profiles

As discussed in the introduction, we choose simple, standard SB profiles for three components: the bulge, disk, and halo. Each profile is given in terms of position on the sky (R, PA) , where R is the projected radius and PA is the position angle measured east of north. Hence, we can fit to an entire SB map, rather than only to data points that happen to fall on the major or minor axis. We assume that the three model components have the same major axis position angle pa .

Bulge For the bulge, we assume a generalized Sérsic profile with Sérsic index n_b , half-light radius R_b and intensity I_b at R_b :

$$\Sigma_b(R_{\text{eff},b}) = I_b \exp \left\{ -A_{n_b} \left[\left(\frac{R_{\text{eff},b}}{R_b} \right)^{1/n_b} - 1 \right] \right\} \quad (5.5)$$

$$R_{\text{eff},b}(R, \Delta PA, \epsilon_b) = R \sqrt{\cos^2 \Delta PA + \left(\frac{\sin \Delta PA}{1 - \epsilon_b} \right)^2} \quad (5.6)$$

$$\Delta PA = PA - pa \quad (5.7)$$

where $A_{n_b} = 1.9992n_b - 0.3271$ (?). The formula is given in terms of the effective major axis (deprojected) coordinate $R_{\text{eff,b}}$, which is a function of the major axis position angle pa and bulge ellipticity ϵ_b . With $n_b = 1$ or $n_b = 4$, the profile reduces to the exponential or de Vaucouleurs profile, respectively.

Disk We assume an exponential SB profile for the disk:

$$\Sigma_d(R_{\text{eff,d}}) = I_d \exp(-R_{\text{eff,d}}/R_d), \quad (5.8)$$

where I_d is the disk surface brightness at the galactic center and R_d is the scale length in the plane of the disk. The formula is given in terms of deprojected radius $R_{\text{eff,d}}$, computed as before in terms of $(R, \Delta PA)$ and disk ellipticity ϵ_d .

Halo Finally, we assume a 2D cored power-law halo SB profile. Though it is also possible to model a halo as a Sérsic function, we adopt a power-law because Courteau et al. (2011) demonstrate that such a model is a more reasonable description of the M31 halo. The halo surface brightness is then given as

$$\Sigma_h(R_{\text{eff,h}}) = \frac{I_h}{(1 + (R_{\text{eff,h}}/R_h)^2)^{\alpha/2}} \quad (5.9)$$

where I_h is the halo surface brightness at the galactic center, R_h is the radius of the core, and the effective major axis coordinate $R_{\text{eff,h}}$ is defined as before in terms of coordinates $(R, \Delta PA)$ and halo ellipticity ϵ_h .

We define

$$\Sigma_T(R, \Delta PA) = \Sigma_b + \Sigma_d + \Sigma_h \quad (5.10)$$

to be the total surface brightness of the model at coordinates $(R, \Delta PA)$.

For each profile, we fit for the central intensity in units of mag arcsec⁻² rather than counts arcsec⁻². So our model parameters describing central magnitudes are μ_b , μ_d , μ_h rather than I_b , I_d , I_h . The three central magnitude parameters μ_k are defined as follows:

$$\mu_b = -2.5 \log_{10} I_b + 25.6 \quad (5.11)$$

$$\mu_d = -2.5 \log_{10} I_d + 25.6 \quad (5.12)$$

$$\mu_h = -2.5 \log_{10} I_h + 25.6 \quad (5.13)$$

Here the zeropoint 25.6 is chosen to match that of the Choi et al. (2002) SB data.

Because the fractional photometric uncertainties e_m on the SB measurements are typically very small, we introduce an uncertainty parameter ϵ_{SB} into the model to allow for additional photometric uncertainty as well as departures from the assumed

functional form of the SB profile. ϵ_{SB} is allowed to have a unique value for each SB dataset j , but may not vary between points i within a given data set. The total fractional uncertainty of a given SB measurement Σ_{ij} is then the quadrature sum of the Poisson uncertainty and the error parameter:

$$\epsilon_{ij} = \Sigma_{ij} \sqrt{e_{m,ij}^2 + \epsilon_{\text{SB},j}^2} \quad (5.14)$$

Disk Fractions

We have measured the disk fraction distribution $\phi_s(f_d)$ in each subregion s , and we want to know what disk fraction our SB model predicts. The fraction of integrated I-band light contributed by the disk is of course Σ_d/Σ_T . However, to compare this to a kinematically-derived disk fraction, we must convert it to a fraction of stars, as sampled by the somewhat arbitrary SPLASH survey, contributed by the disk. This conversion requires knowledge of the intrinsic disk luminosity function \mathcal{L}_d , the total luminosity function \mathcal{L}_T , and the subregion-specific SPLASH selection function $S_s(m_{\text{F814W}})$. The model disk fraction in SPLASH star count units in subregion s is

$$f_{d,s} = \frac{\int S_s n_{d,s} \mathcal{L}_d dm_{\text{F814W}}}{\int S_s n_{t,s} \mathcal{L}_T dm_{\text{F814W}}} \quad (5.15)$$

where the integration only needs to be performed over the magnitude range $m_{\text{F814W}} = [20, 22]$, where the SPLASH selection function is nonzero.

Later in this paper (section 5.5.4), we will use a conversion factor C_s to convert between disk fraction units. C_s is defined as the ratio between $f_{d,s}$ (the disk fraction as measured in SPLASH star counts, as defined in Equation 5.15) and the disk fraction as measured in SB units:

$$C_s \equiv \frac{f_{d,s}}{\Sigma_{d,s}/\Sigma_{T,s}}. \quad (5.16)$$

C_s is constant within a subregion s , but varies from subregion to subregion.

As with the SB data, we introduce a kinematical uncertainty parameter e_k (expressed as a fraction of the empirical skew-normal width σ_s) to account for non-ideal f_d calculation due to, e.g., any non-Gaussianity in the disk line-of-sight velocity distribution. We add this parameter in quadrature to σ_s and re-normalize the widened skew-normal PDF.

Luminosity Functions

Equation 5.15 shows that computing the predicted disk fraction f_d (in star count units) requires knowledge of the disk luminosity function \mathcal{L}_d in the magnitude range $m_{F814W} = [20, 22]$. Therefore, we model the bright end of the disk, bulge, and halo luminosity functions. This magnitude range includes the tip of the red giant branch (TRGB), which lies near $m_{F814W} = 20.5$ at the distance of M31 for stars with $[M/H] \lesssim -0.5$ and fainter for more metal-rich populations (see discussion in § 5.3.3).

Because RGB stars dominate the stellar population faintward of the TRGB,

while younger objects such as AGB stars fill in the brightward portion of the LF, there is reason to expect that the number density of stars should fall at different rates faintward and brightward of the TRGB. Hence, we parameterize the LF of each component $k =$ (disk, bulge, halo) as a broken powerlaw in $\log(\text{number density})$ vs. magnitude space:

$$\log \mathcal{L}_k(m_{F814W}) = c_{0k} + \begin{cases} p_k m & m < l_k \\ q_k m + l_k(p_k - q_k) & m \geq l_k \end{cases} \quad (5.17)$$

where c_{0k} is chosen such that \mathcal{L}_k is normalized to unity over this magnitude range.

Since it is reasonable to expect that the stellar disk may have an age or metallicity gradient, we allow each of the disk LF parameters to depend linearly on radius on the plane of the disk $R_{\text{eff,d}}$ in the radial range of interest $5 \text{ kpc} < R_{\text{eff,d}} < 20 \text{ kpc}$:

$$p_d = p_{d0} + \delta p_d \ln(R_{\text{eff,d}}) \quad (5.18)$$

$$q_d = q_{d0} + \delta q_d \ln(R_{\text{eff,d}}) \quad (5.19)$$

$$l_d = l_{d0} + \delta l_d \ln(R_{\text{eff,d}}) \quad (5.20)$$

However, for simplicity, we require that the power-law slopes p, q and the break magnitude l for the bulge and the halo be constant with radius. This assumption does not affect our results, since the portion of the galaxy covered by the LF and kinematical surveys is almost entirely disk-dominated according to SB-only decompositions

(Courteau et al., 2011) as well as our decomposition. Even if the halo does contribute a significant number of stars in the SPLASH survey region, its range of metallicities is confined to the metal-poor regime (Kalirai et al., 2006b), where the magnitude of the TRGB is insensitive to metallicity. Near the end of Section 6.1, we discuss this point further.

We require for each component that p, q be nonnegative within the radial range $5 \text{ kpc} < R_{\text{eff},d} < 20 \text{ kpc}$, and that l lie within the magnitude range of interest, $m_{F814W} = [20, 22]$. We also require that the bright-end slope of the halo LF to be extremely steep ($p_h = 500$, because an old component should not have stars brighter than the TRGB) and that the break magnitude of each component be fainter than the brightest expected TRGB components ($l_k < 20.3$).

The total normalized LF $\mathcal{L}_{T,s}$ in subregion s is then the weighted sum of those of the disk, bulge, and halo, where the weights correspond to the number density of stars in each subcomponent in the magnitude range of interest. To compute the weights, we assume that in each component, the ratio of the number density n_k of stars in the range $I = [20, 22]$ to the surface brightness Σ_k of that component is a constant y_k . Hence, we introduce four new model parameters y_b , y_h , y_{d0} , and δy_d , defined by the following relationships:

$$y_b \equiv \frac{n_{b,s}}{\Sigma_{b,s}} \quad (5.21)$$

$$y_h \equiv \frac{n_{h,s}}{\Sigma_{h,s}} \quad (5.22)$$

$$y_d(R_{\text{eff,d}}) \equiv \frac{n_{d,s}}{\Sigma_{d,s}} \quad (5.23)$$

$$y_d(R_{\text{eff,d}}) = y_{d0} + R_{\text{eff,d}} \delta y_d \quad (5.24)$$

where $\Sigma_{k,s}$ is the average surface brightness of subcomponent k integrated over the area of subregion s . The ratio y_k is a constant independent of subregion s for the bulge and for the halo, and depends linearly on $R_{\text{eff,d}}$ for the disk. Then the total number density of stars in subregion s is

$$n_{t,s} = (y_{d0} + R_{\text{eff,d}} \delta y_d) \Sigma_{d,s} + y_b \Sigma_{b,s} + y_h \Sigma_{h,s} \quad (5.25)$$

and the normalized model LF is

$$\mathcal{L}_{T,s}(m) = \frac{n_{d,s} \mathcal{L}_d + n_{b,s} \mathcal{L}_b + n_{h,s} \mathcal{L}_h}{n_{t,s}} \quad (5.26)$$

As with the SB and kinematical data, we introduce an uncertainty parameter $\epsilon_{\mathcal{L}}$ to account for differences between the assumed broken power law and the actual shape of the LF. $\epsilon_{\mathcal{L}}$ is a fractional uncertainty whose value is constant across all magnitude

bins and all subregions. The total fractional uncertainty on a given LF bin m and subregion s is the quadrature sum of the Poisson fractional uncertainty and $\epsilon_{\mathcal{L}}$. Then the total uncertainty on that LF bin is

$$\epsilon_{LF,m,s} = \mathcal{L}_{\text{PHAT},m,s} \sqrt{\frac{1}{N_{m,s}} + \epsilon_{\mathcal{L}}^2}. \quad (5.27)$$

5.4.2 Model-Data Comparison: Likelihood Function

The probability that a point in parameter space is a good representation of the data is given by the product of three goodness-of-fit statistics describing the agreement between the model and the three data sets:

$$P = P_{\text{SB}} P_{\text{LF}} P_{\text{fd}} \quad (5.28)$$

We work instead with the log-likelihood:

$$\ln P = \ln P_{\text{SB}} + \ln P_{\text{LF}} + \ln P_{\text{fd}} \quad (5.29)$$

The SB factor is summed over each data point i in each data set j :

$$\ln P_{\text{SB}} = - \sum_{i,j=1}^{N_i, n_i} \left[\frac{(\Sigma_{T,ij} - \Sigma_{\text{obs},ij})^2}{2\pi\epsilon_{ij}^2} \right] \quad (5.30)$$

Meanwhile, the LF and disk fraction factors are summed over each of the 14 subregions. The goodness-of-fit to the disk fraction is the height of the skew-normal function $\phi_s(f_d)$ that describes the probability distribution of the kinematically measured disk fraction in subregion s , evaluated at the model disk fraction $C_s \frac{\Sigma_{d,s}}{\Sigma_{T,s}}$:

$$\ln P_f = \sum_{s=1}^{14} \ln \phi_s \left(f_d = C_s \frac{\Sigma_{d,s}}{\Sigma_{T,s}} \right) \quad (5.31)$$

The LF component of the likelihood is determined by summing the difference between the observed and model total luminosity functions over every magnitude bin m in every subregion s :

$$\ln P_{\mathcal{L}} = - \sum_{s,m=1}^{14,20} \frac{(n_{p,s} \mathcal{L}_{\text{PHAT},s,m} - n_{T,s} \mathcal{L}_{T,s,m})^2}{\epsilon_{LF,s,m}^2}. \quad (5.32)$$

5.4.3 MCMC Sampler

To estimate the probability distribution function of each model parameter, we draw samples from the log-likelihood function (Equation 5.29) using the the Markov chain Monte Carlo sampler `emcee` (Foreman-Mackey et al., 2015). More details on

`emcee` can be found in the Appendix. This section will only summarize the details unique to this paper.

In `emcee`, and more generally MCMC, an ensemble of “walkers” — or points in parameter space — moves through parameter space. During each step, each walker is given the option to move a specified distance along the line in parameter space connecting it to a random other walker. Moves corresponding to increases in the value of the likelihood function are always accepted; moves corresponding to decreases in likelihood are sometimes accepted. After many steps (the “burn-in” phase), the distribution of walkers samples the likelihood function: the density of walkers is highest in high-probability regions of parameter space. In this paper, we allow 256 walkers to burn in for 10,000 steps, and analyze the probability distributions using their positions in their last 100 steps. These probability distributions are shown in Figures 5.19 and 5.20 in the Appendix.

5.5 Results

5.5.1 Confidence Intervals & Correlations

We measure the mean value and 1σ confidence interval for each parameter from its posterior probability distribution. We report these statistics in Table 5.2. In general, the parameters describing the SB profiles are very well constrained, to better than 1%, while the LF parameters are constrained to 10% at best.

Some of the parameters appear in Table 5.2 to be completely unconstrained, but in fact are simply strongly correlated with other parameters. For a given such

pair, the allowed region in 2D space can be quite small. To quantify covariances between parameters, we calculate the Spearman correlation coefficient r between every pair of parameters. Those pairs with $r^2 > 0.6$ are labeled “strongly correlated”, while those with $0.25 < r^2 < 0.6$ are labeled “significantly correlated.” Figure 5.9 shows the 2D probability distributions of 24 of the 28 strongly or significantly correlated pair of parameters.

All six of the strongly correlated pairs, and all but four of the significantly correlated pairs, consist of two parameters describing the same subcomponent. For example, the bulge effective radius R_b is significantly correlated with the bulge Sérsic index n_b , but not with the disk or halo scale radii. Similarly, the disk LF bright-end slope p_{d0} depends on the disk LF break magnitude l_{d0} , but not on the slopes or break magnitudes of the bulge or halo LFs. Additionally, with a few exceptions, the SB parameters tend to be correlated only with other SB parameters, while the LF parameters tend to be correlated only with other LF parameters.

5.5.2 Quality of Profile Fits

Because the degeneracy between *parameters* is generally confined to within a subcomponent, the subcomponent SB and LF *profiles* are well constrained and are not degenerate with one another. The middle panel of Figure 5.10 shows a minor axis projection of the SB decomposition into bulge, disk, and halo subcomponents. One set of profiles (bulge, disk, and halo) is displayed for each of 256 samples drawn from the last step of the walker ensemble. The set of 256 violet lines, then, samples the entire range of bulge profiles allowed by the data; similarly, the set of 256 red lines

represents the entire allowed range of disk profiles. Note that the profiles are relatively well constrained: for example, while the bulge scale radius R_b and Sérsic index n_b each vary significantly, they covary in such a way that the bulge is always small, contributing less SB than the halo at only $R \sim 4$ kpc on the minor axis.

The LF decomposition in each of three representative subregions is shown in Figure 5.11. The disk (red dotted line) dominates in every subregion, so that it is nearly indistinguishable from the total model LF (orange shaded region) or observed PHAT LF (blue line with error bars). For clarity, in this plot we only display the bulge, disk, and halo LFs corresponding to the median values of the parameters. However, in Figure 5.12 we display the three LFs normalized in the magnitude range $20 < m_{F814W} < 22$, with one line drawn for each of 256 samples in the walker ensemble. Since the shape of the disk LF is allowed to change with radius, we display two representative disk LFs: one at the radius of an inner subregion (red lines) and one at the radius of an outer subregion (gold lines). The disk LF is tightly constrained at any given radius despite the significant degeneracy between the slopes and central values of the individual LF parameters. The bulge (violet) and halo (green) LFs are also relatively well constrained, especially considering that neither component dominates in any subregion sampled by our portion of the PHAT survey.

The model fits the SB and LF data sets quite well. The bottom panel of Figure 5.10 displays the difference between the observed and best-fit SB profiles expressed in units of SB uncertainty (a combination of Poisson measurement errors and error parameter). The fit to the SB is very good: the magnitude error is generally less than

10% of the SB uncertainty, corresponding to a median reduced χ^2 of 0.82. The fit to the LF is also acceptable, with a median χ^2 of 0.81. However, the fit to the kinematics is less satisfactory (median reduced $\chi^2 = 1.88$). Figure 5.13 shows that the model overestimates the disk fraction in almost every subregion, and by more than 1σ in four subregions. Though the measured disk fractions do constrain the model, their limited number and relatively large uncertainties mean that they have less effect on the final decomposition than the LF and SB data, which strongly prefer a small Sérsic component and dominant exponential component. The poor fit to some of the kinematical measurements is then a sign of tension between our (very simple) model and the data. In § 5.6.1, we show that allowing for a dynamically hot ($\sigma_v \sim 150 \text{ km s}^{-1}$) component with disklike population and spatial profile (a “kicked-up disk”) can reduce this tension.

The decompositions presented in Figures 5.10 and 5.12 will be discussed in more detail later in this section (§ 5.5.3) and in the Discussion (§ 5.6).

5.5.3 Comparison to previous measurements

The median values of our parameters are presented in Table 5.2. Here, we compare some of the structural parameters to previous measurements.

The disk scale length we measure, $5.76 \pm 0.1 \text{ kpc}$, is consistent with the range of accepted values. Worthey et al. (2005) measure an $R_d = 5.8 \pm 0.2 \text{ kpc}$ from elliptical isophotal fits to the Choi I-band image alone. Seigar et al. (2008) measure $R_d = 5.91 \pm 0.27 \text{ kpc}$ from an IRAC $3.6 \mu\text{m}$ profile. Ibata et al. (2005) measured a scale length of $5.1 \pm 0.1 \text{ kpc}$ by fitting an exponential profile to RGB star counts in the range $R = 20 - 40 \text{ kpc}$ on the major axis — though if the spheroid actually contributes some

additional light in the inner part of that radial range, the disk profile could appear to fall off faster than it actually does, and so the authors would underestimate the disk scale length.

Our R_d is nearly 0.8 kpc (16%) larger than that measured in the same band with an SB-only decomposition (Model U in Courteau et al. (2011)). This offset is likely to arise because we fit for the position angle of the disk, whereas Courteau et al. (2011) simply assume a position angle of 37.7° . The scale length of a disk is largest when measured along its true major axis (that is, when the position angle is correct).

Our median disk ellipticity, 0.725 ± 0.005 , corresponds to an inclination of $74.0^\circ \pm 0.3^\circ$ in the infinitely thin limit. The HI disk at similar radii has a similar inclination (Chemin et al., 2009; Corbelli et al., 2010). Overall, the optical disk appears about three degrees closer to edge-on, though M31's badly warped outer disk (Chemin et al., 2009; Corbelli et al., 2010) implies a strong dependence of inclination on radius. Additionally, the older populations that we sample are more likely to have been heated and thickened via satellite interactions at some point in their lifetimes, making the RGB disk appear more face-on than the younger stellar disk.

Our disk position angle ($44.4^\circ \pm 0.5^\circ$) is 6 – 8 degrees greater than that of the HI disk in the $R \sim 5 - 20$ kpc range (Chemin et al., 2009; Corbelli et al., 2010), but is similar to that seen in the azimuthal number density of all RGB stars in the PHAT dataset (Dalcanton et al., 2012).

Our bulge is nearly identical to that found in the SB-only decomposition of Courteau et al. (2011). The Sérsic index (1.92 ± 0.08) is intermediate between expo-

nential and de Vaucouleurs, characteristic of a disk “pseudobulge.” (Kormendy & Kennicutt, 2004).

Athanassoula & Beaton (2006) and Beaton et al. (2007) have found that NIR photometry and HI/H α kinematics are reproduced by N-body models with at least two bulge components: a small classical bulge, a larger boxy pseudobulge, and a bar that may extend beyond the pseudobulge. We do not distinguish between these subcomponents; instead, our $n = 1.9$ Sérsic bulge includes all of them. As described later (end of §5.6.1), we also try to model a two-component bulge by fitting the sum of an exponential, a power-law, and *two* Sérsic profiles to the composite data set, but the highest-likelihood models are exclusively the single-Sérsic ones.

The core radius of the halo, $R_h = 10.6_{-2.0}^{+2.5}$ kpc, is significantly larger than the value of 3 kpc measured from a fit to the number density profile of resolved blue horizontal branch (BHB) stars (Williams et al., 2012). This quantity is in general difficult to measure, since it relies on tracing the halo in the inner regions of the galaxy, where the faint halo is strongly subdominant to the disk and bulge. Our technique is one of the few that does not involve fixing other structural parameters or using only a single halo tracer such as BHB stars or RGB stars. As shown in Figure 5.9, R_h is degenerate with the power-law halo slope α , but is very unlikely to be shorter than 8 kpc and is not degenerate with any other parameters.

The power-law slope of the halo profile external to the core is -2.5 ± 0.2 , consistent with measurements from SB-only decompositions (Courteau et al., 2011), resolved RGB star counts (Guhathakurta et al., 1911; Gilbert et al., 2012) and BHB

star counts (Williams et al., 2012). A projected surface density power-law slope of -2.5 corresponds to a deprojected density that scales approximately as $r^{-3.5}$, in good agreement with cosmologically motivated simulations in which stellar halos are built up via accretion (e.g., Bullock & Johnston, 2005; Cooper et al., 2010) and accretion plus in situ star formation (e.g., Font et al., 2011). The outer slope of a Navarro et al. (1997) cold dark matter halo also scales as $r^{-3.5}$, suggesting that stars may trace the dark matter profile — but not significantly affect its shape — at large radii.

5.5.4 Conversion between integrated-light and star-count disk fractions

The conversion factor between disk fraction as measured in integrated light and disk fraction as measured in SPLASH star counts (C_s in Equation 5.16) is $0.84 - 0.86$ in the subregions near the major axis and $0.62 - 0.65$ near the minor axis. $C < 1$ implies that the spectroscopic survey is biased towards spheroid stars — not a surprise, since the spectroscopic target selection strategy prioritizes less crowded objects and spheroid stars are less likely to be located in clusters or clumpy structures. In the discussion that follows, as in the analysis, this correction has been applied to the model so that both the model and measured disk fractions represent the fraction of stars, as sampled by the SPLASH survey, that contribute to the disk.

5.6 Discussion

We discuss three classes of new results. First, we argue that a fraction of the stars in M31’s dynamically hot inner spheroid may have originated in the disk. Second, we discuss the transition between the bulge and halo. Third, we discuss the evidence for a radial gradient in the LF of M31’s stellar disk.

5.6.1 Kicked-up disk

Despite the significant dynamically hot population in the kinematical sample, both the SB and LF data are best fit by a model with a bulge too small to contribute light in the SPLASH survey region. Figures 5.10 and Figure 5.11 illustrate the fits. Figure 5.10 shows that the SB profile on the minor axis between projected radii of 1 – 10 kpc — the radial range covered by the overlap of the SPLASH and PHAT surveys — is well fit by an exponential profile. Figure 5.11 shows that the observed LF is also well fit by a decomposition where nearly all of the stars belong to the disk component. Even without trusting the simple toy model decomposition, the shape of the observed LF from PHAT looks disklike in every subregion. The change in slope of the LF at $m_{F814W} \sim 20.5$ corresponds to the TRGB of low- or intermediate-metallicity populations at the distance of M31. In a region dominated by an old population, the LF would drop off steeply brightward of the TRGB; only in a population with a significant young- or intermediate-age fraction (as expected for a disk) can bright stars such as asymptotic giant branch (AGB) stars fill in the bright part of the LF and give it the shallow slope as seen in Figure 5.11.

In summary, the SB and LF data suggest that the region sampled by the SPLASH and PHAT surveys is completely disk dominated, even though the kinematics reveal a dynamically hot population in this region. Figure 5.13 illustrates this tension: even in the best simultaneous fit to the LF, SB, and kinematical data, the kinematically derived disk fractions are systematically lower than the model disk fractions. The inability to simultaneously fit the three data sets is a sign of tension between our (very simple) model and the composite data set. In this section, we propose a dynamically hot (“kicked-up”) disk component as a possible resolution, and walk through some other modifications to the model that cannot resolve the tension.

We have assumed that the population with a disk LF is exclusively dynamically colder than the spheroid, but this may be too restrictive. The tension between the model and the kinematics can be explained if the dynamically hot component identified in the kinematics is inflated by a contribution from stars that were born in the disk but dynamically heated. This population should have a disklike LF and follow the disk SB profile, but have spheroidlike kinematics. Note that such a component is not the same as a thick disk. The velocity dispersion of M31’s thick disk is only about 40% larger than that of the thin disk (Collins et al., 2011). In contrast, a kicked-up disk has a velocity dispersion similar to that of the halo: $\sim 150 \text{ km s}^{-1}$, or more than 300% larger than that of the thin disk (Dorman et al., 2012). In our kinematical decompositions, the “cold” fraction includes contributions from both the thin and thick disks, whereas the “hot” fraction includes contributions from the spheroid and the kicked-up disk. Of course, it is possible that thick disks are created via a heating mechanism similar to

(though less extreme than) that that creates the dynamically hot kicked-up disk.

Kicked-up disk stars have been seen in cosmological simulations for some time. N-body and hydrodynamical simulations predict that minor accretion events can disrupt galactic disks, kicking disk stars enough so that they would be kinematically classified as spheroid members, though they retain some of their angular momentum (Purcell et al., 2010; McCarthy et al., 2012). Purcell et al. (2010) find that a merger with mass ratio 1 : 10 can kick 1% of the disk stars into the halo; this percentage corresponds to a mass similar to the mass accreted from the incoming satellite itself. Tissera et al. (2013) find that 3 – 30% of the halo mass in members of a suite of MW-like galaxies from Aquarius consist of stars that originated in the disk. There is some observational evidence for this as well. Sheffield et al. (2012) identify M giants in the MW halo with velocities and abundances consistent with a kicked-up disk origin. M31 is in the process of merging with the progenitor of the Giant Southern Stream (GSS), which likely was first tidally disrupted less than 1Gyr ago (Fardal et al., 2008), so it is possible that it hosts a nonvirialized kicked-up disk component. [However, note that the GSS progenitor was much smaller than 10% of the mass of M31 (Fardal et al., 2008).]

In Figure 5.14, we map the fraction of the disklike stars that must be dynamically hot to simultaneously fit the LF and kinematics (the “kicked-up fraction”). Overall, the kicked-up fraction decreases with radius, as predicted by McCarthy et al. (2012); Purcell et al. (2010). The kicked-up fraction tends to be higher than the 1% level predicted by Purcell et al. (2010) for a 1:10 merger event, though the fraction in each individual subregion is consistent with 1% at the $1 - 2\sigma$ level. The fraction of

halo stars that come from the disk (a different quantity from the kicked-up fraction) varies from 1 – 30% — identical to the range found in the Aquarius halos by Tissera et al. (2013). Combining all subregions, in Figure 5.15 we show that the distribution in the overall kicked-up fraction is $5.2 \pm 2.1\%$, where the spread in the distribution comes from both the subregion-to-subregion variation and the uncertainty on the mean value. The figure also shows that the kicked-up fraction is greater than 0.6% with 95% confidence. There are two subregions with outlying kicked-up fractions: NE1₃ (with a very high kicked-up fraction) and NE2₁ (where the disklike fraction appears *smaller* than the dynamically cold fraction). Subregion NE1₃ has a very high kicked-up fraction of disk stars: 0.31 ± 0.09 . It happens to be at the inner edge of the distribution of possible locations of the remnant core of the GSS progenitor (Fardal et al., 2013). The broad velocity dispersion in this region could be a recently disturbed portion of the disk, or be biased by the core itself. Subregion NE2₁ sits on top of both the dusty 10 kpc ring and the GSS, so many stars had to be excluded from the disk fraction measurement. While exclusion of stars in dusty regions does not appear to bias the velocity distribution, exclusion of possible GSS debris is much more likely to affect spheroid stars than disk members.

Our decomposition is of course limited by our choice of model, and there are many plausible parameterizations of the bulge, disk and halo SB profiles. For completeness, we consider the possibility that a modified SB model could eliminate the need for a kicked-up disk component. We attempt to fit the data with four different modified SB models, but none result in an acceptable fit to all three data sets.

First, we try fitting a model with a second exponential disk component with its own central intensity, ellipticity, scale length, and LF parameters. However, the most probable models in the posterior distribution are those with a single disk component.

We also try using a model with a single disk whose exponential scale length is allowed to change at a break radius R_{break} (that is, allowing for a Freeman type II or III profile). However, the likelihood values show no preference for a disk with a break.

Bulgeless disk galaxies can have SB profiles that deviate from pure exponential. We try fitting a model with a disk whose Sérsic index is a free parameter. The resulting PDF of the Sérsic index has a mean of 0.94 ± 0.4 — that is, the disk is slightly more cored than a pure exponential. This decomposition does not reduce the required kicked-up fraction.

The BHB profile can be well fit by a halo component alone if the halo core radius is ~ 3 kpc (Williams et al., 2012). We try fixing $R_h = 3.0$ kpc, but this degrades the fit to the kinematics at large radii (regions NE2 and NE3) with minimal to no improvement at smaller radii.

The central region of M31 is complex, with multiple components including a classical bulge that dominates the SB within 0.2 kpc on the major axis, boxy bulge that dominates within 2.7 kpc, and bar that extends to at least 4.5 kpc (Beaton et al., 2007). We allow for the possibility of multiple spheroidal components by fitting a model with two Sérsic bulges with unique central intensities, ellipticities, scale lengths, Sérsic indices, and LFs. However, the fit to the kinematics is not improved.

Finally, we try relaxing the assumption of a Sérsic bulge profile. The kinematics

would be better fit if the bulge were more extended — if it had a shallower slope in the outer regions. However, the outer slope and inner slopes of a Sérsic profile are determined by the same parameter n , which in our case is completely constrained by the inner ~ 0.7 kpc (on the minor axis) where the bulge dominates the SB. We consider a modified Sérsic profile, whose inner and outer slopes are allowed to differ but whose shape reduces to a Sérsic when the slopes are the same. Even with this added flexibility, the data prefer models with near-Sérsic bulge profiles: small bulges nearly identical to those in Figure 5.10.

It is also possible that the excess dynamically hot population belongs to the virialized remnants of tidal debris. Dynamically cold substructure has been seen in M31’s halo and, to a smaller extent, in the portion covered by the PHAT survey; it is not unreasonable to suppose that the remnants from older satellite encounters contribute stars to the central portion of the galaxy. However, note that the debris would have to be old enough to have virialized (because it is dynamically hot), but have a significant fraction of young or intermediate-age stars (to have enough AGB stars to fill out the LF brightward of the TRGB). Analysis of cosmological simulations such as ERIS (Guedes et al., 2011) will provide insight into the relative contributions of young, virialized tidal debris and kicked-up disk stars in the inner halos of large spiral galaxies.

5.6.2 Relationship between bulge and halo

The location of the transition between M31’s bulge and halo has long been unclear. From images and SB decompositions (e.g., Courteau et al., 2011) the bulge appears to be relatively small, with a Sérsic index of around 2 and an effective radius of

around 1 kpc. However, resolved stellar population studies have raised the possibility that the bulge may be much bigger. Deep optical HST CMDs of a minor-axis field at a projected radius of about 15 kpc (~ 12 disk scale radii on the minor axis) revealed a broad spread in $[\text{Fe}/\text{H}]$ ($\sim -1.5 - 0$) and age ($\sim 5 - 13$ Gyr) (Brown et al., 2003, 2006), more similar to the MW’s bulge than to its old, metal-poor inner halo. It appeared that either M31 has a large bulge or else its halo has had a much different formation history than the MW’s.

In Dorman et al. (2012), we showed that a significant fraction ($\sim 10 - 20\%$) of the stars in our kinematical sample belong to a dynamically hot population — presumably either the outer reaches of a centrally concentrated bulge or the inner portion of an extended halo. With kinematics alone, we were unable to distinguish between these two scenarios, but we are now in a position to show that the vast majority of these stars are associated with the halo.

Our decomposition indicates that the SB and LF profiles of the galaxy are much better fit by a small ($R_b = 0.77 \pm 0.03$ kpc) bulge. Figure 5.16 maps the fraction of spheroid SB due to the bulge. This fraction falls below 0.5 — that is, the halo dominates the spheroid SB — exterior to $R_{\text{proj}} \sim 5.5$ kpc on the major axis. Nearly all of the SPLASH survey region falls in this halo-dominated region. As shown in § 5.6.1, some of these stars may have originated in the disk but have been dynamically heated to kinematically resemble halo members.

In Dorman et al. (2012), we found that the dynamically hot population rotates with $v/\sigma \sim 1/3$, significantly more slowly than the typical spiral galaxy bulge (Cappel-

lari et al., 2007). It is now clear that this is not necessarily a useful comparison, as the dynamically hot population is not physically associated with the bulge. A more relevant comparison would be to the “inner halo” of the MW, which rotates much more slowly than our dynamically hot population (Carollo et al., 2007, 2010). It is possible that the accreted halo population in M31 resembles in the inner halo of the MW with almost no rotation, but the portion that originated in the disk has some residual rotation.

We can learn more about the relationship between the bulge and halo by comparing their SB profiles to the density profile of BHB stars, a reliable tracer of metal-poor populations. BHB stars have a mean metallicity of $[Fe/H] \sim -1$, with only a small high-metallicity tail, and so are unlikely to have formed in more metal-rich subcomponents. Williams et al. (2012) combined star counts of CMD-selected BHB stars from the first two years of the PHAT survey. They showed that the density profile of BHB stars increases steeply interior to 10 kpc, but not as steeply as the density profile of RGB stars. Our results are qualitatively consistent with theirs. In Figure 5.17, we compare our halo SB profile and total SB profile (combined halo, bulge, and disk) to the BHB profile of Williams et al. (2012), scaling our SB profiles to match the BHB counts at $R = 20$ kpc. The BHB/SB ratio (which should trace the BHB/RGB ratio) increases with radius, as found in Williams et al. (2012).

However, our halo core radius is large enough that the halo alone cannot account for the BHB density interior to 10 kpc. The disk (and possibly the bulge) must supplement the metal-poor population of the halo. Williams et al. (2012) argued that it was unlikely that a significant portion of the BHB stars belonged to the disk, in part

because the kinematically-derived disk fraction in a field at $R = 9$ kpc on the minor axis – where there were many BHB stars – was no more than 10%. However, we now know two reasons that this disk fraction may have been an underestimate. First, the SPLASH target selection strategy preferentially chooses spheroid stars over disk stars, underestimating the disk fraction by a factor of 1.6 on the minor axis. Second, some fraction of the true disk stars may have been kicked up so that a kinematical decomposition grouped them with the halo. The metallicity distribution of the disk then is likely quite broad: some stars belong to a metal-poor BHB population, while the faint break magnitude of the disk LF indicates a contribution from stars with solar or supersolar metallicities.

5.6.3 Radial Gradient in the Disk LF

As shown in Table 5.2, Figure 5.19 and Figure 5.20, the radial gradient parameters in the disk LF are exclusively nonzero.

To show that a radial gradient is required to simultaneously fit all three data sets, we run a test decomposition with a *constant* disk LF. The resulting fit to the LF is quite poor (median reduced $\chi^2 = 70$ including only observational uncertainties). Figure 5.18 shows the fits in a representative sample of three subregions. Most noticeably, the model predicts a number density that is too high at small radii and too low at large radii. The uncertainty parameter on the LF is driven high in an attempt to resolve the tension between the model and data, inflating the effective uncertainty beyond the Poisson uncertainty. In contrast, when we allow for a radial gradient, the LF uncertainty parameter can be very small, and the structural parameters and the fit to the

kinematics are unaffected (see Figure 5.11).

It is not surprising that there is a radial gradient in the LF of M31's stellar disk. Such a variation could be caused by either an age or metallicity gradient, either of which is plausible: HII region abundance estimates suggest that there may be a small metallicity gradient within M31's (gas) disk (e.g., Sanders et al., 2012; Zurita & Bresolin, 2012), while inside-out disk formation and/or radial migration could induce an age gradient.

As discussed earlier in the paper, the break magnitude of the TRGB can be used as a proxy for metallicity: for populations with $[\text{Fe}/\text{H}] \gtrsim -1$, the TRGB is much more sensitive to metallicity than age, becoming fainter as metallicity increases. The shapes of the LFs in Figure 5.12 illustrate the trend from higher metallicity at low radius (red curves) to lower metallicity at high radius (orange curves). The same trend – metal abundance that decreases with radius – seen in M31 disk PNe (Sanders et al., 2012; Kwitter et al., 2012) and HII regions (Sanders et al., 2012; Zurita & Bresolin, 2012). At large radii, the average metallicity of the disk is similar to that of the bulge, but still higher than that of the halo.

Note that this gradient exists in the old, smooth stellar disk; it does not reflect (and thus is not biased by) the sharp changes in stellar population found in the narrow star-forming rings. Recall that we excluded from the PHAT LF every star that fell within a $25 \times 25 \text{ pc}^2$ pixel with a reddening $A_v > 1.0$; as shown in the middle panel of Figure 5.5, this cut effectively removes stars in the star-forming rings. In addition, we have cut out the few remaining stars blue enough to still be on the young, massive

main sequence.

5.7 Summary

We have presented the first structural decomposition of a large spiral galaxy using simultaneous SB, LF, and kinematical constraints. We have used Bayesian inference to find the probability distribution functions of 32 parameters describing the surface brightness profile and LF of each structural subcomponent. We have found:

1. The structural parameters we measure are consistent with previous measurements. On average, the old stellar disk is more highly inclined and its major axis PA is larger than measured from isophotal SB fitting.
2. A decomposition including a Sérsic bulge, power-law halo, exponential disk and constant shapes to the bulge, halo, and disk LFs cannot simultaneously well fit the SB, LF, and kinematical data. The model poorly predicts the number density of bright stars and overestimates the disk fraction in most subregions.
3. The high kinematically-derived spheroid fractions can be explained if some spatially dependent fraction (between 1% and 30%) of the dynamically hot component is comprised of “kicked-up” disk stars. This kicked-up population has a disklike SB profile and disklike LF, but spheroidlike kinematics ($\sigma_v \sim 150 \text{ km s}^{-1}$).
4. In the I band, the halo is brighter than the bulge exterior to $R = 5.5 \text{ kpc}$ on the major axis.

5. Comparison to BHB stars, a tracer of the metal-poor population, indicates that the disk metallicity distribution has a low-metallicity tail.
6. A SB decomposition including a radially varying disk LF improves the fit to the observed LF and does not affect the structural parameters.




Figure 5.2: Locations of 637 fields from 4 datasets used for SB measurements. The Choi et al. (2002) measurements (teal circles) come from a wedge cut of an I-band image from KPNO; the rest of the SB measurements, including Irwin et al. (2005) (blue triangles), Gilbert et al. (2012) (violet stars), and Pritchett, C.J., van den Bergh (1994) (yellow squares) are derived from RGB star counts. The inner black box has a side length of 32 kpc, approximately the size of the GALEX image of M31. A zoom-in of this region is shown in Figure 5.3.

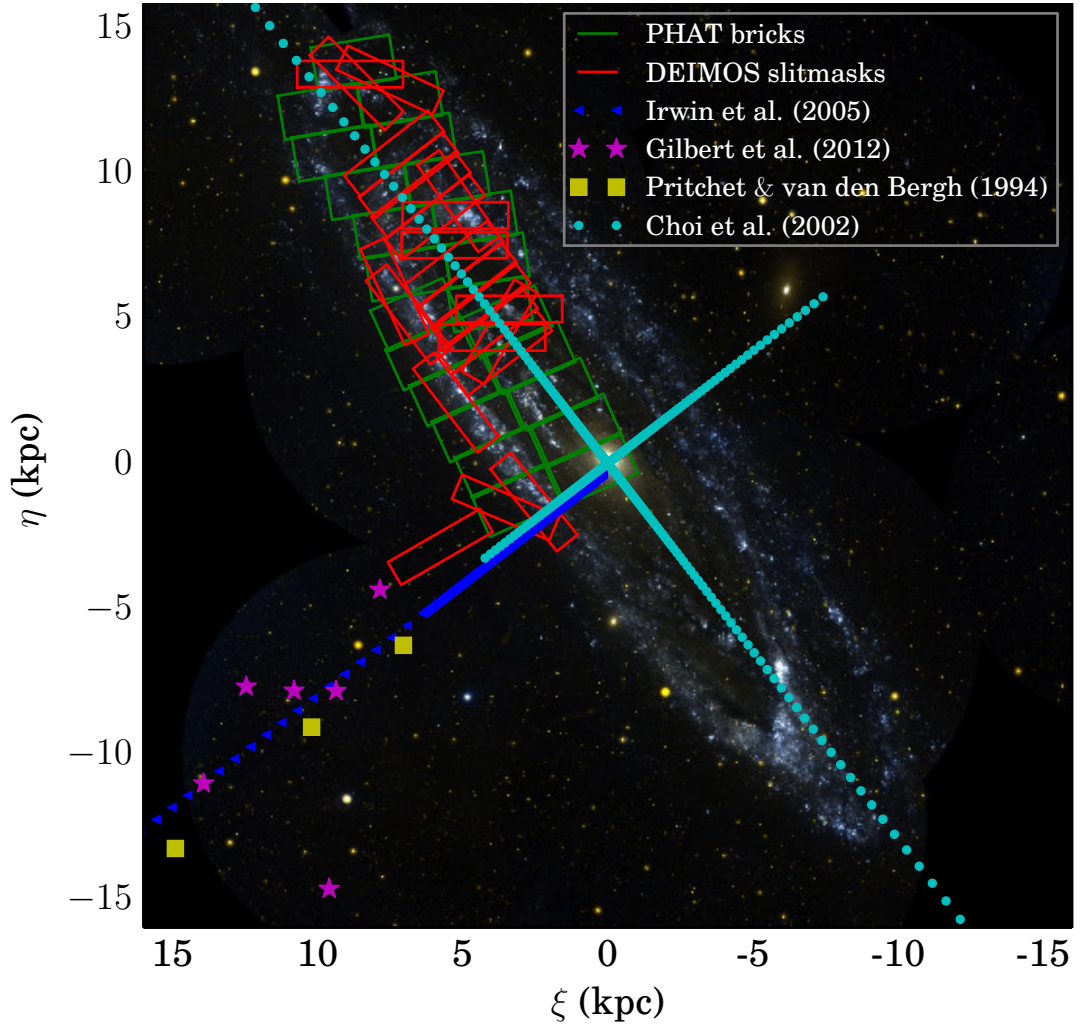


Figure 5.3: Zoom-in on the map of SB fields from Figure 5.2 that fall in the inner regions of M31, overlaid on a GALEX image. Points are color-coded as in Figure 5.2. The Irwin et al. (2005) points (blue triangles) are drawn slightly offset from the minor axis for clarity; they actually overlap the Choi et al. (2002) minor axis points (teal circles). For comparison, bricks in the HST PHAT survey (from which we measure the bright end of the luminosity function) are shown in green and Keck/DEIMOS slitmasks from the SPLASH survey (from which we measure radial velocities of RGB stars) are outlined in red. The kinematical and LF analyses presented in this paper are carried out in the regions where the green bricks and red slitmasks overlap.




Figure 5.4: Major- (left) and minor- (right) axis projections of the SB profiles, assuming a major axis position angle of 37.7° . Samples from opposite sides of the galactic center are collapsed onto a single set of axes. Gilbert et al. (2012) points (violet stars) do not lie on either axis, so for display purposes we project them to the minor axis assuming a circularly symmetric halo. Photometric uncertainties are typically smaller than the point size.




Figure 5.5: Map of target locations in the PHAT and SPLASH surveys. *Left:* Stars in the magnitude range $m_{F814W} = [20, 22]$ in PHAT. (Only a random 4% of the objects are shown.) *Middle:* Same as left panel, but without stars in dusty ($A_V > 1.0$) pixels. The star-forming 10 kpc ring is effectively excluded. *Right:* Same as middle panel, with stars from the SPLASH survey in non-extincted regions overplotted as colored dots. Black lines separate the 4 large “regions,” while colors demarcate the 14 smaller “subregions.” The $F814W$ luminosity functions and kinematically-derived disk fractions are measured in each subregion, where the kinematical and PHAT surveys overlap. Note that the central few kpc are not used in the fits to the kinematics and LF, because this region is too crowded for resolved stellar spectroscopy with DEIMOS.




Figure 5.6: *First four panels:* fits to the velocity distribution of RGB stars from the SPLASH survey in each of four subregions in the NE2 (middle northeast major axis) region, plotted over velocity histograms of stars in each subregion. Each velocity distribution is fit by a sum of two Gaussians corresponding to a dynamically hot spheroid (red) and dynamically cold disk (blue to green). The sum of the two curves is shown in violet. Velocity ranges excluded from the fits due to possible contamination by tidal debris from the Giant Southern Stream (GSS) are shown in two equivalent ways: by the gray shaded regions and by the stars shaded red on the histogram. *Bottom middle panel:* Sum of hot components in the NE2 region (red curve) and sum of cold components (blue curve) overplotted on a histogram of all radial velocities in NE2. *Bottom right panel:* Ellipses represent the mean and uncertainty of each of the parameter pairs (v, σ_v) . The four blue and green ellipses represent the kinematical parameters of the disk components in the four subregions. The red ellipse represents the spheroid component, which has the same (v, σ_v) in every subregion within this region.




Figure 5.7: Probability distribution of the kinematically derived disk fraction in the off-major-axis subregion $NE2_4$. The gray filled histogram is constructed from 50,000 samples from the end of the MCMC chain used to fit the sum of two Gaussians (disk and spheroid) to the observed velocity distribution in this subregion. The black curve represents the skew-normal function ϕ that best fits the histogram. ϕ is later used as the prior on the disk fraction; that is, a model that predicts a disk fraction near the peak (e.g., $f_d \sim 0.8$) will be more favorable than one that predicts a disk fraction far from the peak (e.g., $f_{\text{disk}} \sim 0.3$).




Figure 5.8: *Top panels:* PHAT (blue) and SPLASH (green) luminosity functions (LFs) in two subregions: NE2₁, which lies on the major axis, and SE2₂, which lies on the minor axis. The SPLASH LF has been scaled by a factor of 100. Error bars represent Poisson uncertainties. Note that the slope of the PHAT LF changes twice: at $m_{F814W} \sim 20.4$ mag (the TRGB of a population with $[M/H] \lesssim -0.6$) and around 21 mag (the TRGB of a more metal-rich population). *Bottom panels:* The empirical selection function $N_{\text{SPLASH}}/N_{\text{PHAT}}$ is shown for each subregion. The selection function varies from subregion to subregion.

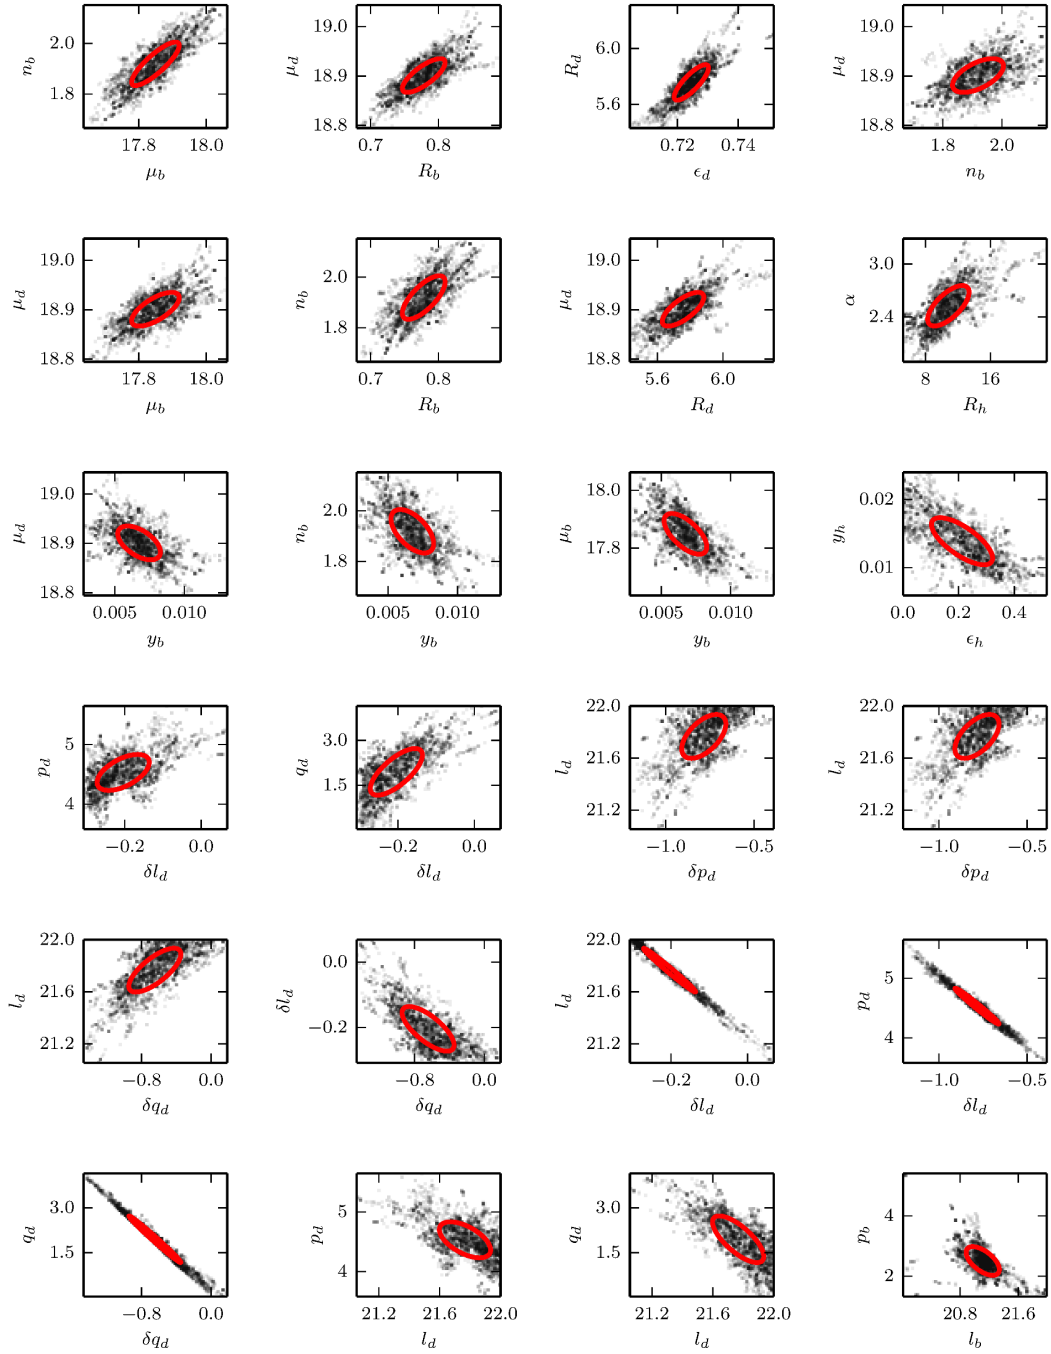


Figure 5.9: 2D posterior distributions of 24 of the strongly (Spearman $r^2 > 0.6$) or significantly ($0.6 \geq r^2 > 0.25$) correlated pairs of parameters. (Only 28 of the 512 pairs of parameters fall into one of these categories.) Red dashed contours are 1σ error ellipses. In general, a parameter is correlated only with other parameters describing the same subcomponent (bulge, disk or halo). LF parameters tend to be correlated only with other LF parameters, and SB parameters only with other SB parameters.

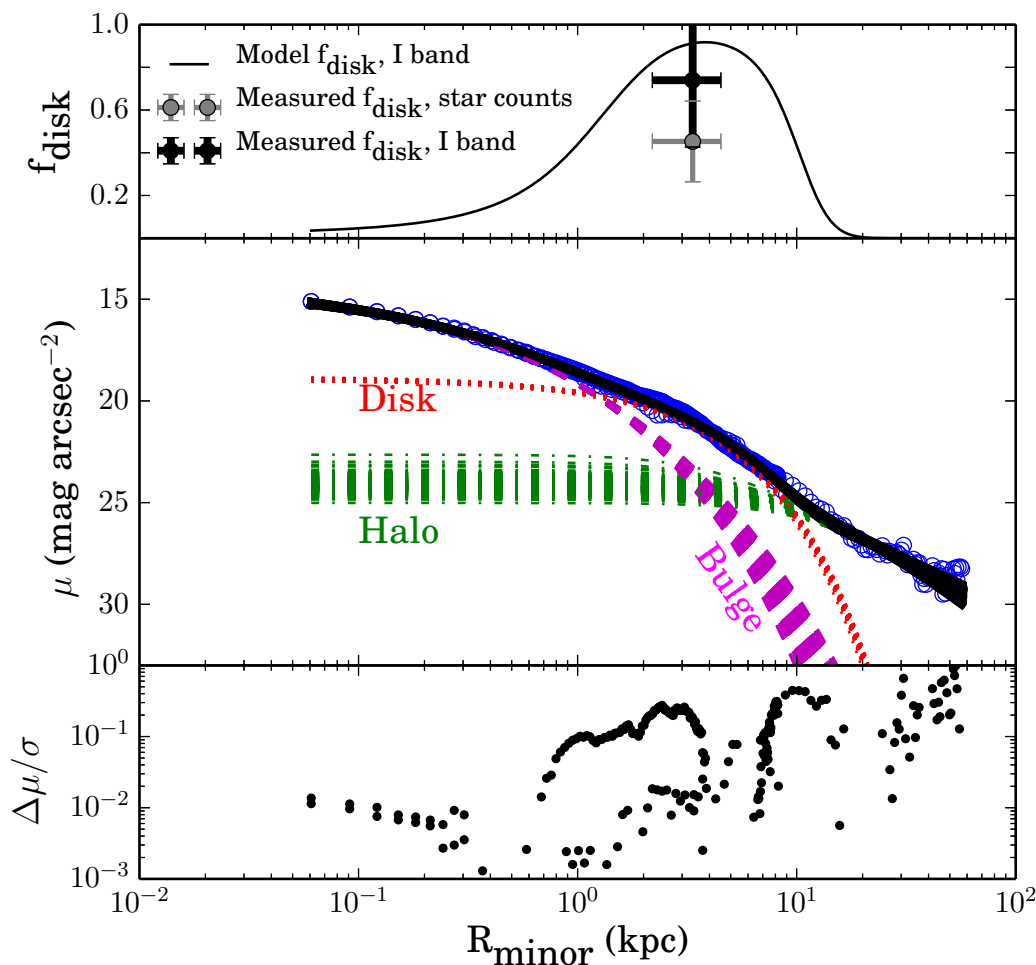


Figure 5.10: Minor axis SB decomposition. *Middle*: Fit of the model SB profile (black lines) to the observed SB profile (blue circles) as a function of projected radius. The model is also shown decomposed into the disk (red dotted), halo (green dot-dashed), and bulge (magenta dashed) components. One line is displayed for each of 256 samples drawn from the posterior distribution of walkers, so that the width of each region encloses the entire uncertainty (not the 1σ uncertainty) associated with that profile. The disk, inner bulge, and outer halo profiles are well constrained, whereas the inner halo and outer bulge profiles are less well constrained. *Top*: Model disk fraction in SB units (black line) slightly overpredicts kinematically measured disk fraction in the minor axis subregion (SE_2) converted to SB units (black cross). The conversion factor from disk fraction in star counts (gray cross) to disk fraction in SB units (black cross) is about 1.2 in this subregion: the kinematical survey oversamples the spheroid population. *Bottom*: Average residual between model and measured SB as a function of projected radius, relative to the measurement uncertainties on the SB. The lower and upper tracks between 0.3 and 2.5 kpc correspond to the Choi and Irwin data sets, respectively.




Figure 5.11: Comparison between observed PHAT (blue) and model (brown) total normalized LFs after a fit to the SB with kinematical constraints, assuming radial gradients in the disk LF shape parameters, in three representative subregions. The shaded portion shows where the SPLASH selection function is at least 30% of its maximum value. The model is shown decomposed into disk (red), bulge (magenta), and halo (green) components. Each LF is weighted by the number density of stars in the magnitude range $F814W = [20, 22]$ in that subcomponent in the subregion of interest. The width of the brown region is given by the variation in the model LF. Error bars on the PHAT LF include contributions from both the \sqrt{N} Poisson uncertainty in each 0.1 mag bin and the LF uncertainty parameter.




Figure 5.12: Normalized median bulge (magenta), disk (red, orange), and halo (green) LFs. For each component, one line is displayed for each of 256 samples drawn from the posterior distribution of walkers, so that the width of each region encloses the entire uncertainty (not the 1σ uncertainty) associated with that profile. The disk LF depends on radius in the plane of the disk R_d ; the disk LFs in the subregions with smallest and largest R_d are displayed in red and orange, respectively, to illustrate the range in disk LF over the PHAT survey region.




Figure 5.13: Comparison between the measured cold and model disk fractions, in units of SPLASH star counts, in the final decomposition. Error bars reflect estimated 1σ uncertainties from the appropriate distributions at the end of the MCMC chain. Points are color-coded by spatial region. Measured cold fractions are systematically lower than model disk fractions, suggesting that there may be a dynamically hot (“kicked-up”) disk component for which our model has not accounted.




Figure 5.14: Fraction of disk-like population that must be dynamically hot in order to simultaneously fit the LF and kinematics. Vertical error bars denote 1σ uncertainties and horizontal error bars denote the entire range of R_{proj} subtended by each subregion. The kicked up fraction in most of the subregions is consistent with, though systematically larger than, that predicted by Purcell et al. (2010) to result from a merger event at low impact angle with mass ratio 1:10.




Figure 5.15: PDF of f_{kicked} , the fraction of disk stars that must be dynamically hot in order to simultaneously fit the LF and kinematics. This probability distribution is constructed from the 14 PDFs corresponding to the individual subregions, weighted by the number of stars in each subregion. The most probable value of f_{kicked} is 5.2%. The width of the distribution, a combination of spatial variation in f_{kicked} and uncertainty on the mean, is $1\sigma = 2.1\%$. The distribution is narrow enough that a model with zero kicked-up fraction is effectively excluded.




Figure 5.16: Map of the relative contribution of the bulge to the spheroid SB. The bulge dominates over the halo in the central 5.5 kpc, but its surface brightness quickly falls off at larger radii. The spectroscopic sample (black dots) falls almost entirely in the region where the spheroid light is dominated by the halo; the dynamically hot stars are more likely to be associated with the halo than the bulge.




Figure 5.17: Comparison between counts of blue horizontal branch (BHB; blue dots) stars, presumed to trace the metal-poor population (Williams et al., 2012), and the most probable halo profile from our decomposition (green dot-dashed lines), scaled to match the BHB counts at 20 kpc. One line is drawn for each of 256 points from the posterior probability distribution, so that the colored regions approximate the entire allowed (not 1σ) region of parameter space. The power-law halo component cannot explain all of the BHB star counts; the bulge and/or the disk must contain a significant metal-poor population. The ratio of BHB counts to SB (red lines) increases with radius.




Figure 5.18: Same as Figure 5.11, assuming a constant disk LF shape. The blue error bars, the quadrature sum between the Poisson uncertainties and the uncertainty parameter, have been significantly inflated relative to those in Figure 5.11 in an attempt to reduce the tension between the data and this model. A radial gradient in the disk LF does a much better job of fitting the LF data without affecting the fits to the SB or kinematical data.

Table 5.2: Model Parameters

x	Description	Units	Prior ^a	Results
μ_d	Disk central SB	mag	$[0, \infty)$	$18.901^{+0.032}_{-0.031}$
R_d	Disk scale length	kpc	$[0, \infty)$	$5.76^{+0.101}_{-0.113}$
ϵ_d	Disk ellipticity	1	$[0, 1)$	$0.725^{+0.005}_{-0.005}$
μ_b	Bulge SB at R_e	mag	$[0, \infty)$	$17.849^{+0.066}_{-0.066}$
R_b	Bulge half-light radius	kpc	$[0, \infty)$	$0.778^{+0.03}_{-0.028}$
n_b	Bulge Sersic index	1	$[0, 10]$	$1.917^{+0.082}_{-0.081}$
ϵ_b	Bulge ellipticity	1	$[0, 1]$	$0.277^{+0.011}_{-0.011}$
μ_h	Halo central intensity	mag	$[0, \infty)$	$24.18^{+0.295}_{-0.328}$
R_h	Halo scale radius	kpc	$[0, \infty)$	$10.631^{+2.459}_{-2.034}$
ϵ_h	Halo ellipticity	1	$[0, < 1)$	$0.215^{+0.109}_{-0.12}$
α_h	Halo power-law slope	1	$[0, \infty)$	$2.508^{+0.232}_{-0.199}$
PA	Major axis PA	degrees	$[0, 180]$	$6.632^{+0.459}_{-0.511}$
y_{d0}	Disk # density/SB at 1 kpc	$\frac{10^5 N}{\text{mag}}$	$[0, \infty)$	$1.93^{+4.00}_{-3.90}$
p_{d0}	Disk LF bright slope at 1 kpc	$\log(N)/\text{mag}$	$[0, \infty)$	$4.519^{+0.313}_{-0.267}$
q_{d0}	Disk LF faint slope at 1 kpc	$\log(N)/\text{mag}$	$[0, \infty)$	$2.012^{+0.826}_{-0.843}$
l_{d0}	Disk LF break mag at 1 kpc	mag	$[20, 22]$	$21.782^{+0.146}_{-0.209}$
δy_d	Gradient in disk # density/SB	$\frac{N}{\text{mag} \ln(\text{kpc})}$	$(-\infty, \infty)$	$0.001^{+0.0}_{-0.0}$
δp_d	Gradient in disk LF bright slope	$\log(N)/\text{mag}/\ln(\text{kpc})$	$(-\infty, \infty)$	$-0.21^{+0.083}_{-0.06}$
δq_d	Gradient in disk LF faint slope	$\log(N)/\text{mag}/\ln(\text{kpc})$	$(-\infty, \infty)$	$-0.777^{+0.111}_{-0.134}$
δl_d	Gradient in disk LF break mag	$\text{mag}/\ln(\text{kpc})$	$(-\infty, \infty)$	$-0.672^{+0.321}_{-0.308}$
y_b	Bulge # density/SB	N/mag	$[0, \infty)$	$0.007^{+0.002}_{-0.001}$
p_b	Bulge LF bright slope	$\log(N)/\text{mag}$	$[0, \infty)$	$2.494^{+0.405}_{-0.38}$
q_b	Bulge LF faint slope	$\log(N)/\text{mag}$	$[0, \infty)$	$0.424^{+0.434}_{-0.307}$
l_b	Bulge LF break mag	mag	$[20, 22]$	$21.125^{+0.144}_{-0.189}$
y_h	Halo # density/SB	$N/\text{arcsec}^2/\text{mag}$	$[0, \infty)$	$0.014^{+0.004}_{-0.003}$
q_h	Halo LF faint-end slope	$\log(N)/\text{mag}$	$[0, \infty)$	$1.592^{+0.301}_{-0.351}$
l_h	Halo LF break magnitude	mag	$[20.3, 22]$	$20.805^{+0.046}_{-0.052}$
ϵ_{LF}	LF uncertainty parameter	1	$(0, \infty)$	$0.129^{+0.008}_{-0.007}$

Table 5.2: Model Parameters

x	Description	Units	Prior ^a	Results
ϵ_{Choi}	Uncertainty on Choi SB data	1	(0, 2]	$0.086^{+0.005}_{-0.005}$
$\epsilon_{\text{Gilbert}}$	Uncertainty on Gilbert SB data	1	(0, 2]	$0.815^{+0.155}_{-0.124}$
ϵ_{PvdB}	Uncertainty on PvdB SB data	1	(0, 2]	$0.577^{+0.7}_{-0.333}$
ϵ_{Irwin}	Uncertainty on Irwin SB data	1	(0, 2]	$0.086^{+0.005}_{-0.005}$

^a We used a flat prior within the range indicated.

5.8 MCMC Sampler

To estimate the marginalized posterior probability function of the model parameters, we draw samples from the distribution P (Equation 5.29) using a Markov chain Monte Carlo sampler (MCMC; Bishop, 2003; Gelman et al., 2003; Macarthur, 2003; Press et al., 2007). MCMC algorithms offer a method of efficiently drawing unbiased samples from any distribution that can be evaluated for a given set of parameters. In this case, the distribution of interest is the posterior probability which is—up to an unimportant normalization constant—the product of the likelihood function (Equation 5.29) and the prior distributions over the parameters (which we choose to be flat for all parameters). Drawing samples from this distribution is equivalent to drawing a representative sampling of physical models that are consistent with the data and the uncertainties.

The most popular class of MCMC algorithms are based on the Metropolis–Hastings (M–H) method (e.g., Macarthur, 2003). For this project we use a more efficient

ensemble sampler called `emcee` that takes advantage of an affine-invariant proposal distribution in order to sample efficiently even in arbitrarily covariant parameter spaces (Goodman & Weare, 2010; Foreman-Mackey et al., 2015). In many cases, `emcee` requires many fewer computations than required by a standard M–H algorithm to draw the same number of independent samples from a distribution. The advantage comes from the fact that `emcee` uses an affine-invariant algorithm, meaning that it performs equally well regardless of the covariance between parameters. If a M–H chain is to perform well, the proposal distribution—which generally has $\sim D^2$ free parameters, where D is the number of dimensions in the problem—must be tuned to match the covariances in the target density using a computationally expensive “burn-in” phase. In contrast, `emcee` has only 2 free parameters and the proposal density adaptively fits the target density without any fine tuning. This is achieved by simultaneously evolving K coupled MCMC chains or “walkers”. Each walker produces unbiased samples from the target distribution but the proposal distribution for each walker is determined by the current positions of all the other walkers. For more information about this algorithm and the implementation details, see the discussion in Foreman-Mackey et al. (2015).

For this project, we use 256 walkers and keep the `emcee` proposal scale—called a in Foreman-Mackey et al. (2015)—at the default value of 2. Each MCMC run is composed of $\sim 10,000$ “burn-in” steps for each walker and then 100 production steps. The results given in the paper are based on the $256 \times 20 = 51,200$ samples produced after burn-in. In particular, the final column in Table 5.2 reports the sample mean and standard deviation of the parameters in the production chain.

5.9 Parameter Distributions

Figures 5.19 and 5.20 show the 1D posterior probability distributions of all model parameters using samples from the last 20 steps of the MCMC chain. Each distribution is normalized to an area of 1 for display purposes. Descriptions of parameters are given in Table 5.2.

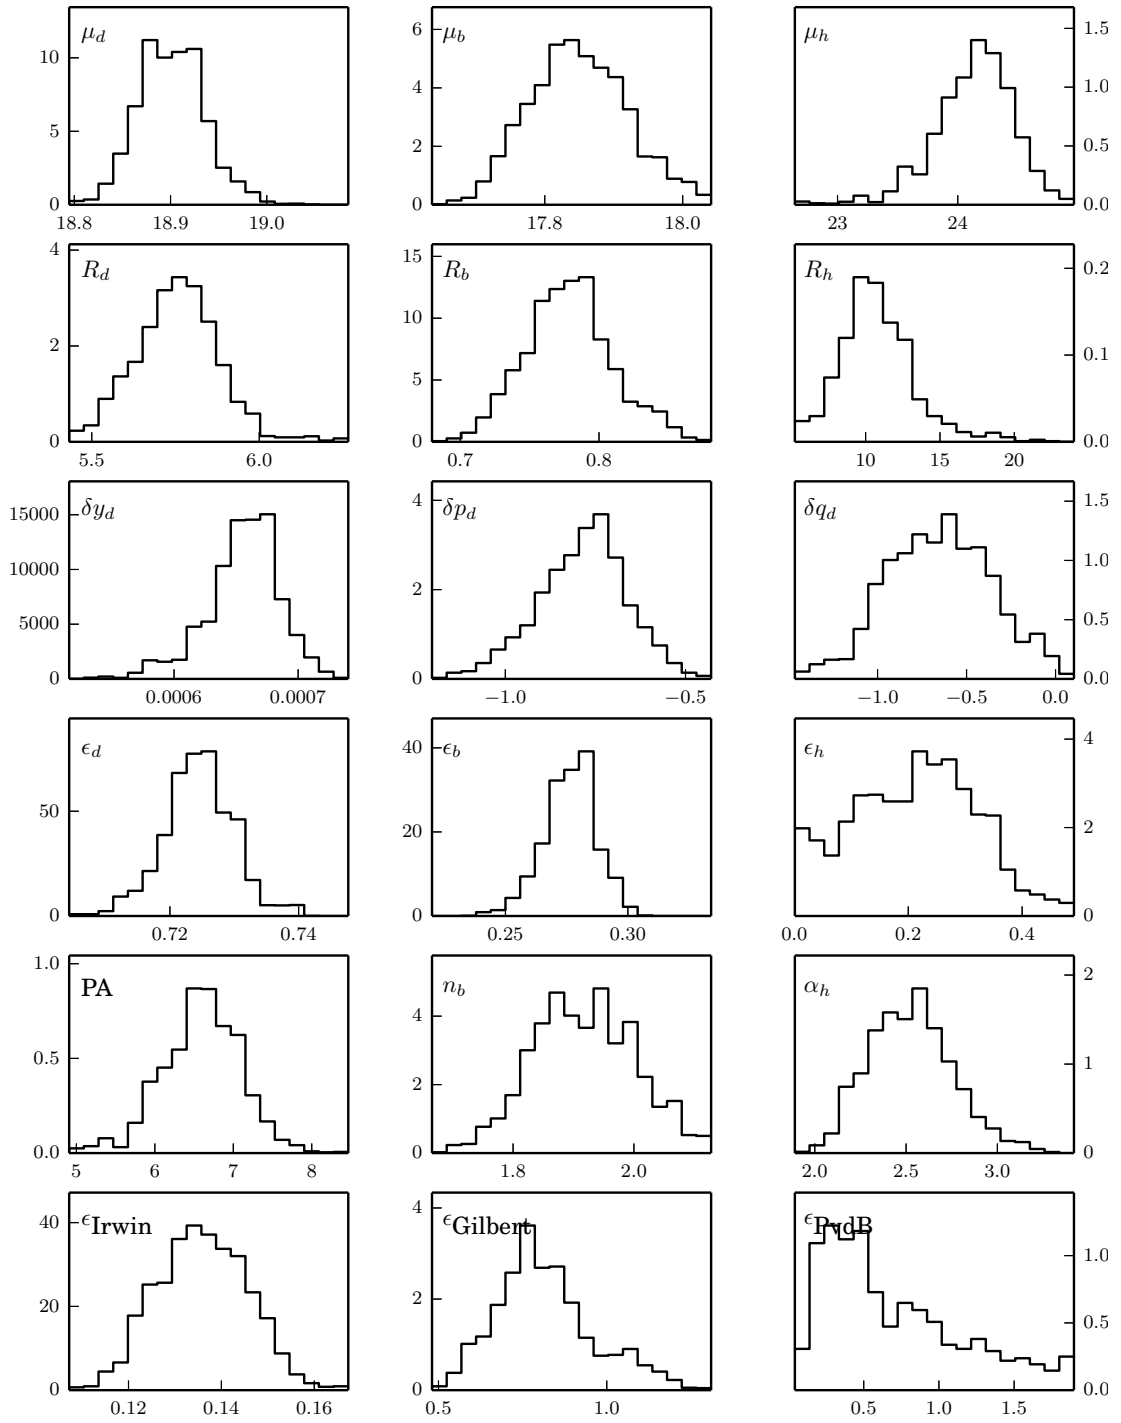


Figure 5.19: Posterior probability distributions of each parameter, as drawn from the last 20 steps of the MCMC chain. Units are given in Table 5.2. With the exception of the third row, parameters describing the disk, bulge, and halo are shown in the left, middle, and right columns, respectively.

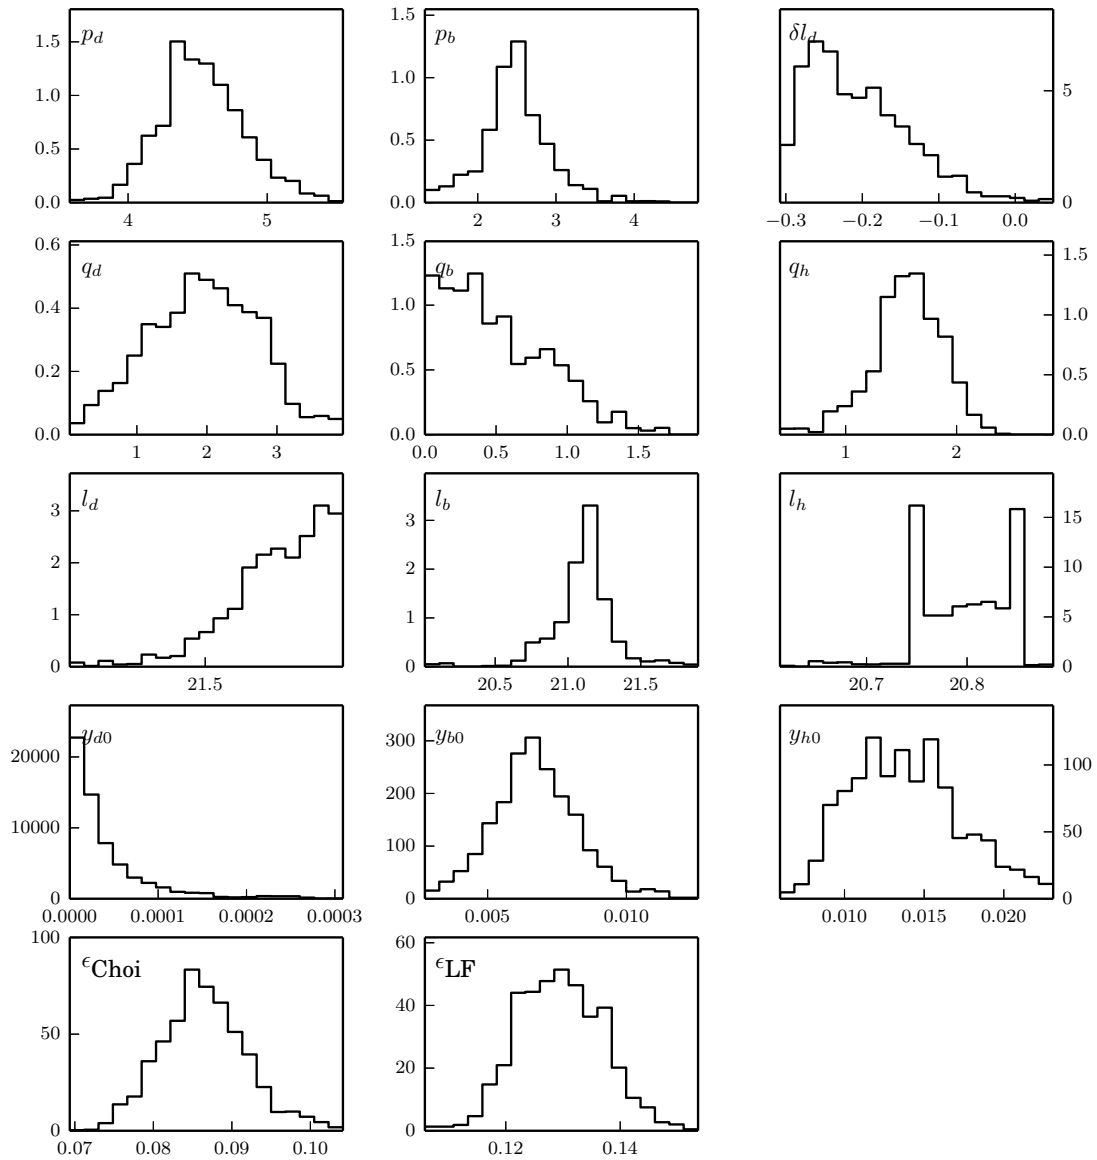


Figure 5.20: Figure 5.19, continued.

Chapter 6

Rotation curves, velocity dispersion profiles, and the asymmetric drift

6.1 INTRODUCTION

The degree of heating of a stellar disk can be parameterized by its velocity dispersion, as in Chapter 3 and Dorman et al. (2015). But an inflated velocity dispersion is not the only effect of perturbations to a disk: at a given location, the *mean* of the stellar velocity distribution is also expected to lag behind the local circular velocity (Binney & Tremaine, 2008). The magnitude of this lag is called the asymmetric drift v_a . Asymmetric drift is a result of conservation of angular momentum $l = mvR$ of a stellar orbit: a star currently at R whose has moved outwards from its guiding center

(average location) $R' < R$ has a lower velocity than the circular velocity v_c that it would have had at its guiding center. Because of the number density gradient in a stellar disk, there are many more stars at a location R that originated at lower radii than higher radii. Therefore, there are more stars with low orbital velocities $v < v_c$ than $v > v_c$, and the mean of the local velocity distribution at R is lower than the circular velocity there.

The Jeans equation for an axisymmetric disk says that, at a given radius, the asymmetric drift is proportional to the square of the stellar velocity dispersion, and also depends on the axis ratio of the velocity ellipsoid and the radius R . Hence, in a steady state system, the asymmetric drift is a probe of the heating history of the galaxy. This relationship appears to hold in the solar neighborhood of the Milky Way (e.g., Golubov et al., 2013). But the velocity dispersions in M31 are much higher, and the dispersion distribution over the entire disk is much less homogenous than over the small solar neighborhood (Dorman et al., 2015), so it is interesting to measure the asymmetric drift in M31 and see if it also behaves as predicted.

Here we use the neutral hydrogen gas (HI) rotation curve from Chemin et al. (2009) as a proxy for the circular velocity profile in M31's disk. HI is not expected to exhibit significant asymmetric drift: unlike stars, it dissipates the dynamical energy from heating events or perturbations, so that its velocity dispersion remains small and thus its asymmetric drift is negligible (assuming enough time has passed since the last significant heating event). We assume that the HI has disk-like orbits everywhere (rather than, say, bar-like orbits) so that we can deproject the line-of-sight velocity distribution to a

circular velocity distribution. However, we will see in Section 6.5 that the assumption that the HI traces the circular velocity may not be warranted everywhere in the disk.

In this chapter, we measure the rotation curve and dispersion profile of the stars in three different age bins. We also measure those same quantities for the HI, and explore the relationship between the stellar velocity dispersion and the asymmetric drift.

6.2 Data

We employ four sets of data in this study: optical HST/ACS photometry of about 8000 individual stars in M31 from the PHAT survey (Dalcanton et al., 2012); Keck/DEIMOS spectroscopy of those same 10,000 stars from the SPLASH survey (Dorman et al., 2012; Howley et al., 2013; Dorman et al., 2013; Dorman et al., 2015); CFHT/MegaCam i' photometry of about 5000 of the SPLASH targets (Gilbert et al., 2009; Dorman et al., 2012; Howley et al., 2013); and 21-cm HI data from the Dominion Radio Astrophysical Observatory (Chemin et al., 2009). All these data have been described in detail elsewhere; here, we briefly summarize the relevant properties of each data set.

6.2.1 Stars: Photometry

Most of the stars in our sample fall within the PHAT survey footprint. For these stars, we have HST photometry in six filters ranging from the ultraviolet to the near-infrared. For this chapter, we make the use of the $F814W$ and $F475W$ filters,

which are roughly comparable to Cousins i and g bands, respectively.

Many stars, including about 20% of those in the PHAT footprint and all of those outside it (those on the SE minor axis or to the south of the galactic center), have single-filter i' photometry from CFHT/MegaCam. This photometry is described in more detail in Howley et al. (2013).

6.2.2 Stars: Spectroscopy

We used the DEIMOS multi-object spectrograph on Keck II between 2007 and 2012 to obtain optical spectra and measure radial velocities of over 10,000 of the stars with PHAT or MegaCam photometry. In 2011 and 2012, when PHAT photometry was available prior to spectroscopy, we targeted stars of a wide range of evolutionary stage: from young massive upper main sequence stars to old red giant branch stars. Each targeted star was brighter than $F814W = 22$ or $F475W = 24$. The target selection scheme is described in Dorman et al. (2013) and Dorman et al. (2015). When PHAT photometry was not available, we targeted stars in the magnitude range $20 < i' < 22$ using a target selection scheme described in Dorman et al. (2012) and Howley et al. (2013).

Raw spectra were processed and then cross-correlated against template spectra to measure radial velocities as described in Dorman et al. (2012, 2013); Dorman et al. (2015). Likely Milky Way foreground contaminants were removed using a combination of photometric and spectroscopic diagnostics as described in Dorman et al. (2015).

6.2.3 Atomic Hydrogen Gas

We use the 21 cm HI velocity field from Chemin et al. (2009). These data were taken with the Synthesis Telescope and the 26 m antenna at the Dominion Radio Astrophysical Observatory. Structures down to the resolution limit of $58'' \times 58'' / \sin(\delta)$ are resolved. For this paper, we use only the lines of sight that align with the locations of the stars in our spectroscopic sample.

There are multiple peaks in the HI spectrum at each location, due in part to a disk warp and the inclination of the system. Sometimes up to five peaks are seen. At each location, we, as in Chemin et al. (2009) choose the component that has the largest velocity relative to v_{sys} , while avoiding isolated faint features. This results in a velocity field without any discontinuities.

6.3 Age Groups

6.3.1 Using PHAT photometry: Summary of previous work

We divide the M31 stars with PHAT photometry into 3 age groups (young MS, intermediate-age AGB, and old RGB) using the photometric method described in Dorman et al. (2015). (Here we combine the younger AGB and older AGB into a single AGB bin.) Figure 6.1 shows the age group boundaries in the optical CMD. As in Dorman et al. (2015), we estimate the average ages of the age bins using a simple simulated CMD that assumes a constant SFR and a MDF derived from the PHAT RGB sample. The average ages of the MS+, AGB, and RGB age groups are 30 Myr, 1 Gyr, 4 Gyr.

6.3.2 Using CFHT-based photometry

We do not have PHAT photometry for the 1/3 of the stars that fall outside the PHAT footprint. For these stars, we only have single-filter ground-based i' CFHT/MegaCam photometry. Most of the stars can be expected to be RGB or AGB stars, since they were red enough to be very bright in the i' image. To use these data, we need to be able separate them into RGB and AGB categories without the aid of PHAT photometry. We do this classification by finding a diagnostic in (i', T_{eff}) space.

In addition to the i' photometry, we also have an estimate of the effective temperature of the stars directly from their spectra. To derive estimates of T_{eff} from our spectra, we first calibrated their shape. This was done using a spectrum of HD 52005, observed with DEIMOS in April 2013 for calibration purposes. The spectrum of HD 52005 was processed identically to the 600line data, and so has the identical instrumental artifacts. We computed the transformation between our uncalibrated spectrum of HD 52005 and the calibrated spectrum in the X-SHOOTER Spectral Library (Chen et al., 2014). Next, we rebinned the sample of science stars observed with the 1200line configuration onto the wavelength grid of the 600line configuration, and applied the aforementioned transformation to the entire science sample. This process is further outlined in Hamren et al. (2015, in preparation).

Using our calibrated spectra, we determined T_{eff} by cross-correlating against a grid of synthetic spectra from Husser et al. (2013). These models use the PHOENIX stellar atmosphere code to generate spectra with a resolution of $R = 500000$, which we rebinned to match the resolution of our sample. We used a subset of the available

models encompassing the following parameters; $2300 \leq T_{eff} \leq 12000$ in increments of 100-200, $0 \leq \log(g) \leq 4.5$ in increments of 0.5, and $-1.5 \leq Z \leq 0.5$ in increments of 0.5. For each of these models, alpha enhancement ($[\alpha/Fe]$) is 0.

We now use the i' photometry and the T_{eff} estimates to classify the stars into AGB and RGB populations. As a training set, we take stars for which we have both PHAT photometry and CFHT photometry: that is, we know from PHAT to what class each star in the training set should belong. We then train a decision tree to use only the i' photometry and the T_{eff} to do the classification, and apply the resulting classification scheme to the stars without PHAT photometry. Figure 6.2 shows the training set (left panel) and classified CFHT-only data (right panel). The left panel shows that the classification scheme is not perfect: while it is possible to get a very clean sample of RGB stars, the AGB class is more contaminated, with a precision of only 82%.

6.4 Kinematics

6.4.1 Projected velocity and dispersion maps

In Figure 6.3, we present three projections of the stellar data: the raw velocity field, the smoothed line-of-sight (LOS) mean velocity field, and the smoothed LOS velocity dispersion field. The mean velocity μ and dispersion σ fields are calculated at each location using stars within circles of radius $200''$, calculated as in Dorman et al. (2015).

Calculating the mean velocity and velocity dispersion for the gas takes extra care. The beam size for the HI data is fairly large, meaning that some stars will project

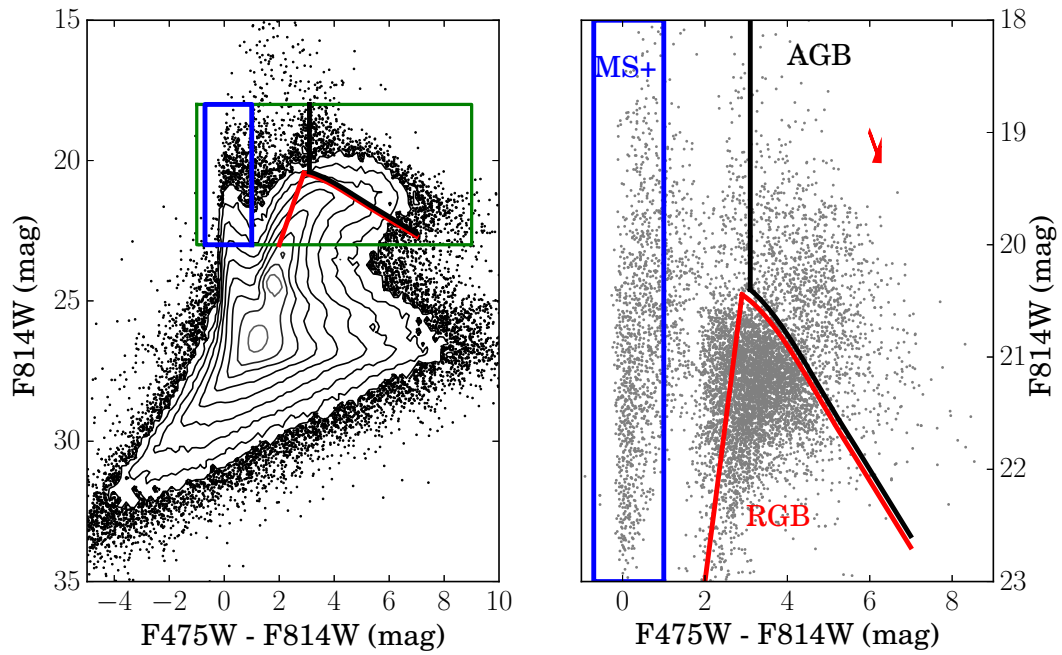


Figure 6.1: *Left*: Optical CMD of a representative region of the PHAT survey. *Right*: Zoom-in of the region of the CMD sampled by the spectroscopic SPLASH survey (green box in left panel). Only stars with reliable photometry *and* radial velocity measurements are shown. The red arrow shows the direction of the MW reddening vector in the direction of M31. The three age bins (young main sequence, intermediate-age AGB, and old RGB) are outlined in both panels. Stars outside the three regions are not included in the kinematical analysis.

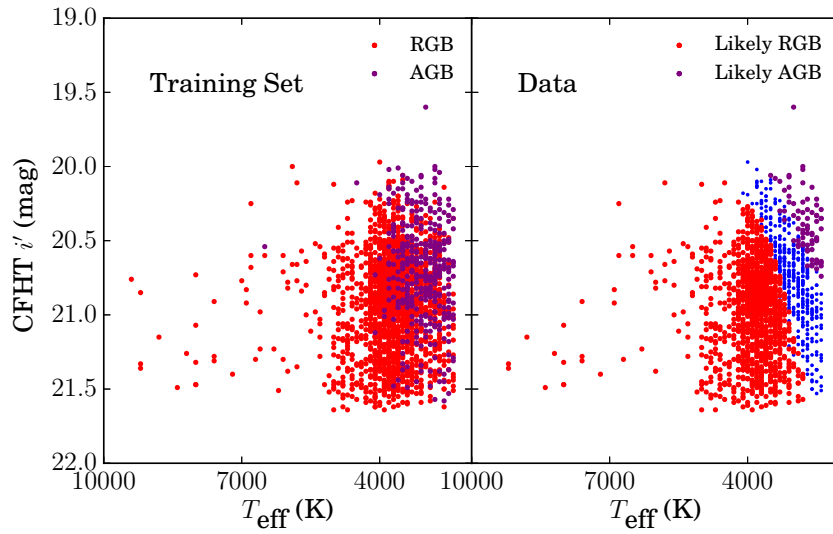


Figure 6.2: Distinguishing between AGB and RGB stars in the sample without multi-band PHAT photometry. *Left:* Training set of stars with both optical PHAT photometry and CFHT/MegaCam i' photometry. PHAT photometry are used to distinguish between RGB (red) and AGB (violet) members. A decision tree algorithm is used to find the decision boundary in (i', T_{eff}) space. *Right:* Set of data without PHAT photometry. The logistic regression curve from the training set is applied to distinguish between evolutionary stages. Stars in the blue region are not used because their membership class is very uncertain.

onto the same pixel and thus end up with the same velocity value, decreasing the measured dispersion and biasing the mean velocity. So we include only distinct velocity measurements when calculating average properties of the gas. The result is that in regions that are densely sampled by the SPLASH survey, the number of data points included in the stellar velocity dispersion and mean velocity calculations is somewhat larger than the number of data points included in the gas dispersion and mean velocity calculations.

In Figure 6.4, we present the mean velocity and velocity dispersion maps for the neutral hydrogen data, which are created in the same way as the stellar maps except that we do not separate the data into age bins because those are not meaningful for the gas.

6.4.2 Deprojection

Assigning P.A. and inclination

In order to deproject the LOS velocity and dispersion to the tangential components (the components in the direction of rotation), we need to know the major axis position angle and the inclination of the disk. The HI disk is known to have a warp: that is, the PA and inclination are not constant over the disk (Chemin et al., 2009; Corbelli et al., 2010). Instead of fitting for the PA and inclination of the stellar disk as a function of position, we adopt the HI PA and inclination values from Chemin et al. (2009), which were measured in rings of constant radius in the plane of the disk. For each star, we assign the PA and inclination of the nearest HI ring.

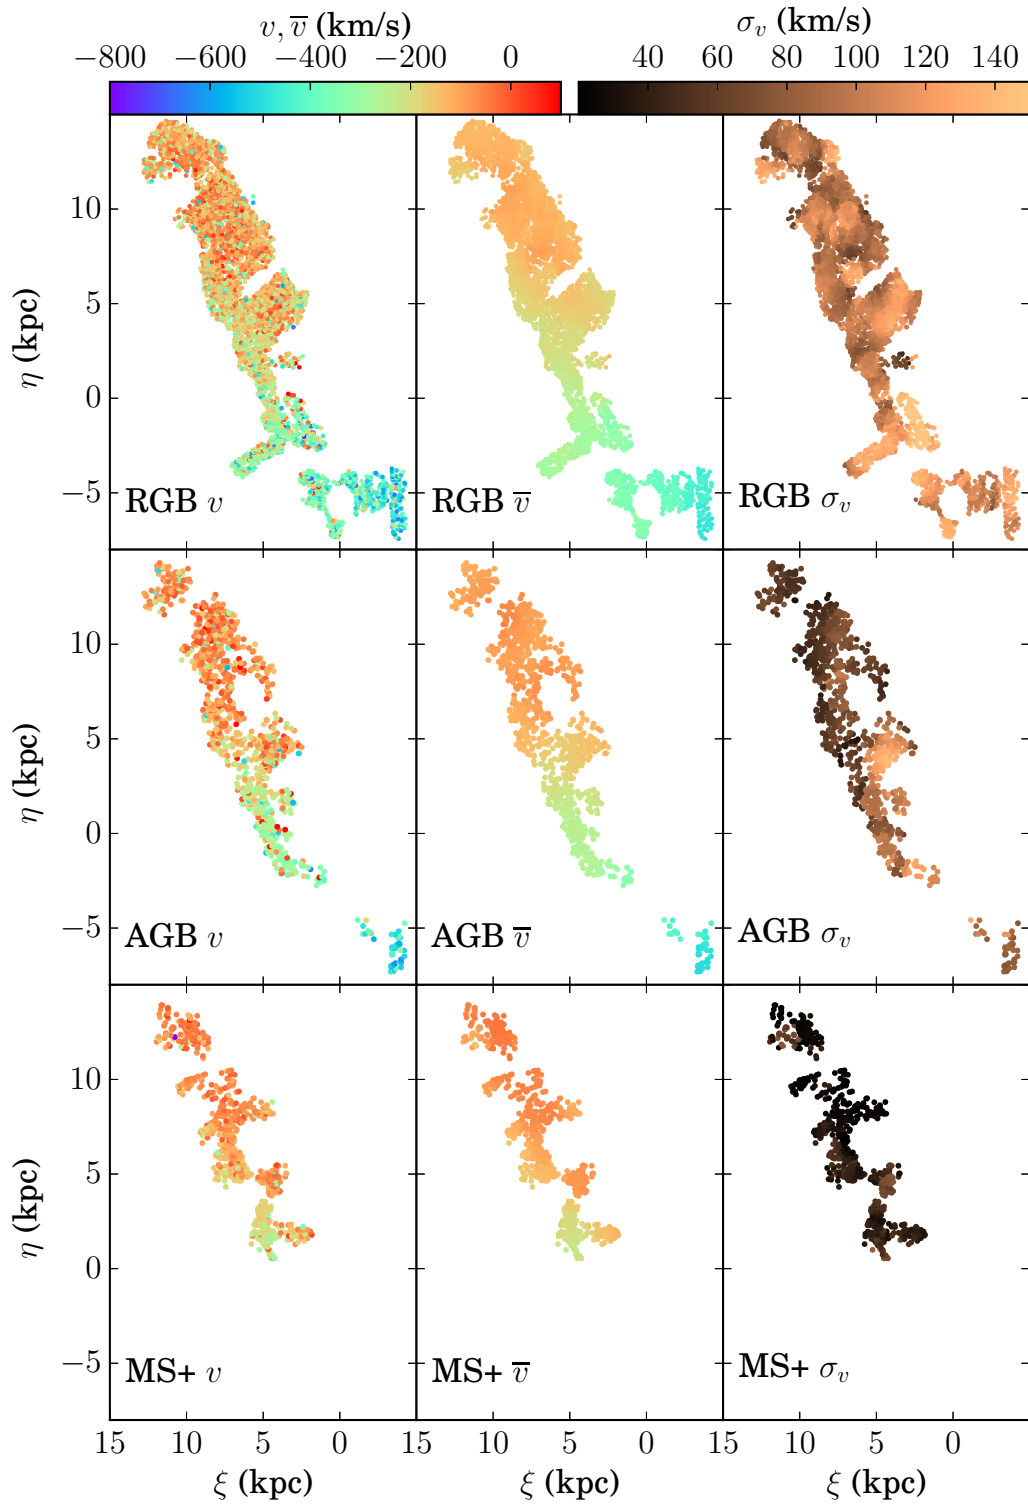


Figure 6.3: Left to right: Raw velocity field, mean velocity field, and velocity dispersion field. Each row corresponds to stars from a different age bin. Smoothing is performed over a $200''$ radius for RGB and MS+ stars and over a $275''$ for the AGB stars.




Figure 6.4: Same as Figure 6.3, for neutral hydrogen data from Chemin et al. (2009).

Rotation curve

To deproject the LOS velocity v_j of each star j to the component in the direction of rotation v_ϕ , we use the following equation:

$$v_{\phi,j} = \pm(v_j - v_{\text{sys}}) \sqrt{\frac{1 + \tan^2(\text{PA} - \text{PA}_j)}{\cos^2 i}} \quad (6.1)$$

where the sign in front of the second term is positive and negative for points to the NE and SW of the minor axis, respectively. $v_{\text{sys}} = -300$ km/s is the systemic velocity of M31's center relative to the MW. The position angle PA_j is the position angle of data point j measured east of north.

We also apply this equation to deproject the mean smoothed velocity μ_j centered on each star j to $\bar{v}_{\text{rot},j}$.

Velocity ellipsoid

When presenting velocity dispersion profiles, we want to take out geometric effects and only present the tangential component, σ_ϕ . To do so we need to know the axis ratio σ_ϕ/σ_R . The velocity ellipsoid of stars in a disk is expected to be anisotropic (Sellwood, 2014, and references therein), and so in general the line-of-sight component of the disk velocity dispersion, σ_{LOS} , will vary with position across the disk. In theory, we can use this variation to measure the shape of the velocity ellipsoid. However, in reality, M31's disk is inclined so close to edge-on that we cannot constrain the vertical component of the dispersion at all, and the kinematical survey does not have high enough target density on the minor axis to accurately measure the tangential velocity ellipsoid axis ratio as a function of radius.

Because we cannot accurately measure the velocity ellipsoid axis ratio, we assume one instead in order to deproject σ_{LOS} to σ_ϕ . In the Milky Way, this axis ratio varies between 0.5 – 0.7, with the younger/colder disk less isotropic (Nordstrom et al., 2004). For this chapter, we assume an axis ratio of 0.7. In Figure 6.5, we show that for the stars within 13° of the major axis, this introduces an uncertainty in σ_ϕ of only $< 5\%$, even if the true axis ratio is actually 0.5. The uncertainty is a monotonic function of distance from the major axis, increasing towards the minor axis.

In §6.6, we perform a by-eye fit to the distribution of points in asymmetric drift-velocity dispersion space to estimate the tangential component of the velocity distribution.




Figure 6.5: Fractional offset in σ_ϕ estimate when we assume a velocity ellipsoid axis ratio of $\sigma_\phi/\sigma_R = 0.5$ vs. 0.7 , as a function of angular offset from the major axis. The choice of axis ratio has no effect on the major axis but increases towards the minor axis. We do not include data with offset $> 13^\circ$ ($> 5\%$ uncertainty in σ_ϕ).

6.5 Rotation curve and dispersion profiles

Figure 6.6 shows the rotation curves (v_ϕ vs. radius) and dispersion profiles (σ_ϕ vs. radius) for the stars (colored points) and gas (gray points). Average values are presented in Tables 6.1-6.4. The AGB and RGB stellar disks typically have a lower rotational velocity than the gas (“lag” the gas disk) and a higher velocity dispersion. The MS+ stellar disk also has a higher dispersion than the gas disk along the same lines of sight, but its rotational velocity is comparable to that of the gas. A few features are present in all stellar age bins: the stellar velocity dips and the dispersion increases around $R_{\text{disk}} = 13$ kpc. The gas dispersion increases a bit at this point, but the gas velocity does not decrease in tandem with the stars. The gas velocity and dispersion and stellar velocity and dispersion all increase around $R = 8$ kpc, which on the major axis coincides with the end of the long bar (Athanasoula & Beaton, 2006; Dorman et al., 2015; Gregersen et al., 2015).

The stellar rotation curves and dispersion profiles are complex in a way that is not matched the gas disk. The means of the stellar profiles vary by up to 25% over scales of ~ 4 kpc, while the HI profiles remain relatively smooth. The variations are greater for older stellar populations. These jagged profiles may be due to the fact that the disk geometry adopted from the gas disk is not the correct geometry for the stellar disk.




Figure 6.6: Deprojected rotation curves (*Left*) and tangential dispersion profiles (*Right*) for stars (colored dots) and gas (gray dots). The top, middle, and bottom panels correspond to RGB, AGB, and MS+ stars (or HI gas along those stars' lines of sight), respectively. For the dispersion profile, the tangential anisotropy parameter $\beta \equiv 1 - (\sigma_\phi/\sigma_R)^2$ was assumed to be 0.5.

6.6 Asymmetric drift

Figure 6.6 shows a clear lag between the stellar and gas rotation curves. The lag clearly increases with age bin (average velocity dispersion). The Jeans equation predicts that the asymmetric drift of a stellar population with a uniform scale length, evaluated at a single radius, is proportional to the square of the stellar velocity dispersion. This has been observationally verified in the solar neighborhood of the Milky Way, where the v_a vs. σ_R^2 relationship is linear for a single stellar population and becomes steeper for populations with increasingly higher metallicities. (Golubov et al., 2013).

We measure the asymmetric drift at the location of each star as the difference of the deprojected mean smoothed velocities of the stars and the gas at that location:

$$v_a = \bar{v}_{\phi,\text{gas}} - \bar{v}_{\phi,\text{stars}} \quad (6.2)$$

Tables 6.1-6.3 report the average value of v_a in each 1 kpc-wide ring. Figure 6.7 compares the asymmetric drift to the squared velocity dispersion for each of our three age groups. While we do see the expected correlation, there is significant scatter. Here we show that the scatter can be explained by the range in radii covered by our sample, closely following the derivation in Binney & Tremaine (2008).

The Jeans equation for an axisymmetric stellar disk, evaluated at the midplane, is

$$v_a^2 - 2v_c v_a = \frac{\sigma_R^2}{2v_c} \left[\frac{\sigma_\phi^2}{\sigma_R^2} - 1 - \frac{\partial \ln(\nu \sigma_R^2)}{\partial \ln R} - \frac{R}{\sigma_R^2} \frac{\partial(\overline{v_R v_Z})}{\partial z} \right]. \quad (6.3)$$

where ν is the stellar density at radius R . If we assume that ν is an exponentially declining function of R with scale length R_d , and that σ_R is either constant or exponentially declining with scale length R_d , then this simplifies to

$$v_a^2 - 2v_c v_a \simeq \frac{\sigma_R^2}{2v_c} \left[\frac{\sigma_\phi^2}{\sigma_R^2} + k \frac{R}{R_d} - 1.4 \right] \quad (6.4)$$

where $k = 2$ when σ_R is exponential and $k = 1$ when σ_R is constant. Here, as in Binney & Tremaine (2008), we have assumed that $\sigma_Z^2/\sigma_R^2 \simeq 0.2$, as it is in the Milky Way. This assumption is as good as any, since M31's inclination is such that have no observational constraint on the vertical axis ratio of the velocity ellipsoid.

We solve Equation 6.4 for the asymmetric drift $v_a(\sigma_\phi, R)$, assuming $R_d = 5.76$ kpc (Dorman et al., 2015) and $\sigma_\phi/\sigma_R = 0.8$. We allow for an exponentially declining dispersion profile interior to $R = 10$ kpc, but require a constant dispersion profile external to that, as inferred from Figure 16 in Dorman et al. (2015). (It is clear from Figure 6.6 that this is a strong oversimplification, but it serves as a first order approximation.) The colored points in Figure 6.8 show the distribution of predicted v_a at the observed σ_ϕ in our data set. The points are color coded by radius, since the varying radius is the primary source of scatter about the $v_a \propto \sigma_\phi^2$ relationship. We have also

added Gaussian noise, proportional to the observed velocity errors, to the calculated lag in order to better simulate our data. From a by-eye fit, we find best fit velocity ellipsoid axis ratios of $\sigma_\phi/\sigma_R = 0.8$ for the RGB stars and 1 for the AGB stars. It is clear from the figure that the observed scatter in lag-dispersion space can be adequately modeled simply from the Jeans equation.

There are a few discrepancies between the model and data that can be explained by dynamical effects in the HI disk. For the AGB stars, the model over-predicts the asymmetric drift at low radii relative to the observations. However, this may be due to the fact that the HI does not trace the circular velocity at low radii. Figure 6.6 shows that the HI velocity and dispersion both increase at large radii. The inflated gas dispersion probably means that the gas is not tracing a relaxed disk: it may be influenced by bar dynamics, as this region appears to coincide with the end of the bar (Athanasoula & Beaton, 2006; Dorman et al., 2015; Gregersen et al., 2015). If the gas is presently being stirred up by the bar, then its velocity *also* lags the circular velocity, and thus a measurement of $\bar{v}_{\text{gas}} - \bar{v}_{\text{stars}}$ is an underestimate of the asymmetric drift $v_c - \bar{v}_{\text{stars}}$. So while there is offset between the model and measured velocity lag at low radii, there may not actually be a significant offset between the model and the true asymmetric drift.

For the RGB stars, the same over prediction phenomenon happens at high radii ($R \sim 20$ kpc). Again, the gas velocity dispersion at high radius is high (possibly because of the warp in the disk), suggesting that the mean gas rotation velocity is less than the true circular velocity and that our measurement is simply a lower limit on the




Figure 6.7: Asymmetric drift as a function of stellar velocity dispersion squared, for stars in different age bins.

true lag.

Table 6.1: RGB Stellar Disk

R_{disk} (kpc)	PA ($^{\circ}$)	i	\bar{v} (km/s)	δv	$\bar{\sigma}$	$\delta\sigma$	$\bar{v} - \bar{v}_{\text{gas}}$
4.5	34.6	67.0	154.42	4.24	112.39	1.77	-81.53
5.5	34.7	67.0	162.24	0.15	123.85	0.1	-71.53
6.5	36.5	72.0	154.62	0.09	121.85	0.06	-82.76
7.5	36.8	74.5	173.25	0.09	98.23	0.09	-68.27

8.5	37.8	74.3	179.98	0.12	100.27	0.1	-63.53
9.5	38.7	74.8	193.5	0.16	100.0	0.11	-48.36
10.5	39.0	75.9	207.85	0.08	84.54	0.07	-34.19
11.5	38.4	76.2	197.76	0.06	92.05	0.05	-47.13
12.5	37.1	75.6	194.1	0.09	99.14	0.06	-57.61
13.5	36.1	74.6	194.61	0.06	99.85	0.04	-63.4
14.5	36.6	74.2	213.36	0.11	91.16	0.07	-43.51
15.5	37.2	74.1	213.27	0.28	74.24	0.29	-45.07
16.5	37.5	75.4	180.41	0.06	102.19	0.06	-79.36
17.5	37.5	74.4	180.58	0.09	95.14	0.08	-76.32
18.5	37.5	73.4	176.68	0.1	95.64	0.09	-76.79
19.5	37.6	73.4	185.97	0.33	107.76	0.52	-72.34

Table 6.2: AGB Stellar Disk

R_{disk} (kpc)	PA ($^{\circ}$)	i	\bar{v} (km/s)	δv	$\bar{\sigma}$	$\delta\sigma$	$\bar{v} - \bar{v}_{\text{gas}}$
5.5	34.6	67.0	173.31	0.58	115.29	0.78	-56.57
6.5	36.5	72.0	167.52	0.48	122.05	0.35	-68.31
7.5	36.9	74.5	181.16	0.53	112.25	0.35	-65.75
8.5	37.8	74.3	185.86	1.03	99.66	1.4	-60.82
9.5	38.6	74.8	204.24	0.99	92.28	1.16	-42.81
10.5	39.0	75.8	227.95	0.44	83.0	0.46	-14.18
11.5	38.4	76.2	244.85	0.35	65.01	0.17	1.23
12.5	37.0	75.6	223.12	0.38	69.55	0.11	-26.45
13.5	36.1	74.6	234.66	0.19	64.75	0.1	-20.84
14.5	36.6	74.1	247.79	0.11	55.84	0.09	-8.82
15.5	37.2	74.1	256.13	0.38	42.92	0.35	-0.14
16.5	37.5	75.4	231.02	0.31	52.65	0.22	-27.07
17.5	37.5	74.3	223.35	0.16	54.6	0.13	-35.45

18.5	37.5	73.5	213.12	0.42	55.99	0.42	-42.46
19.5	37.7	73.4	211.48	0.48	66.51	0.63	-47.9

Table 6.3: MS+ Stellar Disk

R_{disk} (kpc)	PA ($^{\circ}$)	i	\bar{v} (km/s)	δv	$\bar{\sigma}$	$\delta\sigma$	$\bar{v} - \bar{v}_{\text{gas}}$
3.5	34.6	67.0	251.31	0.06	37.67	0.0	49.68
6.5	36.5	72.0	238.49	0.21	37.98	0.06	6.17
7.5	36.8	74.5	255.81	0.8	35.77	0.22	16.15
8.5	37.6	74.4	241.41	2.85	32.71	0.6	3.4
9.5	38.7	74.8	253.23	0.87	39.25	0.29	8.95
10.5	39.0	75.8	243.5	0.25	31.51	0.25	-6.07
11.5	38.6	76.3	255.6	0.28	29.97	0.16	8.08
12.5	37.0	75.6	259.14	0.64	30.4	0.25	5.05
13.5	36.1	74.6	281.6	1.67	41.67	0.13	42.25
14.5	36.5	74.3	265.19	1.11	36.76	0.64	27.12
15.5	37.2	74.1	275.29	0.19	27.59	0.1	23.04
16.5	37.5	75.4	285.42	0.15	24.78	0.06	30.55
17.5	37.5	74.5	294.66	1.72	20.1	1.14	42.93

Table 6.4: HI Disk

R_{disk} (kpc)	\bar{v} (km/s)	δv	$\bar{\sigma}$	$\delta\sigma$
19.5	198.72	0.4	47.4	0.09
3.5	186.43	0.15	44.42	0.19
4.5	218.8	0.32	26.36	0.64

5.5	233.63	0.03	22.66	0.02
6.5	236.46	0.01	24.8	0.01
7.5	243.06	0.02	26.84	0.02
8.5	244.44	0.04	26.95	0.04
9.5	236.8	0.03	15.89	0.02
10.5	241.68	0.02	12.31	0.01
11.5	247.66	0.01	11.75	0.01
12.5	255.11	0.01	9.85	0.01
13.5	257.54	0.01	11.72	0.01
14.5	258.91	0.01	11.57	0.01
15.5	259.48	0.01	6.86	0.01
16.5	260.47	0.01	5.47	0.0
17.5	259.23	0.02	9.15	0.01
18.5	253.45	0.01	13.91	0.03
19.5	247.67	0.11	16.92	0.04

6.7 Summary

In this Chapter, we have used the line-of-sight velocity distribution of stars in Andromeda’s disk in tandem with the previously measured geometry of the HI disk to measure the stellar rotation curve and azimuthal velocity dispersion profile at radii between 3 and 20 kpc from the galactic center. We have made these measurements for stars in each of three age bins: young, massive upper main sequence stars; intermediate-age asymptotic giant branch stars; and old red giant branch stars.

Using HI gas kinematics from (Chemin et al., 2009) as a tracer of the circular




Figure 6.8: Predicted lag-dispersion relation as computed via Equation 6.4. Points are color-coded by deprojected radius in the plane of the disk. Lines are the average lag-dispersion relationship from Figure 6.7. The model generally reproduces the observed lag-dispersion relationship. It overpredicts the lag a bit at very high and very low radii; these may be a result of the HI not tracing the circular velocity in the dynamically disturbed warp and bar regions, respectively.

velocity, we have measured the asymmetric drift of each of the three stellar populations as a function of velocity dispersion. For each population, the $v_a - \sigma_\phi^2$ relationship has more scatter than is observed in the Milky Way, but the scatter can be easily accounted for by the large range of radii observed in M31 compared to the solar neighborhood in our own galaxy. At large radii, the relationship breaks down; this is likely a result of the warp in the gas disk beyond 18 kpc (Chemin et al., 2009).

Chapter 7

Conclusions and Future

Directions

In this thesis, I have presented data and results from a comprehensive resolved stellar spectroscopic survey of the disk-dominated region of the Andromeda galaxy. The survey, an extension of the SPLASH survey of M31's halo, includes radial velocities of over 10,000 individual stars of a variety of ages and metallicities. The ages and metallicities are measured by high-resolution, six-filter HST imaging from the PHAT survey.

Using this new data set, I have explored the structure and dynamics of Andromeda's disk-dominated region, and found evidence for dynamical heating from multiple sources. I measured the overall properties of the stars in the disk-dominated region in two ways: first by measuring the stellar age-velocity dispersion correlation (Chapter 3), and second by reporting the overall rotation curve and tangential velocity dispersion

profile as a function of stellar age (Chapter 6). Both show strong evidence for recent dynamical effects. The continuous stellar age-velocity dispersion correlation implies that a continuous heating or cooling process has occurred over at least the last 6 Gyr. This heating or cooling mechanism, though, is at least twice as strong as anything seen in either the Milky Way or existing models (which is partly by construction, since many models have been tuned to match the Milky Way). The stellar rotation curve and dispersion profiles are far less smooth – even in their medians – than the corresponding HI profiles, indicating that the stars no longer trace the geometry and/or kinematics of the gas disk.

I identified a spatially inhomogeneous population of RGB stars with halo-like (high velocity dispersion) kinematics. This population manifests itself in several ways. It inflates the overall velocity dispersion of RGB stars in some patches relative to that of the underlying disk (Chapter 3). It appears as a low-velocity tail, with mean velocity $\sim 50 \text{ km s}^{-1}$ (in the same direction as the disk’s rotation) and velocity dispersion $\sim 150 \text{ km s}^{-1}$, in the stellar line-of-sight velocity distribution (Chapter 4). The number of stars in the low-velocity tail is much higher than the number of stars with a bulge- or halo-like luminosity function, implying that many of the hot component members originated in the disk and were dynamically heated later by satellites or the bar (Chapter 5). The dynamically hot population is scattered all over the disk. The portion far from the galactic center may have been “kicked out” from satellite impacts. But there is also a concentration of velocity outliers in a patch about 5 – 6 kpc from the galactic center. This patch coincides with the end of the long bar, suggesting that

the inflated velocity dispersion is due to a superposition of disk and bar members. The neutral hydrogen gas kinematics are also disturbed in this location, further implying that a current (rather than past) perturbation (such as a bar) is disturbing the orbits of disk stars (Chapter 6).

I also see evidence for an accreted population in the form of tidal stream debris: the first kinematical detection of the northern extension of the Giant Southern Stream.

I have also studied the structure of the old (RGB + AGB) stellar disk, bulge, and halo separately via a detailed structural decomposition that incorporates kinematical, resolved photometric, and integrated-light data sets (Chapter 5). I have presented average central surface brightnesses, scale lengths, ellipticities, and major axis position angles for the bulge, disk, and halo. I have found that the disk metallicity decreases with radius from high-metallicity at the center to bulge-like at $R = 20$ kpc .

The high velocity dispersion and existence of a spatially inhomogenous metal-rich, dynamically hot population support the existing evidence that M31's merger history may have been more violent than the Milky Way's, and more in line with Λ CDM cosmological predictions. However, more modeling is needed to confirm this: for example, can cosmological simulations of interactions between satellites and disks such as Purcell et al. (2010) that produce velocity dispersion similar to that of M31 also produce the observed continuous age-velocity dispersion relationship? If the age-velocity dispersion relationship is the result of disk settling during formation, rather than heating, then the initial conditions of the disk in simulations like Purcell et al. (2010) need to be modified; is the resulting post-merger simulated disk still similar to M31, or is it

too hot?

There is much work that needs to be done in order to better understand the detailed structure of M31's disk as well as the constraints that these observations place on the disk formation, evolution, and merger history. Possible directions include:

- Repeat the structural decomposition described in Chapter 5, while modeling the bar in addition to the bulge, disk, and halo. The concentration of velocity outliers in every age group in the “Brick 9” region that coincides with the end of the bar suggests that many of the “kicked-out” stars in this region may be either bar members or disk members whose orbits have been disturbed by the bar, rather than stars heated by satellite impacts. To explore this effect, one could add a fourth component, corresponding to a long bar, to the structural decomposition model. This addition would require detailed modeling of the expected LOS kinematics in a subregion that includes both bar stars and unperturbed disk stars: the distribution likely cannot be modeled as a single Gaussian.
- Repeat the asymmetric drift analysis of Chapter 6 including ionized gas kinematics, in order to better understand the very young disk. The SPLASH data set contains serendipitous detections of ionized gas ($H\alpha$, SII, and NII) in places where part of a slit happened to cross the line of sight of a diffuse gas cloud. One could measure radial velocities for these serendipitous detections and measure the lag between ionized gas and neutral hydrogen, and between stars and ionized gas. Preliminary studies of HII regions in Andromeda suggest that the kinematics of the ionized gas are much less regular than those of the neutral hydrogen (Nelson

Caldwell, private communication).

- Compare the observed M31 disk velocity distribution — including the age-dispersion relationship, the clumpy spatial distribution of the dynamically hot population, the velocity dispersion map, and the velocity lag — to a suite of cosmological models to constrain the formation and accretion history of the inner parts of the galaxy. In addition to the velocity distribution, we can require that the simulations also reproduce the disk scale height as measured from the PHAT RGB population in Dalcanton et al (2015, submitted). A challenge will lie in choosing the initial conditions for the disk: disks that form thin will require more mergers to attain their present-day velocity dispersion than disks that formed from a thick, settling gas disk.
- Search for evidence for or against the existence of a distinct thick disk (that with a velocity dispersion higher than that predicted by the overall age-dispersion relationship). This exercise will require greater age resolution than we have from the PHAT photometry alone; the age-metallicity degeneracy in the PHAT data prevent us from isolating a clean, extremely old RGB sample. One option is to obtain much cleaner spectra with higher S/N to better estimate spectroscopic metallicities and better constrain stellar ages.

All of these projects will enable a more detailed understanding of the structure and evolution of M31 and the Local Group. But, just as the Milky Way's history is not representative of that of both the large spirals in the Local Group, our Local Group is not necessarily representative of galaxy formation overall. As more powerful telescopes

come online, studying the resolved stellar kinematics of a larger set of galaxy disks will allow more meaningful statistical comparisons to cosmological simulations and place stronger constraints on the physical processes that influence the evolution of galactic disks.

Bibliography

Abadi, M. G., & Navarro, J. F. 2003, *The Astrophysical Journal*, 21

Athanassoula, E., & Beaton, R. L. 2006, *Monthly Notices of the Royal Astronomical Society*, 370, 1499

Beaton, R. L., et al. 2007, *The Astrophysical Journal*, 658, 91

Bernard, E. J., et al. 2012, *Monthly Notices of the Royal Astronomical Society*, 420, 2625

Binney, J., & Tremaine, S. 2008, *Galactic Dynamics: Second Edition* (Princeton University Press)

Bird, J. C. 2015, in *American Astronomical Society Meeting Abstracts*, Vol. 225, American Astronomical Society Meeting Abstracts, 227.03

Bird, J. C., Kazantzidis, S., Weinberg, D. H., Guedes, J., Callegari, S., Mayer, L., & Madau, P. 2013, *The Astrophysical Journal*, 773, 43

Bishop, C. M. 2003, *Pattern Recognition and Machine Learning* (Springer)

- Bournaud, F., Elmegreen, B. G., & Martig, M. 2009, *The Astrophysical Journal*, 707, L1
- Bovy, J., Rix, H.-W., & Hogg, D. W. 2012, *The Astrophysical Journal*, 751, 131
- Bressan, A., Marigo, P., Girardi, L., Salasnich, B., Dal Cero, C., Rubele, S., & Nanni, A. 2012, *Monthly Notices of the Royal Astronomical Society*, 427, 127
- Brown, T. M., Ferguson, H. C., Smith, E., Kimble, R. A., Sweigart, A. V., Renzini, A., Rich, R. M., & Vandenberg, D. A. 2003, *The Astrophysical Journal*, 592, 17
- Brown, T. M., Smith, E., Ferguson, H. C., Rich, R. M., Guhathakurta, P., Renzini, A., Sweigart, A. V., & Kimble, R. A. 2006, *The Astrophysical Journal*, 652, 323
- Büdenbender, A., van de Ven, G., & Watkins, L. L. 2014, *ArXiv e-prints*
- Bullock, J. S., & Johnston, K. V. 2005, *The Astrophysical Journal*, 635, 931
- Caldwell, N., Morrison, H., Kenyon, S. J., Schiavon, R., Harding, P., & Rose, J. a. 2010, *The Astronomical Journal*, 139, 372
- Cappellari, M., et al. 2007, *Monthly Notices of the Royal Astronomical Society*, 379, 418
- Carollo, D., et al. 2007, *Nature*, 450, 1020
- . 2010, *The Astrophysical Journal*, 712, 692
- Chapman, S. C., Ibata, R., Lewis, G. F., Ferguson, A. M. N., Irwin, M., Mconnachie, A., & Tanvir, N. 2006, *The Astrophysical Journal*, 653, 255

- Chapman, S. C., et al. 2008, *Monthly Notices of the Royal Astronomical Society*, 1452, 1437
- Chemin, L., Carignan, C., & Foster, T. 2009, *The Astrophysical Journal*, 705, 1395
- Chen, Y., Trager, S., Peletier, R., Lancon, A., Vazdekis, A., Prugniel, P., Silva, D. R., & Gonneau, A. 2014
- Chiba, M., & Beers, T. I. C. 2000, *The Astronomical Journal*, 119, 2843
- Choi, P., Guhathakurta, P., & Johnston, K. V. 2002, *The Astronomical Journal*, 124, 310
- Collins, M. L. M., et al. 2011, *Monthly Notices of the Royal Astronomical Society*, 413, 1548
- Cooper, a. P., et al. 2010, *Monthly Notices of the Royal Astronomical Society*, 406, 744
- Corbelli, E., Lorenzoni, S., Walterbos, R., Braun, R., & Thilker, D. 2010, *Astronomy and Astrophysics*, 511, A89
- Courteau, S. 1996, *Astrophysical Journal*, 103, 363
- Courteau, S., Widrow, L. M., McDonald, M., Guhathakurta, P., Gilbert, K. M., Zhu, Y., Beaton, R. L., & Majewski, S. R. 2011, *The Astrophysical Journal*, 739, 20
- Dalcanton, J. J., et al. 2012, *The Astrophysical Journal Supplement Series*, 200, 18
- Davis, M., et al. 2002, 12

- Deason, a. J., Belokurov, V., Koposov, S. E., & Rockosi, C. M. 2014, *The Astrophysical Journal*, 787, 30
- Dorman, C. E., et al. 2012, *The Astrophysical Journal*, 752, 147
- . 2013, *The Astrophysical Journal*, 779, 103
- Dorman, C. E., et al. 2015, *The Astrophysical Journal*, 803, 24
- Fardal, M. A., Babul, A., Guhathakurta, P., Gilbert, K. M., & Dodge, C. 2008, *The Astrophysical Journal*, 682, 33
- Fardal, M. a., Guhathakurta, P., Babul, a., & McConnachie, a. W. 2007, *Monthly Notices of the Royal Astronomical Society*, 380, 15
- Fardal, M. A., et al. 2013, *MNRAS*, 28, 1
- Ferguson, A. M. N., Irwin, M. J., Ibata, R. A., & Tanvir, N. R. 2002, *The Astronomical Journal*, 124, 1452
- Font, A. S., McCarthy, I. G., Crain, R. A., Theuns, T., Schaye, J., Wiersma, R. P. C., & Dalla Vecchia, C. 2011, *MNRAS*, 416, 2802
- Forbes, J., Krumholz, M., & Burkert, A. 2012, *The Astrophysical Journal*, 754, 48
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2015, 125, 306
- Geha, M., Guhathakurta, P., Rich, R. M., & Cooper, M. C. 2006, *The Astronomical Journal*, 131, 332

- Geha, M., van der Marel, R. P., Guhathakurta, P., Gilbert, K. M., Kalirai, J., & Kirby, E. N. 2010, *The Astrophysical Journal*, 711, 361
- Gelman, A., Carlin, J. B., Stern, H. S., Rubin, D. B., & Dunson, D. B. 2003, *Bayesian Data Analysis* (Chapman & Hall)
- Gilbert, K. M., Font, A. S., Johnston, K. V., & Guhathakurta, P. 2009, *The Astrophysical Journal*, 701, 776
- Gilbert, K. M., et al. 2006, *The Astrophysical Journal*, 652, 1188
- Gilbert, K. M., et al. 2007, *The Astrophysical Journal*, 668, 245
- Gilbert, K. M., et al. 2012, *The Astrophysical Journal*, 760, 76
- Girardi, L., et al. 2010, *The Astrophysical Journal*, 724, 1030
- Golubov, O., Just, A., Bienaymé, O., Gibson, B. K., Grebel, E. K., & Munari, U. 2013, *Astronomy & Astrophysics*, 92, A92
- Goodman, J., & Weare, J. 2010, *Communications in Applied Mathematics and Computational Science*, 5
- Gregersen, D., et al. 2015, *Astronomical Journal*, submitted
- Guedes, J., Callegari, S., Madau, P., & Mayer, L. 2011, *The Astrophysical Journal*, 742, 76
- Guhathakurta, P., van Gorkom, J. H., Kotanyi, C. G., & Balkowski, C. 1988, *Astronomical Journal*, 96, 851

- Guhathakurta, P., et al. 1911, 1
- . 2006, *The Astrophysical Journal*, 2497
- Howley, K. M., Geha, M., Guhathakurta, P., Montgomery, R. M., Laughlin, G., & Johnston, K. V. 2008, *The Astrophysical Journal*, 683, 722
- Howley, K. M., et al. 2013, *The Astrophysical Journal*, 765, 65
- Husser, T., Berg, S. W.-v., Dreizler, S., Homeier, D., & Reiners, A. 2013, *Astronomy & Astrophysics*, 6, 1
- Ibata, R., Chapman, S., Ferguson, a. M. N., Irwin, M., Lewis, G., & McConnachie, a. 2004, *Monthly Notices of the Royal Astronomical Society*, 351, 117
- Ibata, R., Chapman, S., Ferguson, A. M. N., Lewis, G., Irwin, M., & Tanvir, N. 2005, *The Astrophysical Journal*, 287
- Ibata, R., Irwin, M., Lewis, G., Ferguson, A., & Tanvir, N. 2001, 1
- Ibata, R., Martin, N. F., Irwin, M., Chapman, S., Ferguson, a. M. N., Lewis, G. F., & McConnachie, a. W. 2007, *The Astrophysical Journal*, 671, 1591
- Ibata, R. A., Gilmore, G., & Irwin, M. J. 1994, *Nature*, 370, 194
- Ibata, R. a., et al. 2014, *The Astrophysical Journal*, 780, 128
- Ida, S., Kokubo, E., & Makino, J. 1993, *MNRAS*, 263, 875
- Irwin, M. J., Ferguson, A. M. N., Ibata, R. A., Lewis, G. F., Tanvir, N. R., & Al, I. E. T. 2005, *The Astrophysical Journal*, 5, 105

- Johnson, L. C., et al. 2012, *The Astrophysical Journal*, 752, 95
- Kalirai, J. S. 2012, *Nature*, 486, 90
- Kalirai, J. S., Guhathakurta, P., Gilbert, K. M., Reitzel, D. B., Majewski, S. R., Rich, R. M., & Cooper, M. C. 2006a, *The Astrophysical Journal*, 641, 268
- Kalirai, J. S., et al. 2006b, *The Astrophysical Journal*, 648, 389
- Kormendy, J. 1982, *Astrophysical Journal*, 257, 75
- Kormendy, J., & Illingworth, G. 1982, *Astrophysical Journal*, 256, 460
- Kormendy, J., & Kennicutt, R. C. 2004, *Annual Review of Astronomy and Astrophysics*, 42, 603
- Kwitter, K. B., Lehman, E. M. M., Balick, B., & Henry, R. B. C. 2012, *The Astrophysical Journal*, 753, 12
- Macarthur, L. A. 2003, 689
- Majewski, S. R., Skrutskie, M. F., Weinberg, M. D., & Ostheimer, J. C. 2003, *Astrophysical Journal*, 599, 1082
- Martinsson, T. P. K., Verheijen, M. A. W., Westfall, K. B., Bershad, M. A., Schechtman-rook, A., Andersen, D. R., & Swaters, R. A. 2013, *Astronomy & Astrophysics*, 557, A130
- Mathewson, D. S., Cleary, M. N., & Murray, J. D. 1974, *The Astrophysical Journal*, 190, 291

- McCarthy, I. G., Font, a. S., Crain, R. a., Deason, a. J., Schaye, J., & Theuns, T. 2012, Monthly Notices of the Royal Astronomical Society, 420, 2245
- McConnachie, A. W., Irwin, M. J., Ferguson, A. M. N., Ibata, R. A., Lewis, G. F., & Tanvir, N. 2005, Monthly Notices of the Royal Astronomical Society, 356, 979
- McConnachie, A. W., et al. 2009, Natur, 461, 66
- Mendez, B., Davis, M., Madore, B. F., & Freedman, W. L. 2002, The Astronomical Journal, 124, 213
- Merrett, H. R., et al. 2006, Monthly Notices of the Royal Astronomical Society, 369, 120
- Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, The Astrophysical Journal, 1, 493
- Newberg, H. J., et al. 2003, 3, 191
- Nordstrom, B., et al. 2004, Astronomy and Astrophysics, 418, 989
- Peng, C. Y., Ho, L. C., Impey, C. D., & Rix, H.-W. 2002, The Astronomical Journal, 124, 266
- Pillepich, A., Madau, P., & Mayer, L. 2015, The Astrophysical Journal, 799, 184
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 2007, Numerical Recipes (Cambridge University Press)
- Pritchett, C.J., van den Bergh, S. 1994, The Astronomical Journal, 107, 1730
- Pryor, C., & Meylan, G. 1993, PASP, 50

- Purcell, C. W., Bullock, J. S., & Kazantzidis, S. 2010, *MNRAS*, 404, 1711
- Purcell, C. W., Kazantzidis, S., & Bullock, J. S. 2009, *The Astrophysical Journal*, 694, L98
- Quinn, P. J., Hernquist, L., & Fullagar, D. P. 1993, *The Astrophysical Journal*, 403, 74
- Reitzel, D. B., Guhathakurta, P., & Gould, A. 1998, *The Astronomical Journal*, 707
- Richardson, J. C., et al. 2008, *The Astronomical Journal*, 135, 1998
- Robin, A. C., Reyl, C., & Derri, S. 2003, *Astronomy & Astrophysics*, 540, 523
- . 2004, *Astronomy & Astrophysics*, 157, 20040968
- Rocha-Pinto, H. J., Majewski, S. R., Skrutskie, M. F., & Crane, J. D. 2003, *The Astrophysical Journal*, 594, 115
- Saglia, R. P., et al. 2010, *Astronomy & Astrophysics*, 524, A6
- Salaris, M., & Cassisi, S. 1997, *Monthly Notices of the Royal Astronomical Society*, 289, 406
- Sanders, N. E., Caldwell, N., McDowell, J., & Harding, P. 2012, *The Astrophysical Journal*, 758, 133
- Sarajedini, A. 2005, *The Astronomical Journal*, 1627
- Schiavon, R. P., Barbury, B., Rossi, S. C. F., & Milone, A. 1997, *The Astrophysical Journal*, 20, 902

- Schönrich, R., & Binney, J. 2012, *Monthly Notices of the Royal Astronomical Society*, 419, 1546
- Seigar, M. S., Barth, A. J., & Bullock, J. S. 2008, *Monthly Notices of the Royal Astronomical Society*, 389, 1911
- Sellwood, J. A. 2014, *Reviews of Modern Physics*, 86, 1
- Sersic, J. L. 1968, *Atlas de galaxias australes*
- Seth, A. C., & Dalcanton, J. J. 2005, *The Astronomical Journal*, 130, 1574
- Sheffield, A. a., et al. 2012, *The Astrophysical Journal*, 761, 161
- Simard, L., et al. 2002, *Astrophysical Journal Supplement*, 142, 1
- Simon, J. D., & Geha, M. 2007, *Astrophysical Journal*, 670, 313
- Sohn, S. T., et al. 2007, *Astrophysical Journal*, 663, 960
- Stetson, P. B. 1994, *PASP*, 106, 250
- Stewart, K. R., Bullock, J. S., Wechsler, R. H., Maller, A. H., & Zentner, A. R. 2008, *The Astrophysical Journal*, 683, 597
- Tanaka, M., Chiba, M., Komiyama, Y., Guhathakurta, P., Kalirai, J. S., & Iye, M. 2010, *The Astrophysical Journal*, 708, 1168
- Tissera, P. B., Scannapieco, C., Beers, T. C., & Carollo, D. 2013, *Monthly Notices of the Royal Astronomical Society*, 432, 3391
- Tollerud, E. J., et al. 2012, *Astrophysical Journal*, 752, 45

- Vanhollebeke, E., Groenewegen, M. A. T., & Girardi, L. 2009, MNRAS, 107, 95
- Velazquez, H., & White, S. D. M. 1999, MNRAS, 270, 254
- Williams, B. F., et al. 2012, The Astrophysical Journal, 759, 46
- . 2014, The Astrophysical Journal Supplement Series, 215, 9
- Worthey, G., Espan, A., & Macarthur, L. A. 2005, The Astrophysical Journal, 631, 820
- Yanny, B., et al. 2003, The Astrophysical Journal, 588, 824
- Yoachim, P., & Dalcanton, J. J. 2006, The Astronomical Journal, 131, 226
- Zolotov, A., Willman, B., Brooks, A. M., Governato, F., Hogg, D. W., Shen, S., & Wadsley, J. 2010, Astrophysical Journal, 721, 738
- Zurita, A., & Bresolin, F. 2012, Monthly Notices of the Royal Astronomical Society, 427, 1463