NOTE TO USERS

This reproduction is the best copy available.

UMI®

ALIGNMENT OF GALACTIC COMPONENTS IN MODELS OF GALAXY FORMATION

by

Jeremy Bailin

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF ASTRONOMY

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

2004

UMI Number: 3158068

INFORMATION TO USERS

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleed-through, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

UMI®

UMI Microform 3158068

Copyright 2005 by ProQuest Information and Learning Company.
All rights reserved. This microform edition is protected against
unauthorized copying under Title 17, United States Code.

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

The University of Arizona ®
Graduate College

As members of the Final Examination Committee, we certify that we have read the

dissertation prepared by JEREMY BAILIN

entitled ALIGNMENT OF GALACTIC COMPONENTS IN MODELS OF GALAXY FORMATION

and recommend that it be accepted as fulfilling the dissertation requirement for the

Degree of Doctor of Philosophy

	9/10/04 date
DANIEL EISENSTEIN	9/10/04 date
MATTHIAS STEINMETZ	date
EDWARD OLSZEWSKI	9/10/04 date
PHILIP PINTO	date
MARCIA RIEKE	9/10/04

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

Dissertation Director: DANIEL EISENSTEIN MATTHIAS STEINMETZ — 10/12/04 date

STATEMENT BY AUTHOR

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at The University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgment of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the head of the major department or the Dean of the Graduate College when in his or her judgment the proposed use of the material is in the interests of scholarship. In all other instances, however, permission must be obtained from the author.

SIGNED: ______

ACKNOWLEDGMENTS

First, I would like to thank my three loves: Sharon Bailin, Alyson Ford, and Trajan, for motivating me, inspiring me, and keeping me sane.

This work could not have been done without the incredible support of my advisor, Matthias Steinmetz, whose suggestions, comments, and questions were crucial ingredients in the creation of this work, and whose financial support enabled me to perform this research and discuss it with many other workers in the field. I am also extremely grateful to Daniel Eisenstein, who agreed to co-chair my thesis committee, and has provided invaluable insights into galaxy formation.

I would like to thank Volker Springel, for writing the GADGET code and providing an advance copy of version 2, Lexi and John Moustakas for creating the RED cosmography package, and Michelle Cournoyer, Catalina Diaz-Silva, Joy Facio, and Sharon Jones for their help dealing with the inevitable bureaucracy.

I can barely begin to scratch the surface of the list of people I would like to thank for their friendship, for provoking stimulating conversations, and for prompting useful discussions, but I would like to mention Morten Andersen, Eric Bell, Michael Brown, James Bullock, Martin Bureau, Lise Christensen, Richard Cool, Romeel Davé, Roelof de Jong, Ruben de Kemp, Walter Dehnen, Rich Dice, Christian Drouet d'Aubigny, Mark Dusings, Vincent Eke, Kris Eriksen, Keith Felberg, Ken Freeman, Brad Gibson, Chris Gottbrath, Chris Groppi, Mike Habicher, Chris Hartig, Eric Hayashi, Knud Jahnke, Gyula Jozsa, Rob Kennicutt, Matt Kenworthy, Anatoly Klypin, Karen Knierman, Craig Kulesa, Georg Lamer, Beka Lash, Andrea Leistra, Jim Liebert, Wilson Liu, Joe Lobianco, Nicolas Lodieu, Bishop Lord, Kerry Lord, Eric Mamajek, Andy Marble, Richard Martin, Lucio Mayer, Casey Meakin, Michael Meyer, Doug Miller, Lissa Miller, Jackie Monkiewicz, Julio Navarro, Erika Nurmsoo, Ed Olszewski, Tim Pickering, Phil Pinto, Chris Power, Jane Rigby, Greg Rudnick, Erin Ryan, Ryan Shook, Kim Sharabura, Scott Sharabura, Sanjib Sharma, Josh Simon, James Taylor, Sarah Ternoway, Toshio Tsuchiya, Octavio Valenzuela, Marc Verheijen, Simon White, Patrick Young, and Dennis Zaritsky.

The computations described in this thesis were performed on the Astrophysikalisches Institut Potsdam beowulf cluster *octopus*, the University of Arizona High Performance Computing Center alpha supercomputer *aura*, and the Celeron 400 workstation *clavelina*.

This work was supported by grants from the National Aeronautics and Space Administration (NAG 5-10827), the David and Lucile Packard Foundation, the Bundesministerium für Bildung und Forschung (FKZ 05EA2BA1/8), the National Science Foundation (9807151 and PHY99-0749), the National Science and Engineering Research Council (PGSB-233028-2000), and the American Astronomical Society (International Travel Grant).

DEDICATION

In memory of my father, Ailey Bailin, who wanted to be an astronomer but could never stay awake late.

TABLE OF CONTENTS

LIST OF FIGURES 8

LIST OF TABLES 13

ABSTRACT 14

CHAPTER 1 Introduction 16

CHAPTER 2 Torques and Misaligned Disk Galaxies 23

- 2.1 Introduction 23
- 2.2 Torques from misaligned halos 25
  - 2.2.1 Framework for torqued orbits 25
    - 2.2.1.1 Introduction to torque 25
    - 2.2.1.2 Tilting timescale 26
  - 2.2.2 Expected torques from misaligned flattened halos 27
- 2.3 Reaction of self-gravitating disks to an external torque 30
  - 2.3.1 The simulations 32
  - 2.3.2 Warped disks 34
  - 2.3.3 Warp evolution 36
- 2.4 Conclusions 45

CHAPTER 3 Warps from Satellite-Disk Interactions 46

- 3.1 Evidence for coupling between Sgr and the Milky Way warp 46
  - 3.1.1 Introduction 46
  - 3.1.2 The angular momentum of satellite galaxies 47
  - 3.1.3 The angular momentum associated with the Milky Way warp 52
  - 3.1.4 Discussion 55
  - 3.1.5 Summary 58
- 3.2 Satellite-disk simulations 59
  - 3.2.1 Introduction 59
  - 3.2.2 Summary of simulations 59
  - 3.2.3 Precession and sinking 66
  - 3.2.4 Tilting and warping disks 68
    - 3.2.4.1 Tilting due to L_{precess} 68
    - 3.2.4.2 Tilting and warping due to L_{shrink} 68
    - 3.2.4.3 Warping of the disk due to a slewing potential . . 69
    - 3.2.4.4 Differential precession 69
  - 3.2.5 Simulation results 69
    - 3.2.5.1 Orbital inclination 72
    - 3.2.5.2 Satellite mass 81

TABLE OF CONTENTS — *Continued*

3.2.5.3 Orbital eccentricity 91
3.2.6 Discussion 94

CHAPTER 4 Figure Rotation of Cosmological Dark Matter Halos 97
4.1 Introduction 97
4.2 The simulations 101
4.3 Methodology 104
4.3.1 Introduction 104
4.3.2 Halo matching 108
4.3.3 Error in axis orientation 109
4.3.4 Figure rotation 114
4.3.4.1 Plane method 115
4.3.4.2 Quaternion method 121
4.4 Results 121
4.5 Conclusions 144

CHAPTER 5 Internal and External Alignment of the Shapes and Angular Momenta of ΛCDM Halos 147
5.1 Introduction 147
5.2 Methodology 155
5.2.1 The simulation 155
5.2.2 Measuring the axes 156
5.3 Shapes 161
5.4 Internal alignment 167
5.4.1 Principal axes 167
5.4.2 Angular momentum 172
5.4.3 Alignment between the angular momentum and the halo shape 174
5.5 External alignment 183
5.5.1 Introduction 183
5.5.2 Axis alignments 186
5.5.3 Axis correlations 194
5.5.4 Angular momentum 198
5.6 Discussion 208
5.7 Summary 211

CHAPTER 6 Conclusions 213
6.1 Summary 213
6.2 Outlook 216

REFERENCES 221

LIST OF FIGURES

2.1 Diagram of a tilted circular orbit with angular velocity ω 25
2.2 Magnitude of torques experienced by disks in misaligned flattened halos 29
2.3 x-z projection of a simulated torqued disk galaxy after 1 Gyr 33
2.4 Radial dependence of the tilt of the disk shown in Figure 2.3 34
2.5 Evolution of warp radius for simulations with $\beta = -2.0$ 37
2.6 Evolution of warp radius for simulations with $\beta = -2.5$ 38
2.7 Evolution of warp radius for simulations with $\beta = -3.0$ 39
2.8 Disk surface density at the warp radius for disks of different mass . 41
2.9 Disk surface density at the warp radius for simulations with different β 42

3.1 Angular momentum of each satellite and the warp 51
3.2 Diagram demonstrating the precession of a satellite in a flattened potential 66
3.3 Evolution of the angular momentum of the inner and outer disk in the undisturbed simulation 70
3.4 Images of the undisturbed disk at $t = 1.8$ Gyr 71
3.5 Evolution of the angular momenta of the inner and outer disk in prograde, polar, and retrograde simulations with satellite S1 73
3.6 Evolution of the warp angular momentum in the prograde, polar, and retrograde simulations with satellite S1 74
3.7 Images of the disks in the prograde, polar, and retrograde simulations after the first perigalacticon 76
3.8 Images of the disks in the prograde, polar, and retrograde simulations after the second perigalacticon 77
3.9 Evolution of the angular momenta of the inner and outer disk in prograde, polar, and retrograde simulations with satellite S2 79

LIST OF FIGURES — *Continued*

3.10 Evolution of the warp angular momentum in the prograde, polar, and retrograde simulations with satellite S2 80
3.11 Images of the disk in the simulations with a low mass satellite after the second perigalacticon 82
3.12 Evolution of the angular momenta of the inner and outer disk in prograde, polar, and retrograde simulations with satellite S3 84
3.13 Evolution of the warp angular momentum in the prograde, polar, and retrograde simulations with satellite S3 85
3.14 Images of the disk in the simulations with a high mass satellite after the first perigalacticon 86
3.15 Evolution of the specific angular momenta of the inner and outer disk in simulations with low, intermediate, and high mass satellites 88
3.16 Evolution of the warp specific angular momentum in simulations with low, intermediate, and high mass satellites 89
3.17 Images of the disks in simulations with low, intermediate, and high mass satellites after the first perigalacticon 90
3.18 Evolution of the specific angular momenta of the inner and outer disk in low-eccentricity orbit simulations 92
3.19 Evolution of the warp specific angular momentum in low-eccentricity orbit simulations 93
3.20 Images of the disk in a low-eccentricity orbit simulation 93

4.1 Fraction of the final mass that comes from the progenitor halo . . . 108
4.2 Major axis angular bootstrap error as a function of the number of particles 111
4.3 Major axis angular bootstrap error as a function of b/a 112
4.4 Diagram demonstrating the plane method 115
4.5 Projected density of halos with a range of subhalo fractions 118
4.6 Images of the sample halo at each snapshot, and the phase of the major axis rotation as a function of time 120
4.7 Pattern speed of the major axis rotation as a function of its error . . 122
4.8 Histogram of $\log \Omega_p$ 123

LIST OF FIGURES — *Continued*

4.9 Fractional histogram of Ω_p 124
4.10 Pattern speed of the major axis rotation as a function of its error, only using snapshots b096 through b102 127
4.11 Alignment of the figure rotation axis and the minor axis 129
4.12 Comparison between the figure rotation axis determined using the plane method and the quaternion method 130
4.13 Alignment of the figure rotation axis and the major axis, for halos where the plane and quaternion methods disagree 131
4.14 Alignment of the figure rotation axis and the major axis, for halos where the plane method gives only an upper limit 132
4.15 Pattern speed of the minor axis rotation as a function of its error, for potential major axis rotators 133
4.16 Alignment between the figure rotation axis and the angular momentum vector 135
4.17 Pattern speed as a function of halo mass 137
4.18 Pattern speed as a function of λ 138
4.19 Alignment between the figure rotation axis and the angular momentum vector, as a function of λ 139
4.20 Comparison of the pattern speed at different radii 141
4.21 Comparison of the figure rotation axis at different radii 142
4.22 Fractional rate of change of the halo axis ratios 144

5.1 Measured versus input c/a axis ratios of randomly-sampled halos . 159
5.2 Histogram of the minor-to-major axis ratio for each halo 160
5.3 Median axis ratios as a function of radius within the halo 161
5.4 Histogram of the triaxiality T of all halos 164
5.5 Relationship between b/a and c/a 165
5.6 Histogram of projected ellipticities 166
5.7 Histogram of the internal alignment of the major axes 169
5.8 Median internal alignment of the principal axes as a function of radius within the halo 170

LIST OF FIGURES — *Continued*

5.9 Median internal alignment of the major axes for halos of different mass . . . 171
5.10 Histogram of the internal alignment of the angular momentum . . 173
5.11 Median internal alignment of the angular momentum as a function of radius within the halo . . . 175
5.12 Median internal alignment of the angular momentum, for halos of different mass . . . 176
5.13 Histograms of the alignment of the angular momentum and the major axis . . . 177
5.14 Histograms of the alignment of the angular momentum and the intermediate axis . . . 178
5.15 Histograms of the alignment of the angular momentum and the minor axis . . . 179
5.16 Median alignment of the angular momentum and the principal axes as a function of radius within the halo . . . 181
5.17 Median alignment of the angular momentum and the principal axes for halos of different mass . . . 182
5.18 Orientation of the principal axes and angular momenta of all halos in one quarter of the simulation volume . . . 184
5.19 Orientation of principal axes and angular momenta of high mass halos in one quarter of the simulation volume . . . 185
5.20 Histograms of the alignment of the major axes of halos and the location of surrounding structure . . . 187
5.21 Mean alignment $\xi_{|ar|}$ of the major axes of halos and the surrounding structure . . . 189
5.22 Mean alignment $\xi_{|br|}$ of the intermediate axes of halos and the surrounding structure . . . 191
5.23 Mean alignment $\xi_{|cr|}$ of the minor axes of halos and the surrounding structure . . . 192
5.24 Mean correlation $\xi_{|aa|}$ of the major axes of neighbouring halos . . . 195
5.25 Mean correlation $\xi_{|bb|}$ of the intermediate axes of neighbouring halos 196

LIST OF FIGURES — *Continued*

5.26 Mean correlation $\xi_{|cc|}$ of the minor axes of neighbouring halos . . . 197
5.27 Mean vector alignment ξ_{Lr} of the angular momenta of halos and the location of surrounding structure 199
5.28 Mean vector alignment ξ_{Lr} as a function of local density 200
5.29 Mean alignment $\xi_{|Lr|}$ of the angular momenta of halos and the location of surrounding structure 202
5.30 Alignment $\omega(r)$ of the angular momenta of halos and the location of surrounding structure 203
5.31 Mean vector correlation ξ_{LL} of the angular momenta of neighbouring halos 205
5.32 Mean correlation $\xi_{|LL|}$ of the angular momenta of neighbouring halos 206
5.33 Correlation $\eta(r)$ of the angular momenta of neighbouring halos . . 207

LIST OF TABLES

3.1 Properties of the Sagittarius dwarf 48
3.2 Angular momenta of the Milky Way warp and some Milky Way satellites 50
3.3 Disk parameters 53
3.4 Mass components of the disk galaxy models. 61
3.5 Initial position and velocity of satellites in the simulations. 63
3.6 Initial orbital parameters of the satellites in the simulations. 64

4.1 Parameters of the cosmological simulation. 102
4.2 Boundaries of the slabs in the cosmological simulation. 105
4.3 Snapshots used to calculate figure rotations. 106

5.1 Number of halos with axes determined to within 0.2 radians 168
5.2 Number of halos with angular momentum directions determined to within 0.4 radians 172

ABSTRACT

In this thesis, we study the relationship between the angular momentum and shape of galactic disks, satellite galaxies, dark matter halos, and large scale structure using N-body simulations in the context of current models of galaxy formation.

In warped galactic disks, the angular momenta of the inner and outer disk are misaligned. We have calculated the torques a misaligned halo imparts on an embedded galactic disk. They are very strong, with tilting timescales significantly less than a Hubble time. N-body simulations of disks subject to torques of this strength indicate that the disk tilts in response and develops a trailing warp of the same magnitude as that of the Milky Way. The warp starts where the disk surface density drops below $70\ M_\odot\ \mathrm{pc}^{-2}$ and moves out through the disk on a timescale of 500 Myr.

We have investigated whether the warp of the Milky Way's disk is caused by nearby satellite galaxies. The misaligned angular momentum that causes the warp is almost perfectly anti-aligned with that of the orbital angular momentum of the Sgr dSph, and is of the same magnitude. We estimate the probability that this is a chance configuration to be 1.4%. This suggests that Sgr is responsible for the warp. However, we have performed N-body simulations of such disk-satellite interactions and found that the warps excited by Sgr with its current mass and orbit are much smaller than the warp of the Milky Way. We have investigated a range of satellite masses and orbits, and found that inclined orbits, eccentric orbits, and especially more massive satellites are more efficient at exciting warps.

The source of misaligned angular momentum is ultimately the surrounding

dark matter halo and large scale environment. Therefore, several aspects of halo alignment in cosmological dark matter simulations have been studied. We have analyzed several late snapshots of such a simulation and found rotation of the triaxial figure of galaxy and group mass dark matter halos. The figure is found to rotate about the minor axis in most cases. The rate of figure rotation follows a log-normal distribution centred on $\Omega_p = 0.148\,h\,\mathrm{km\,s^{-1}\,kpc^{-1}}$.

The halos have triaxial shapes that become more spherical at larger radii, with minor-to-major axis ratios distributed around $c/a = 0.6$. The principal axes of individual halos show strong internal alignment. The angular momentum also shows relatively good internal alignment, and is usually oriented along the minor axis. However, this alignment is not perfect, and the median misalignment is large enough to cause galactic warps as described earlier. Comparison of different halos reveals that the minor axes of halos show a strong tendency to point perpendicular to filaments and sheets. Major axes show a weaker tendency to point along filaments. These alignments are much stronger for group mass halos than for galaxy mass halos. The angular momenta of galaxy mass halos tend to point along filaments and sheets, while those of group mass halos point perpendicular to the surrounding mass distribution.

CHAPTER 1

INTRODUCTION

In this thesis, we study the relationship between the angular momentum and shape of galactic disks, satellite galaxies, dark matter halos, and large scale structure using N-body simulations in the context of current models of galaxy formation.

The formation of galaxies has long been a subject of interest and debate in astrophysics. The current model, which is well supported by observations on scales from individual galaxies at the small end to the Cosmic Microwave Background (CMB) at the large end, envisions a hierarchical assembly of halos made of collisionless dark matter. Inside of each halo lies the condensed baryonic component, presumed to be the luminous galaxy. The scenario has been laid out in detail by White & Rees (1978). As the universe expands, each density enhancement grows linearly through gravitational instability (Zeldovich, 1970). Eventually the halo reaches a sufficient density that its self-gravity overwhelms the cosmological expansion, and the halo turns around and begins to collapse. Starting at the center and moving outward, shells of material that were initially distinct cross, leading to mixing of the halo material and virialization. Baryons, which are initially well-mixed with the dark matter, cool if the cooling time is sufficiently short, forming a condensed baryonic core where star formation can occur.

The initial spectrum of density perturbations for the Cold Dark Matter (CDM) model has sufficient power at small scales that the growth of structure is "bottom-up" or hierarchical: small structures form first around short-wavelength peaks in the density field. Larger structures form later from the material surrounding longer-wavelength peaks in the density field. Some of the material surrounding these larger-scale peaks is smoothly accreted, but much of the material has itself previously collapsed into smaller structures, leading to the picture of larger galaxies formed from the agglomeration of smaller galaxies. Recent N-body simulations indicate that the central concentrations of pre-existing smaller halos maintain their distinct identity at least temporarily, even after they are accreted into larger halos, resulting in a large population of subhalos within each virialized halo (Klypin et al., 1999; Moore et al., 1999).

The role of angular momentum within this framework is a topic of considerable interest, but one that is not yet well understood. In a universe with no primordial vorticity, angular momentum originates from cosmological torques due to the surrounding matter (Strömberg, 1934; Hoyle, 1949), and manifests itself as shear flows during the linear evolution of density perturbations (Peebles, 1969; Doroshkevich, 1970; White, 1984). As each halo collapses and virializes, its angular momentum gets evenly redistributed. The baryons initially share the same phase space and specific angular momentum distribution as the dark matter. However, as they cool, angular momentum conservation requires that they settle into rotationally-supported disks oriented perpendicular to their angular momentum vector. When two halos with previously-condensed baryonic disks merge, the disks spiral together through dynamical friction, transferring some angular momentum from the gas to the dark matter. In addition, angular momentum can be transferred between the disks and dark matter by tidal torques.

This basic picture provides a good overview of the origin of angular momentum in galaxies, but it does not answer a large number of important questions. In particular, the alignment between the angular momentum of the baryonic disk, the angular momentum of the smooth component of the dark matter halo, the angular momentum of the subhalos within each dark matter halo, the shape of the dark matter halo, and the location of surrounding large scale structure is strongly affected by non-linear processes, and therefore requires non-linear techniques such as full N-body simulations to study (Porciani et al., 2002a).

In this thesis, I investigate two aspects of angular momentum alignment in galaxy formation: the origin of galactic warps, and the orientation and internal alignment of the shape, angular momentum, and figure rotation of galaxy and group mass dark matter halos.

Part I: Warped Galactic Disks

The disk of the Milky Way is warped like an integral sign, rising above the plane on one side and falling below the plane on the other. The Sun lies along the line of nodes of the warp, where tilted outer rings cross the inner plane. This warp is seen both in maps of neutral hydrogen (e.g., Diplas & Savage, 1991) and in the stellar distribution (Reed, 1996; Drimmel et al., 2000; López-Corredoira et al., 2002b). The Milky Way is far from unique in this regard; the planes of most external galaxies are warped to some degree (Bosma, 1981; Briggs, 1990; Christodoulou et al., 1993; Reshetnikov & Combes, 1998; Schwarzkopf & Dettmar, 2001), despite the tendency of warps to disperse when isolated (Kahn & Woltjer, 1959; Hunter & Toomre, 1969). This has driven many authors to search for universal mechanisms to excite or maintain warps (see Binney, 1992, for a review). Many of these proposed mechanisms rely on the dark halo to either stabilize warps as discrete bending modes within the halo (Sparke & Casertano 1988; Kuijken 1991; but see also Binney et al. 1998), or to provide the torque necessary to create the warp (Ostriker & Binney, 1989; Debattista & Sellwood, 1999; Ideta et al., 2000).

Other proposed mechanisms include the infall of intergalactic gas (Kahn & Woltjer, 1959; López-Corredoira et al., 2002a), magnetic fields (Battaner et al., 1990), and interactions with satellite galaxies (Huang & Carlberg, 1997).

Each of these mechanisms can, in particular circumstances, produce warps that appear similar to observed galactic warps. Although no single mechanism appears universal enough to account for all warps, the evolution toward a bending mode (even when no discrete mode exists) appears sufficiently similar to observed warps (Hofner & Sparke, 1994) that warping may be a generic response of disks to the individual perturbations they experience. In this case, we should examine individual warped galaxies for specific evidence of particular perturbations that explain their warps rather than search for a universal mechanism that may not exist.

A warp is fundamentally a misalignment between the angular momenta of the inner and outer disk. Angular momentum is transferred by torques, so the evolution of warped galaxies depends essentially on the torques they experience, and how the disks react to those torques. In Chapter 2, we calculate the torques a galactic disk might experience due to misalignment of the disk with the dark matter halo, and use N-body simulations to study the warping response of the disk to torques of this magnitude. In Chapter 3, we study another source of misaligned angular momentum: infalling satellite galaxies. We compare the angular momentum of the Milky Way warp to that of all the nearby satellite galaxies in search of interesting correlations, and perform N-body simulations of a variety of satellite-disk interactions to investigate the warping caused by satellite galaxies.

Part II: Properties of Cosmological Dark Matter Halos

It is suggested above that galactic warps may be caused by a misalignment between the angular momentum of the disk and the orientation of its triaxial halo, or a misalignment between the angular momentum of the disk and the angular momentum of the lumpy component of the dark matter halo in the form of satellite galaxies. It is therefore important to study the alignment of angular momentum within halos and its relation to the three-dimensional structure of the halos, as misalignments may be responsible for torques that affect the embedded galactic disk. In addition, the results of these studies can provide insight into the galaxy formation process, and provide important constraints on other studies of galaxy formation and evolution.

The ellipticity of the dark halo of the Milky Way and other galaxies can be measured using a variety of methods: the orbits of tidal streams (Law et al., 2004), the anisotropy of stellar velocities (Olling & Merrifield, 2000), the flaring of the gas disk (Olling & Merrifield, 2000), the projected shape of X-ray gas (Buote et al., 2002), the kinematics of polar ring galaxies (Sackett et al., 1994), and weak gravitational lensing (Hoekstra et al., 2004). These results suggest that galaxy halos have a wide range of flattenings.

Halos formed in cosmological simulations are also generally not spherical, but have an ellipsoidal shape due to the anisotropic velocity dispersion imparted by the external tidal field on protohalo material (Eisenstein & Loeb, 1995). While there have been several studies of the shapes of halos in numerical simulations (Frenk et al., 1985; Dubinski & Carlberg, 1991; Katz, 1991; Warren et al., 1992; Cole & Lacey, 1996; Bullock, 2002; Jing & Suto, 2002), they have often suffered from poor resolution, non-cosmological initial conditions, or initial conditions based on cosmologies that are no longer consistent with observations. There has been relatively little work done on how those shapes evolve with time; "figure rotation", where the orientation of the halo principal axes rotate, may have a considerable influence on the structure of embedded galactic disks (Bureau et al., 1999; Tremaine & Yu, 2000; Bekki & Freeman, 2002; Masset & Bureau, 2003).

Cosmological simulations have also been used to study the alignment of angular momentum within dark matter halos (Dubinski, 1992; Warren et al., 1992; Bullock et al., 2001). It has generally been found that the angular momentum is relatively well aligned within halos, and that it is most often aligned with the minor axis of the mass distribution, with considerable halo-to-halo variation. However, these studies have been performed with relatively low resolution, and in background cosmologies that do not match current observations; this has motivated us to test these results with modern simulations.

The alignment of the shape and angular momenta of nearby halos act as a test of the theory of structure formation, inform our understanding of the origin of halo shapes and angular momenta, and are important systematic uncertainties in weak lensing measurements that must be calibrated. While there have been several theoretical studies of the alignment of halos with each other, those based on linear theory (Lee & Pen, 2000, 2001; Pen et al., 2000) are suspect due to the importance of non-linear effects in determining the orientation of halo angular momentum (Porciani et al., 2002a), while those based on N-body simulations have not had sufficient resolution to detect effects for any but the most massive clusters (Barnes & Efstathiou, 1987; West et al., 1991; van Haarlem & van de Weygaert, 1993; Splinter et al., 1997; Onuora & Thomas, 2000; Hatton & Ninin, 2001; Faltenbacher et al., 2002; Kasun & Evrard, 2004).

We have performed a large high-resolution ΛCDM N-body simulation to study these effects. In Chapter 4, we examine the figure rotation of triaxial dark matter halos formed in this simulation. In Chapter 5, we present an extensive study of the shapes of the dark matter halos, the internal alignment of their principal axes and angular momenta, and the alignment of these properties with external halos.

CHAPTER 2

TORQUES AND MISALIGNED DISK GALAXIES

2.1 Introduction

There is observational and theoretical evidence that dark matter halos are flattened with minor-to-major c/a axis ratios ranging as low as 0.3. Cosmological simulations show that while the angular momentum and minor axis of a flattened halo are usually correlated, the alignment is not perfect for a large number of halos (see § 5 and references therein for a full discussion of the alignment of the shape and angular momentum of dark matter halos). Recent high-resolution simulations suggest that the angular momentum of the baryons may not even be aligned with that of the dark matter (van den Bosch et al., 2002; Sharma & Steinmetz, 2004), and that therefore the disks that form from those baryons are usually misaligned with the minor axis of the halo matter distribution.

It has been suggested that warped disks may result from such misalignment between the disk and a flattened dark matter halo (Dekel & Shlosman, 1983; Toomre, 1983). A warped disk is a stable solution inside an oblate halo if it is a discrete normal mode of the system (Toomre, 1983; see also Lynden-Bell, 1965). Calculations of such normal modes, assuming that the disk can be approximated by a series of spinning concentric rings, have been performed in the linear (Sparke & Casertano, 1988) and non-linear (Kuijken, 1991) regimes. These discrete modes exist if the disk is abruptly truncated; however, if the surface density declines more gradually, in accordance with observations, there are no discrete modes. Such discrete modes do exist if the halo is prolate rather than oblate, elongated perpendicular to the disk (Ideta et al., 2000). However, alignment between a halo's angular momentum and its major axis is uncommon (§ 5); even if the angular momenta of the baryons that form a disk are misaligned with the angular momentum of the halo by 20° (Sharma & Steinmetz, 2004), it is unlikely that there exists a large population of galaxies with disks lying near the symmetry plane of a prolate halo. Furthermore, all of the aforementioned studies assume that the halo potential can be assumed to be static. Binney et al. (1998), following on an idea of Nelson & Tremaine (1995), have found that the halo responds to the presence of a warped disk, destroying the normal modes and causing the warp to disperse.

These studies have all assumed that warps are steady-state uniformly processing phenomena. However, this may not be the case; Hipparcos measurements of the Galactic warp indicate that it is not uniformly precessing (Drimmel et al., 2000). This should not be unexpected: the orientations of halo shapes change with time (§ 4), and the orientations of galactic disks change with time as they accrete angular momentum (Quinn & Binney, 1992). Therefore, even if the equilibrium configuration of the disk is not warped, the disk may spend a large fraction of its time out of equilibrium, in a warped state.

An important ingredient of any non-equilibrium explanation for galactic warps is an understanding of the dynamical response of disks to the torques they experience. In this chapter, we present a simple model of the torques misaligned halos exert on disk galaxies and the reaction of disks to these torques, to determine if they appear similar to observed warped galaxies.

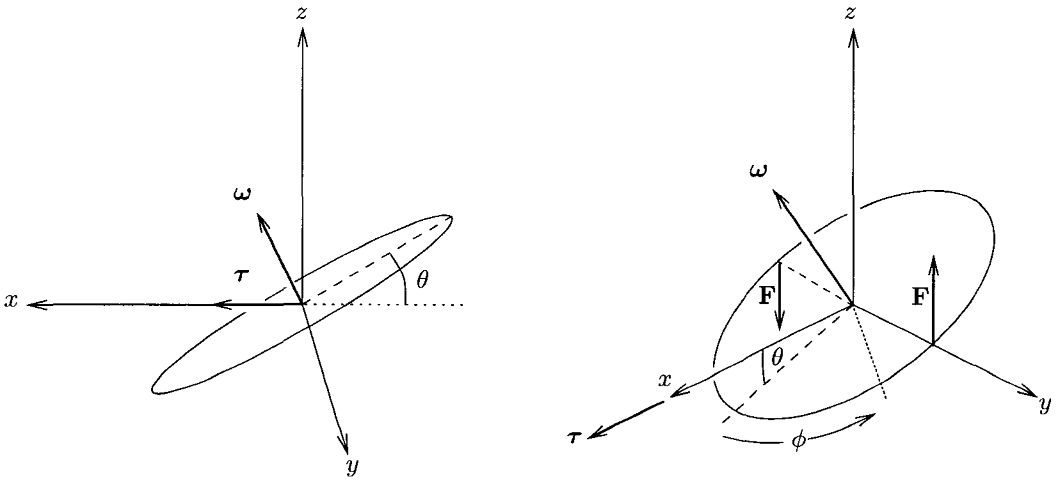

Figure 2.1: Two views of a circular orbit with angular velocity vector $\boldsymbol{\omega}$ tilted from the xy-plane toward the x-axis by an angle θ. ϕ is the azimuthal angle around the orbit. A torque $\boldsymbol{\tau}$ directed along the x-axis produces forces $\boldsymbol{F}$, as in equation (2.1).

2.2 Torques from misaligned halos

2.2.1 Framework for torqued orbits

2.2.1.1 Introduction to torque

We introduce a simplified model to establish a framework in which the reaction of galactic disks to torques can be studied. Consider a star in the disk on a circular orbit in the xy-plane about the origin with angular velocity $\omega(r) = v_c(r)/r$ (see Figure 2.1). We apply a torque along the x-axis by accelerating the star in the y and z directions around the x-axis with an antisymmetric pseudo-tensor T_{ij}. If the acceleration in the i direction is a_i and the j-position of the star is r_j then the applied acceleration is

$$a_i = \sum_j T_{ij} r_j \tag{2.1}$$

where

$$\boldsymbol{T} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -\tau(r) \\ 0 & \tau(r) & 0 \end{pmatrix}. \tag{2.2}$$

This is an angular acceleration (or equivalently a specific torque) around the x-axis of magnitude $a/r = \tau(r)$. Since the angular acceleration is the time derivative of the angular velocity, the x-component of the angular velocity increases in response to the torque, and the angular momentum of the star slews toward the x axis, tilting the orbit around the y-axis (Bardeen & Petterson, 1975).

2.2.1.2 Tilting timescale

For an orbit in the xy-plane, the angular momentum is initially aligned with the z axis ($\boldsymbol{\omega} = \omega_0(r)\,\hat{\boldsymbol{z}}$). The torque adds angular momentum in the x direction at a constant rate:

$$\frac{d\boldsymbol{\omega}}{dt} = \tau(r)\,\hat{\boldsymbol{x}}. \tag{2.3}$$

If the radius of the orbit stays constant, then the angular velocity grows as

$$\omega(t) = t\tau(r)\,\hat{\boldsymbol{x}} + \omega_0(r)\,\hat{\boldsymbol{z}}. \tag{2.4}$$

As the torque adds angular momentum along the x axis, the orbit tilts toward the x axis by an angle θ, where

$$\tan\theta = \frac{\omega_x}{\omega_z} = \frac{t\tau(r)}{\omega_0(r)}. \tag{2.5}$$

The rate of tilting at small angles is

$$\frac{d\theta}{dt} = \frac{\tau(r)}{\omega_0(r)}, \tag{2.6}$$

with a characteristic timescale to tilt one radian of

$$t_{\text{tilt}} = \frac{\omega_0(r)}{\tau(r)}. \tag{2.7}$$

This timescale is a useful quantity for characterizing the strength of the torque. The form of equation (2.7) can also be motivated by noting that ω_0 is equivalent to the specific angular momentum and τ is the specific torque, which is the rate of change of the specific angular momentum. Note that for a flat rotation curve where $\omega_0 = v_c/r$, t_{tilt} is a monotonically increasing function of r if $\tau(r)$ falls off faster than r^{-1}; i.e. for a torque that decreases with radius faster than r^{-1}, the inner regions of the disk should always tilt faster than the outer regions, resulting in a trailing warp.

2.2.2 Expected torques from misaligned flattened halos

We calculate the magnitude of the torque for a disk misaligned in a flattened halo with the radial dependence of an NFW profile (Navarro et al., 1996) but flattened along the z axis:

$$\rho(x, y, z) = \frac{\rho_0}{(m/r_s)\,(1 + m/r_s)^2} \tag{2.8}$$

for modified radius $m^2 = x^2 + y^2 + z^2/q^2$ where q is the c/a axis ratio. The force from this distribution on a given point is calculated using equation (2-88) of Binney & Tremaine (1987). For reference, we also calculate the torques inside flattened isothermal profiles, as in equation (2-54a) of Binney & Tremaine (1987). The torques are calculated by evaluating the forces at opposite points along a fictitious disk centred at the origin and inclined by an angle θ to the xy plane (see Figure 2.1). The acceleration at radius r due to a torque of magnitude τ is orthogonal to the symmetry plane, is of opposite sign on opposite sides of the disk, and is of magnitude $F = \tau r$. Given the forces on two test points F_1 and F_2 at opposite sides of the disk, we calculate the component of the force perpendicular to the plane of the disk $F_{\perp 1}$ and $F_{\perp 2}$. The torque is

$$\tau = \frac{F_{\perp 1} - F_{\perp 2}}{2r}. \tag{2.9}$$

The angular velocity ω is determined at each point from the radial force: $\omega = \sqrt{F_r/r}$ where the radial force F_r is defined to be positive when it is directed toward the center of the halo. The average value of ω for the two points across the disk is used. The tilting timescale is given by equation (2.7).

Figure 2.2 illustrates the torques derived using this method for halos with a cross-section of properties expected for galactic halos. The fiducial NFW model has flattening $q = 0.7$, virial velocity $v_c = 175\,\mathrm{km\,s^{-1}}$ (corresponding to $v_{\mathrm{rot}} = 205\,\mathrm{km\,s^{-1}}$ at a galactocentric radius of 10 kpc), and a disk inclined 20° to the symmetry plane of the halo. All NFW halos have concentration parameters $c_{200} = 15$ (Navarro et al., 1997). The different panels show the effect of varying the disk angle by 10°, the axis ratio from 0.5 to 0.9, and the virial velocity by 75 km s^{-1} compared to the fiducial model. The bottom-right panel of Figure 2.2 shows the torques inside an isothermal profile with a range of axis ratios. The magnitude of the torque in the isothermal case is similar to that in the NFW case, but falls off less rapidly with radius.

Figure 2.2 demonstrates that the torque from a misaligned flattened halo is significant to its evolution on a cosmological timescale, as the tilting timescale t_{tilt} is less then the Hubble time over the entire length of the disk for all halos modelled. The torques scale approximately as the density, and therefore fall off with radius as $\tau(r) \propto r^{\beta}$ with β ranging from -1 at small radii to -3 at large radii for NFW halos and ranging from 0 to -2 for isothermal halos. The typical torque follows the relation

$$\tau(r) = \tau_0 \left(\frac{r}{r_0}\right)^{\beta}, \tag{2.10}$$

where $\tau_0 = 10^{-30}\ \mathrm{s}^{-2}$, $\beta = -2.5$, and $r_0 = 1$ kpc. The angular velocity $\omega(r)$ varies as $\omega(r) \propto r^{\alpha}$ where α ranges from -0.5 to -2 in NFW halos, and is -1 at most radii in isothermal halos. Therefore, the tilting timescale t_{tilt}, as calculated

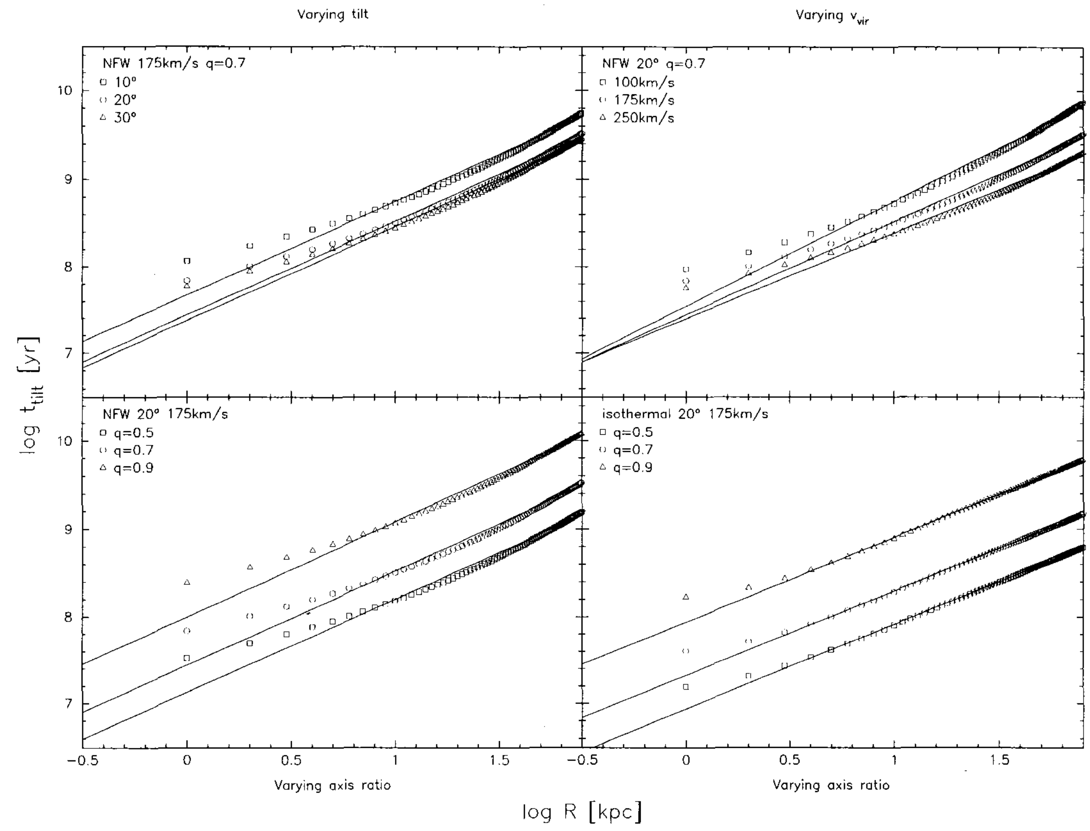

Figure 2.2: The symbols indicate the torques experienced by disks misaligned inside flattened halo profiles, expressed in terms of the tilting timescale t_{tilt}. The lines are power law fits for $\log R > 0.1$. Note that stronger torques have *shorter* timescales and appear lower in these graphs. *Top-left:* Torques inside NFW halos with virial velocities $v_{200} = 175\,\mathrm{km\,s^{-1}}$, flattenings $q = 0.7$, and concentration parameters $c_{200} = 15$, for disks misaligned by 10°, 20°, and 30° from the symmetry plane. *Top-right:* Torques inside NFW halos with virial velocities $v_{200} = 100,\ 175,$ and $250\,\mathrm{km\,s^{-1}}$, flattenings $q = 0.7$, and concentration parameters $c_{200} = 15$, for disks misaligned by 20° from the symmetry plane. *Bottom-left:* Torques inside NFW halos with virial velocities $v_{200} = 175\,\mathrm{km\,s^{-1}}$, flattenings $q = 0.5,\ 0.7,$ and 0.9, and concentration parameters $c_{200} = 15$, for disks misaligned by 20° from the symmetry plane. *Bottom-right:* Torques inside isothermal halos with virial velocities $v_{200} = 175\,\mathrm{km\,s^{-1}}$, flattenings $q = 0.5,\ 0.7$, and 0.9, and core radii $R_c = 1$ kpc, for disks misaligned by 20° from the symmetry plane.

using equation (2.7), rises approximately linearly with radius. For example, in an isothermal halo $\tau(r) \propto r^{-2}$ while $\omega(r) \propto r^{-1}$ so $t_{\mathrm{tilt}} = \omega(r)/\tau(r) \propto r$. The timescale is shorter in the inner regions of the disk, and therefore the inner disk tilts faster than the outer disk in response to torques from misaligned halos: the torque increases more rapidly toward the center than does the disk's ability to resist the torque due to its angular momentum.

The further the halos and disks are misaligned, the stronger the torque is; however, the torque profile is mostly unchanged for angles beyond 20°. Increasing the virial velocity, and therefore the mass of the halo, increases the magnitude of the torque. Because the concentration c_{200} is held constant amongst these models, changing the virial velocity also changes the scale radius r_s, which can be seen as the shift in the radius of the knee in the torque profiles of the top-right panel of Figure 2.2. The flattening of the halo has a large effect on the magnitude of the torque, with torques strengthening as the halo departs further from spherical symmetry. The torques fall off more slowly in the isothermal halos than in the NFW halos, but are of similar magnitude over most of the disk radius.

2.3 Reaction of self-gravitating disks to an external torque

While there have been many studies of the effects of torques on massive disks, almost all are based on the assumption that the prevalence of warps is due to their stability, and therefore seek to solve the winding problem (Kahn & Woltjer, 1959; Hunter & Toomre, 1969) by finding a solution that precesses uniformly (Lynden-Bell, 1965; Sparke & Casertano, 1988; Kuijken, 1991; Ideta et al., 2000) or reaches a uniform configuration in the presence of a given steady torque (Ostriker & Binney, 1989; Debattista & Sellwood, 1999; López-Corredoira et al., 2002a). The other possible solution, that warps are transient but frequently excited, has not been thoroughly investigated. A lone exception is Hofner & Sparke (1994), who found that in the linear regime of tilt angles and static oblate halos, disks which are initially inclined evolve toward discrete warp modes (whether or not they exist over the entire disk) from the inside out, settling on a timescale that depends on the group velocity of vertical bending waves.

The evolution of a massless disk is easily found from equation (2.5). This leads to a disk with a curved inner region, becoming flat in the outer regions where t_{tilt} is large. Real warped galaxies appear quite the opposite, with a disk that lies flat in its inner regions and warps beyond the Holmberg radius (Briggs, 1990). It is therefore clear that the self-gravity of galactic disks is very important to their dynamics. This is not surprising, as "maximal disk" models, where the mass-to-light ratios of the disks (M/L) are as large as is consistent with the rotation curves, provide good fits to most rotation curves (Broeils & Courteau, 1997; Bell & de Jong, 2001; Bolatto et al., 2002; Simon et al., 2003). Sackett (1997) demonstrates that the Milky Way itself is consistent with having a maximal disk. Although Courteau & Rix (1999) have used the residuals of the Tully-Fisher relation (Tully & Fisher, 1977) to estimate that disks provide only 40% of the dynamical mass at 2.2 exponential scale lengths, and the rotation curves of the majority of low surface brightness galaxies imply that their disks are substantially sub-maximal (de Blok & McGaugh, 1997; Zavala et al., 2003), the prevalent phenomena of spiral arms and bars, which are disk instabilities, require that self-gravity be important to the dynamics of many disks (Athanassoula et al., 1987).

The self-gravity of the disk acts to keep the disk flat. Ostriker & Binney (1989) examined the effect of a slewing disk potential on a set of self-gravitating rings and found that regions of high surface density react like a solid body. Since the surface density of the disk is highest in the central regions, the central parts of the disk will be kept locally flat, resulting in disks that closer resemble observed warped galaxies. López-Corredoira et al. (2002a) take this into account by developing a form for the internal torque $\tau_{\rm int}$. Here we do not use their $\tau_{\rm int}$, but rather integrate the orbits numerically. While this does not provide us with an explicit set of differential equations for $\theta(r, t)$, it frees us from such assumptions of theirs as circular orbits, the constancy of $\omega(r)$, the existence of an equilibrium configuration, and the lack of a dark matter halo.

In this section, we perform N-body simulations of self-gravitating disks subject to torques from the surrounding halo, as calculated in § 2.2.2, and investigate the nature of the transient warps that develop.

2.3.1 The simulations

The initial conditions for the simulations were constructed using the method of Hernquist (1993), for disk masses of $1 \times 10^{10}\,\mathrm{M}_\odot$, $3 \times 10^{10}\,\mathrm{M}_\odot$, and $5.6 \times 10^{10}\,\mathrm{M}_\odot$, scale length $r_d = 3.5\,\mathrm{kpc}$ and vertical scale height $h_z = 325\,\mathrm{pc}$ in a static spherically-symmetric NFW halo potential with concentration parameter $c_{200} = 15$ and virial velocity $v_{200} = 175\,\mathrm{km\,s^{-1}}$. These disks were then allowed to relax under the force of gravity until they appeared to be in equilibrium. The halo masses were varied such that the mass of the disk plus halo was $4.1 \times 10^{11}\ M_\odot$ in each case.

The models were evolved using the GRAPESPH code (Steinmetz, 1996), where interparticle gravitational forces are computed using direct summation with GRAPE-3 hardware (Okumura et al., 1993). A Plummer softening (Plummer, 1911; Aarseth, 1963) of 0.3 kpc has been used. The models were evolved for 2 Gyr, which took 5000-7000 timesteps depending on the model. The torque was applied as an external force of the form (2.10), with $\tau_0 = 10^{-30}\ \mathrm{s}^{-2}$ at $r_0 = 1\ \mathrm{kpc}$. Simulations were performed with torque slopes of $\beta = -2.0$, -2.5, and -3.0. The torque was

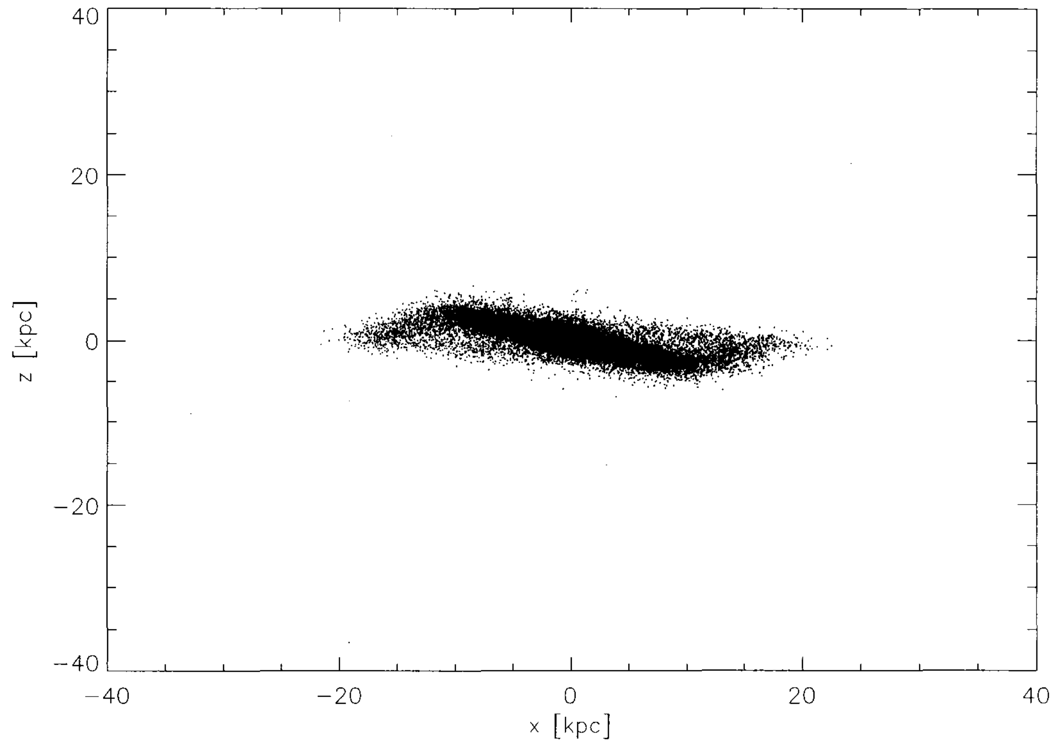

Figure 2.3: x-z projection of a simulated disk galaxy with mass $3 \times 10^{10}\ M_{\odot}$ after 1 Gyr in a torque of $\tau_0 = 10^{-30}\ \mathrm{s}^{-2}$ at $r_0 = 1$ kpc and $\beta = -2.5$.

capped at $\tau = \tau_0$ inside $r < r_0$ in order to prevent extremely small timesteps for particles at very small radii where the torquing force diverges.

The disks contained 16384 particles. The $3 \times 10^{10}\ M_{\odot}$ simulations were also performed with 32767 particles to see if the resolution was sufficient. The results for the high resolution simulations were identical to those for the lower resolution simulations to well within the angular errors computed in § 2.3.3 except for in the spherical shells which contained very few particles and had already been neglected from the analysis on those grounds.

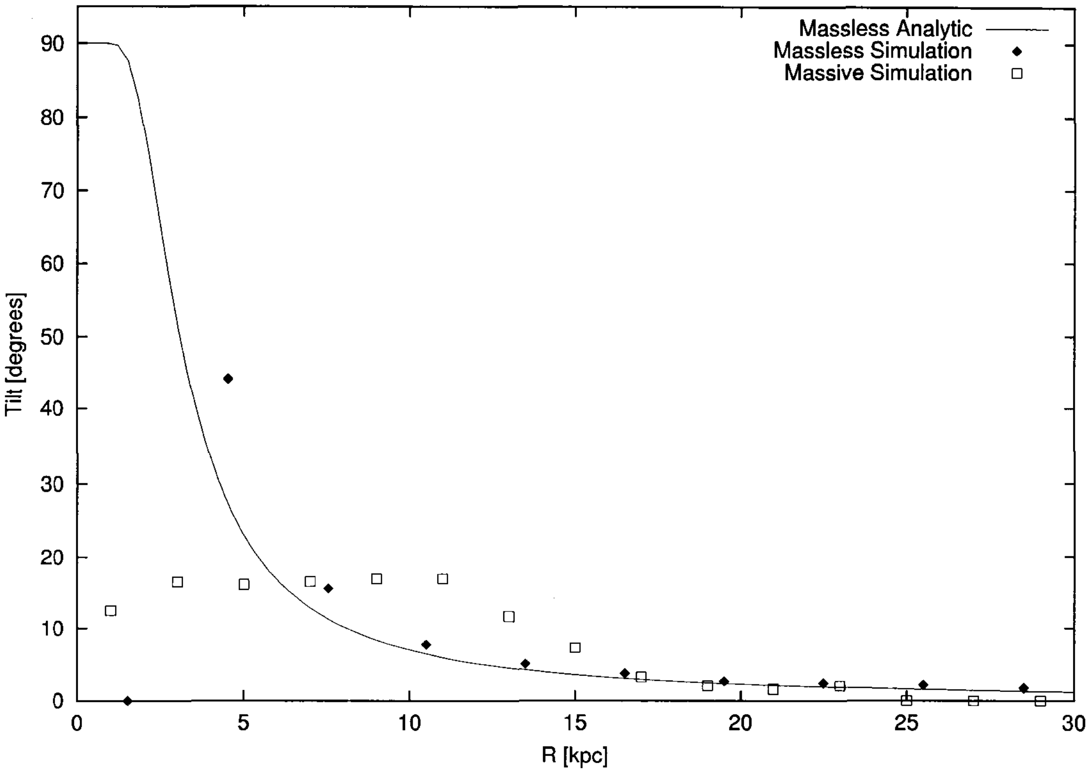

Figure 2.4: Open squares show the radial dependence of the tilt of the disk shown in Figure 2.3, computed in spherical shells 2 kpc wide. The solid line is an analytic estimate of the warp of a massless disk, and the solid diamonds are an N-body simulation of the same disk but with the self-gravity of the disk turned off.

2.3.2 Warped disks

Figure 2.3 shows the simulation of a $3 \times 10^{10}\,M_\odot$ disk with 32767 particles subject to a $\beta = -2.5$ torque for 1 Gyr. The main plane of the disk is flat and clearly tilted toward the positive x-axis. Beyond 10 kpc, the disk no longer remains flat but warps back toward the original plane; this appears similar to many observed warped disks. The particles which appear to be filling in the area between the main disk and the warp are projections and are actually in front of or behind the galaxy at large radii.

We bin the simulation particles into spherical shells of thickness 2 kpc and calculate their degree of tilt using the moment of inertia tensor. Figure 2.4 shows a plot of the radial profile of the tilt of the disk shown in Figure 2.3. Also shown are a non-gravitating disk subject to the same torque and an analytic estimate of the behaviour of a massless disk in this torque. There are three important regimes: the central flat disk, the warp, and the effectively massless outer region. The central 11 kpc is tilted uniformly 17° from the original plane (we ignore the innermost point because the torque is capped inside r_0) and corresponds to the visually flat part of the disk. The effect of the disk's self-gravity is to prevent the inner regions of the disk from tilting as much as they would otherwise, while pulling the outer regions of the disk to more inclined orbits. At the warp radius r_w=11 kpc, the disk is no longer flat, and the disk warps back toward the original plane. Finally, from 17 kpc to 23 kpc (the final extent of the disk), the disk follows the analytic prediction, acting like a massless disk. This result agrees qualitatively with that of Ostriker & Binney (1989), who found that regions of high surface density remain flat when torqued, and the position of a warp is determined in part by a drop in surface density.

A useful diagnostic for the strength of a warp is the total amount of angular momentum which is misaligned with that of the flat central disk. We have calculated this warped angular momentum for the highest-mass disk simulation (which is most similar to the Milky Way) at 1 Gyr, and found that the misaligned component has a magnitude of 2.0×10^{12} $M_\odot$ kpc km s^{-1}. To facilitate comparison between galaxies with different mass distributions, it is useful to normalize the misaligned angular momentum by the total angular momentum in the warped region of the disk; we find that a fraction 0.17 of the angular momentum is misaligned.

2.3.3 Warp evolution

We now investigate how the disk evolves over time under the influence of the torque. The general form of the warp is as above, but as the disk evolves under the influence of the torque, the warp radius moves out through the disk at a rate that depends on the mass of the disk. At early times, typically before 200 Myr, the inner disk does not act like a solid body, and the tilt profile lies below the prediction for the massless disk.

We developed an algorithm to automatically detect the warp radius. The particle positions were stored every 40 timesteps. For each of these outputs, the tilt angles were calculated in spherical shells as in Figure 2.4. We did a bootstrap analysis to estimate the error in these angles: for each spherical shell, we drew 10 random sets from the particles in that shell, calculated the tilt of each bootstrap set, and then used the standard deviation of those angles as an estimate of the error in the tilt of that bin. The error depended on the number of particles in the shell, but was typically less than 1°.

Shells with errors greater than 0.5° were not analyzed, as these generally had few particles and were dominated by numerical noise. For each remaining shell with radius r and tilt $\theta_{\rm sim}$, we calculated the deviation $\Delta\theta$ between the simulation tilt $\theta_{\rm sim}$ and the tilt expected for a massless disk at that radius. The central radius r of the radial bin with the highest $\Delta\theta$ was then considered to be the warp radius r_w. This is the radius at which the difference between the tilt of the massive disk and the tilt of an equivalent massless disk is maximized. For example, in Figure 2.4, the highest $\Delta\theta$ occurs at $r = 11$ kpc where the data point $\theta_{\rm sim} = 17^\circ$ and the predicted line $\theta_{\rm pred} = 7^\circ$. This corresponds exactly to the radial bin at the end of the plateau in Figure 2.4. Examining a few randomly chosen outputs from each simulation showed that in each case this automated r_w agreed with our

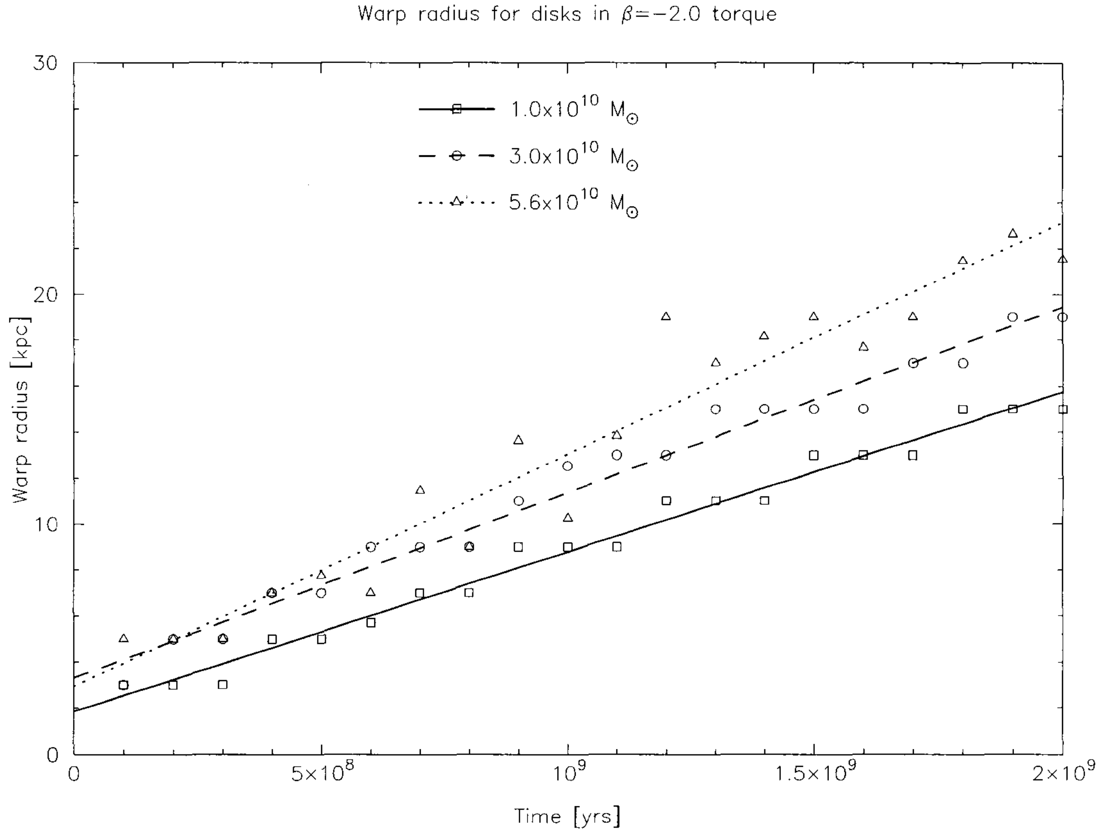

Figure 2.5: Radial bins at which the disks exhibit their warps as the simulations evolve in time. Points are plotted every 10^8 years. Because the particles were binned into spherical shells 2 kpc wide, the warp radii appear quantized. The lines are linear least square fits. This figure shows the simulations with $\beta = -2.0$ torques and disk masses $1.0 \times 10^{10}\ M_\odot$, $3.0 \times 10^{10}\ M_\odot$, and $5.6 \times 10^{10}\ M_\odot$ represented by the squares, circles, and triangles respectively.

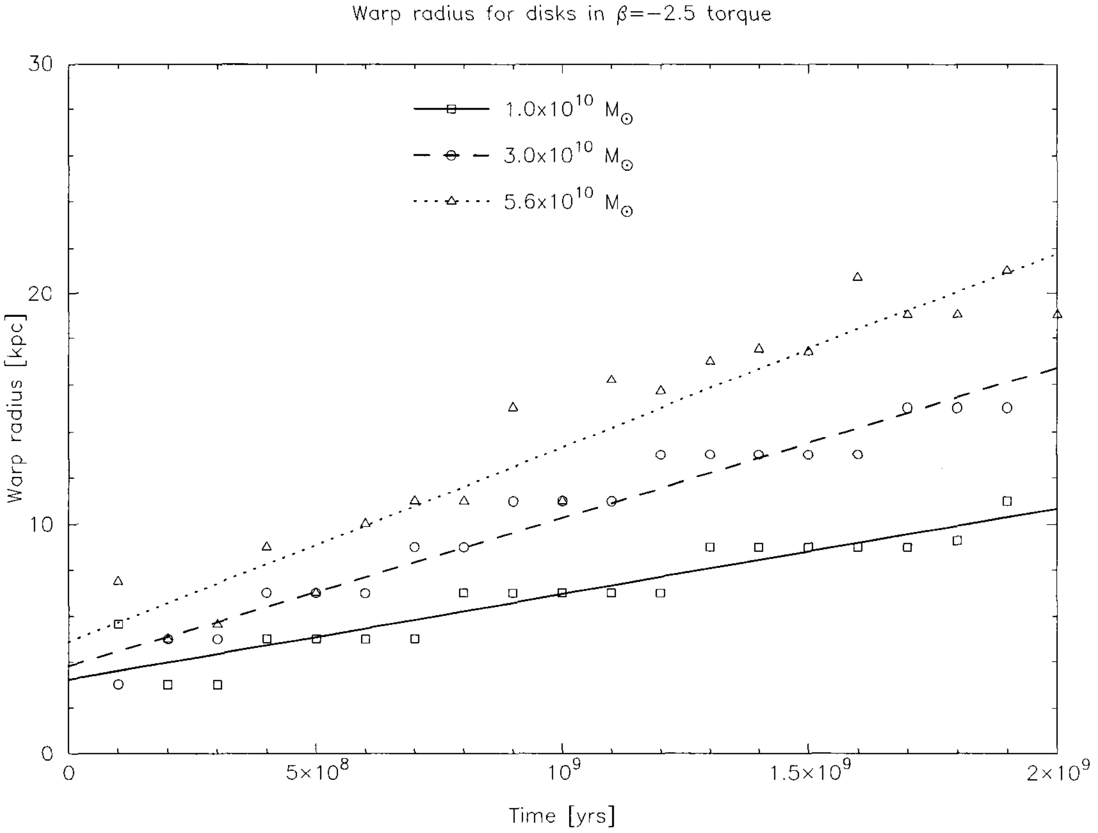

Figure 2.6: As in Figure 2.5 but for simulations with $\beta = -2.5$ torques.

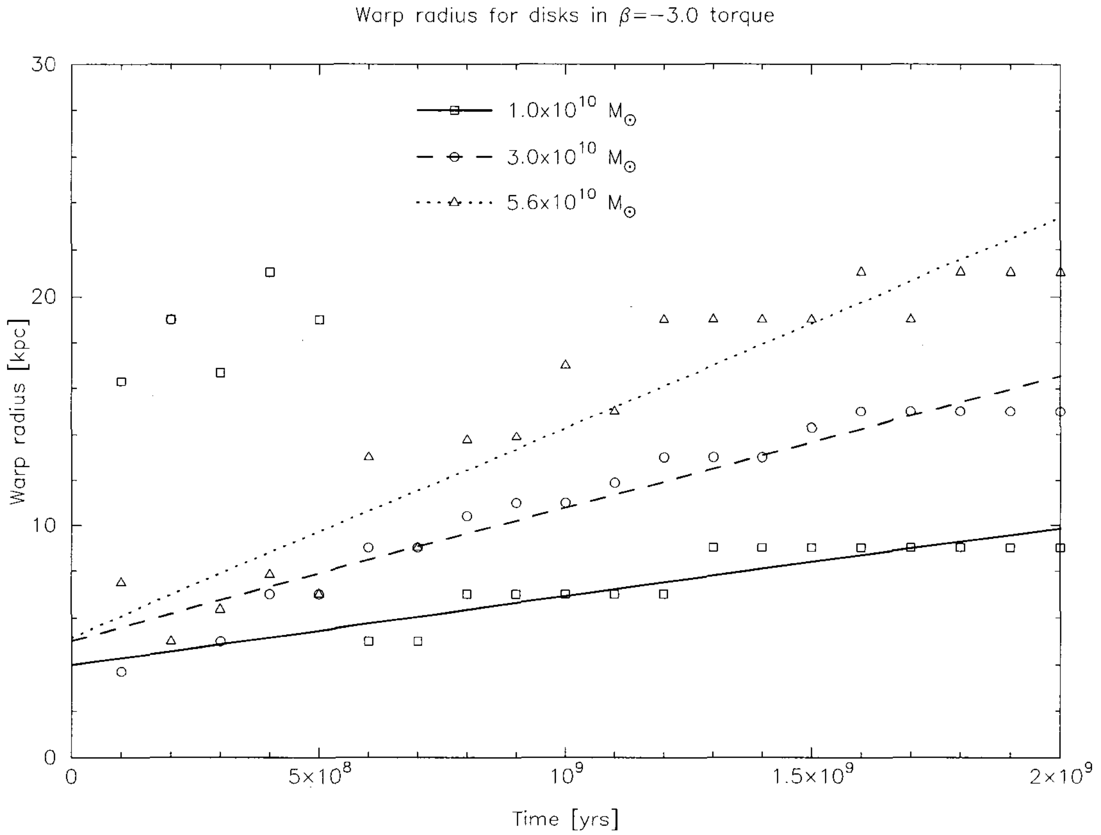

Figure 2.7: As in Figure 2.5 but for simulations with $\beta = -3.0$ torques. The warp takes time to develop, resulting in the noise apparent early in the simulations, which were not included in the straight line fits.

intuition, except at early times when there was no solid disk and the automated r_w was dominated by noise.

Figures 2.5, 2.6, and 2.7 show the evolution of the warp radius for simulations with differing torque slopes β. The squares, circles, and triangles correspond to simulations with disk masses of 1, 3, and 5.6×10^{10} $M_\odot$ respectively. The warp radii appear "quantized" because of the 2 kpc-wide radial bins used to calculate the tilt (see, e.g. Figure 2.4). One point was plotted every 10^8 years, linearly interpolating r_w from the neighbouring simulation outputs. The early times before the warp has developed show up as noise in the upper-left region of some of the figures, particularly Figure 2.7. To quantify the growth of the warp, straight lines were fit through the profiles. While one might think that the lines ought to be constrained to pass through the origin, since there is no warp at $t = 0$, the data do not support this. This is because the warp instability requires that the surface density of the disk be below a certain critical surface density, as we demonstrate later. We examined Figures 2.5–2.7 by eye and excluded the obviously noisy early times from the fits; in practice, this meant excluding $t \leq 5 \times 10^8$ yr for the $\beta = -3.0$, $M_d = 1.0 \times 10^{10}$ $M_\odot$ simulation and $t \leq 2 \times 10^8$ yr for the $\beta = -3.0$, $M_d = 3.0 \times 10^{10}$ $M_\odot$ simulation. We ended the simulations after 2 Gyr, the time when the warp radius reached the edge of the particle disk in some simulations, at which point the algorithm to find r_w simply detects the edge of the disk. The fits are shown as the solid, dashed, and dotted lines for the simulations with disk masses of 1, 3, and 5.6×10^{10} $M_\odot$ respectively. The warp radius moves out through high mass disks faster than low mass disks, while the effect of the torque slope is somewhat ambiguous.

In a massive disk, particles at different radii are coupled to each other gravitationally and act to keep the disk locally flat. This suggests that the crucial factor

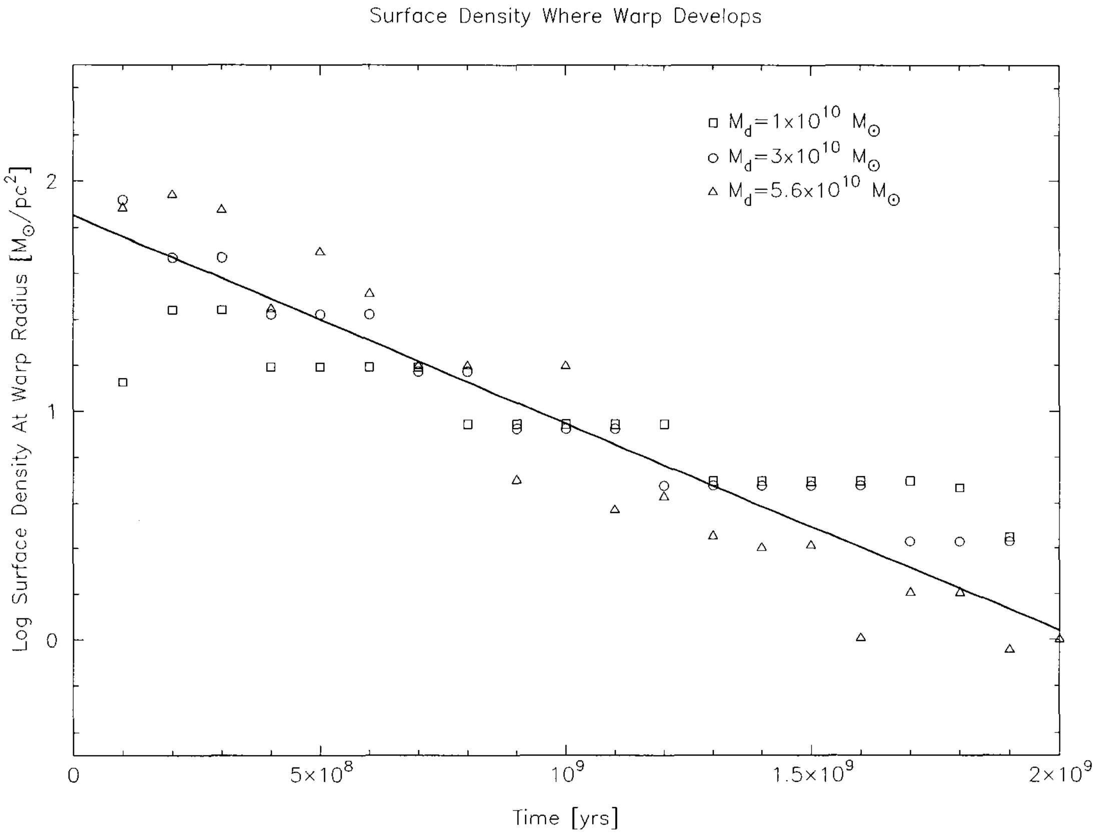

Figure 2.8: Surface densities at the warp radii as the simulations evolve. As the warp moves out through the disk, the local surface density at the warp radius falls according to equation (2.12). The squares, circles, and triangles show simulations with disk masses of 1, 3, and 5.6×10^{10} $M_\odot$ respectively in a $\beta = -2.5$ torque. The solid line is the fit given by equation (2.13). Compare this with Figure 2.6, which shows the radial evolution of the warp for the same set of simulations. The local surface density at the warp is similar for all models at all times.

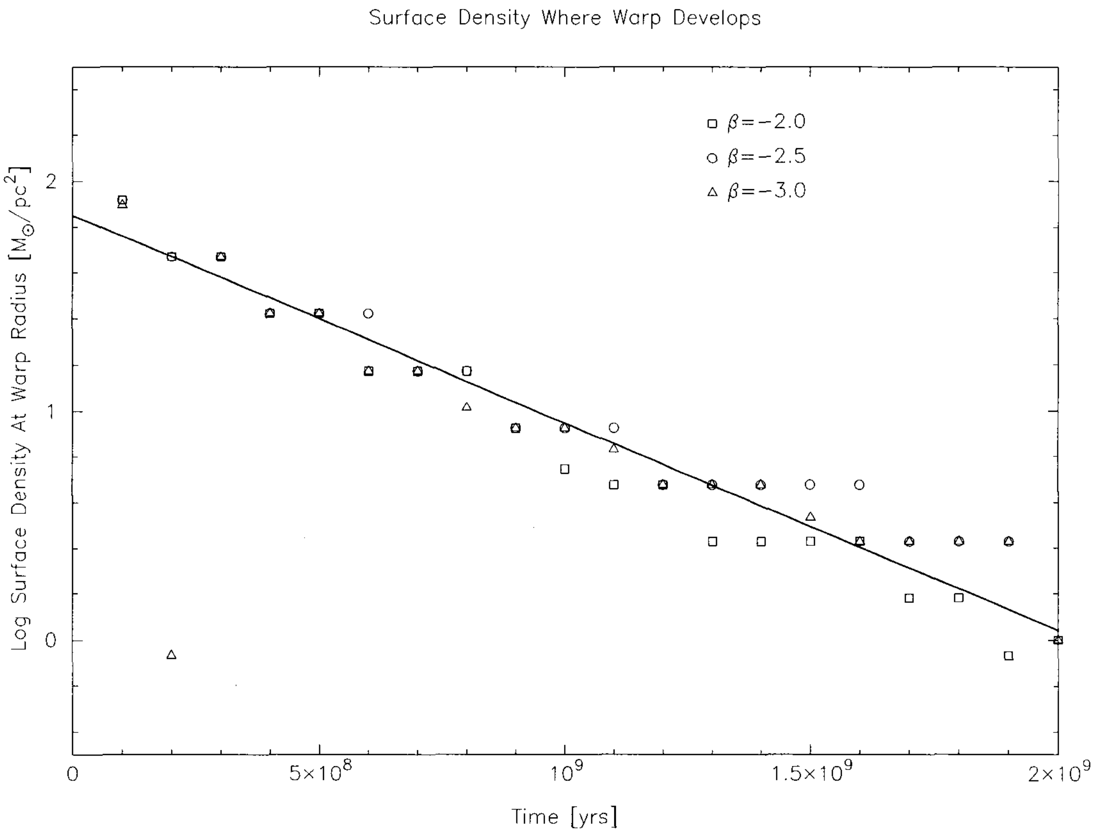

Figure 2.9: As in Figure 2.8, but for 3.0×10^{10} $M_{\odot}$ disk simulations with different torque slopes β.

that determines where the warp develops is the local surface density of the disk. Ostriker & Binney (1989) examined the effect of a slewing disk potential on a set of self-gravitating rings and found qualitatively that regions of high surface density react like a solid body, but that warps can occur where the surface density is lower. Hofner & Sparke (1994) noted that in most cases the group speed of bending waves in a disk of surface density $\Sigma(r)$ and angular rotation velocity $\omega(r)$ is

$$c_g = \frac{\pi G \Sigma(r)}{\omega(r)}, \tag{2.11}$$

and therefore the time for a warp to settle at a given radius is inversely proportional to the surface density at that radius.

To test this, we translate the warp radii of the simulations into local surface densities using

$$\Sigma(r_w) = \frac{M_d}{4\pi r_e^2} e^{-r_w/r_e} \tag{2.12}$$

for a disk of mass M_d and exponential scale length r_e. Figure 2.8 shows the surface density at the warp radius as a function of time for the three disks of different mass in a $\beta = -2.5$ torque (i.e. for the same simulations shown in Figure 2.6), while Figure 2.9 shows the same relation for a $3.0 \times 10^{10}\ M_\odot$ disk in torques with varying β. The surface density at the warp radius falls as the warp moves out through the disk. The local surface densities at the warp *at a given time* are quite similar for all disk masses, and are almost indistinguishable for all torque slopes. It appears that the local surface density is the most important parameter for determining how far the warp settles in a given time.

The evolution of the warp is well described by a decaying exponential

$$\Sigma(r_w, t) = \Sigma_{w0} e^{-t/t_w} \tag{2.13}$$

with $\Sigma_{w0} = 70\ M_\odot\ \mathrm{pc}^{-2}$ and $t_w = 480$ Myr, shown as the straight line in Figures 2.8 and 2.9. For such a result to hold, the timescale t_w should depend only on global properties of the galaxy independent of disk mass or torque; in particular, it can only depend on Σ_{wo}, the virial velocity of the galaxy $v_{200} = 175\,\mathrm{km\,s^{-1}}$, the exponential scale length of the disk $r_d = 3.5\,\mathrm{kpc}$, and/or the vertical scale height $h_z = 325\,\mathrm{pc}$. An interesting timescale that matches this is the characteristic timescale for bending waves (Hofner & Sparke, 1994) one disk scale length away from Σ_{wo}:

$$t_g = \frac{r}{c_g} = \frac{v_{200}}{\pi G \Sigma_{wo} e^{-1}} \tag{2.14}$$

which is 490 Myr in these cases.

Σ_{wo} is the extrapolation of equation (2.13) to $t = 0$, and is the critical surface density above which the disk does not develop a warp. At higher surface densities, the self-gravity of the disk is always sufficient to keep the disk flat. It is interesting that the extrapolation to $t = 0$, when there is no torque, gives a finite well-defined value for the warp surface density. This is why the linear $r_w(t)$ fits do not go through the origin and suggests that at this surface density the disk is marginally unstable to warping. Therefore, we do not expect to see warps occurring at surface densities higher than $70\ M_\odot\ \mathrm{pc}^{-2}$. In the Milky Way, the warp begins at or slightly beyond the solar circle (Binney, 1992), while Hipparcos determinations of the Milky Way surface density in the solar neighbourhood give $\Sigma_0 = 40\,M_\odot\,\mathrm{pc}^{-2}$ (Creze et al., 1998). It is encouraging that this is less than Σ_{wo}. In external galaxies, warps begin between 25 and 26.5 mag arcsec^{-2} (Briggs, 1990), which for a mass to light ratio in B of $1\,M_\odot/L_\odot$ gives surface densities of between 1.8–7.0 $M_\odot\,\mathrm{pc}^{-2}$, ignoring projection effects and extinction. This is also consistent with a picture in which warps only can occur at surface densities below Σ_{wo}.

2.4 Conclusions

When studying disks in misaligned halos, almost all previous work has assumed that it is possible to study the system in isolation and find the equilibrium configuration. Even Hofner & Sparke (1994) and Ideta et al. (2000), who studied the time dependence of such configurations, studied the evolution from isolated initial conditions. However, we have argued that in a cosmological context, the orientations of both halos and disks change on short enough timescales that there is often misalignment between halos and disks regardless of the tendency for isolated disk-halo systems to align themselves. We have studied the effects of such torques on disks, and found that the disks warp in a manner similar to the observed Milky Way warp. We have discovered that such warps begin at the radius where the surface density of the disk drops below $70\ M_{\odot}\ \mathrm{pc}^{-2}$, regardless of the disk mass, and move linearly through the disk at a rate of one scale length per 500 Myr, a timescale characteristic of the timescale of bending waves in the outer (warped) region of the disk. Therefore, if disks and halos are often misaligned, as we have suggested, then such misalignments will produce realistic-looking warps.

CHAPTER 3

WARPS FROM SATELLITE-DISK INTERACTIONS

Some of the contents of this chapter were previously published in Bailin (2003).

3.1 Evidence for coupling between Sgr and the Milky Way warp

3.1.1 Introduction

The closest warped galaxy to us is our own, the Milky Way. Ideally, a description of warped galaxies should at least be able to explain the Milky Way warp, about which we have more information than about any other galaxy. In particular, if warping is a generic response of disks to the individual perturbations they experience, the likelihood of identifying the cause of the Milky Way warp is higher than of identifying the cause of other galactic warps due to the proximity of the satellite galaxies that are likely to perturb the disk.

The Magellanic Clouds have often been proposed as the perturbation responsible for the Milky Way's warp. While Hunter & Toomre (1969) found that the tidal distortion from the clouds alone is not sufficient to cause the observed warp, Weinberg (1998) proposed that orbiting satellites could set up wakes in the Milky Way's halo which could provide the necessary torque. Tsuchiya (2002) performed self-consistent simulations of such a system and confirmed that for a sufficiently massive halo ($2.1 \times 10^{12}\ M_{\odot}$), the magnitude of the torque can be increased enough to cause a warp of the same magnitude as the Milky Way's.

The Magellanic Clouds orbit about the center of the Galaxy in a direction orthogonal to the line of nodes, i.e., near the line of maximum warp. García-Ruiz et al. (2002) demonstrated that the warp caused by a satellite will rock back and forth with its line of nodes oriented *along* the satellite's orbit. A simple way of understanding this result is to recognize that a torque is a transfer of angular momentum, and therefore the disk will acquire angular momentum along the same axis as the orbital angular momentum of the satellite which is providing the torque, and tilt toward that axis. Therefore, the Magellanic Clouds are not suitable candidates for producing the Milky Way warp.

The orbital plane of the Sagittarius dwarf galaxy (Ibata et al., 1994) does intersect the line of nodes, suggesting that it may be a good candidate for producing the Milky Way warp (Lin, 1996). It is located behind the Galactic bulge and is on a nearly polar orbit (Ibata et al., 1997). Ibata & Razoumov (1998) have performed simulations which suggest that the passage of a sufficiently massive Sgr ($5 \times 10^9\ M_\odot$) through the disk could produce a warp. Alternatively, its gravitational tides or the tides of a wake it produces in the dark halo could exert a warp-inducing torque on the disk.

If Sgr is responsible for the warp, its angular momentum will be coupled to that of the warp. In this section, we calculate the orbital angular momentum of Sgr, along with the warped component of the Milky Way disk's angular momentum. We show that they have the same magnitude and are anti-aligned. As there is no *a priori* reason to expect them to be within orders of magnitude of each other, this is evidence that Sgr may be coupled to the warp, and therefore responsible for it.

3.1.2 The angular momentum of satellite galaxies

The position, distance, and motion of Sagittarius are given in Table 3.1, along with estimates of its mass and orbital angular momentum. Estimates of the angular momentum range between 1.7 and $8.6 \times 10^{12}\ M_\odot\ \mathrm{kpc\,km\,s^{-1}}$, and are directed

Table 3.1. Properties of the Sagittarius dwarf

Parameter	Value	
Data from Ibata et al. 1997		
Galactic coordinates	$l = 5.6°, b = -14°$	
Galactocentric distance	16 ± 2 kpc	
Space motion (U, V, W)	$(232, 0, 194) \pm 60$ km s^{-1}	
Galactocentric radial velocity	150 ± 60 km s^{-1}	
Galactocentric tangential velocity	270 ± 100 km s^{-1}	
Derived angular momentum:		
Assumed mass (10^9 $M_\odot$)	0.4	2.0
Angular momentum (10^{12} $M_\odot$ kpc km s^{-1})	1.7 ± 0.6	8.6 ± 3.4
Direction	$l = 276°, b = -7°$	

toward $l = 276°, b = -7°$.

The major uncertainty in this calculation is the determination of the mass. Ibata & Lewis (1998) argued that in order for the satellite to have survived to the present day, it must have a massive extended dark matter halo and a total $M/L \sim 100$ in solar units (Ibata et al., 1997). However, Helmi & White (2001) found viable models with more moderate masses ranging from $4.66 \times 10^8\ M_\odot$ for a purely stellar model to $1.7 \times 10^9\ M_\odot$ for their model with an extended dark matter envelope (see also Jiang & Binney, 2000, who find that if the original mass of Sgr was large enough for dynamical friction to be important, the majority of the mass would have been stripped off after a Hubble time leaving a current mass of 1–3 $\times\ 10^9\ M_\odot$). The properties of Helmi & White (2001)'s models seem most in agreement with the expected properties of dwarf spheroidal galaxies, and therefore we adopt 0.4–$2.0 \times 10^9\ M_\odot$ as the range of possible masses of the Sagittarius dwarf; the lower end of this range is favoured by the tidal stream models of Law et al. (2004).

Table 3.2 shows the magnitude and direction of the angular momenta of Sgr and the 6 other Galactic satellites with measured proper motions, along with that of the Milky Way warp which is calculated in § 3.1.3. The masses and galactocentric radii are also shown for reference. The orbital angular momenta of the Large Magellanic Cloud (LMC) and the Small Magellanic Cloud (SMC) were calculated using data from Kroupa & Bastian (1997). For the remaining satellites, the mass was taken from Mateo (1998) and the velocity vector from the tabulated reference. Two recent measurements of the proper motion of the Fornax dSph give quite different values (Piatek et al., 2002; Dinescu et al., 2004). The reason for this discrepancy is unclear, so both values are tabulated in Table 3.2.

The results of Table 3.2 can be appreciated visually in Figure 3.1. Each line

Table 3.2. Angular momenta of the Milky Way warp and some Milky Way satellites

Object	Mass $(M_\odot)$	Galactocentric radius (kpc)	Angular momentum $(M_\odot \text{ kpc km s}^{-1})$	Direction (deg)	Reference
Milky Way warp	...	...	$1.7\text{–}8.6 \times 10^{12}$	$l = 90, b = 0$	§ 3.1.3
Sgr dSph	$0.4\text{–}2.0 \times 10^{9}$	17	$2.0\text{–}9.8 \times 10^{12}$	$l = 276, b = -7$	1, 2
LMC	2×10^{10}	48	2×10^{14}	$l = 184, b = 9$	3
SMC	2×10^{9}	54	3×10^{13}	$l = 200, b = -1$	3
Fornax dSph	7×10^{7}	140	3×10^{12}	$l = 189, b = 20$	4, 5
			2×10^{12}	$l = 127, b = -7$	4, 6
Ursa Minor dSph	2×10^{7}	66	3×10^{11}	$l = 187, b = -17$	4, 7
Carina dSph	1×10^{7}	101	1×10^{11}	$l = 314, b = 52$	4, 8
Sculptor dSph	6×10^{6}	78	1×10^{11}	$l = 74, b = -7$	4, 9

References. — (1) Ibata et al. 1997; (2) Helmi & White 2001; (3) Kroupa & Bastian 1997; (4) Mateo 1998; (5) Piatek et al. 2002; (6) Dinescu et al. 2004; (7) Schweitzer 1996; (8) Piatek et al. 2003; (9) Schweitzer et al. 1995

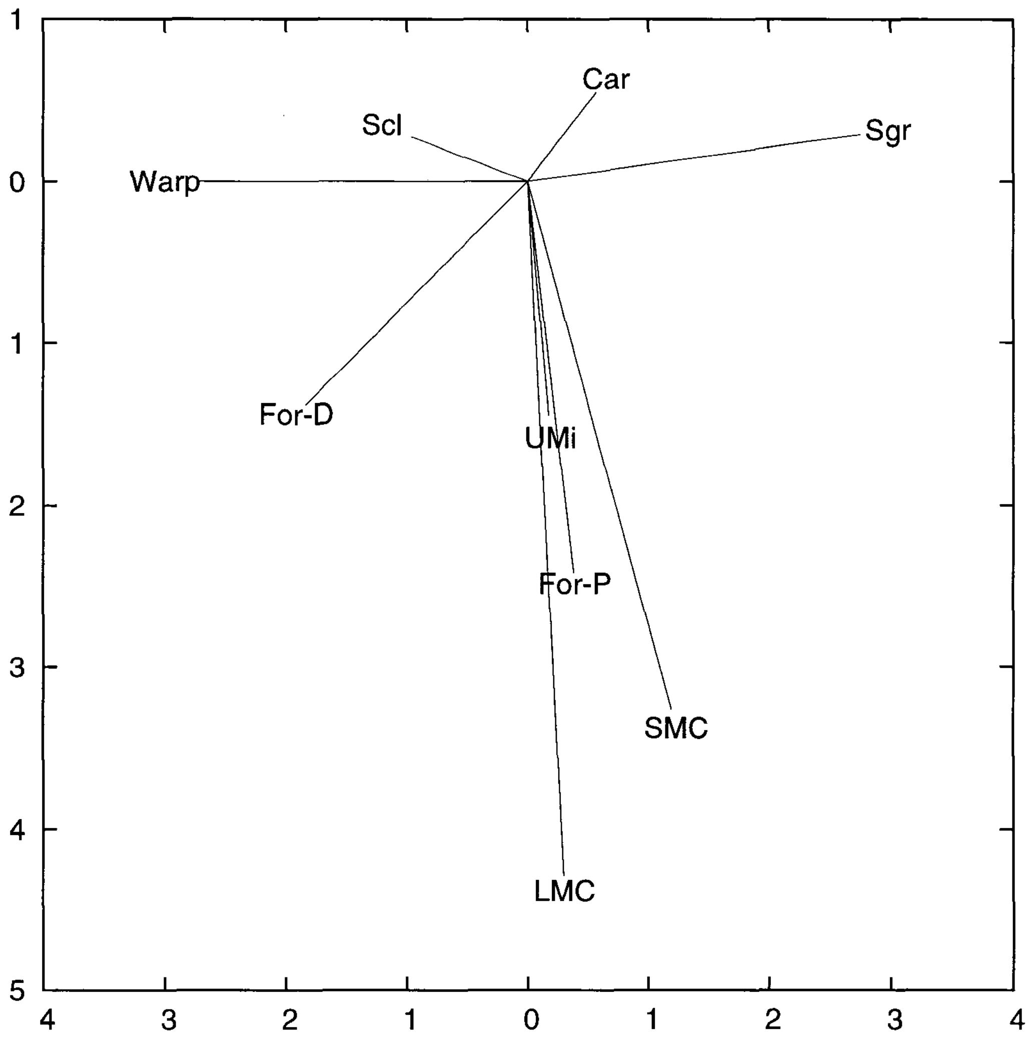

Figure 3.1: The galactocentric angular momentum of each satellite listed in Table 3.2, projected onto the galactic plane, as viewed from the North Galactic Pole (NGP). The length of each vector (including the unseen z component) is $\log L - 10$, where L is the magnitude of the angular momentum in units of $M_\odot$ kpc km s^{-1}. The galactic center is at the origin, while the sun lies below the origin. For-P and For-D refer to the Fornax values derived by Piatek et al. (2002) and Dinescu et al. (2004) respectively. The Milky Way warp and Sgr are both plotted at the mean of the minimum and maximum values tabulated in Table 3.2.

shows the galactocentric direction of the angular momentum vector of a satellite, projected onto the Galactic plane, as viewed from the North Galactic Pole (NGP). The Galactic center is at the origin, while the sun lies below the origin. The length of each vector (including the unseen z component) is $\log L - 10$, where L is the magnitude of the angular momentum in units of $M_\odot$ kpc km s^{-1}. This diagram clearly shows that Ursa Minor and possibly Fornax (if the Piatek et al. (2002) value for its proper motion is more accurate) are on similar orbits to the Magellanic Clouds (Lynden-Bell, 1976; Lynden-Bell & Lynden-Bell, 1995; Palma et al., 2002).

3.1.3 The angular momentum associated with the Milky Way warp

We calculate the component of the disk angular momentum which is due to the warp in the Milky Way's disk, i.e., that which is not directed toward the South Galactic Pole (SGP). Because the line of nodes is straight, this component has the same orientation at all radii. If the disk rises a height $h(R)$ above the plane at cylindrical radius R, then the total angular momentum in the disk which is due to the warp is

$$L_w = \int_{R_w}^{\infty} 2\pi R^2 v_c \Sigma(R) \frac{h(R)}{\sqrt{h(R)^2 + R^2}}\, dR. \tag{3.1}$$

The mass distribution of the disk is taken from Dehnen & Binney (1998) (hereafter DB98). The disk surface density for a given component in these models is given by

$$\Sigma(R) = \Sigma_d \exp\left(-\frac{R_m}{R} - \frac{R}{R_d}\right), \tag{3.2}$$

where Σ_d is the normalization, R_d is the scale length of the component, and R_m is introduced to allow the ISM to have a central depression[1]. $R_m = 4$ kpc for the gas disk and $R_m = 0$ for the stellar disk. The relative contributions to the

---

[1] Note that equation (1) of DB98 has a typo which is fixed above (W. Dehnen 2002, private communication)

Table 3.3. Disk parameters

Parameter	Model			
	1	2	3	4
Data from DB98:				
Stellar disk scale length $R_{d,*}$ (kpc)	2.0	2.4	2.8	3.2
ISM disk scale length $R_{d,\mathrm{ISM}}$ (kpc)	4.0	4.8	5.6	6.4
Surface density at solar circle Σ_0 ($M_\odot$ pc^{-2})	43.3	52.1	52.7	50.7
Derived warp angular momenta[a] (10^{12} $M_\odot$ kpc km s^{-1})				
Stellar disk	1.10	2.16	3.36	4.73
ISM disk	0.52	1.10	1.78	2.57
Total	1.62	3.26	5.14	7.30

[a] Direction is $l = 90° \pm 10°, b = 0°$

surface density Σ_0 at the solar circle $R_0 = 8$ kpc are 0.25 for the ISM and 0.75 for the stars. DB98 distinguish between thin and thick disk components of the stellar disk. However, the difference between their vertical scale heights does not affect the angular momentum, and while the radial scale length of the thick disk is poorly known, the fraction of disk mass in the thick disk is sufficiently small (less than 10% in the DB98 models) that any difference between the radial distribution of the thin and thick disks will have a much smaller effect on the angular momentum than the model-to-model differences. Therefore, we treat the stellar disk as a single component. DB98's models 1–4, which differ primarily in disk scale length, R_d, are all acceptable fits to the observations, and therefore provide a reasonable range of mass distributions with which to estimate the angular momentum. Table 3.3 gives the essential parameters for the four models.

The circular velocity, v_c, of the disk from 3 kpc to the solar circle is ≈ 200 km s^{-1} (e.g., Merrifield, 1992). While most measurements at $R > R_0$ show a rising rotation curve, Binney & Dehnen (1997) argue that a constant rotation curve is consistent with the data when the correlations between errors are taken into account. We adopt $v_c = 200$ km s^{-1} at all radii. The uncertainty in the angular momentum due to uncertainties in the mass models dominates over any error in the circular velocity.

The height of the warp above the plane as a function of radius, $h(R)$, appears to differ for the stars and for the gas. Drimmel et al. (2000) fit Hipparcos measurements of OB stars and find

$$h(R) = \begin{cases} (R - R_w)^2/R_h & R > R_w \\ 0 & R \leq R_w \end{cases}, \tag{3.3}$$

with the warp starting at $R_w = 6.5$ kpc and scaled by $R_h = 15$ kpc. Binney &

Merrifield (1998) approximate the $m = 1$ mode of the ISM warp as

$$h(R) = \begin{cases} (R - R_w)/a & R > R_w \\ 0 & R \leq R_w \end{cases}, \tag{3.4}$$

where $R_w = 10.4$ kpc and $a = 5.6$ when converted to $R_0 = 8$ kpc (Tsuchiya, 2002). Binney & Merrifield (1998) also fit an $m = 2$ mode, but the net angular momentum of any even m mode is aligned with the angular momentum of the flat disk, so it does not contribute.

We use equation (3.3) for the stellar disk and equation (3.4) for the gas disk. The results are shown in Table 3.3. The majority of the warp angular momentum is contained in the range $10 \lesssim R \lesssim 25$ kpc in all models. The Sun lies within $10°$ of the line of nodes, so L_w is directed toward $80° \lesssim l \lesssim 100°$, $b = 0°$.

To facilitate comparison with simulations of warped disks that do not necessarily have the same mass or angular momentum distribution, it is convenient to normalize L_w by the total angular momentum in the disk between 10 and 30 kpc, $L_{\rm tot} = \int 2\pi R^2 v_c \Sigma(R)\, dR$ (for $L_{\rm tot} \gg L_w$). $L_w/L_{\rm tot}$ varies from 0.13 for Model 1 to 0.19 for Model 4. We note that both the total magnitude of the warp angular momentum and the misaligned fraction are comparable to the values found for the Milky Way mass disk in § 2.3.2.

3.1.4 Discussion

Table 3.2 shows the angular momenta of the Milky Way and the seven dwarfs (Grimm & Grimm, 1857). These are also shown visually in Figure 3.1. The magnitude of the angular momentum of Sagittarius is strikingly similar to that of the Milky Way warp, while the direction of the angular momentum is almost exactly opposite. There is no *a priori* reason to expect any relation; the angular momenta of the other satellites with known orbits span three orders of magnitude and a full range of galactic longitude, although there is a strong tendency for the satellites to have polar orbits with low values of $|b|$ (Lynden-Bell, 1976; Lynden-Bell & Lynden-Bell, 1995; Hartwick, 2000; Palma et al., 2002). The only other satellite whose angular momentum is of a similar magnitude is Fornax, but its orientation is unrelated to that of the warp for both measurements of its proper motion. Therefore, the coincidence of the two angular momentum vectors is probable evidence that they are dynamically coupled, i.e., that Sagittarius is the perturber responsible for the Galactic warp.

In order to determine how much of a coincidence this is, we calculate the chance of a randomly-oriented vector lying within 10° of being anti-aligned with the warp angular momentum (as the error in the orientation of the Milky Way line of nodes is 10°, and the orbital angular momentum of Sgr lies within this 10° cone), and find that it is 1.5%. It is more difficult to assess the probability of finding a satellite with the same magnitude of angular momentum as the warp. In order to do this, we take all of the masses, radii, and perpendicular velocities of the satellites tabulated in Table 3.2, and calculate the angular momentum from each combination. We use the mean of the high and low masses for Sgr, and the Piatek et al. (2002) value for Fornax velocity. Out of the 343 combinations, 23 give values of the angular momentum that lie between the high and low values calculated for the Milky Way warp, giving a probability of 7%. Therefore, the probability that the angular momentum of any one satellite is both anti-aligned with that of the Milky Way warp and of the same magnitude is 0.1%, and the the probability of finding at least one such satellite in a sample of 7 is 0.7%. However, it is difficult to objectively interpret such *a posteriori* probabilities. There may be other configurations that would also have been suggestive of a coupling. For example, if the angular momentum of Sgr were perfectly aligned with with that of the warp, rather than anti-aligned, we would also have concluded that they were coupled. We can take this particular case directly into account by doubling the calculated probability; we therefore conclude that there is a 1.4% chance that the relationship we have found between the angular momenta of Sgr and of the warp arises from chance. However, we urge caution when interpreting this number due to the possible existence of other suspicious configurations which we have not thought of.

While it is clear that there should be a relationship between the orientation of the angular momentum of a perturbing satellite and of the excited warp (García-Ruiz et al., 2002), the relationship between the magnitude of the angular momenta is less clear. However, the satellite cannot donate more than its total angular momentum to the disk. Therefore, the satellite can only donate an amount of angular momentum comparable to its orbital angular momentum before its orbit must decay dramatically and be accreted by the Galaxy. Assuming that most of the donated angular momentum goes into the warp, the equilibrium configuration is for the satellite and warp to have comparable amounts of angular momentum due to the coupling.

There are three possibilities for the nature of the coupling. The first is a direct gravitational tidal torque by the satellite itself (Hunter & Toomre, 1969), the second is the gravitational torque of a wake in the Galactic dark matter halo (Weinberg, 1998; Tsuchiya, 2002), and the third is an impulsive deposition of momentum to the gas disk by passage through it (Ibata & Razoumov, 1998). The direct tidal torque for a satellite of mass m and distance r scales as m/r^3. Therefore, the direct tidal effect of Sgr is no stronger than that of the LMC, whose direct tidal torque is not sufficient to induce the warp (Hunter & Toomre, 1969). This means that the gravitational torque of Sgr itself cannot be the coupling mechanism.

If the primary perturber is instead a wake in the halo, the strength of the torque scales as $m_{\rm wake}/r^3_{\rm wake}$. The mass of the wake scales as the mass of the satellite and as the density of the halo at the wake radius (Weinberg, 1998). The wake develops at half the satellite's orbital radius (Tsuchiya, 2002). Therefore, for an isothermal halo, the strength of the torque scales as m/r^5. In this case, the effect of Sagittarius is 10–50 times stronger than that of the LMC. It is plausible that in Tsuchiya (2002)'s lower mass simulation, in which the LMC did not excite a warp, a satellite with Sagittarius's parameters would have. Further simulations which better reproduce the observed Sgr-MW system could confirm or falsify this suggestion.

Ibata & Razoumov (1998) suggest that the impulsive deposition of momentum to the gas disk could excite the warp. The mass they use for Sgr, $5 \times 10^9\ M_\odot$, is quite large, and they find very little warping in their $1 \times 10^9\ M_\odot$ simulation. However, they only model a single interaction. In order for the angular momenta to reach equilibrium, as they appear to have done, there must be repeated or continual encounters. Helmi & White (2001) find orbital periods of $\sim$ 1 Gyr for Sagittarius, indicating that it has passed through the disk several times. Further simulations that follow the evolution of the system over many encounters are necessary to better understand the predictions of this model; meanwhile, it cannot be ruled out.

3.1.5 Summary

The orbital angular momentum of the Sagittarius dwarf galaxy and the component of the Milky Way disk angular momentum due to the Galactic warp are almost perfectly anti-aligned, and both have a magnitude of $2\text{–}8 \times 10^{12}\ M_\odot$ kpc km s^{-1}. We estimate the probability that this configuration arises by chance to be 1.4%. Such a coincidence suggests that they are a coupled system, i.e., that Sgr is responsible for the warp. The direct gravitational tidal torque of Sgr cannot cause the warp. Interaction via a gravitational wake in the Milky Way's dark matter halo, and impulsive deposition of momentum into the disk by passing through it are still both possible coupling mechanisms. More simulations of each of these models are necessary to discriminate between their effects.

3.2 Satellite-disk simulations

3.2.1 Introduction

§ 3.1 presents circumstantial evidence that the Sagittarius dSph galaxy may be responsible for the warp of the Milky Way. However, the simulations of satellite-disk interactions that have been performed to date cannot test this hypothesis. In this section, we present a set of N-body simulations designed to study the interactions between satellites and warps, and test the scenario presented in § 3.1. We first present the details of the simulations, then discuss the physical principles that guide the evolution of the system, and finally examine the results of the simulations to discover the relationship between the nature and orbit of satellites and the warps they cause.

3.2.2 Summary of simulations

The simulations were run using the parallel version of the GADGET N-body code (Springel et al., 2001), and only take into account gravitational forces. Observed galactic disks contain not only collisionless stars but also gas that experiences hydrodynamic forces. The warp mechanism is purely gravitational, so hydrodynamic forces are not required to accurately model the transfer of angular momentum from the satellite to the disk. However, the internal reaction of a gaseous disk may differ from that of a disk made up of collisionless particles due to pressure forces and viscous damping. To estimate the effect of the pressure forces, we note that the sound speed in an isothermal gas at 10^4 K is $c_s \sim 10 \text{ km s}^{-1}$.

The radial range over which the warp occurs is from 10 kpc to 30 kpc. Therefore, the timescale over which the outer disk comes to pressure equilibrium is $\sim 10^9$ yr, which, as we shall see, is several times longer than the timescale over which the disk warps, and so neglecting pressure forces should not change the magnitude of the excited warps. The other possible hydrodynamic effect is viscous damping. However, the vertical shear due to the warp is certainly smaller than the shear due to differential rotation of the disk. Therefore, the lack of significant accretion flows in the outer regions of galactic gas disks implies that the the dissipation due to viscosity is much too small to affect the results.

The simulations were run using the parallel version of the GADGET N-body code (Springel et al., 2001). The galaxy model is based on the Milky Way models of DB98, particularly model 3 which has a disk scale length of 2.8 kpc. The disk vertical scale height and the bulge scale length were both set to 0.1 times the disk scale length, or 280 pc. The models were created using the method of Hernquist (1993), with a Toomre Q parameter in the disk of 1.5 (Toomre, 1964). The models were oriented such that the angular momentum was directed along the $-\hat{z}$ axis, i.e. the rotation is clockwise as seen from the NGP, along the $+\hat{z}$ axis. The location of the sun is along the $\hat{x}$ axis at Cartesian coordinates (-8,0,0) kpc. The angular momentum of Sgr in this model is directed along the $-\hat{y}$ axis, i.e. at $b = 0°$, $l = 270°$.

The models consist of 1048576 halo particles, 969438 disk particles and 79138 bulge particles (see Table 3.4). The satellite consists of a single particle containing the entire satellite mass. Because the primary interest of this work was to see whether the Sagittarius dSph galaxy could excite the Galactic warp, we chose two satellite models that bracket the range of estimated masses and sizes of Sgr. The first, labelled "S1", is based on Helmi & White (2001)'s model which includes a

Table 3.4. Mass components of the disk galaxy models.

Component	Total mass ($10^{10} M_\odot$)	N	Particle mass ($10^4 M_\odot$)	Softening pc
Halo	89.565	1048576	85.4158	550
Disk	4.45166	969438	4.5920	200
Bulge	0.36489	74138	4.6108	200
Total	94.38155	2097152		
Medium satellite (S1)	0.2	1	2×10^5	2000
Small satellite (S2)	0.04	1	4×10^4	5000
Large satellite (S3)	1.0	1	1×10^6	2000

dark halo, and has a mass of $2 \times 10^9 M_\odot$ and a radius (represented by its softening length) of 2 kpc, while the second, labelled "S2", is based on the Helmi & White (2001) constant mass-to-light model, and has a mass of $4 \times 10^8 M_\odot$ and a radius of 5 kpc. We have also created a more massive satellite model, S3, with a mass of $1 \times 10^{10} M_\odot$. This could represent the LMC or Sgr before it lost significant mass through its tidal interaction with the Milky Way, and also helps elucidate how the satellite mass affects the strength of the coupling between satellites and the warp.

Because the satellite is represented by a single particle, its internal degrees of freedom are not taken into account. This is only expected to have small effects on the results because even if the satellite were rotating at break-up speed, its internal angular momentum would be more than an order of magnitude smaller than its orbital angular momentum, so the internal rotation cannot be an important reservoir of angular momentum in the system unless the satellite is disrupted. Observed satellite galaxies do lose mass and eventually disrupt in the tidal field of the parent galaxy, losing approximately 30–40% of their mass per orbit (Taylor & Babul, 2001, 2004). Even then, the stripped stars continue in their original orbital path about the parent galaxy, so their gravitational effects are still felt by the disk, though no longer in phase with the body of the satellite. Treating the satellite as a single particle simplifies the interpretation of the physical processes in our simulation; however, it also means our results are only valid during the time before which a fully self-consistent satellite would disrupt or lose most of its mass, i.e. for only a couple of orbits.

The initial orbital parameters were chosen to mimic those of Sgr; however, a wide range of additional orbits were simulated in order to understand the effect of the orbital eccentricity, and inclination. The initial conditions of the simulations are listed in Table 3.5. Orbits were either polar, inclined 45° prograde to the disk rotation, or inclined 45° retrograde to the disk rotation, and either had the same apo- and perigalactica as Sgr, or were on relatively circular orbits with the same angular momentum (the "r20" models). In addition, one model was run with no satellite ("undist"). The orbital parameters of the satellites, based on the first perigalactic and apogalactic passages, are given in Table 3.6.

Some of the simulations were performed on the Astrophysikalisches Institut Potsdam beowulf cluster *octopus*, and some on the University of Arizona High Performance Computing Center alpha supercomputer *aura*, each using 16 processors. Each simulation was evolved from $t = 0$ to $t = 4$ Gyr. The time step criterion was

$$\Delta t = \sqrt{2\eta\epsilon/|\mathrm{a}|}, \tag{3.5}$$

where ϵ is the force softening length, a is the acceleration, and η was set to a con-

Table 3.5. Initial position and velocity of satellites in the simulations.

Name	Satellite	Initial position (kpc)	Initial velocity (km s^{-1})	Inclination
undist	...	...	...	...
S1-i0	S1	(-70,0,0)	62	polar
S1-i+45	S1	(-70,0,0)	62	prograde
S1-i-45	S1	(-70,0,0)	62	retrograde
S2-i0	S2	(-70,0,0)	62	polar
S2-i+45	S2	(-70,0,0)	62	prograde
S2-i-45	S2	(-70,0,0)	62	retrograde
S3-i0	S3	(-70,0,0)	62	polar
S3-i+45	S3	(-70,0,0)	62	prograde
S3-i-45	S3	(-70,0,0)	62	retrograde
S1-r20-i0	S1	(-20,0,0)	220	polar
S2-r20-i0	S2	(-20,0,0)	220	polar

Table 3.6. Initial orbital parameters of the satellites in the simulations.

Name	Apogalacticon (kpc)	Perigalacticon (kpc)	Semi-major axis (kpc)	Eccentricity	Period (Gyr)
S1-i0	70.0	12.7	41.4	0.693	1.15
S1-i+45	70.0	12.8	41.4	0.690	1.15
S1-i-45	70.0	12.9	41.5	0.688	1.15
S2-i0	70.0	13.1	41.5	0.685	1.20
S2-i+45	70.0	13.2	41.6	0.683	1.20
S2-i-45	70.0	13.2	41.6	0.682	1.20
S3-i0	70.0	11.6	40.8	0.717	1.00
S3-i+45	70.0	11.7	40.9	0.713	1.00
S3-i-45	70.0	11.6	40.8	0.717	1.00
S1-r20-i0	28.5	20.0	24.3	0.175	0.60
S2-r20-i0	29.2	20.0	24.6	0.186	0.65

servative value of 0.01 to ensure that vertical oscillations were not artificially amplified by integration errors. This resulted in an average timestep of 8×10^{-5} Gyr. The typical simulation therefore had $\sim 5 \times 10^4$ timesteps, and took approximately 27 days on the beowulf cluster or 5 days on the alpha. Snapshots were output every 0.05 Gyr. Because mass loss and tidal disruption, which do not occur in our simulations, would have a significant impact on the satellite after a couple of orbits, we only analyze each simulation up to the apogalacticon after the second perigalactic passage.

These simulations are similar in spirit to the simulations Tsuchiya (2002) uses to examine the effects of the LMC on the Milky Way. However, there are some important differences. First, as we are most concerned with the effect of Sgr on the Milky Way disk, our fiducial model has a satellite with a lower mass on an orbit closer to the parent galactic disk; however, we also explore a wide range of satellite masses and orbits (including satellites of LMC mass) to fully probe the physics of the interaction. Secondly, we have been very conservative with our numerical parameters to ensure that our simulations are accurate representations of the physical system; our models contain twice as many disk and halo particles as Tsuchiya (2002), ensuring that we have the 10^6 particles in the halo required to resolve the wake (Weinberg, 1998), and our average timestep is an order of magnitude smaller than that used by Tsuchiya (2002). In addition, we use a single N-body algorithm, a tree code, rather than the mixture of tree and Self-Consistent Field algorithms used by Tsuchiya (2002). Finally, we note that our halo, based on the Milky Way models of DB98, is slightly more massive than the low halo mass model of Tsuchiya (2002) (his "model S").

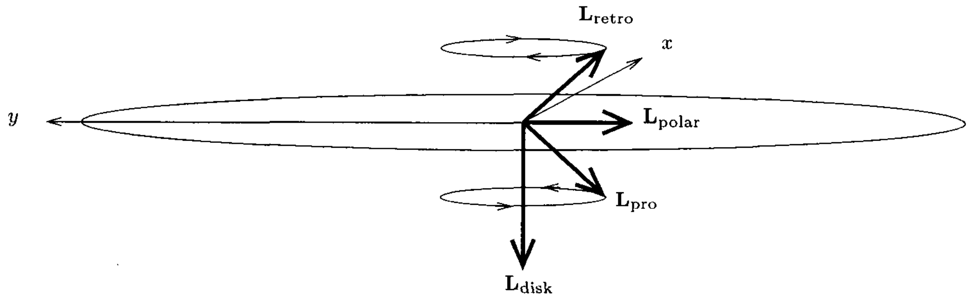
$\mathbf{L}_{\text{retro}}$
x
y
$\mathbf{L}_{\text{polar}}$
$\mathbf{L}_{\text{pro}}$
$\mathbf{L}_{\text{disk}}$

Figure 3.2: Diagram demonstrating the precession of a satellite on a prograde, retrograde, or polar orbit around the flattened disk potential.

3.2.3 Precession and sinking

The satellite interacts with the galaxy, transferring angular momentum. There are two main processes by which the satellite and galaxy transfer angular momentum. The first is the sinking of the satellite. Dynamical friction transfers orbital energy (and consequently angular momentum) from the satellite into the halo and disk (Byrd et al., 1986; Bontekoe & van Albada, 1987; Cora et al., 1997; Colpi et al., 1999; Taylor & Babul, 2001, 2004). This transfer of angular momentum causes the disk to tilt (Huang & Carlberg, 1997). The angular momentum transferred has the same direction as the satellite's angular momentum, and magnitude L_{shrink}. L_{shrink} depends on the rate at which the satellite sinks and on the magnitude of the satellite's initial angular momentum. The rate at which the satellite sinks is a function of the mass of the satellite, its velocity, the local halo density, and the velocity dispersion of the halo (Chandrasekhar, 1943; Cora et al., 1997). For a given halo model, it therefore depends on the satellite mass and its galactocentric radius.

The second process that affects the angular momentum of the satellite is precession of the satellite in the non-spherical potential of the disk. The orbit precesses clockwise for retrograde orbits, counter-clockwise for prograde orbits, and not at all for polar orbits, as viewed from the NGP. Figure 3.2 demonstrates this for the orbits in our simulations. We refer to the magnitude of the change in angular momentum due to precession as $L_{\rm precess}$. The angular momentum acquired by the galaxy is opposite in direction to $L_{\rm precess}$, and is therefore directed along $-\hat{x}$ ($+\hat{x}$) for satellites on prograde (retrograde) orbits that are initially directed toward $-\hat{y}$. Through two-body interactions, the disk also donates some of its $-L_z$ angular momentum to the satellite, so after their first disk passage, the satellites which were initially on polar orbits become very slightly prograde, resulting in very slight precession for the lowest mass satellites. $L_{\rm precess}$ depends on the inclination of the satellite orbit, the flattening of the potential (which is a strong function of radius, as the relatively spherical halo dominates at large radius while the disk dominates at small radius), and the magnitude of the satellite's initial angular momentum.

These two processes transfer angular momentum from the satellite to different galactic components. In particular, the dynamical friction force comes from both the disk and the halo, while the precession of the satellite's orbit is due to the flattened potential of the disk. Therefore, by Newton's Third Law (Newton, 1687), $L_{\rm shrink}$ is initially transferred partially to the halo and partially to the disk, while $L_{\rm precess}$ is transferred entirely to the disk.

When examining the simulations, we found an additional source of angular momentum: the halo. Because the method of Hernquist (1993) involves sampling the velocity distribution randomly, it is possible to end up with a small net component of angular momentum which is misaligned with the disk due to shot noise. In the case of our initial conditions, 0.26% of the halo angular momentum is misaligned with the disk. However, due to the large mass of the halo, this amounts to 3.6×10^{12} $\mathrm{M}_{\odot}$ kpc km s^{-1}, comparable to that in the satellite. In most of the simulations, the effects of the satellite were dominant, but in some simulations the coupling between the disk and the satellite was small, and therefore the disk tilted slightly and warped due to the presence of this additional halo component. We use the undisturbed simulation to calibrate out the effect of this, as discussed in § 3.2.5.

3.2.4 Tilting and warping disks

There are four physical effects that govern the evolution of the disk in the presence of the satellite: tilting of the disk due to L_{precess}, tilting and warping of the disk due to L_{shrink}, warping of the disk due to a slewing potential, and differential precession of the warped outer disk. The direction that the disk tilts can be deduced from the direction of the angular momentum it acquires. As demonstrated by Figure 3.2, the initial orientation of the disk angular momentum is $-L_z$. We now examine in detail each of these physical effects.

3.2.4.1 Tilting due to L_{precess}

The satellite precesses in the flattened potential of the disk. The potential of the entire disk is responsible, and therefore the exchange of angular momentum occurs uniformly across the disk. For prograde satellites, the satellite acquires $+L_x$ and the disk acquires $-L_x$, while for retrograde satellites, the satellite acquires $-L_x$ and the disk acquires $+L_x$. Therefore, to first order, the disk tilts uniformly about the $+\hat{y}$ axis for prograde satellites and about the $-\hat{y}$ axis for retrograde satellites.

3.2.4.2 Tilting and warping due to L_{shrink}

Dynamical friction transfers angular momentum from the satellite to the disk and halo. In all cases, this is directed toward $-L_y$, and therefore the disk and halo both tilt about the $-\hat{x}$ axis. Dynamical friction is a local process, so the effects of the satellite are strongest at radii through which it orbits. In none of our simulations does the satellite pass into the inner disk; therefore, the outer disk acquires a proportionally larger fraction of $L_{\rm shrink}$ and tilts further.

3.2.4.3 Warping of the disk due to a slewing potential

Ostriker & Binney (1989) determined that if the potential of the disk slews about an axis $\hat{\Omega}$, the angular momenta of rings within the disk tilt toward $\hat{\Omega}$. The effective torque of the tilting disk potential is proportional to its surface density, and therefore falls off exponentially with radius. As a result, the inner disk responds faster and the disk warps.

This acts as a secondary response for both processes discussed above. For example, if the disk tilts about the $+\hat{y}$ $(-\hat{y})$ axis due to precession of a prograde (retrograde) satellite, the inner disk acquires $+L_y$ $(-L_y)$ due to this effect. Similarly, if the disk and halo tilt about the $-\hat{x}$ axis due to $L_{\rm shrink}$, the inner disk acquires $-L_x$ due to this effect.

3.2.4.4 Differential precession

When the disk is warped, the outer disk precesses about the plane of the inner disk counter-clockwise when viewed from the NGP. The rate of precession is a function of radius and of warp amplitude. This differential precession quickly damps the warp in the absence of the above processes (Kahn & Woltjer, 1959; Hunter & Toomre, 1969).

3.2.5 Simulation results

Here we examine the angular momentum in the simulations. For most purposes, we split the disk up into two radial regions: the inner disk, with radii between 0–10 kpc, and the outer disk, with radii between 10–30 kpc. In § 3.1.3, we found that most of the warped angular momentum in the Milky Way is contained between

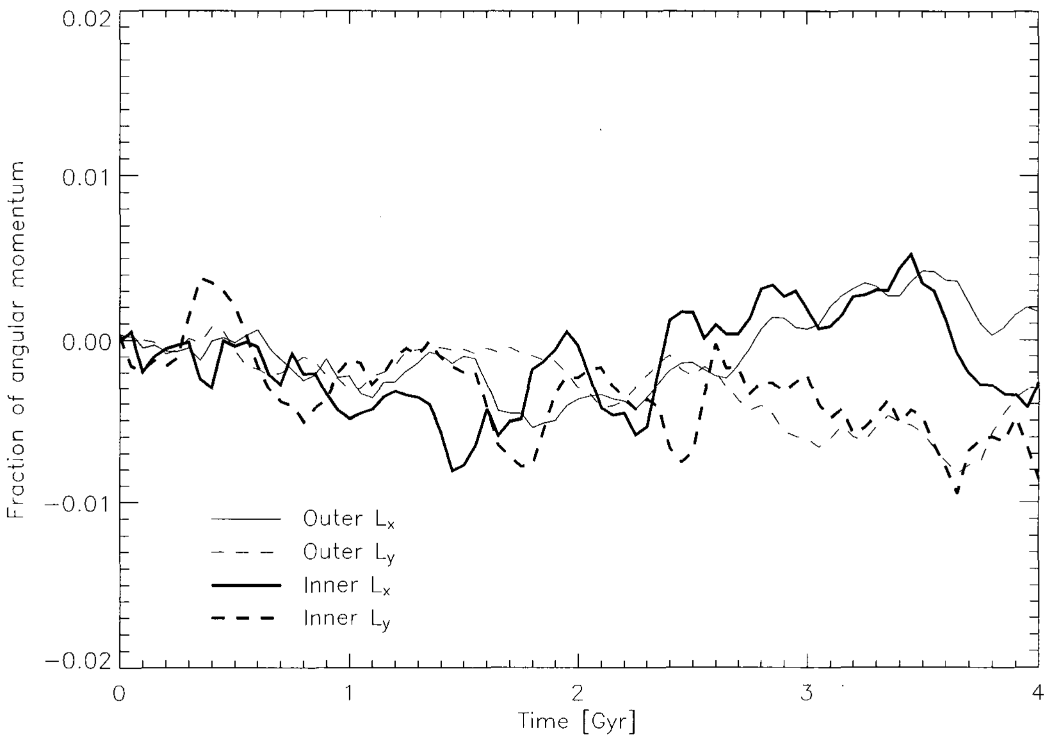

Figure 3.3: Evolution of the x (solid) and y (dashed) components of the angular momentum of the undisturbed simulation, as a fraction of the total angular momentum. Thick lines represent the innermost 10 kpc of the disk, while the thin lines represent the outer 10 – 30 kpc of the disk.

10 and 25 kpc, so we expect the difference between our outer and inner regions to represent the difference between the warped and unwarped regions of the disk.

The disk in the undisturbed simulation tilts slightly due to small amounts of misaligned angular momentum in the halo. In Figure 3.3, we plot the evolution of the x and y components of the angular momentum of the inner and outer regions of the disk in the undisturbed simulation, as a fraction of the total angular momentum in each region of the disk. This is equivalent to the angle in radians

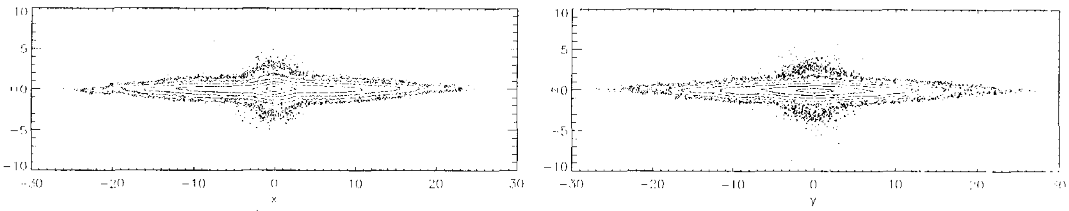

Figure 3.4: Log-spaced contours of the projected density of the bulge and disk in the undisturbed simulations, viewed edge-on from two orthogonal directions at time $t = 1.8$ Gyr, when the angular momenta of the inner and outer disk are most misaligned (see Figure 3.3). A very slight warp is apparent on the left edge of the x-z projection beyond 20 kpc, and in the y-z projection between 12 and 18 kpc.

that each region of the disk is tilted. These changes produce only an extremely mild warp; Figure 3.4 shows a contour plot of the disk at $t = 1.8$ Gyr, when the difference between the angular momentum in the inner and outer disk is particularly large (and is in the time span analyzed in most of the simulations with satellites). The global tilt of the disk in this and all following projected images (which is extremely small in the undisturbed simulation, but can be quite large in the simulations with satellites) has been removed to enhance our ability to discern any warp that may exist. These images show a very small warp due to the exchange of angular momentum between the disk and the halo (Debattista & Sellwood, 1999). When examining all future simulations, we subtract out the undisturbed simulation. In particular, we define

$$\tilde{L}_i^{\text{inner}} \equiv \frac{L_i^{\text{inner}}}{L^{\text{inner}}} - \frac{L_{\text{undist},i}^{\text{inner}}}{L_{\text{undist}}^{\text{inner}}} \tag{3.6}$$

and

$$\tilde{L}_i^{\text{outer}} \equiv \frac{L_i^{\text{outer}}}{L^{\text{outer}}} - \frac{L_{\text{undist},i}^{\text{outer}}}{L_{\text{undist}}^{\text{outer}}}, \tag{3.7}$$

for $i = x, y$.

The magnitude of the warp is well described by the difference in magnitude between $\tilde{\mathbf{L}}^{\text{inner}}$ and $\tilde{\mathbf{L}}^{\text{outer}}$. We define the warp angular momentum as $\tilde{\mathbf{L}}_w$, which has magnitude

$$\tilde{L}_w \equiv \sqrt{(\tilde{L}_x^{\text{outer}} - \tilde{L}_x^{\text{inner}})^2 + (\tilde{L}_y^{\text{outer}} - \tilde{L}_y^{\text{inner}})^2}, \tag{3.8}$$

and is oriented toward galactic latitude $b = 0°$ and galactic longitude l, where

$$\tan l = \frac{\tilde{L}_y^{\text{outer}} - \tilde{L}_y^{\text{inner}}}{\tilde{L}_x^{\text{outer}} - \tilde{L}_x^{\text{inner}}}. \tag{3.9}$$

Note that $\tilde{L}_w$ is the fraction of the angular momentum which is misaligned, so in order to compare $\tilde{L}_w$ with the warp angular momentum L_w calculated in § 3.1.3, we must multiply by the total outer disk angular momentum L^{outer}, which is initially $1.6 \times 10^{13}\ M_\odot$ kpc km s^{-1}, but increases to $2.2 \times 10^{13}\ M_\odot$ kpc km s^{-1} during the course of the simulation due to redistribution of mass throughout the disk.

3.2.5.1 Orbital inclination

We compare the S1-i0, S1-i+45, and S1-i-45 simulations to study the effects of orbital inclination on the propensity for satellites to excite warps. In Figure 3.5, we plot the x and y components of the angular momenta $\tilde{\mathbf{L}}^{\text{inner}}$ and $\tilde{\mathbf{L}}^{\text{outer}}$. The primary effects of L_{precess} are clearly seen: in the prograde simulation, the disk acquires a significant amount of $-L_x$ whenever the satellite passes through the disk, while in the retrograde simulation the disk acquires $+L_x$. As a secondary effect, the slewing potential induces the angular momentum to tilt toward the slewing vector, acquiring $+L_y$ in the prograde case and $-L_y$ in the retrograde case, with the inner disk tilting further than the outer disk. The evolution in the polar simulation is entirely due to L_{shrink}, which does not cause as much evolution of $\tilde{\mathbf{L}}^{\text{inner}}$ and $\tilde{\mathbf{L}}^{\text{outer}}$. The disk acquires a small amount of $-L_y$ directly, particularly in its outer region, and also feels the effect of the slewing potential, resulting in a

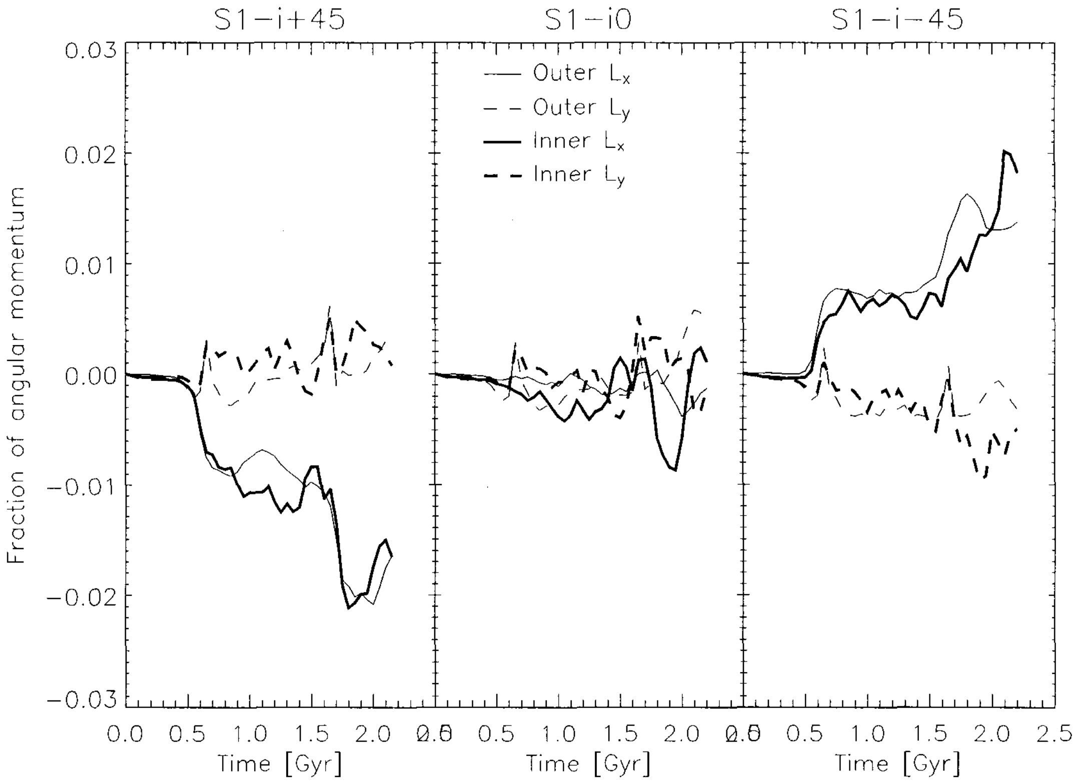

Figure 3.5: Evolution of the the x (solid) and y (dashed) components of the angular momenta $\tilde{\mathbf{L}}^{\rm inner}$ (thick lines) and $\tilde{\mathbf{L}}^{\rm outer}$ (thin lines) of the galactic disks in the prograde (left), polar (middle), and retrograde (right) simulations with intermediate-mass satellite S1. Each simulation is plotted until the apogalacticon after its second perigalactic passage. Perigalactic passages occur at $t = 0.60$ and $t = 1.70$ Gyr.

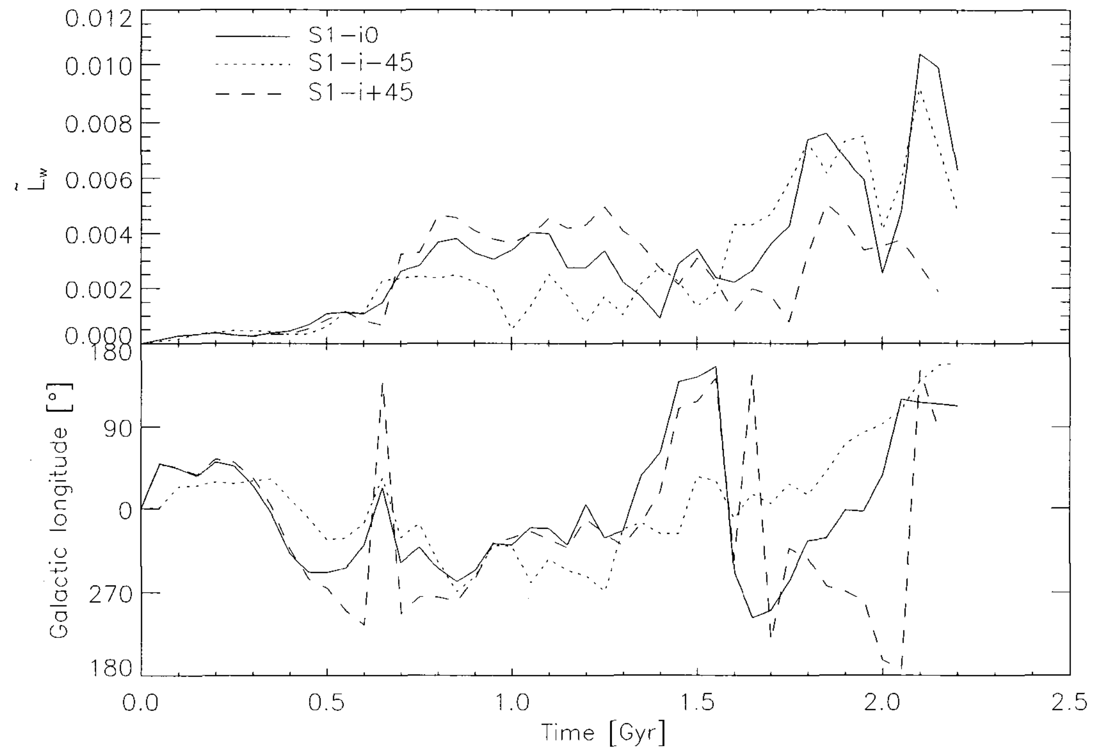

Figure 3.6: Evolution of the warp angular momentum $\tilde{L}_w$ of the galactic disks in the prograde (dashed), polar (solid), and retrograde (dotted) simulations with satellite S1. In the top panel, we plot the magnitude of $\tilde{L}_w$, while in the bottom panel we plot the galactic longitude l that it is directed toward. Note that these can wrap around the $l = 180°$ point. Each simulation is plotted until the apogalacticon after its second perigalactic passage. Perigalactic passages occur at $t = 0.60$ and $t = 1.70$ Gyr.

small decrease in L_x, particularly for the inner disk. These effects are also seen in the prograde and retrograde simulations.

In Figure 3.6, we have plotted the magnitude of the warp angular momentum $\tilde{L}_w$, and its orientation. The warp is excited at the first perigalactic passage at $t = 0.60$ Gyr, reaching its maximum ~ 200 Myr later, or one disk rotation period at $R = 10$ kpc. Afterwards, it oscillates and slowly declines until the satellite approaches a second time at $t = 1.70$ Gyr, when the pattern repeats. The orientation of the warp is similar for all satellites after the first perigalactic passage. This is a direct result of the processes discussed in § 3.2.4: L_{shrink} causes the outer disk to acquire $-L_y$ with respect to the inner disk in all cases, which on its own would result in a warp orientation of $l = 270°$. The secondary effect of L_{shrink} is that the inner disk acquires $-L_x$ with respect to the outer disk, which results in an increase in l, i.e. it trails slightly behind the satellite orbital angular momentum with respect to the galactic rotation. The direct effect of L_{precess} is a uniform tilt of the disk, which causes no warp at all, while the secondary effect is that the inner disk acquires $+L_y$ $(-L_y)$ for the prograde (retrograde) satellite, which slightly decreases (increases) l, and therefore the warp angular momentum leads (trails) the polar value with respect to the galactic rotation (although the effect on the orientation is small, since the primary effect of L_{shrink} already dominates the L_y component). However, it also results in a larger value of $\tilde{L}_w$ for the prograde simulation, where the effects of L_{precess} and L_{shrink} have the same sign, than for the retrograde simulation, where they have opposite signs. Between perigalactic passages, differential precession damps the warp causing the rise in galactic longitude and decline in amplitude of $\tilde{L}_w$. The second passage of the satellite is qualitatively similar to the first; however, the satellite spends more time at smaller radii due to its shrunken orbit, and therefore the effects are significantly

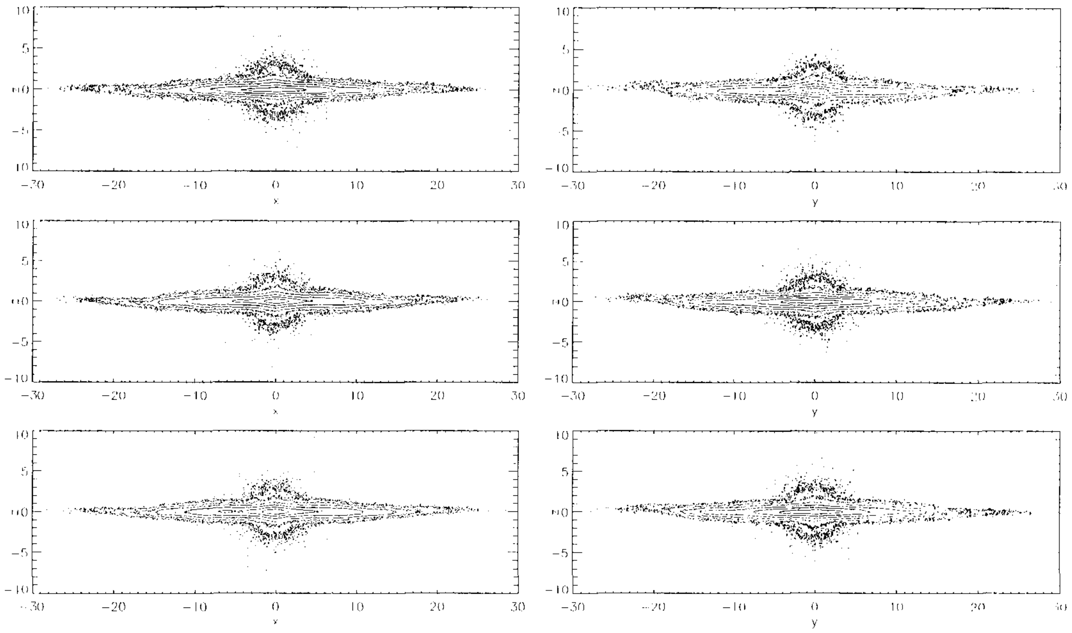

Figure 3.7: Log-spaced contours of the projected density of the bulge and disk in the prograde (top), polar (middle), and retrograde (bottom) simulations for satellite S1 at the time when $\tilde{L}_w$ reaches its first maximum. This occurs at $t = 0.80$ Gyr for S1-i+45 (the prograde simulation), and $t = 0.85$ Gyr for S1-i0 and S1-i-45. The left panels are the x-z projections, while the right panels are the y-z projections. The global tilt of the disk has been removed.

stronger.

After the first perigalactic passage, Figure 3.6 indicates that the peak of $\tilde{L}_w$ is around $t = 0.80$ Gyr for all three simulations. Projected images of the three simulations at the snapshot where $\tilde{L}_w$ is maximized are shown in Figure 3.7. The orientation of the warp is $l = 265°$, $282°$, and $271°$ for the prograde, polar, and retrograde simulations respectively, suggesting that the disk warps up at negative y/down at positive y in all cases with a very small warp up at positive x/down at negative x for the polar and possibly the retrograde simulation (which has a

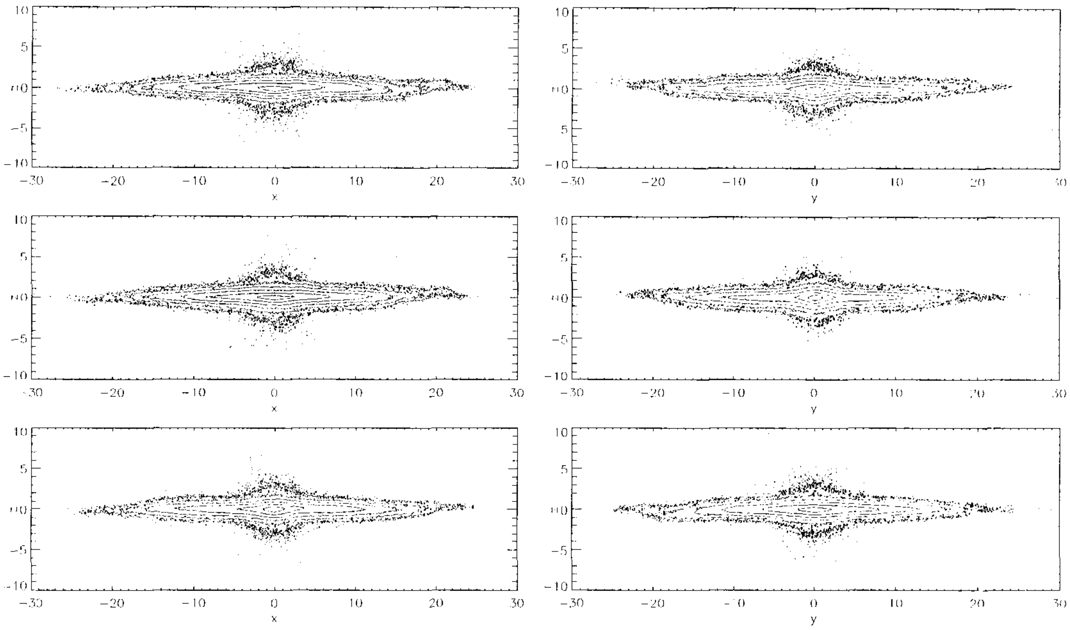

Figure 3.8: Log-spaced contours of the projected density of the bulge and disk in the prograde (top), polar (middle), and retrograde (bottom) simulations for satellite S1 at the time when $\tilde{L}_w$ peaks after the second perigalacticon. This occurs at $t = 1.95$ Gyr for S1-i-45 (the retrograde simulation), and $t = 1.85$ Gyr for S1-i0 and S1-i+45. The left panels are the x-z projections, while the right panels are the y-z projections. The global tilt of the disk has been removed.

larger value of l for most other snapshots around the same time). A warp up at negative y is clearly seen in all simulations, with a maximum amplitude of about 1 kpc. There are also very slight warps down at negative x in all simulations and up at positive x in the polar and retrograde simulations. The amplitude is much smaller than in the y-z projection, never larger than 0.5 kpc.

The warp amplitudes rise to a much higher level after the second perigalactic passage. There are several peaks in $\tilde{L}_w$ for each simulation after the second passage, as differential precession and vertical bending waves transfer the angular momentum through the disk. We examine the first snapshot where images of the simulation show a clear warp. This is $t = 1.95$ Gyr for the retrograde simulation, and $t = 1.85$ Gyr for the polar and prograde simulations. Projected images of the disks are shown in Figure 3.8. The qualitative discussion regarding the warp orientation after the second perigalacticon is the same as after the first: the warp is directed toward lower galactic longitude for prograde satellites and higher galactic longitude for retrograde satellites. However, the effects are strongly amplified due to the smaller orbital radius of the satellite the second time around, and precession modifies the orientation on quite short timescales. In particular, the orientation of the warp is $l = 276°, 329°$, and $84°$ for the prograde, polar, and retrograde simulations respectively. Therefore, we expect a warp up at negative y/down at positive y for the prograde and polar simulations, an exact opposite orientation for the y-z projection of the warp in the retrograde simulation, and a warp up at positive x/down at negative x for the polar simulation. The y-z projections bear out this picture (though the polar simulation appears more like a U-shaped warp, rising on both sides), while all three simulations warp up slightly at positive x (and, for the polar and retrograde satellites, down at negative x). The amplitude of the warp appears larger than after the first satellite passage, although it still never rises much beyond 1 kpc, and appears more like the moderate warps seen in most disk galaxies (Schwarzkopf & Dettmar, 2001). Note that at late times, the warp of the S1-i0 simulation, which was constructed to be similar to the true orbit of Sgr, settles to $l \approx 90°$, similar to the warp of the Milky Way. However, the magnitude of the warp angular momentum peaks at $2 \times 10^{11}\ M_\odot$ kpc km s^{-1}, an order of magnitude smaller than the magnitude of the warp angular momentum L_w found in § 3.1.3.

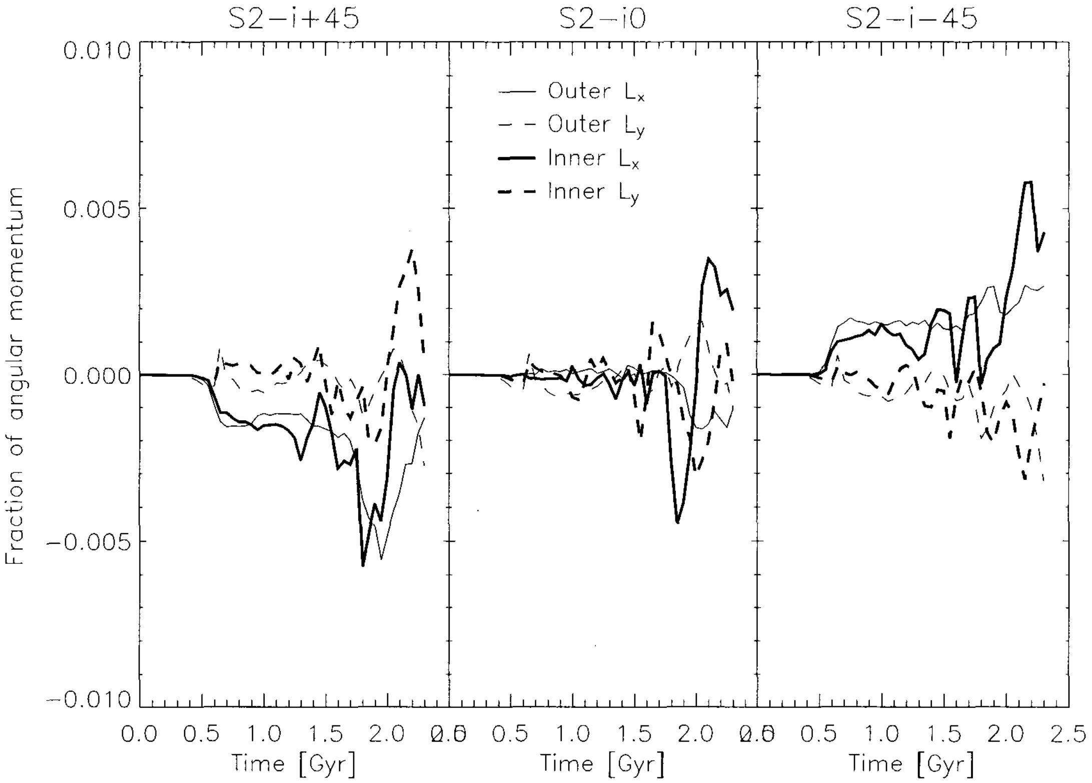

Figure 3.9: As in Figure 3.5, but for the low-mass satellite S2. Perigalactica occur at $t = 0.60$ and $t = 1.75$ Gyr.

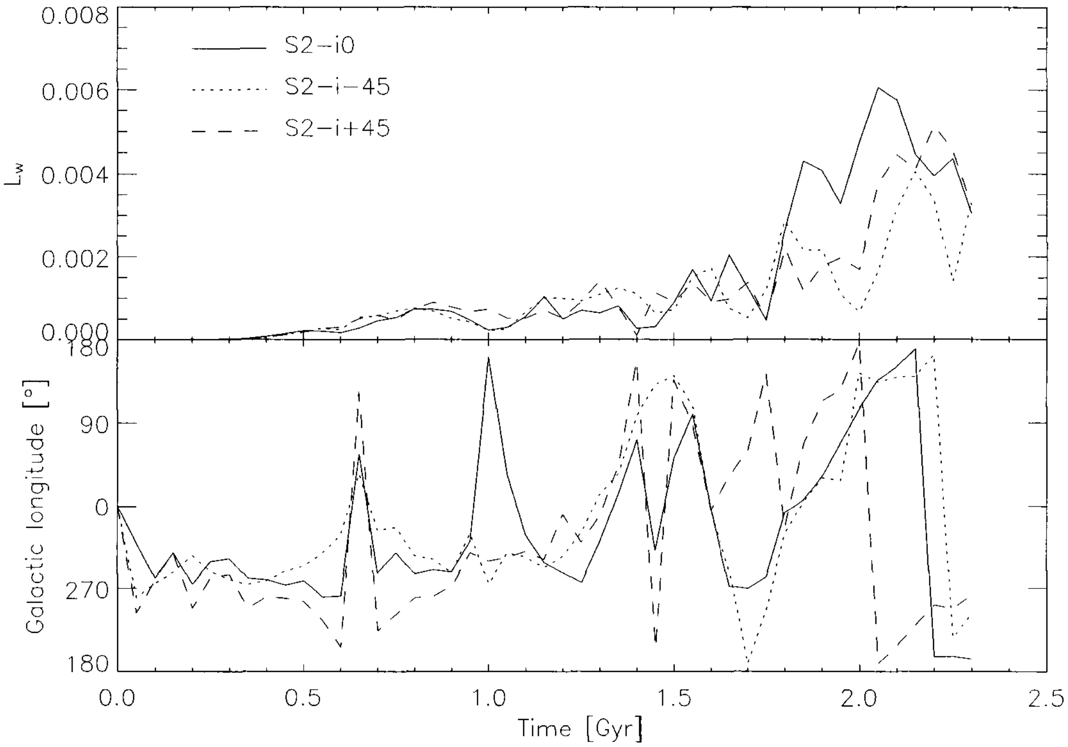

Figure 3.10: As in Figure 3.6, but for the low-mass satellite S2. Perigalactica occur at $t = 0.60$ and $t = 1.75$ Gyr.

3.2.5.2 Satellite mass

To determine the effect of the satellite mass on its ability to warp the disk, we compare the low mass simulations S2-i0, S2-i+45, and S2-i-45, and the high mass simulations S3-i0, S3-i+45, and S3-i-45 to the analogous S1 simulations discussed above. First we focus on S2, which is representative of the low end of the mass estimates for Sgr. We plot the x and y components of $\tilde{\mathbf{L}}^{\mathrm{inner}}$ and $\tilde{\mathbf{L}}^{\mathrm{outer}}$in Figure 3.9 and the magnitude and orientation of $\tilde{\mathbf{L}}_w$ in Figure 3.10. A comparison with Figures 3.5 and 3.6. reveals that the behaviour is very similar to that of the higher mass satellite, but scaled down by a factor of roughly 5, the factor by which the masses differ. We also find that the difference between the inclined simulations and the polar simulation is larger in the S1 simulations than in the S2 simulations. This suggests that L_{precess} is more efficiently transferred to the disk for higher mass satellites.

The angular momentum transferred in a single perigalactic passage is not sufficient to warp the disk for the S2 models, and no warp is visible in projections of the S2 models during the first orbit. Because the S2-i0 simulation was constructed as a plausible model of Sgr if it is at the lower end of the currently discussed mass range (which appears to be favoured by recent detailed models of the Sgr stream; Law et al., 2004), it is interesting to ask whether this simulation ever produces a visible warp in the disk. After the second perigalactic passage, Figure 3.10 suggests that $\tilde{L}_w$ rises to the level of the S1 simulations after their first perigalacticon, where a small warp is visible. We therefore show in Figure 3.11 projected images of the low mass simulations when $\tilde{L}_w$ peaks after the second perigalacticon. This occurs at $t = 2.10$, 2.05, and 2.15 Gyr for the prograde, polar, and retrograde simulations respectively. The warp angular momentum is directed toward $l = 207°$, $136°$, and $140°$ for the prograde, polar, and retrograde simulations re-

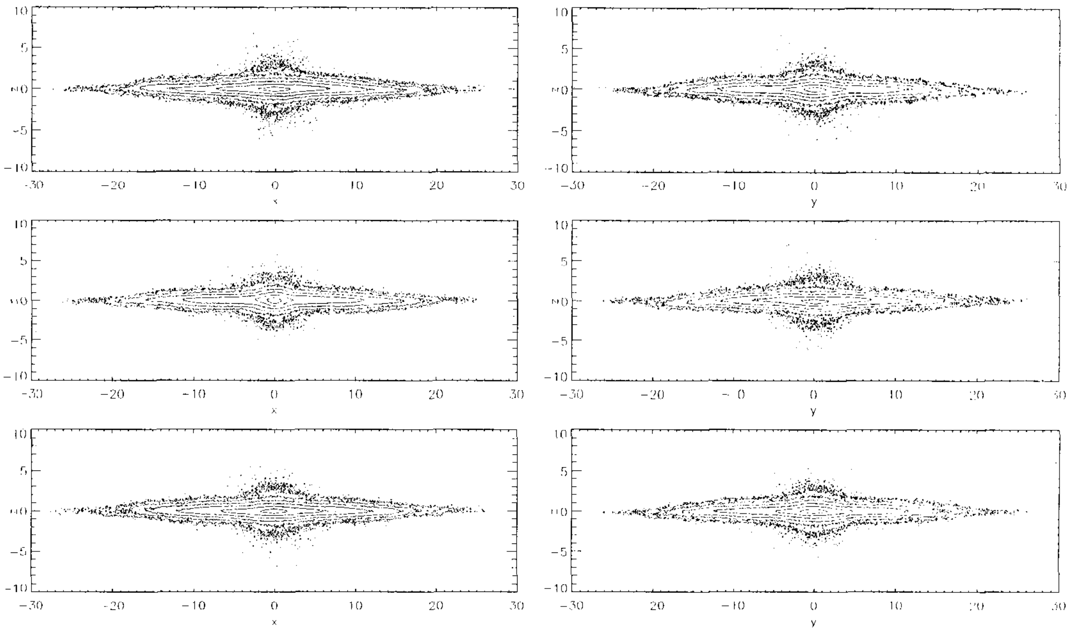

Figure 3.11: Log-spaced contours of the projected density of the bulge and disk in the simulations with low mass satellite S2 on a prograde (top), polar (middle), and retrograde (bottom) orbit when $\tilde{L}_w$ peaks after the second perigalactic passage ($t = 2.10$, 2.05, and 2.15 Gyr for the prograde, polar, and retrograde simulations respectively). The left panels show the x-z projection, while the right panels show the y-z projection. The global tilt of the disk has been removed.

spectively, suggesting that the disk warps up at negative x/down at positive x in all cases, up at positive y/down at negative y in S2-i0 and S2-i-45, and up at negative y/down at positive y in S2-i+45. Instead, S2-i+45 warps up at negative x, S2-i0 warps down at negative x, and S2-i-45 has regions that warp each direction at negative x. All of the disks warp down at negative y. It appears that for small warp amplitudes, the strong differential precession hinders the use of the $\tilde{L}_w$ as an effective measure of the warp orientation; it is also likely that for such small amplitudes, subtraction of the undisturbed simulation may leave artifacts comparable in magnitude to the physical effects of the satellite as suggested by comparing the amplitudes of L^{inner} and L^{outer} in the undisturbed simulation (Figure 3.3) and in the S2 simulations (Figure 3.9). However, we find that it may still be possible for Sgr to excite a small warp even if it is on the lower end of the discussed mass range, providing it has orbited the Galaxy a sufficient number of times.

We have also performed simulations with the very massive S3 satellite. Figures 3.12 and 3.13 plot the x and y components of $\tilde{\mathbf{L}}^{\text{inner}}$ and $\tilde{\mathbf{L}}^{\text{outer}}$, and the magnitude and orientation of $\tilde{L}_w$ respectively. Projections of the disk at the peak of $\tilde{L}_w$ are shown in Figure 3.14. The evolution of the disk is quite violent, and the approximation that the satellite is not disrupted would clearly be violated during the second perigalactic passage, so we only analyze the first orbit. The physical effects discussed in § 3.2.4 are particularly clear for the high mass satellite: L_{shrink} adds $-L_y$ to the disk, especially the outer disk; the slewing potential due to L_{shrink} adds $-L_x$ to the disk, especially the inner disk; L_{precess} adds $-L_x$ $(+L_x)$ to the disk uniformly in the prograde (retrograde) case; the slewing potential due to L_{precess} adds $+L_y$ $(-L_y)$ to the disk, especially the inner disk, in the prograde (retrograde) case; and precession causes a gradual increase in the galactic longitude

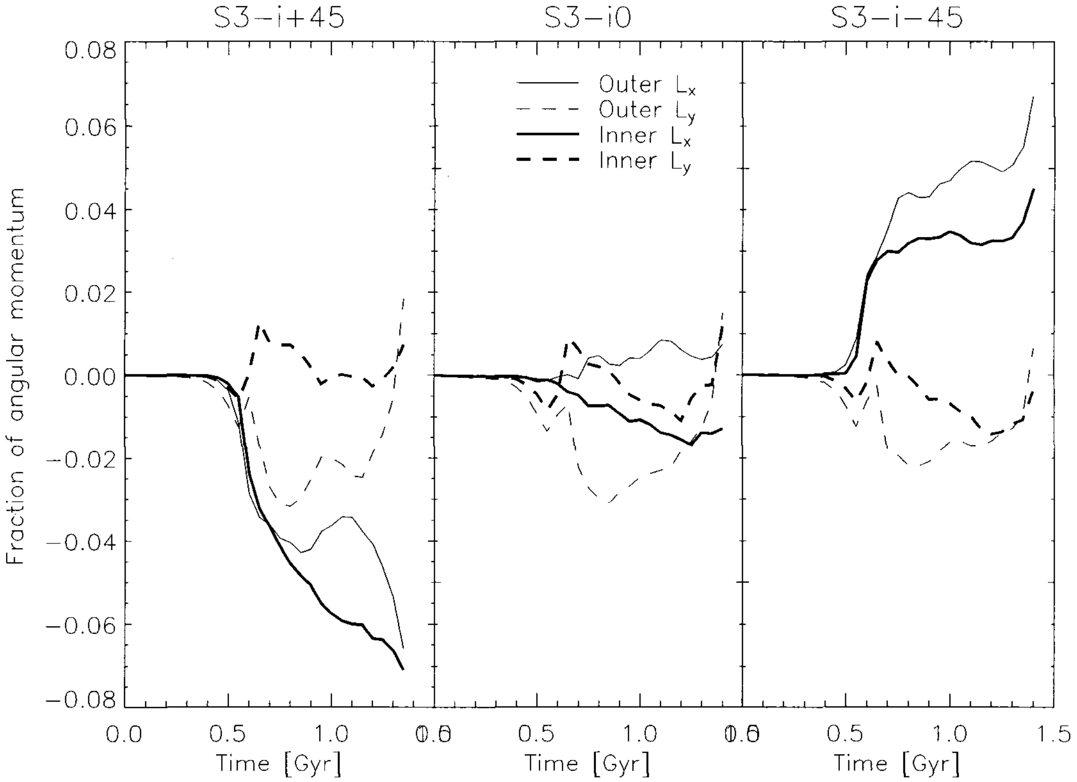

Figure 3.12: As in Figure 3.5, but for the high-mass satellite S3. Perigalactica occur at $t = 0.60$ and $t = 1.45$ Gyr.

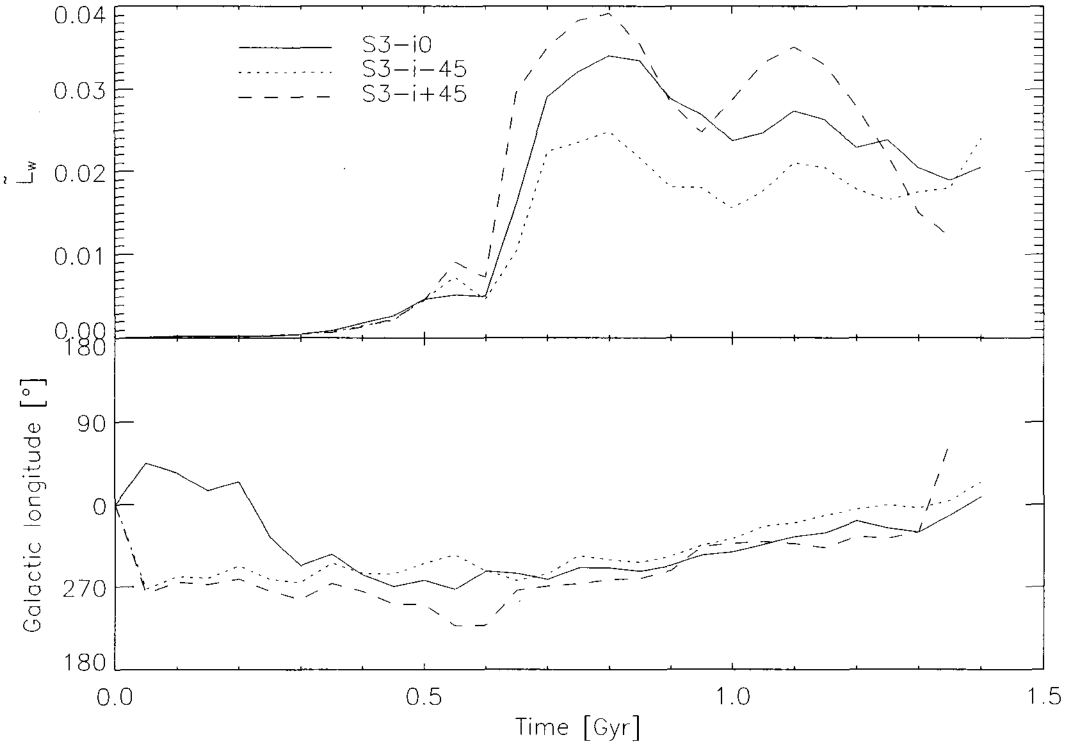

Figure 3.13: As in Figure 3.6, but for the high-mass satellite S3. Perigalactica occur at $t = 0.60$ and $t = 1.45$ Gyr.

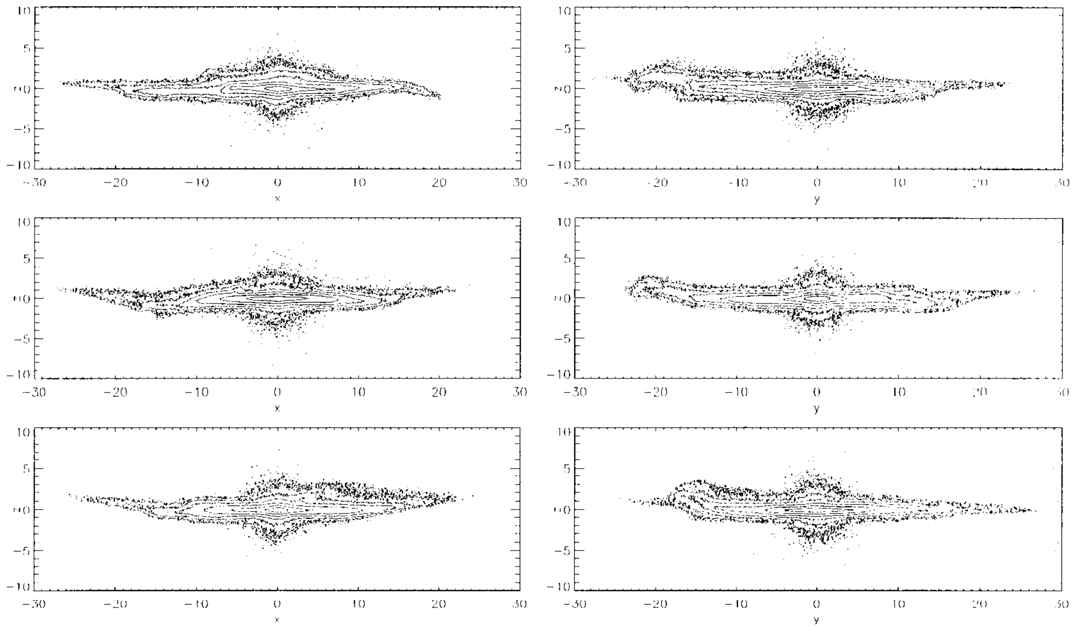

Figure 3.14: Log-spaced contours of the projected density of the bulge and disk in the simulations with high mass satellite S3 on a prograde (top), polar (middle), and retrograde (bottom) orbit when $\tilde{L}_w$ peaks after the first perigalactic passage (t = 0.75 Gyr for the prograde and retrograde simulations and t = 0.80 Gyr for the polar simulation). The left panels show the x-z projection, while the right panels show the y-z projection. The global tilt of the disk has been removed.

of the warp angular momentum. The end result is that the warp angular momentum is directed toward somewhere between 270 and 360° (higher for retrograde satellites, lower for prograde satellites), i.e. it trails slightly behind the satellite orbital angular momentum with respect to galactic rotation, trailing further for retrograde satellites and less for prograde satellites. At the snapshots shown in Figure 3.14, the warp angular momentum is directed toward $l = 273^\circ$, 291°, and 304° for the prograde, polar, and retrograde simulations respectively. The y-z projections show warps up at negative y/down at positive y, as expected. The disks in the prograde and retrograde simulations both warp up at positive x and down at negative x, as expected. The prograde simulation, which is not expected to show any warp in the x-z projection, warps up at positive x between 8 and 15 kpc, followed by a sharp warp down beyond 15 kpc, which combine to give no net $\hat{x}$ component of $\tilde{\mathbf{L}}_w$. However, in all cases the impression of the disk is that it is quite disturbed, rather than warped. The total magnitude of the warp angular momentum peaks at $8 \times 10^{11}\ M_\odot\ \mathrm{kpc\,km\,s^{-1}}$, still smaller than the angular momentum calculated for the Milky Way warp in § 3.1.3 despite the fact that the simulated disks appear significantly more disturbed than the Milky Way disk. This may suggest that the mass and angular momentum distribution of the simulated disks do not match that distributions adopted for the Milky Way in § 3.1.3. However, $\tilde{L}_w$, which is the fraction of the angular momentum which is misaligned, should be much less dependent on the distributions, as should the value of L_w/L_{tot} calculated for the Milky Way. Even when we compare these numbers, we find that warp of the Milky Way ($L_w/L_{\mathrm{tot}} = 0.16 \pm 0.03$, see § 3.1.3) is several times stronger than in any of our high mass simulations (which peak at $\tilde{L}_w = 0.04$).

It is interesting to directly compare the S1-i0, S2-i0, and S3-i0 simulations. In

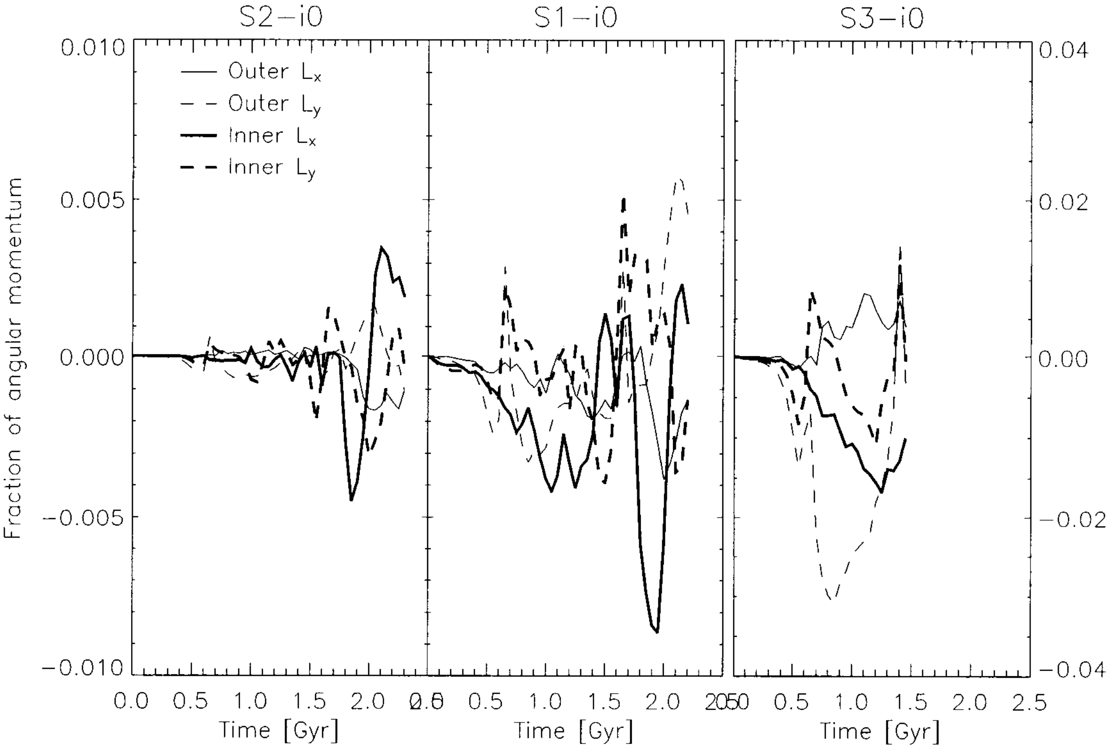

Figure 3.15: As in Figure 3.5, but for the simulations with a low mass (left), intermediate mass (middle), and high mass (right) satellite on a polar orbit. Note that the right-most panel has a different vertical scale. The first perigalacticon occurs at $t = 0.60$ Gyr for all simulations, while the second occurs at $t = 1.75$, 1.70, and, 1.45 for simulations S2-i0, S1-i0, and S3-i0 respectively.

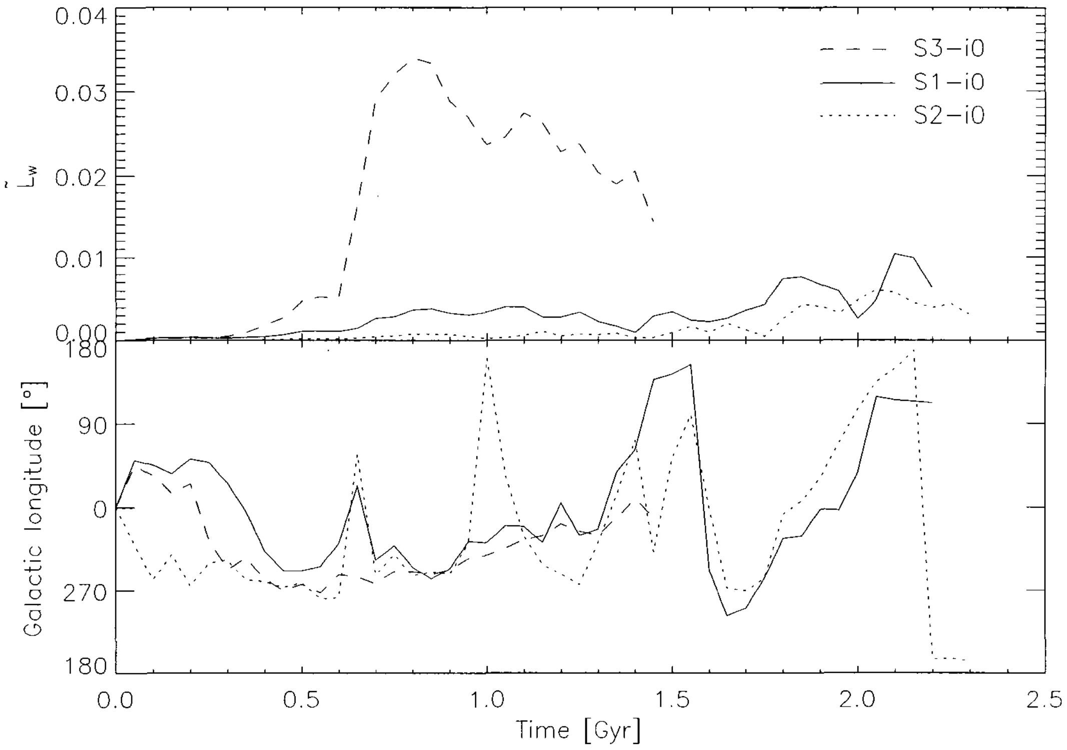

Figure 3.16: As in Figure 3.6, but for the simulations with a low mass (dotted), intermediate mass (solid), and high mass (dashed) satellite on a polar orbit. The first perigalacticon occurs at $t = 0.60$ Gyr for all simulations, while the second occurs at $t = 1.75, 1.70$, and, 1.45 for simulations S2-i0, S1-i0, and S3-i0 respectively.

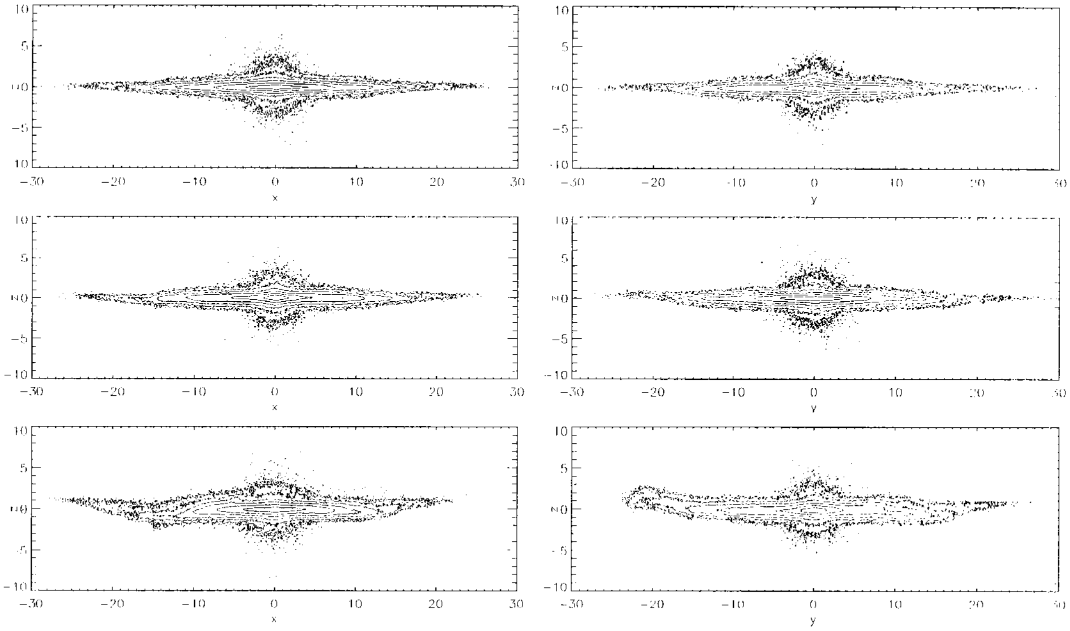

Figure 3.17: Log-spaced contours of the projected density of the bulge and disk in the low mass (top), intermediate mass (middle), and high mass (bottom) simulations with polar orbits at the peak of $\tilde{L}_w$ after the first perigalactic passage (t = 0.80 Gyr for S2-i0 and S3-i0, t = 0.85 Gyr for S1-i0). The left panels are the x-z projections, while the right panels are the y-z projections. The global tilt of the disk has been removed.

Figure 3.15, we plot the x and y components of $\tilde{\mathbf{L}}^{\text{inner}}$ and $\tilde{\mathbf{L}}^{\text{outer}}$, while we plot the magnitude and orientation of $\tilde{L}_w$ in Figure 3.16. All satellites are on polar orbits, so the disks are only affected by the primary and secondary effects of L_{shrink}. The amount of angular momentum transferred to the disk scales roughly as the mass of the satellite — note the difference in vertical scale for the rightmost panel of Figure 3.15.

In Figure 3.17 we show projections of the three simulations at the time when $\tilde{L}_w$ peaks after the first perigalactic passage. The galactic longitude of $\tilde{L}_w$ for all of these simulations at this snapshot is $l \approx 285°$, so we expect the disk to warp up for negative y/down for positive y, and a slight warp up at positive x/down at negative x. The low mass satellite has no discernible effect on the disk. The intermediate mass satellite excites a small warp in the expected direction. The disk interacting with the high mass satellite displays a warp in the expected direction as part of its overall disturbed appearance.

3.2.5.3 Orbital eccentricity

Using the "r20" simulations, we examine how the orbital eccentricity affects the disk warping. The satellites in the S1-r20-i0 and S2-r20-i0 simulations have the same orbital angular momenta as in the S1-i0 and S2-i0 simulations, but are on low eccentricity orbits compared to the high eccentricity of the fiducial S1-i0 simulation. Therefore, they spend a larger fraction of their time at relatively small radii, but do not penetrate as far into the galaxy.

In Figure 3.18, we plot the x and y components of $\tilde{\mathbf{L}}^{\rm inner}$ and $\tilde{\mathbf{L}}^{\rm outer}$ for these simulations. In all cases, the three effects seen are the transfer of $L_{\rm shrink}$ from the satellite to the disk (in particular, the outer disk), the inner disk acquiring $-L_x$ due to the slewing potential, and precession rotating the $-L_y$ component into $+L_x$ (see Figure 3.2, remembering that the disk is prograde, by definition). Figure 3.19 shows the magnitude and orientation of the warp angular momentum $\tilde{L}_w$. Once the warp is initially excited, it is maintained at roughly the same level, which is comparable to that seen after the first perigalactic passage of a satellite on the corresponding high eccentricity orbit. The variation between apo- and perigalacticon is not nearly as large, so the orientation of the angular momentum, which changes somewhat with orbital phase, may be strongly dependent on the particular phase of our initial conditions. Still, we note that the orientation is most often toward $l \approx 270°$.

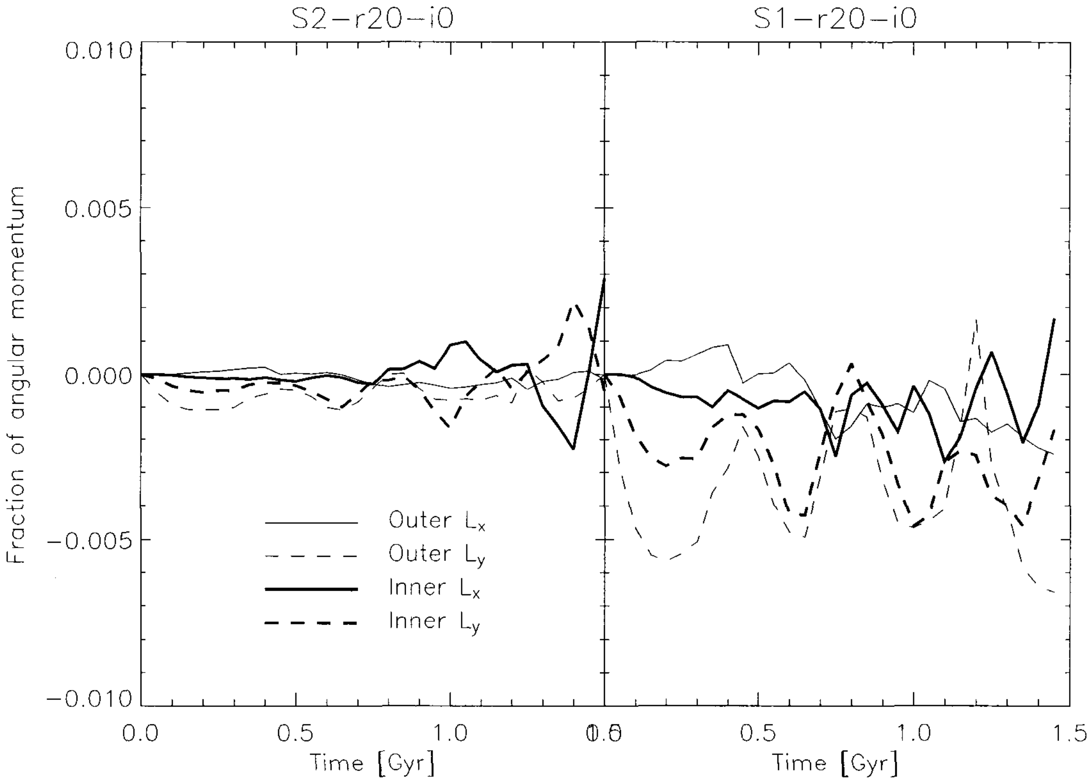

Figure 3.18: As in Figure 3.5, but for satellites on low-eccentricity polar orbits. The left panel represents the low mass satellite S2, while the right panel represents the intermediate mass satellite S1. The perigalactica are at t = 0.00, 0.60, and 1.20 Gyr for S1-r20-i0 and t = 0.00, 0.65, and 1.25 Gyr for S2-r20-i0.

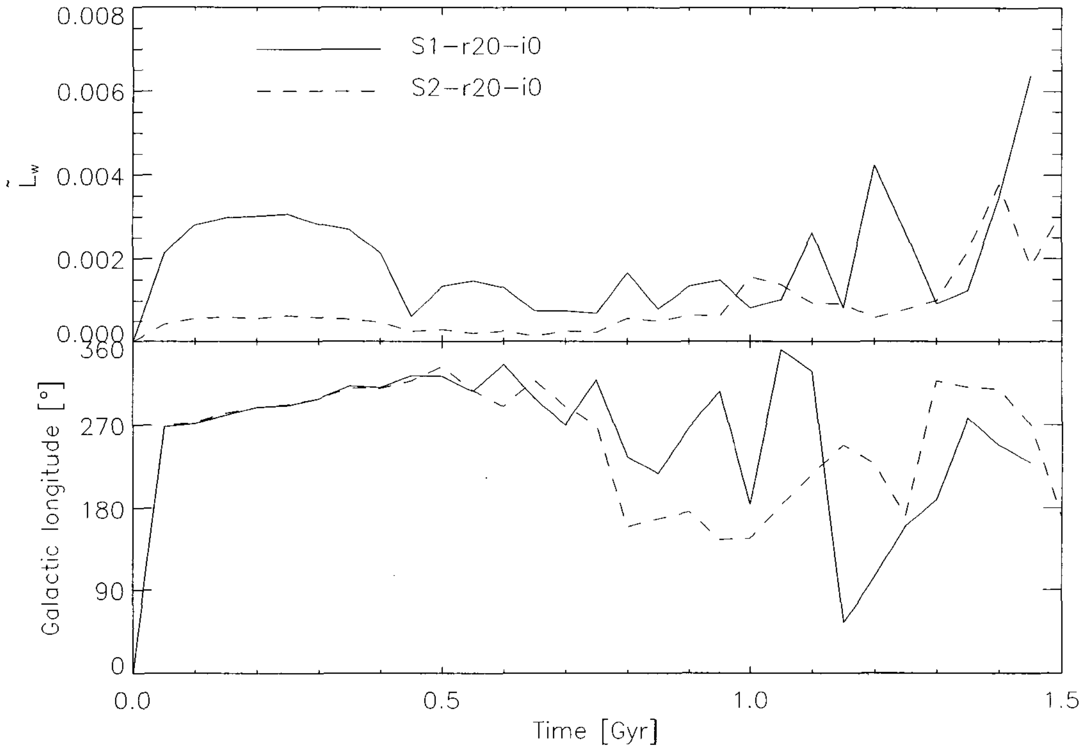

Figure 3.19: As in Figure 3.6, but for satellites on low-eccentricity polar orbits. The perigalactica are at t = 0.00, 0.60, and 1.20 Gyr for S1-r20-i0 and t = 0.00, 0.65, and 1.25 Gyr for S2-r20-i0.

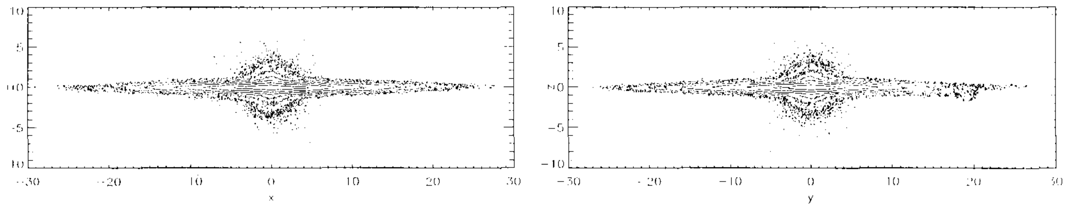

Figure 3.20: Log-spaced contours of the projected density of the bulge and disk in the S1-r20-i0 simulation at time $t = 0.15$ Gyr. The left panels are the x-z projections, while the right panels are the y-z projections. The global tilt of the disk has been removed.

As before, the low mass satellite does not excite a visible warp. The warp angular momentum $\tilde{L}_w$ initially peaks at $t = 0.15$ Gyr in the S1-r20-i0 simulation; Figure 3.20 shows a plot of the simulation at this point. Based on the orientation $l \approx 270^\circ$, we expect the disks to warp up at negative y/down at positive y. A clear warp down at positive y is visible.

3.2.6 Discussion

The relationship between the properties of satellite galaxies and the warps they excite is complicated. In our simulations, we have identified the following trends:

- Angular momentum is transferred between the satellite and the disk because the satellite orbit shrinks and because the satellite precesses if it is on an inclined orbit. The angular momentum due to precession is transferred evenly throughout the disk, tilting it uniformly, while the angular momentum due to the shrinking orbit is transferred more efficiently to the outer disk, tilting it further and warping the disk.
- A secondary effect of the tilting potential is a tilt of the disk angular momentum toward the slewing vector (Ostriker & Binney, 1989). This effect is stronger for the inner disk, causing the disk to warp.
- The warp reaches its maximum amplitude approximately one disk rotation period after the perigalactic passage. Differential precession then dissipates the warp within a couple of disk rotation periods, but it is renewed with each subsequent perigalactic passage. The warp is often asymmetric, with a much larger amplitude on one side than the other, as seen in the Milky Way and about one third of external warped galaxies (Schwarzkopf & Dettmar, 2001).

- The orientation of the warp has a complicated dependence. As a general rule, the warp angular momentum trails slightly behind the satellite angular momentum with respect to the galactic rotation when the satellite is on a polar orbit ($l \approx 285°$ if the satellite angular momentum is directed toward $l = 270°$), while the warp angular momentum leads (trails) this value when the satellite is on a prograde (retrograde) orbit. However, precession quickly scrambles these relationships, and the disk in the polar S1-i0 simulation has a warp directed toward $l \approx 90°$ at late times.

- Satellites on highly eccentric orbits excite warps at perigalacticon. These warps slowly dissipate, but can be enhanced by further perigalactic passages. Satellites on low eccentricity orbits quickly reach an equilibrium warp amplitude that is comparable to the warp excited by the same satellite after one perigalactic passage on a highly eccentric orbit with the same angular momentum.

- The warps occur in the outer regions of the disk, typically between 15 and 20 kpc, while the Milky Way warp starts at much smaller radius.

- Larger satellites are much more efficient at exciting warps than smaller satellites. In particular, a satellite with a mass of $1 \times 10^{10}\ M_\odot$ has a large effect on a Milky Way mass disk (though this might more properly be considered disruption of the disk than warping), a satellite with a mass of $2 \times 10^{9}\ M_\odot$ can excite a small warp in the disk, and a satellite with a mass of $4 \times 10^{8}\ M_\odot$ can only excite a very mild warp after several perigalactica. The degree of warping is roughly proportional to the satellite mass; the total magnitude of the warp angular momentum is typically $\approx 1\%$ of the orbital angular momentum of the satellite. Therefore, Sgr, with its current mass and orbit, is too small to be responsible for the Milky Way warp. The LMC is significantly larger, and we find that the Milky Way warp angular momentum is $\approx 1\%$ of the orbital angular momentum of the LMC. However, the orientation of the warp leads the LMC, which is only expected if the LMC is on a prograde orbit; in fact, the LMC is on a retrograde orbit. Therefore, it is not yet clear if any of the Milky Way satellites are responsible for its warp.

- The magnitude of the warp angular momentum in all of our simulations is less than the angular momentum calculated for the Milky Way's warp in § 3.1.3. This is even true for our high-mass S3 models, in which the disk appears severely disrupted, far more disrupted than the observed plane of the Milky Way. This may suggest that the mass models adopted in § 3.1 and § 3.2 are mismatched. We attempt to compensate for this by comparing the simulation's $\tilde{L}_w$ to the Milky Way's $L_w/L_{\rm tot}$, which should be much less dependent on the adopted mass and angular momentum distributions. Even when we do this, the warp of the Milky Way appears several times larger than any warp seen in our simulations.

CHAPTER 4

FIGURE ROTATION OF COSMOLOGICAL DARK MATTER HALOS

This chapter has been accepted for publication as Bailin & Steinmetz (2004a).

4.1 Introduction

Although there have been many theoretical studies of the shapes of cosmological dark matter halos (e.g. Dubinski & Carlberg, 1991; Warren et al., 1992; Cole & Lacey, 1996; Jing & Suto, 2002), there has been relatively little work done on how those figure shapes evolve with time, in particular, whether the orientation of a triaxial halo stays fixed, or whether the figure rotates. While the orientation of the halo can clearly change during a major merger, it is not known whether the orientation changes in between cataclysmic events. Absent any theoretical prediction one way or the other, it is usually assumed that the figure orientation of triaxial halos remain fixed when in isolation (e.g. Subramanian, 1988; Johnston et al., 1999; Lee & Suto, 2003)

Early work at detecting figure rotation in simulated halos was done by Dubinski (1992) (hereafter D92). While examining the effect of tidal shear on halo shapes, he found that in all 14 of his 1–2 × 10^{12} $M_\odot$ halos, the direction of the major axis rotated uniformly around the minor axis. The period of rotation varied from halo to halo, and ranged from 4 Gyr at the fast end to 50 Gyr at the slow end, or equivalently the angular velocities ranged between 0.1 and 1.6 km s^{-1} kpc^{-1}. [1] It is difficult to draw statistics from this small sample size, especially since the initial conditions of this simulation were not drawn from cosmological models, but were performed in a small isolated box with the linear tidal field of the external matter prescriptively superimposed (Dubinski & Carlberg, 1991). Given that the main result of D92 is that the tidal shear may have a significant impact on the shapes of halos, it is clearly important to do such studies using self-consistent cosmological initial conditions.

---

[1] It may be more intuitive to think of angular velocity in units of radians Gyr^{-1} rather than the common unit of pattern speed, km s^{-1} kpc^{-1}. Fortunately, the two units give almost identical

Recent studies of figure rotation come from Bureau et al. (1999) (BFPM99) and Pfitzner (1999) (P99). P99 compared the orientation of cluster mass halos in two snapshots spaced 500 Myr apart in an SCDM simulation ($\Omega = 1$, $\Lambda = 0$, $h = 0.5$). He detected rotation of the major axis in $\sim 5\%$ of them, and argued that the true fraction with figure rotation is probably higher. BFPM99 presented one of these halos which was extracted from its cosmological surroundings and left to evolve in isolation for 5 Gyr. During that time, the major axis rotated around the minor axis uniformly at all radii at a rate of 60° Gyr^{-1}, or about 1 $\mathrm{km\ s^{-1}\ kpc^{-1}}$.

There may be observational consequences to a dark matter halo whose figure rotates. BFPM99 suggested that triaxial figure rotation is responsible for the spiral structure of the blue compact dwarf galaxy NGC 2915. Outside of the optical radius, NGC 2915 has a large H I disk extending to over 22 optical disk scale lengths (Meurer et al., 1996). The gas disk shows clear evidence of a bar, and a spiral pattern extending over the entire radial extent of the disk. BFPM99 argue that the observed gas surface density is too low for a bar or spiral structure to form by gravitational instability, and that there is no evidence of an interacting companion to trigger the pattern. They propose that the pattern may instead be triggered by a rotating triaxial halo.

Bekki & Freeman (2002) followed this up with Smoothed Particle Hydrody-

---

numerical values.

namics (SPH) simulations of a disk inside a halo whose figure rotates, and showed that a triaxial halo with a flattening of $b/a = 0.8$ and a pattern speed of $3.84\ \mathrm{km\ s^{-1}}$ $\mathrm{kpc^{-1}}$ could trigger spiral patterns in the disk, or warps when the figure rotation axis is inclined to the disk symmetry axis. Masset & Bureau (2003) (hereafter MB03) found that in detail, the observations of NGC 2915 are better fit by increasing the disk mass by an order of magnitude (for example, if most of the hydrogen is molecular, e.g. Pfenniger, Combes, & Martinet, 1994), but that a triaxial halo with $b/a \approx 0.85$ and a pattern speed of between 6.5 and $8.0\ \mathrm{km\ s^{-1}\ kpc^{-1}}$ also provides an acceptable fit.

MB03 concluded that if the halo were undergoing solid body rotation at this rate, its spin parameter would be $\lambda \approx 0.157$, which is extremely large (only 5×10^{-3} of all halos have spin parameters at least that large). However, this argument may be flawed because the figure rotation is a pattern speed, not the speed of the individual particles which constitute the halo, and so it is in principle independent of the angular momentum; in some cases the figure may even rotate retrograde to the particle orbits (Freeman, 1966). Schwarzschild (1982) discusses in detail the orbits inside elliptical galaxies with figure rotation. He finds that models can be constructed from box and X-tube orbits, which have no net streaming of particles with respect to the figure (though they have prograde streaming at small radius and retrograde streaming at large radius), and so result in figures and particles with the same net rotation. He also constructs models that include prograde-streaming Z-tube orbits, which result in a figure that rotates slower than the particles. Stable retrograde Z-tube orbits also exist, but Schwarzschild (1982) did not attempt to include them in his models, so it may also be possible for the figure to rotate faster than the particles. While these results demonstrate the independence of the figure and particle rotation, it is not clear if they can be translated directly to dark matter halos. Dark matter halos may have different formation mechanisms and may be subject to different tidal forces than elliptical galaxies, and the different density profile may also have a large effect on the viable orbital families (Gerhard & Binney, 1985).

There are other consequences of triaxial figure rotation. A rotating potential introduces an oscillating force on particles moving within the potential. Disk stars that have orbital frequencies in resonance with this oscillating force may experience very large changes in their orbit due to the figure rotation. For instance, Tremaine & Yu (2000) examined the behaviour of disks in halos with retrograde figure rotation. In these disks, stars can get trapped in the Binney resonance, where $\Omega_3 - \Omega_2 = \Omega_p$, for vertical and azimuthal frequencies Ω_2 and Ω_3 respectively, and a halo pattern speed of Ω_p (Binney, 1981). If the pattern speed falls slowly toward zero, stars trapped in this resonance are pulled out of the disk and into polar orbits, while if the figure rotation smoothly proceeds from retrograde to prograde, the stars trapped in this resonance are flipped 180° and end up on retrograde orbits. Figure rotation may also erase or modify any intrinsic alignments between the orientation of neighbouring halos (see § 5).

If there are observational consequences to dark halo figure rotation, such as those found by Bekki & Freeman (2002) and Tremaine & Yu (2000), they can be used as a direct method to distinguish between dark matter and models such as MOdified Newtonian Dynamics (MOND) that propose to change the strength of the force of gravity (Milgrom, 1983; Sanders & McGaugh, 2002). Many of the traditional methods of deducing dark matter cannot easily distinguish between the presence of a roughly spherical dark matter halo and a modified force or inertia law. However, a major difference between dark matter and MOND is that dark matter is dynamical, and so tests that detect the presence of dark matter in motion are an effective tool to discriminate between these possibilities. Among the tests that can make this distinction are the ellipticities of dark matter halos as measured using X-ray isophotes, the Sunyaev-Zeldovich effect (Sunyaev & Zeldovich, 1980), and weak lensing (Buote et al., 2002; Lee & Suto, 2003, 2004; Hoekstra et al., 2004); the presence of bars with parameters consistent with being stimulated by their angular momentum exchange with the halo (Athanassoula, 2002; Valenzuela & Klypin, 2003); and spatial offset between the baryons and the mass in infalling substructure measured using weak lensing (Clowe, Gonzalez, & Markevitch, 2004). Rotation of the halo figure requires that dark matter is dynamic, and therefore observable structure triggered by figure rotation potentially provides another test of the dark matter paradigm.

In this chapter, we use cosmological simulations to determine how the figures of ΛCDM halos rotate. The organization of the chapter is as follows. § 4.2 presents the cosmological simulations. § 4.3 describes the method used to calculate the figure rotations, which are presented in § 4.4. Finally we discuss our conclusions in § 4.5.

4.2 The simulations

The halos are drawn from a large high resolution cosmological N-body simulation performed using the GADGET2 code (Springel et al., 2001). We adopt a "concordance" cosmology (e.g. Spergel et al., 2003) with $\Omega_0 = 0.3$, $\Omega_\Lambda = 0.7$, $\Omega_{\rm bar} = 0.045$, $h = 0.7$, and $\sigma_8 = 0.9$. The only effect of $\Omega_{\rm bar}$ is on the initial power spectrum, since no baryonic physics is included in the simulation. The simulation contains $512^3 = 134,217,728$ particles in a periodic volume $50\ h^{-1}$ Mpc comoving on a side. All results are scaled into h-independent units when possible. The full list of parameters is given in Table 4.1.

Table 4.1. Parameters of the cosmological simulation.

Parameter	Value
N	512^3
Box size (h^{-1} Mpc comoving)	50
Particle mass ($10^7\ h^{-1}\ \mathrm{M}_\odot$)	7.757
Force softening length (h^{-1} kpc)	5
Hubble parameter h ($H_0 = 100\ h$ km s^{-1} Mpc^{-1})	0.7
Ω_0	0.3
Ω_Λ	0.7
σ_8	0.9
Ω_{bar}	0.045

Because of the large size of the files and limited available shared-memory, the simulation was split into 8 overlapping slabs segregated in the Cartesian z coordinate. Slab A0 was still too large, so it was split again into 4 overlapping slices segregated in x. The boundaries of the regions are listed in Table 4.2. To overcome problems at the edges of the slabs, only those groups whose centres of mass lie within the central 50% of each sliced dimension are used. For example, in slab A1, whose particles span a range in z from -12500 to $0\ h^{-1}$ kpc, we only use those halos whose centres of mass lie in the range $-9375 < z \leq -3125$. Columns 4 and 5 of Table 4.2 list the boundaries of this central region for each slab. The overlapping slabs provide complete coverage of the simulation volume.

The minimum distance between the edge of a slab and the center of mass of a group is $3125\ h^{-1}$ kpc comoving, much larger than the radius of any collapsed structure expected; the mass of a structure with virial radius $r_{\rm vir}$ is

$$M = \frac{4}{3}\pi \Delta_c \rho_{\rm crit} r_{\rm vir}^3, \tag{4.1}$$

where Δ_c is the enclosed overdensity for a virialized halo and scales as $\Delta_c = 178\Omega_m^{0.45}$ if $\Omega_m + \Omega_\Lambda = 1$ (Eke, Navarro, & Frenk, 1998). For the adopted cosmology, $\Delta_c = 103.5$. Therefore, in order to encounter the edge of its slab, a halo would need to have a mass of at least $3 \times 10^{15}\ h^{-1}\ \mathrm{M}_{\odot}$, over an order of magnitude larger than the largest halo found. If there were a halo with an intrinsic mass larger than this whose center of mass lay as close as possible to a slab boundary, it would still be detected but would be assigned a mass of $\approx 3 \times 10^{15}\ h^{-1}\ \mathrm{M}_{\odot}$ rather than its true mass. Therefore, the lack of any detected halo whose mass is within an order of magnitude of this is proof that there are no halos affected by slicing up the simulation volume.

A friends-of-friends algorithm is used to identify halos (Davis et al., 1985). We use the standard linking length of

$$b = 0.2\bar{n}^{-1/3}, \tag{4.2}$$

where $\bar{n} = N/V$ is the global number density.

Measuring the figure rotation requires comparing the same halo at different times during the simulation. We analyze snapshots of the simulation at lookback times of approximately 1000, 500, 300, and 100 h^{-1} Myr with respect to the $z = 0$ snapshot. The scale factor a of each snapshot, along with its corresponding redshift and lookback time, is listed in Table 4.3.

4.3 Methodology

4.3.1 Introduction

The basic method is to identify individual halos in the final $z = 0$ snapshot of the simulation, find their respective progenitors in slightly earlier snapshots, and measure the rotation of the axes through their common plane as a function of time.

Precisely determining the direction of the axes is crucial and difficult. When merely calculating axial ratios or internal alignment, errors on the order of a few degrees are tolerable. However, if a pattern speed of 1 km s^{-1} kpc^{-1}, as observed in the halo of BFPM99, is typical, then a typical halo will only rotate by 4° in between the penultimate and final snapshots of the simulation. Therefore, the axes of a halo must be determined more precisely than this in order for the figure rotation to be detectable. In fact, we should strive for even smaller errors to see if slower-rotating halos exist. It would have been difficult for P99 to detect halos rotating much slower than the halo presented in BFPM99; although the error varies from halo to halo (for reasons discussed in section 4.3.3), Figure 5.23 of P99

Table 4.2. Boundaries of the slabs in the cosmological simulation.

Slab	Boundary in x (h^{-1} kpc)	Boundary in z (h^{-1} kpc)	Central Region in x (h^{-1} kpc)	Central Region in z (h^{-1} kpc)
b*xxx*-A0a	-25000 – 0	-25000 – -12500	-18750 – -6250	-21875 – -15625
b*xxx*-A0b	0 – 25000	-25000 – -12500	6250 – 18750	-21875 – -15625
b*xxx*-A0c	-12500 – 12500	-25000 – -12500	-6250 – 6250	-21875 – -15625
b*xxx*-A0d	12500 – -12500	-25000 – -12500	18750 – -18750	-21875 – -15625
b*xxx*-A1	-25000 – 25000	-12500 – 0	-25000 – 25000	-9375 – -3125
b*xxx*-A2	-25000 – 25000	0 – 12500	-25000 – 25000	3125 – 9375
b*xxx*-A3	-25000 – 25000	12500 – 25000	-25000 – 25000	15625 – 21875
b*xxx*-B0	-25000 – 25000	-18750 – -6250	-25000 – 25000	-15625 – -9375
b*xxx*-B1	-25000 – 25000	-6250 – 6250	-25000 – 25000	-3125 – 3125
b*xxx*-B2	-25000 – 25000	6250 – 18750	-25000 – 25000	9375 – 15625
b*xxx*-B3	-25000 – 25000	18750 – -18750	-25000 – 25000	21875 – -21875

Note. — The periodic box extends from -25000 h^{-1} kpc to 25000 h^{-1} kpc in each dimension. Slabs A0d and B3 span the periodic boundary in x and z respectively. All units are comoving.

Table 4.3. Snapshots used to calculate figure rotations.

Snapshot Name	Scale Factor (a)	Redshift (z)	Lookback Time (h^{-1} Myr)
b090	0.89	0.1236	1108
b096	0.95	0.0526	496
b098	0.97	0.0309	296
b100	0.99	0.0101	98
b102	1.0	0.0	0

shows that most of his halos had orientation errors of between 8° and 15°, corresponding to a minimum resolvable figure rotation of ~ 0.6 km s^{-1} kpc^{-1} for a 2σ detection in snapshots spaced 500 Myr apart.

A major difficulty in determining the principal axes so precisely is substructure. The orientation of a mass distribution is usually found by calculating the moment of inertia tensor $I_{ij} = \sum_k m_k r_{k,i} r_{k,j}$, and then diagonalizing I_{ij} to find the principal axes. However, this procedure weights particles by r^2. Therefore, substructure near the edge of the halo (or of the subregion of the halo used to calculate the shape) can exert a large influence on the shape of nearly spherical halos, especially if a particular subhalo is part of the calculation in one snapshot but not in another, such as when it has just fallen in. This is particularly problematic because subhalos are preferentially found at large radii (Ghigna et al., 2000; De Lucia et al., 2004; Gill et al., 2004; Gao et al., 2004). Moving substructures can also induce a false measurement of figure rotation due to their motion within the main halo at approximately the circular velocity.

To mitigate this, we first use particles in a spherical region of radius 0.6 r_{vir} surrounding the center of the halo, rather than picking the particles from a density dependent ellipsoid as in Warren et al. (1992) or Jing & Suto (2002). We find that those methods allow substructure at one particular radius to influence the overall shape of the ellipsoid from which particles are chosen for the remainder of the calculation, and therefore bias the results even when other measures are adopted to minimize their effect. The choice of a spherical region biases the derived axis ratios toward spherical values, but does not affect the orientation. Secondly, the particles are weighted by $1/r^2$ so that each mass unit contributes equally regardless of radius (Gerhard, 1983). Both D92 and P99 take similar approaches, but using radii based on ellipsoidal shells. Therefore, we base our analysis on the principal axes of the reduced inertia tensor

$$\tilde{I}_{ij} = \sum_k \frac{r_{k,i} r_{k,j}}{r_k^2}. \tag{4.3}$$

In the majority of halos, the substructure is a small fraction of the total mass, usually less than 5% of the total mass within 60% of the virial radius (De Lucia et al., 2004, Figure 8), so its effect is much reduced. There are still some halos which have not yet relaxed from a recent major merger, in which case the "substructure" constitutes a significant fraction of the halo. To find these cases, the reduced inertia tensor is separately calculated enclosing spheres of radius 0.6, 0.4, 0.25, 0.12, and 0.06 times the virial radius to search for deviations as a function of radius (see § 4.3.4.1 for details). These radii are always with respect to the $z = 0$ value of r_{vir}.

We find that only halos with at least 4×10^3 particles, or masses of at least $\sim 3 \times 10^{11}\ h^{-1}\ \mathrm{M}_{\odot}$ have sufficient resolution for the orientation of the principal axes to be determined at sufficient precision (see § 4.3.3). There are 1432 halos in

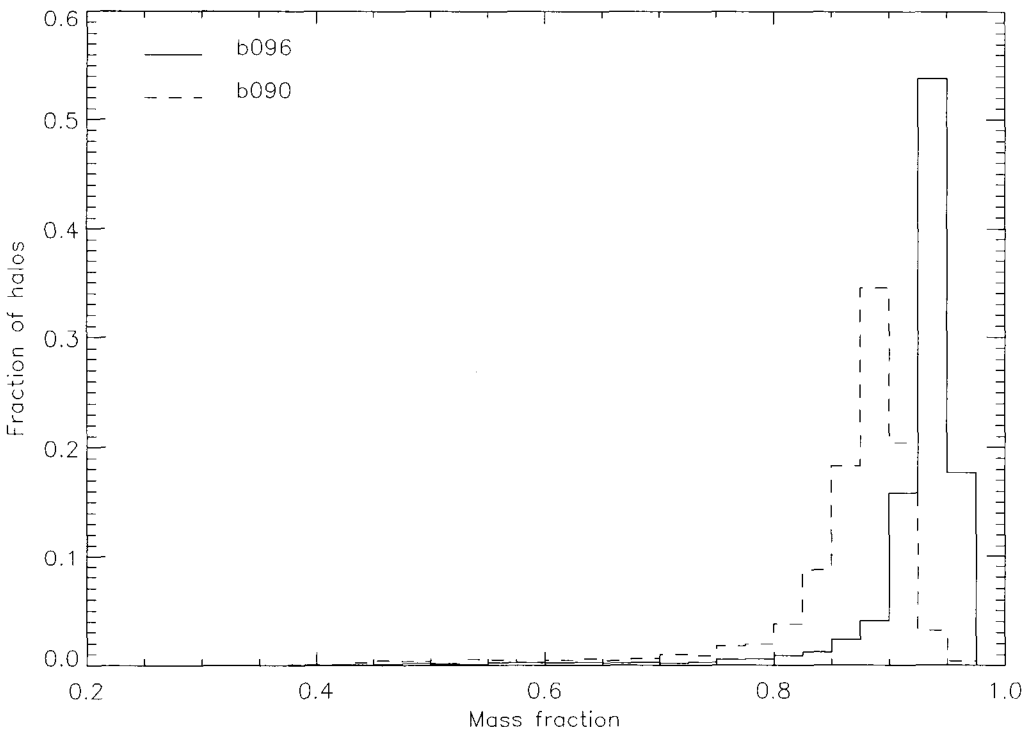

Figure 4.1: Histogram of the fraction of the final mass that comes from the b096 ($z \approx 0.05$, solid line) and b090 ($z \approx 0.12$, dashed line) halo which contributes the most mass.

the $z = 0$ snapshot satisfying this criterion, with masses extending up to $2.8 \times 10^{14}\, h^{-1}\, \mathrm{M}_{\odot}$.

4.3.2 Halo matching

To match up the halos at $z = 0$ with their earlier counterparts, we use the individual particle numbers provided by GADGET which are invariant from snapshot to snapshot, and find which halo each particle belongs to in each snapshot. The progenitor of each $z = 0$ halo in a given $z > 0$ snapshot is the halo that contributes $\geq 90\%$ of the final halo mass. Sometimes no such halo exists; in these cases, the halo has only just formed or underwent a major merger and so is not useful for our purposes. Figure 4.1 shows a histogram of the fraction of the final halo mass that comes from the b096 ($z \approx 0.05$) halo which contributes the most mass. There are also some cases where two nearby objects are identified as a single halo in an earlier snapshot but as distinct objects in the final snapshot. We therefore impose the additional constraint that the mass contributed to the final halo must also be $\geq 90\%$ of the progenitor's mass. In the longer time between the earliest snapshot b090 and the final snapshot b102, a halo typically accretes a greater fraction of its mass, and so a more liberal cut of 85% is used for this snapshot (see the dashed histogram in Figure 4.1). 492 of the halos that satisfied the mass cut did not have a progenitor which satisfied these criteria in at least one of the $z > 0$ snapshots and so were eliminated from the analysis, leaving a sample of 940 matched halos.

4.3.3 Error in axis orientation

There are two sources of errors that enter into the determination of the axes: how well the principal axes of the particle distribution can be determined, and whether that particle distribution has a smooth triaxial figure. Here we estimate the error assuming that it is not biased by substructure. The halos for which this assumption does not hold will become apparent later in the calculation.

For a smooth triaxial ellipsoid represented by N particles, the error is a function of N and of the intrinsic shape: as the axis ratio b/a or c/b approaches unity, the axes become degenerate. To quantify this, we have performed a bootstrap analysis of the particles in a sphere of radius 0.6 $r_{\rm vir}$ of each $z = 0$ halo (Heyl et al., 1994). If the sphere contains N particles then we resample the structure by randomly selecting N particles from that set, with no constraint against selecting duplicate particles, and determine the axes from this bootstrap set. We do this 100 times for each halo. The dispersion of these estimates around the calculated axis is taken formally as the "1σ" angular error, and is labelled $\theta_{\rm boot}$.

As expected, the two important parameters are N and the axis ratio. We focus here on the major axis, for which the important axis ratio is b/a (the results for the minor axis are identical with the minor-to-intermediate axis ratio c/b replacing b/a). The top panels of Figures 4.2 and 4.3 show the dependence of the bootstrap error on N and b/a respectively for all halos with $M > 10^{11}\ h^{-1}\ {\rm M}_{\odot}$. The solid lines are empirical fits:

$$\theta_{{\rm err},N} = \frac{2}{\sqrt{N}}, \tag{4.4}$$

and

$$\theta_{{\rm err},b/a} = 0.005\frac{\sqrt{b/a}}{1-b/a}. \tag{4.5}$$

The form of equation (4.4) is not surprising; if a smooth halo was randomly sampled, we would expect the errors to be Poissonian with an $N^{-1/2}$ dependence. However, the cosmological halos are not randomly sampled. Individual particles "know" where the other particles are, because they have acquired their positions by reacting in the potential defined by those other particles. Therefore, the errors may be less than expected from a randomly sampled halo. To test this, we construct a series of smooth prolate NFW halos (Navarro, Frenk, & White, 1996) with b/a axis ratios ranging from 0.5 to 0.9, randomly sampled with between 3×10^3 and 3×10^5 particles, and perform the bootstrap analysis identically for each of these halos as for the cosmological halos. Because the method for calculating axis ratios outlined in § 4.3.1 biases axis ratios toward spherical, the recovered b/a of these randomly sampled halos is larger than the input value, and ranges from 0.65 to 0.95. The errors for these randomly sampled smooth halos are shown as asterisks in Figures 4.2 and 4.3.

The top panel of Figure 4.2 shows a rise in the dispersion of the error for $N \lesssim 4000$, with many halos having errors greater than the 0.1 radians necessary

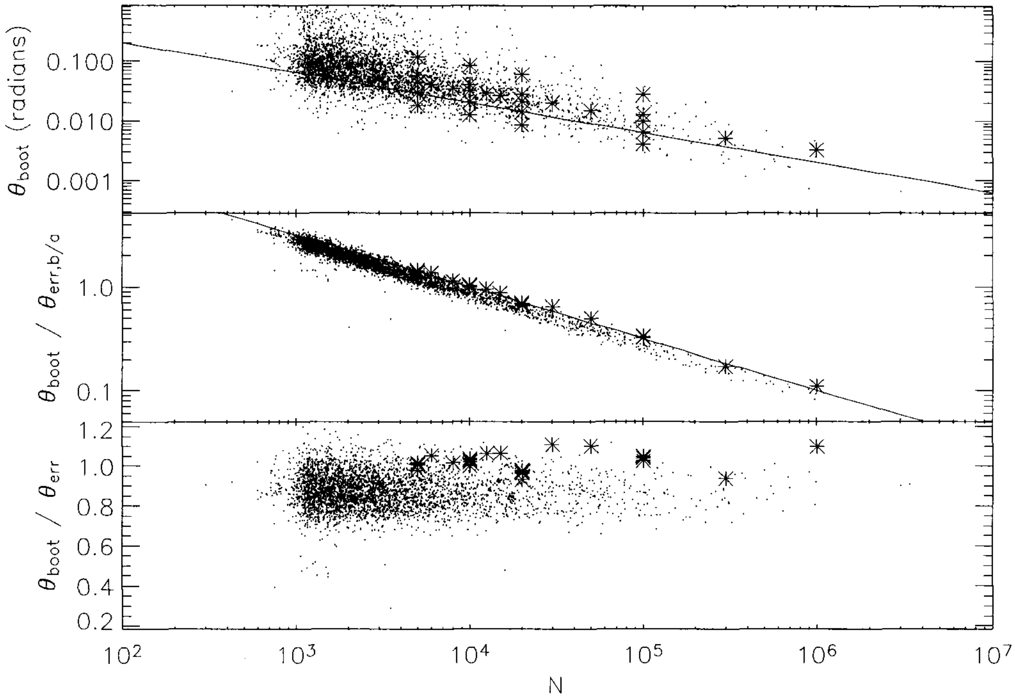

Figure 4.2: Angular bootstrap error $\theta_{\rm boot}$ as a function of the number of particles N within the central 0.6 $r_{\rm vir}$ of each halo. Points are the cosmological halos, and asterisks are randomly sampled smooth NFW halos. *(Top):* Angular error $\theta_{\rm boot}$. The solid line is the fit $\theta_{{\rm err},N}$ from equation (4.4). *(Middle):* Ratio between the angular error and the error expected for the halo given its axis ratio b/a, i.e. $\theta_{\rm boot}/\theta_{{\rm err},b/a}$. The solid line is $\theta_{{\rm err},N}$ from equation (4.4) renormalized to the typical error of 0.02 radians. *(Bottom):* Ratio between the angular error and the analytic estimate $\theta_{\rm err}$ from equation (4.6).

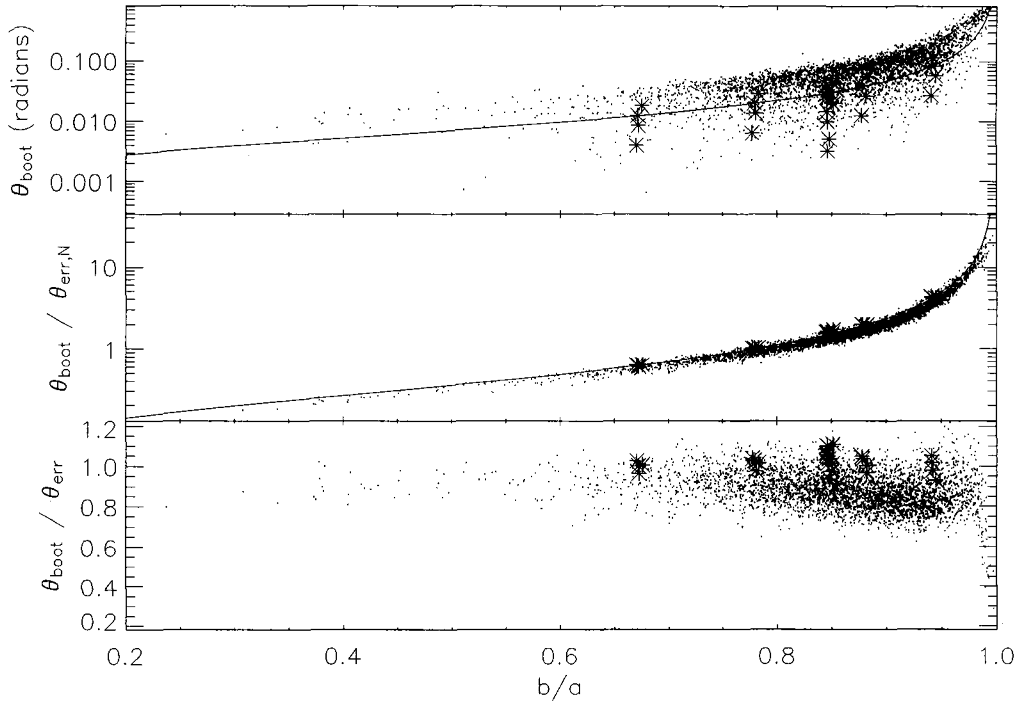

Figure 4.3: Angular bootstrap error $\theta_{\rm boot}$ as a function of the axis ratio b/a of each halo. Points are the cosmological halos, and asterisks are randomly sampled smooth NFW halos. *(Top):* Angular error $\theta_{\rm boot}$. The solid line is the fit $\theta_{{\rm err},b/a}$ from equation (4.5). *(Middle):* Ratio between the angular error and the error expected for the halo given the number of particles N, i.e. $\theta_{\rm boot}/\theta_{{\rm err},N}$. The solid line is $\theta_{{\rm err},b/a}$ from equation (4.5) renormalized to the typical error of 0.02 radians. *(Bottom):* Ratio between the angular error and the analytic estimate $\theta_{\rm err}$ from equation (4.6).

to detect the figure rotation of the halo presented in BFPM99. Therefore, we only use halos with $N > 4000$, or equivalently masses $M > 3.1 \times 10^{11}\, h^{-1}\,\mathrm{M}_{\odot}$.

The bootstrap error appears to be completely determined by N and b/a. The residuals of θ_{boot} with respect to $\theta_{\mathrm{err},N}$ are due to $\theta_{\mathrm{err},b/a}$ and vice versa. This is shown in the middle panels of Figures 4.2 and 4.3. In the middle panel of Figure 4.2 we have divided out the dependence of θ_{boot} on the axis ratio, making apparent an extremely tight relation between the residual and N, while in the middle panel of Figure 4.3 we have divided out the dependence of θ_{boot} on N, showing the equally tight relation between the residual and b/a. It is apparent from comparing the points and asterisks that the errors in the cosmological halos are slightly smaller than for randomly sampled smooth halos.

Combining equations (4.4) and (4.5), and noting that the typical error is $\theta_{\mathrm{boot}} \approx 0.02$ radians, we find the bootstrap error is well fit by

$$\theta_{\mathrm{err}} = \frac{1}{2\sqrt{N}} \frac{\sqrt{b/a}}{1 - b/a}. \tag{4.6}$$

The bottom panels of Figures 4.2 and 4.3 show the residual ratio between the bootstrap error θ_{boot} and the analytic estimate θ_{err}. The vast majority of points lie between 0.8 and 1.0, indicating that θ_{err} overestimates the error by $\sim 10\%$. Equation (4.6) breaks down as b/a approaches unity; these halos are nearly oblate and so do not have well-defined major axes. It also becomes inaccurate at very low b/a due to low-mass poorly-resolved halos. Even in these cases, the error estimate is conservative, but to be safe we have eliminated axes with $b/a < 0.35$ or $b/a > 0.95$ from the subsequent analysis, regardless of the nominal error. The randomly-sampled smooth halos follow equation (4.6) extremely well, so the non-Poissonianity of the sampling in simulated halos reduces the errors by 10%.

Calculating the bootstraps is computationally expensive, so equation (4.6) is used for the error in all further computation. Because this estimate is expected to be correct for smooth ellipsoids, cases where the error is anomalous are indications of residual substructure.

4.3.4 Figure rotation

Ideally one would fit the figure rotation by comparing the orientation of each of the axes at each snapshot to that of a unit vector rotating uniformly along a great circle, and minimize the χ^2 to find the best fit uniform great circle trajectory. However, this requires minimizing a non-linear function in a four-dimensional parameter space, a non-trivial task.

We adopt two simpler and numerically more robust methods for measuring the figure rotation. The first method, referred to as the "plane method", involves fitting the major or minor axis measurements at all five snapshots to a plane, and then measuring the rotation of the axis along the plane. This fully takes the errors and measurements at all snapshots into account. However, it presupposes that the figure rotation axis is perpendicular to the plane containing the major or minor axis. The second method, referred to as the "quaternion method", involves comparing all of the axes at two snapshots to find the axis through which the figure has rotated. This method gives a figure rotation axis that is not constrained to have any particular relation to the major or minor axis. However, by construction it measures the rotation from a single reference frame to another single reference frame, and therefore can only include information from two snapshots at a time. It is also not possible to take into account the errors in the axis determinations — in particular, for prolate halos, where the error in the determination of the intermediate and minor axes are much larger than the error in the major axis, physical rotation of the major axis can be masked by spurious fluctuations in the two degenerate axes.

The strengths and weaknesses of these two methods complement each other

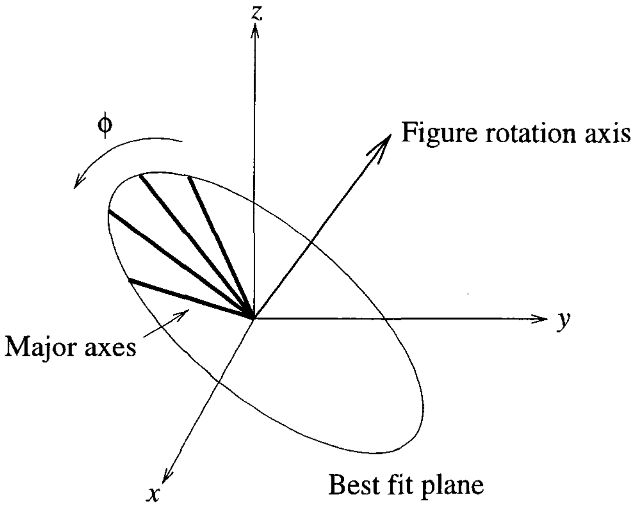

Figure 4.4: Diagram that demonstrates how we fit a plane to the major axis measurements at all snapshots (thick lines) and then find the increase of phase ϕ as a function of time. The figure rotation axis is perpendicular to the best fit plane, and defined such that ϕ increases around it counter-clockwise with time.

well. We adopt the plane method as our primary method of measuring the figure rotation. The quaternion method is used to check for bias in the derived figure rotation axes.

4.3.4.1 Plane method

For the plane method, we first solve for the plane $z = ax + by$ that fits the major axis measurements of the halo best at all timesteps, assuming the error is negligible. The change of the phase of the axes in this plane as a function of time are then fit by linear regression. A schematic diagram of this process is shown in Figure 4.4. We follow the same procedure for the minor axes when appropriate, as discussed in § 4.4.

The degree to which the axes are consistent with lying in a plane is checked by calculating the out-of-plane χ^2:

$$\chi^2_{\rm oop} \equiv \frac{1}{\nu}\sum_i \frac{\Delta\theta_i^2}{{\theta_{\rm err}}_i^2}, \tag{4.7}$$

where ν is the number of degrees of freedom and $\Delta\theta_i$ is the minimum angular distance between the major axis at timestep i and the great circle defined by the best fit plane.

Because the axes have reflection symmetry, it is impossible to measure a change in phase of more than $\pi/2$. The phases are adjusted by units of π such that the difference in phase between adjacent snapshots is always less than $\pi/2$. If the figure were truly rotating by 90° or more in between the snapshots, it would be impossible to accurately measure this rotation since the angular frequency would be larger than the Nyquist frequency of our sampling rate (Nyquist, 1928). Any faster pattern speeds would be aliased to lower angular velocities, with an aliased angular velocity of $\Omega_{\rm Nyq} - (\Omega_p - \Omega_{\rm Nyq})$, where Ω_p is the intrinsic angular velocity of the pattern and $\Omega_{\rm Nyq}$ is the Nyquist frequency. For snapshots spaced 500 h^{-1} Myr apart, the maximum time between the snapshots we analyze, the maximum detectable angular velocity is 3.8 h km s^{-1} kpc^{-1}. We do not expect the figure to change so dramatically as we have excluded major mergers. However, this can be checked *post facto* by checking whether the distribution of measured angular velocities extends up to the Nyquist frequency; if so, then there are likely even more rapidly rotating figures whose angular frequency is aliased into the detectable range, fooling us into thinking they are rotating slower. If the measured distribution does not extend to the Nyquist frequency, then it is unlikely that there are any figures rotating too rapidly to be detected (see § 4.4).

The best fit linear relation for the phase as a function of time is found by linear regression. Because the component of an isotropic angular error projected onto a plane is half of the isotropic error, we divide the error of equation (4.6) by two before we perform the regression. The slope of the linear fit gives the pattern speed Ω_p of the figure rotation. The error is the 1σ limit on the slope.

Once we have calculated the pattern speed for each halo, we can eliminate the cases where substructure has severely impacted the results. In these cases, the signal is dominated by a large subhalo at a particular radius, so the derived pattern speed will be significantly different when the sphere is large enough to include the subhalo from when the subhalo is outside the sphere. We have calculated the pattern speed using enclosing spheres of radius 0.6, 0.4, 0.25, 0.12, and 0.06 of the virial radius. The fraction of mass in subhalos can be estimated via the change in the pattern speed Ω_p at adjacent radii. Because the reduced inertia tensor is mass-weighted, the figure rotation of a sphere with a smooth component rotating at $\Omega_{p,\mathrm{smooth}}$ plus a subhalo containing a fraction f_s of the total mass moving at the circular velocity v_c at radius $R\,r_{\mathrm{vir}}$ is approximately

$$\Omega_p \approx (1 - f_s)\Omega_{p,\mathrm{smooth}} + f_s \frac{v_c}{R\,r_{\mathrm{vir}}}, \tag{4.8}$$

where the difference due to the presence of the subhalo is

$$\Delta\Omega_p = f_s \frac{v_c}{R\,r_{\mathrm{vir}}} = f_s \sqrt{\frac{GM(< R\,r_{\mathrm{vir}})}{R^3\,r_{\mathrm{vir}}^3}}. \tag{4.9}$$

If the density profile is roughly isothermal, the enclosed mass is

$$M(< R\,r_{\mathrm{vir}}) = \frac{4}{3}\pi \Delta_c \rho_{\mathrm{crit}} R\,r_{\mathrm{vir}}^3. \tag{4.10}$$

Solving equations (4.9) and (4.10) gives an expression for the fraction of the mass in substructure given a jump of $\Delta\Omega_p$ in the measured pattern speed when crossing radius $R\,r_{\mathrm{vir}}$:

$$f_s = \frac{\Delta\Omega_p R}{\sqrt{\frac{4}{3}\pi G \Delta_c \rho_{\mathrm{crit}}}}. \tag{4.11}$$

The term in the square root is equal to $0.72\ h\ \mathrm{km\ s^{-1}\ kpc^{-1}}$. For each halo, we compute the value of f_s due to the jump $\Delta\Omega_p$ between each set of adjacent radii,

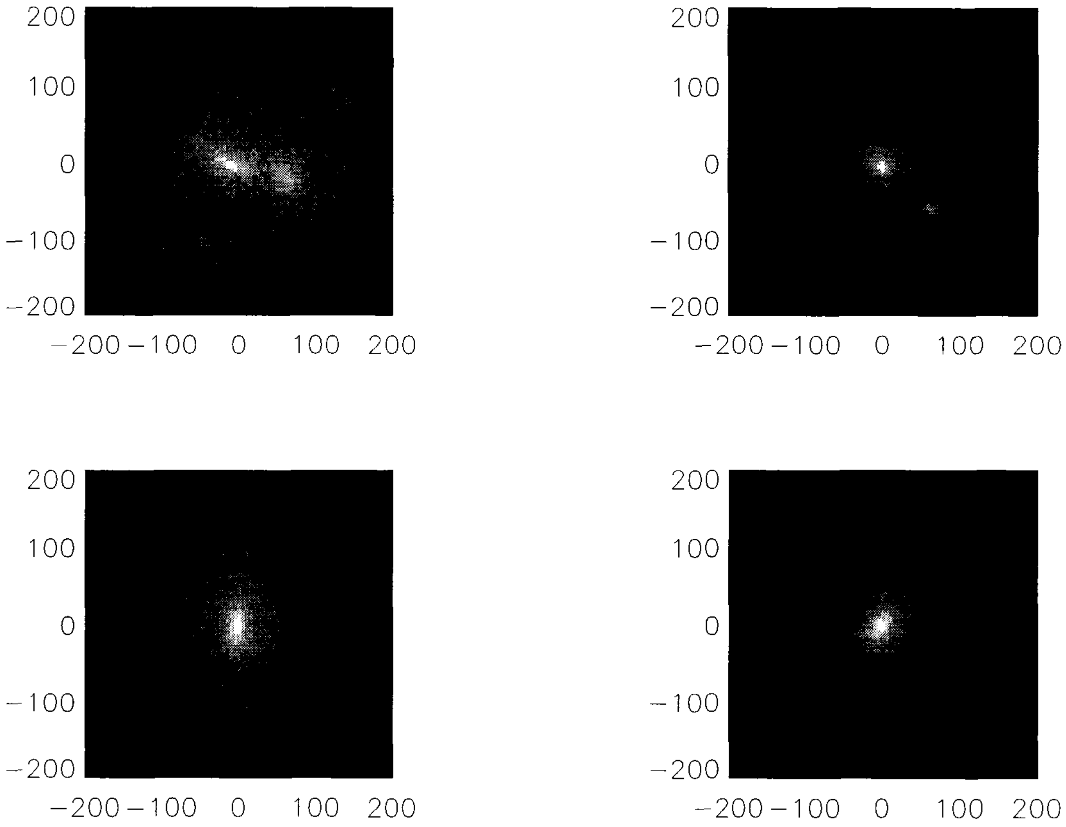

Figure 4.5: Log-weighted projected density of 4 halos with a range of subhalo fractions f_s. The subhalo fractions are 0.166 (top-left), 0.065 (top-right), 0.045 (bottom-left), and 0.016 (bottom-right). Axes are in units of h^{-1} kpc from the halo center. All halos have masses in the range $2 - 3 \times 10^{12}$ h^{-1} $M_{\odot}$.

i.e. for $R = 0.4, 0.25, 0.12$, and 0.06, and adopt the largest value of f_s as the substructure fraction of the halo. Figure 4.5 shows log-weighted projected densities of 4 halos in the mass range $2 - 3 \times 10^{12}\ h^{-1}\ \mathrm{M}_{\odot}$ with a variety of values of f_s, ranging from 0.166 at the top-left to 0.016 at the bottom right. After examining a number of halos spanning a range of f_s, we adopt a cutoff of $f_s < 0.05$ for undisturbed halos. This eliminates 289 of the 940 halos, leaving 651 undisturbed halos.

A further 158 halos were eliminated because the angular error approached $\pi/2$ in at least one of the snapshots. This includes the halos with $b/a < 0.35$ or $b/a > 0.95$ discussed in § 4.3.3. We also eliminate cases where the reduced χ^2 from the linear fit of phase versus time indicates that the intrinsic error of the direction determination is much lower than suggested by equation (4.6), indicating that the model of the halo as a smooth ellipsoid is violated (10 halos with $\chi^2_\nu < 0.1$), and those cases where the phase does not evolve linearly with time (134 halos with $\chi^2_\nu > 10$). Finally, we eliminate halos where the axes do not lie on a common plane, i.e. the 32 halos where $\chi^2_{\mathrm{oop}} > 10$. Therefore, the final sample consists of 317 halos.

A sample halo is shown in the first five panels of Figure 4.6. It was chosen randomly from the halos with relatively low errors and typical pattern speeds. It has a mass of $1.9 \times 10^{12}\ h^{-1}\ \mathrm{M}_{\odot}$, and a pattern that rotates at $0.32 \pm 0.01\ h\ \mathrm{km\ s^{-1}\ kpc^{-1}}$. It has a spin parameter $\lambda = 0.047$, and axis ratios of $b/a = 0.86$ and $c/a = 0.77$ at $z = 0$. The derived substructure fraction is $f_s = 0.045$, and the out-of-plane $\chi^2_{\mathrm{oop}} = 8.5$. The solid line shows the measured major axis in each snapshot, which rotates counterclockwise in this projection. The phase of its figure rotation as a function of time is shown in the bottom-right panel of Figure 4.6. The zero point is arbitrary, but is consistent from snapshot to snapshot. The linear fit is also

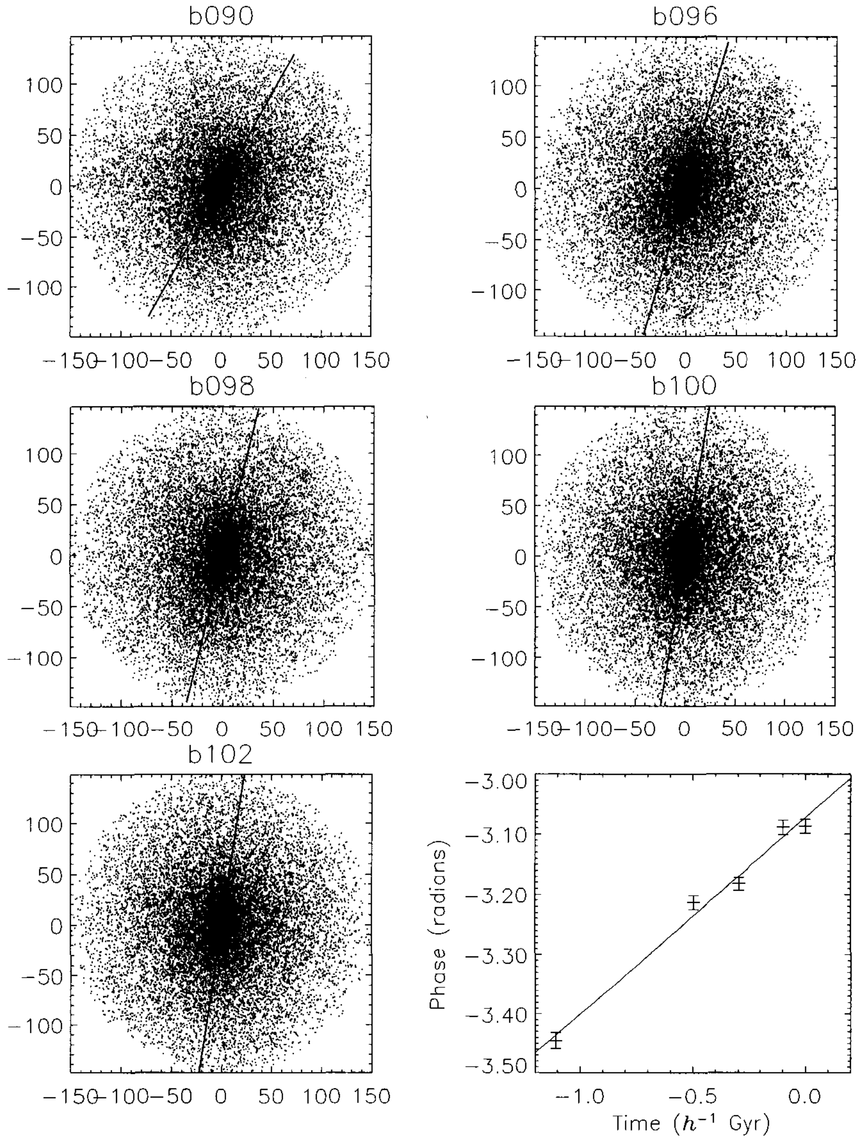

Figure 4.6: *(Upper-left five panels):* Projection onto the best fit plane of the inner 0.6 r_{vir} of a sample halo at the five snapshots we analyze. Axes are in units of h^{-1} kpc from the halo center. From left to right, top to bottom, the snapshots are at 1108, 496, 296, 98, and 0 h^{-1} Myr before $z = 0$. The solid line is the major axis, which rotates counterclockwise by 20° from beginning to end. *(Bottom-right):* Phase of the major axis in the rotational plane of the sample halo. The zero point is arbitrary, but identical in all snapshots. The solid line is the linear fit, with a slope of 0.33 h radians Gyr^{-1}.

shown, which has a reduced χ^2 of 2.9.

4.3.4.2 Quaternion method

For the quaternion method, we directly measure the rotation between the axes at two snapshots. If the major, intermediate, and minor axes for snapshot j are $\hat{\mathrm{a}}_j$, $\hat{\mathrm{b}}_j$, and $\hat{\mathrm{c}}_j$ respectively, we construct a rotation matrix R_j that transforms vectors into the principal axis frame:

$$R_j = \begin{pmatrix} a_{j,x} & a_{j,y} & a_{j,z} \\ b_{j,x} & b_{j,y} & b_{j,z} \\ c_{j,x} & c_{j,y} & c_{j,z} \end{pmatrix} . \tag{4.12}$$

The matrix expressing the rotation from snapshot j to snapshot k is $A = R_j R_k^T$. This rotation matrix can be represented as a quaternion[2]. Defining $T \equiv 2\sqrt{1 + \mathrm{Tr}\ A}$, the components of the quaternion q are given by

$$\begin{aligned} q_w &= T/4 \\ q_x &= (A_{32} - A_{23})/T \\ q_y &= (A_{13} - A_{31})/T \\ q_z &= (A_{21} - A_{12})/T. \end{aligned} \tag{4.13}$$

The angle of rotation between the snapshots is then $\theta = 2 \arccos q_w$, and the unnormalized figure rotation axis is (q_x, q_y, q_z).

4.4 Results

Figure 4.7 shows the measured figure rotation speeds of the major axes for all of the halos in the sample, as a function of their error using the plane method. Halos with measured pattern speeds less than twice as large as the estimated error (the dashed line) are taken as non-detections. 278 of the 317 halos have detected figure

---

[2] see, for example, the Matrix and Quaternion FAQ at `http://vamos.sourceforge.net/matrixfaq.htm`

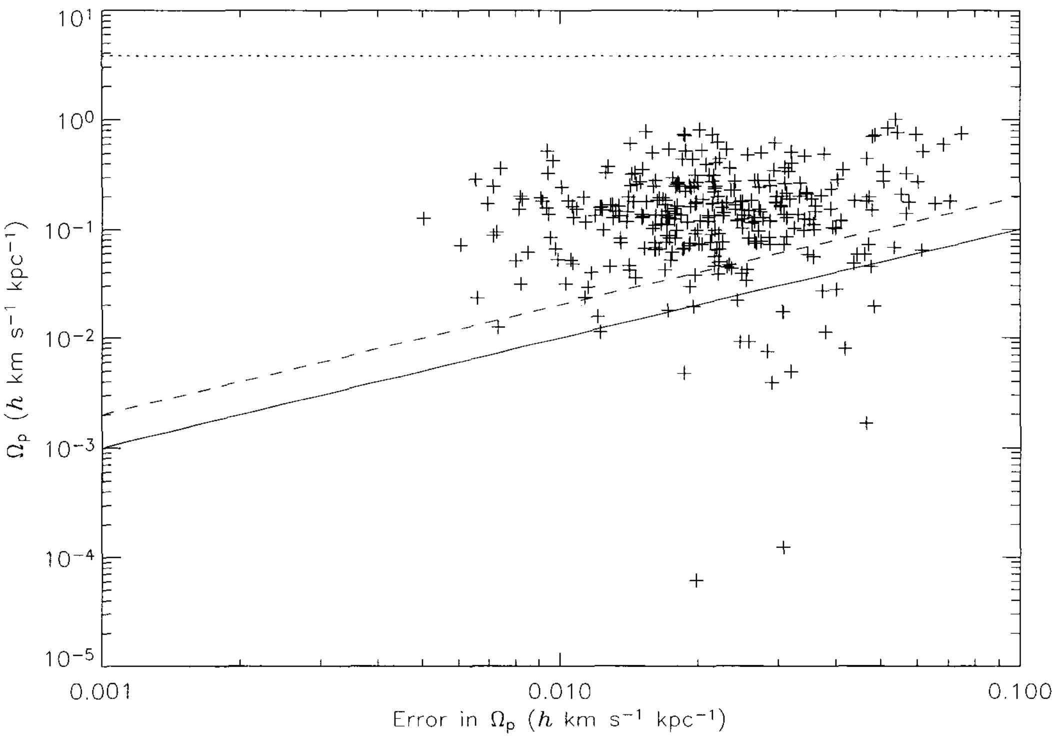

Figure 4.7: Measured pattern speed of the rotation of the major axis. The x-axis is the error in the pattern speed. The solid line indicates where the measured pattern speed is equal to the estimated error, while the dashed line is the 2σ limit. The horizontal dotted line indicates the Nyquist limit.

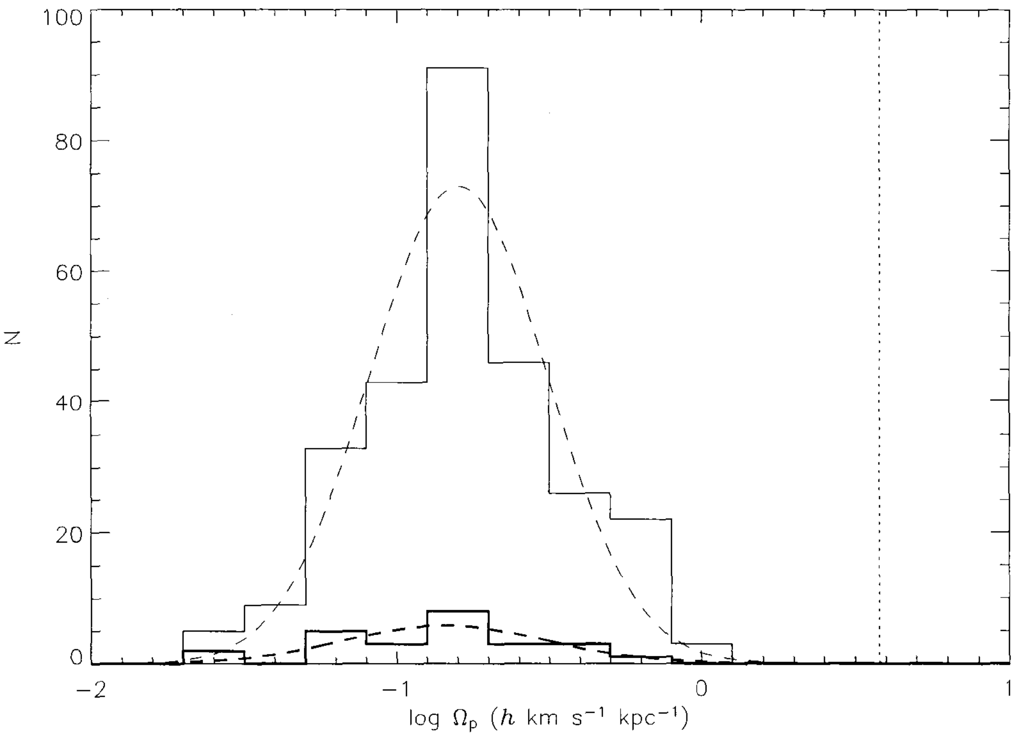

Figure 4.8: Histogram of the pattern speeds of the figure rotation, expressed in $\log \Omega_p$. The thin histogram contains all halos that have 2σ detections of figure rotation, i.e. those above the dashed line of Figure 4.7, and is incomplete at $\Omega_p < 0.126\ h\ \mathrm{km\ s^{-1}\ kpc^{-1}}$ or equivalently $\log \Omega_p < -0.9$. The thick histogram contains only those halos with errors less than $0.01\ h\ \mathrm{km\ s^{-1}\ kpc^{-1}}$, and is incomplete at $\Omega_p < 0.015\ h\ \mathrm{km\ s^{-1}\ kpc^{-1}}$ or $\log \Omega_p < -1.8$. The dashed curves are Gaussian fits to the histograms. The vertical dotted line indicates the Nyquist limit.

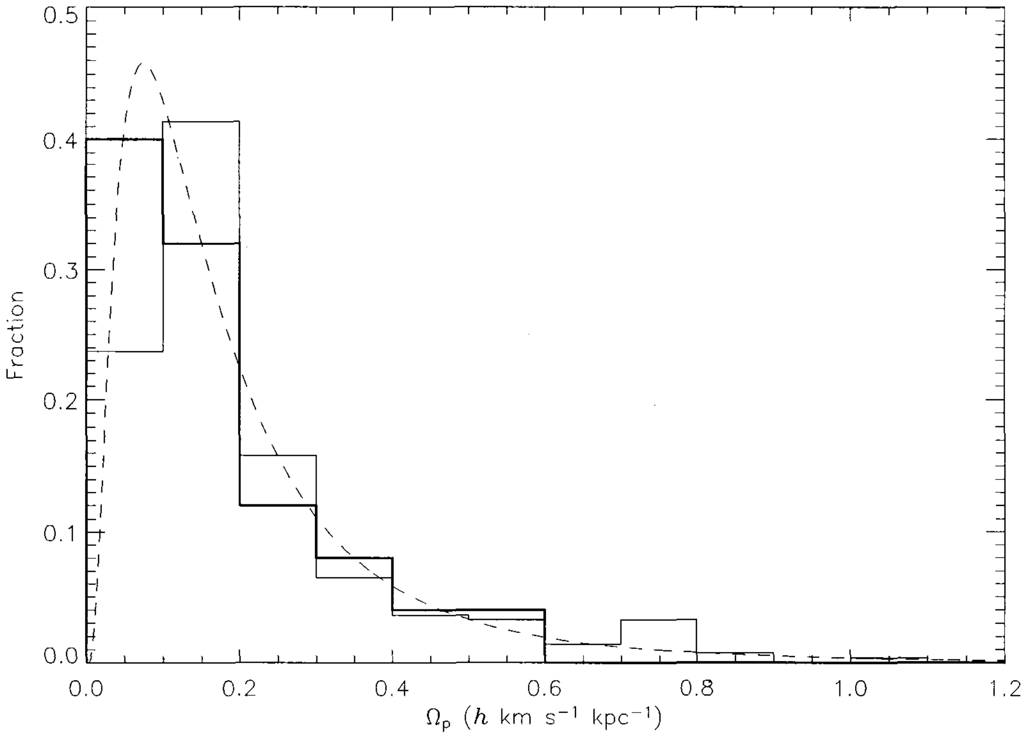

Figure 4.9: Fractional histogram of pattern speeds of the figure rotation, expressed linearly in Ω_p. The thin histogram contains all halos that have 2σ detections of figure rotation (incomplete at $\Omega_p < 0.126$ h km s^{-1} kpc^{-1}), while the thick histogram contains halos with errors less than 0.01 h km s^{-1} kpc^{-1}and is complete down to $\Omega_p = 0.015$ h km s^{-1} kpc^{-1}. The dashed curve is the log normal fit given by equation (4.14).

rotation. A histogram of the pattern speeds is presented in Figure 4.8, expressed in $\log \Omega_p$. The thin histogram contains all halos with 2σ detections, while the thick histogram contains those with the smallest errors, less than 0.01 h km s^{-1} kpc^{-1}. The largest upper limit due to a non-detection in the main sample is $\Omega_p < 0.126$ ($\log \Omega_p < -0.9$), so the thin histogram is incomplete below this level, while the low error sample contains only one upper limit, $\Omega_p < 0.015$ ($\log \Omega_p < -1.8$), so the thick histogram is complete down to this level. The dashed curves are Gaussian fits to the histograms. The fit to the thin histogram, which has the largest sample size but is incomplete at low Ω_p, peaks at $\log \Omega_p = -0.80$ and has a standard deviation of 0.29, while the thick curve, which contains fewer halos but is less biased toward large values of Ω_p, peaks at $\log \Omega_p = -0.84$ and has a standard deviation of 0.34. We give more weight to the thick histogram, whose points all have very small errors, and propose that the true distribution peaks at $\log \Omega_p = -0.83$ with a standard deviation of 0.36. Expressed as a log normal distribution, the probability is

$$P(\Omega_p) = \frac{1}{\Omega_p \sigma \sqrt{2\pi}} \exp\left(-\frac{\ln^2(\Omega_p/\Omega_{p_0})}{2\sigma^2}\right), \tag{4.14}$$

where $\Omega_{p_0} = 10^{-0.83} = 0.148\, h$ km s^{-1} kpc^{-1} and the natural width $\sigma = 0.36 \ln 10 = 0.83$. This fit is shown in Figure 4.9, compared to the fractional distribution of halos in the full (thin) and low error (thick) samples, and encompasses both the large number of halos with low Ω_p seen when the errors are sufficiently small, and the tail at high Ω_p seen when the sample size is sufficiently large.

For comparison, the halo in BFPM99 has a pattern speed of 2 h km s^{-1} kpc^{-1}. This lies slightly above the top end of our distribution; the maximum pattern speed in our sample is 1.01 h km s^{-1} kpc^{-1}. Based on the log normal fit of equation (4.14), we estimate the fraction of halos with $\Omega_p \geq 2$ h km s^{-1} kpc^{-1} to be $\sim 10^{-3}$. Therefore, this halo is unusual, but it is not unreasonable to find a halo with such a pattern speed in a large simulation. Given the size of the errors in P99, and that he found very few halos with figure rotation, it should not be surprising that P99 could only detect pattern speeds at the upper end of the overall distribution. The different adopted cosmologies may also influence the results (note, however, that this comparison is performed in h-independent units). Our results are also mostly consistent with D92, who finds pattern speeds of between 0.1 and 1.6 $\mathrm{km\ s^{-1}\ kpc^{-1}}$ in a sample of 14 halos. We have trouble reproducing the most rapidly rotating halos in D92, but this may be a product of the heuristic initial conditions in D92 compared to the cosmological initial conditions we use. Another useful reference value is v_{200}/r_{200}, the orbital frequency at r_{200}, which has a value of $1.0\ h\ \mathrm{km\ s^{-1}\ kpc^{-1}}$. This is the same for all halos, because the only timescale in a purely gravitational system is $1/\sqrt{\bar{\rho}}$, and the mean density $\bar{\rho}$ inside r_{200} is the same for all halos by definition.

In order to account for the spiral structure in NGC 2915, a triaxial figure would need to rotate at $7 \pm 1\ \mathrm{km\ s^{-1}\ kpc^{-1}}$ (MB03). This is almost an order of magnitude faster than the fastest of the halos in our sample, and the log normal fit from equation (4.14) suggests that the fraction of halos with $\Omega_p \geq 6\ \mathrm{km\ s^{-1}\ kpc^{-1}}$ is 5×10^{-7}. Therefore, the figure rotation of undisturbed ΛCDM halos can not explain the spiral structure of NGC 2915. SPH simulations of gas disks inside triaxial halos with pattern speeds of 0.77 $\mathrm{km\ s^{-1}\ kpc^{-1}}$, comparable to the fastest pattern speeds in our sample, show very weak if any enhancement of spiral structure compared to a static halo (Bekki & Freeman, 2002, Figure 2f). Therefore, it is unlikely that triaxial figure rotation can be detected in extended gas disks.

The dotted line in both Figures 4.7 and 4.8 shows the Nyquist frequency of $3.8\ h\ \mathrm{km\ s^{-1}\ kpc^{-1}}$. If the measured distribution of pattern speeds extended up to the Nyquist frequency, the intrinsic distribution would likely extend above the

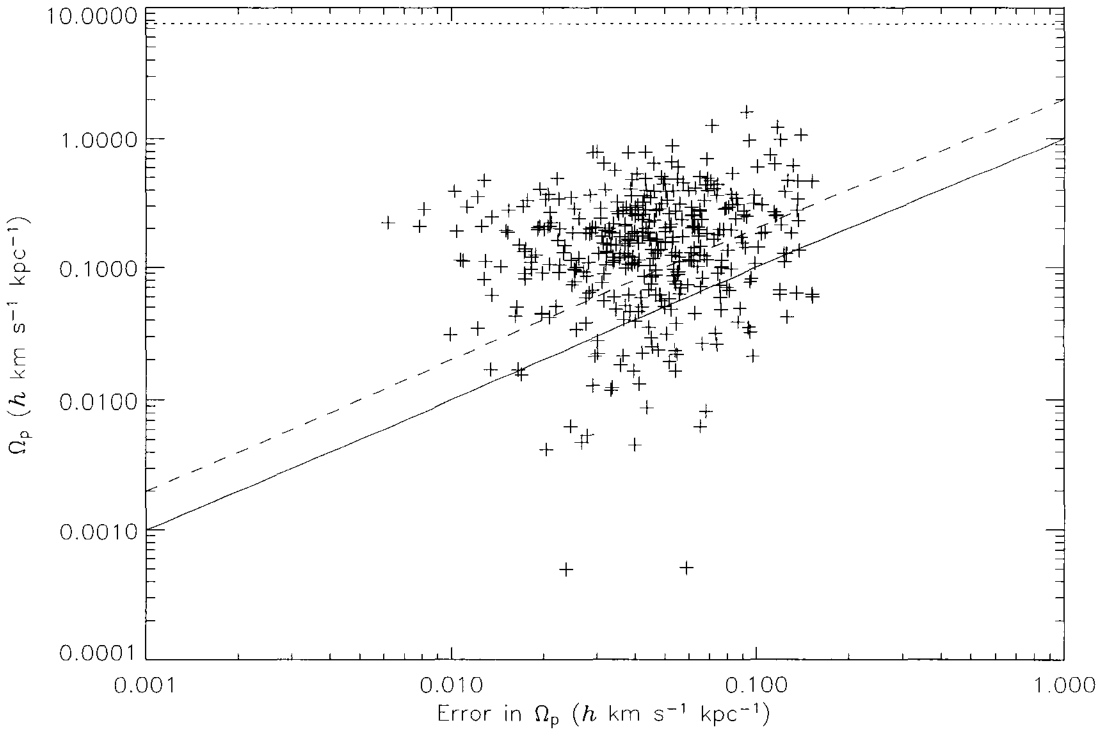

Figure 4.10:

As in Figure 4.7, but only including snapshots b096 through b102.

Nyquist frequency, and the results would be affected by frequency aliasing. However, the measured distribution does not approach the Nyquist frequency. Therefore, any halo whose figure rotation is aliased would need to be wildly anomalous, with a pattern speed many times faster than any other halo in our sample. We consider this unlikely. Figure 4.10 shows the pattern speeds as a function of the error, as in Figure 4.7, except that it only uses snapshots b096 through b102, so the maximum time between snapshots is 200 h^{-1} Myr, and the corresponding Nyquist frequency is 7.6 h km s^{-1} kpc^{-1}, shown again as the dotted line. The top of the distribution does not change between Figures 4.7 and 4.10, demonstrating that the results are not affected by aliasing.

We investigate how the figure rotation axis relates to two other important axes. Both D92 and P99 claim that the major axis rotates around the minor axis. The direction cosine between the rotation axis and the minor axis is plotted both as a function of the pattern speed and as a histogram in Figure 4.11. We confirm that the axis about which the major axis rotates aligns very well with the minor axis.

Measuring the rotation of the major axis using the plane method precludes finding halos with figure rotation about the major axis, which is a theoretically stable configuration. Therefore, the first and last snapshots are compared using the quaternion method. A comparison between the derived rotation axes, for those halos with 2σ detections of figure rotation, is shown in Figure 4.12. The axes agree in the majority of the halos. There are three halos where the axes are anti-aligned. Examination of these halos reveals that in these cases, the figure rotation from snapshots b096 through b102 is smooth, but snapshot b090 is anomalous (these halos have high values of χ^2_ν, though not high enough to have been excluded from the sample). More interesting are the halos where the alignment

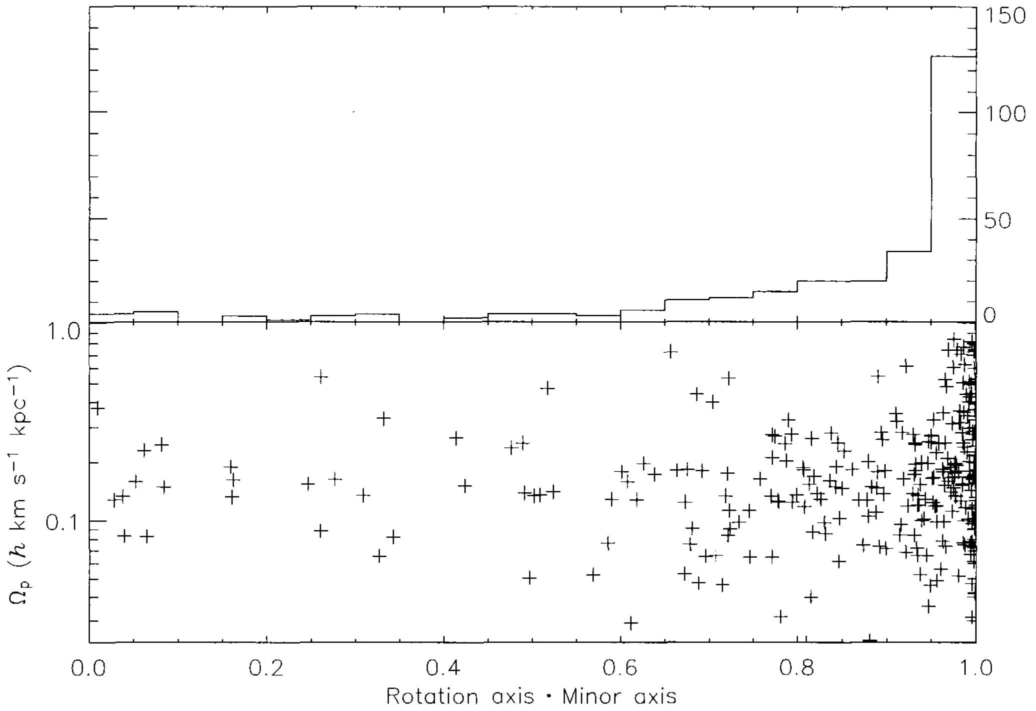

Figure 4.11: *(Top)*: Histogram of the direction cosine between the minor axis of the halo and the figure rotation axis, for halos with 2σ detections of figure rotation. Because the minor axis has reflection symmetry, this is always positive. *(Bottom)*: Direction cosine as a function of the magnitude of figure rotation.

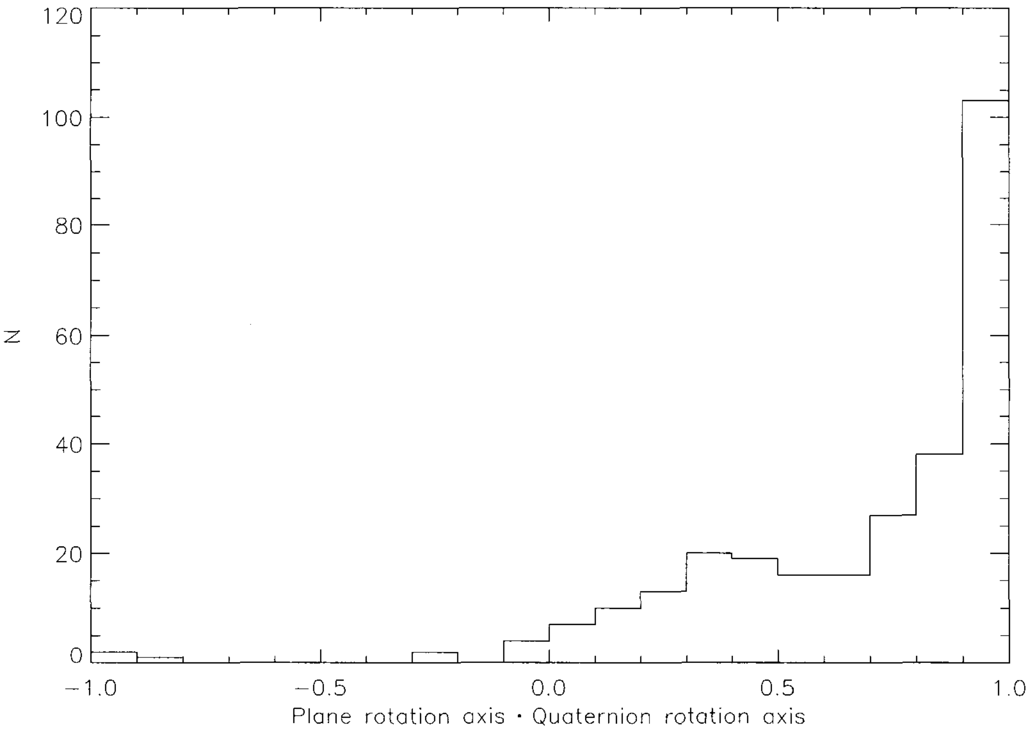

Figure 4.12: Histogram of the alignment between the figure rotation axis determined using the plane method and the figure rotation axis determined using the quaternion method (see § 4.3.4.1), for all of the halos with 2σ detections of figure rotation using the plane method.

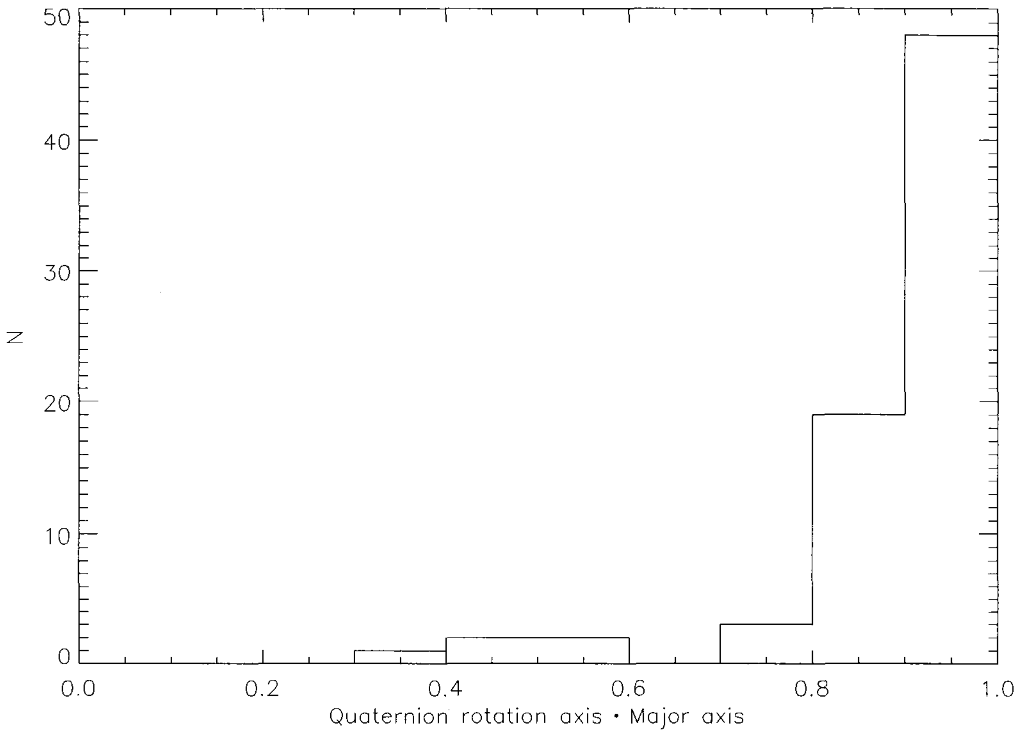

Figure 4.13: Histogram of the alignment between the figure rotation axis determined using the quaternion method and the halo major axis, for the halos where the agreement between the plane and quaternion methods is poor (the alignment of the axes is between -0.1 and 0.5).

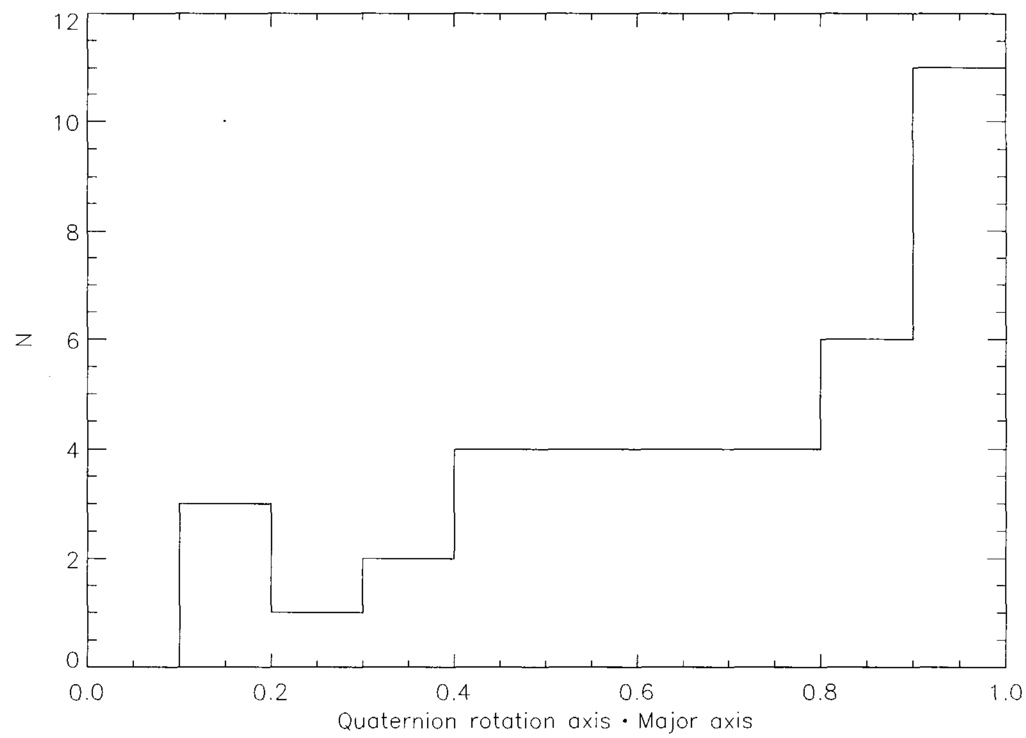

Figure 4.14: Histogram of the alignment between the figure rotation axis determined using the quaternion method and the major axis, for halos where we found no detectable figure rotation using the plane method.

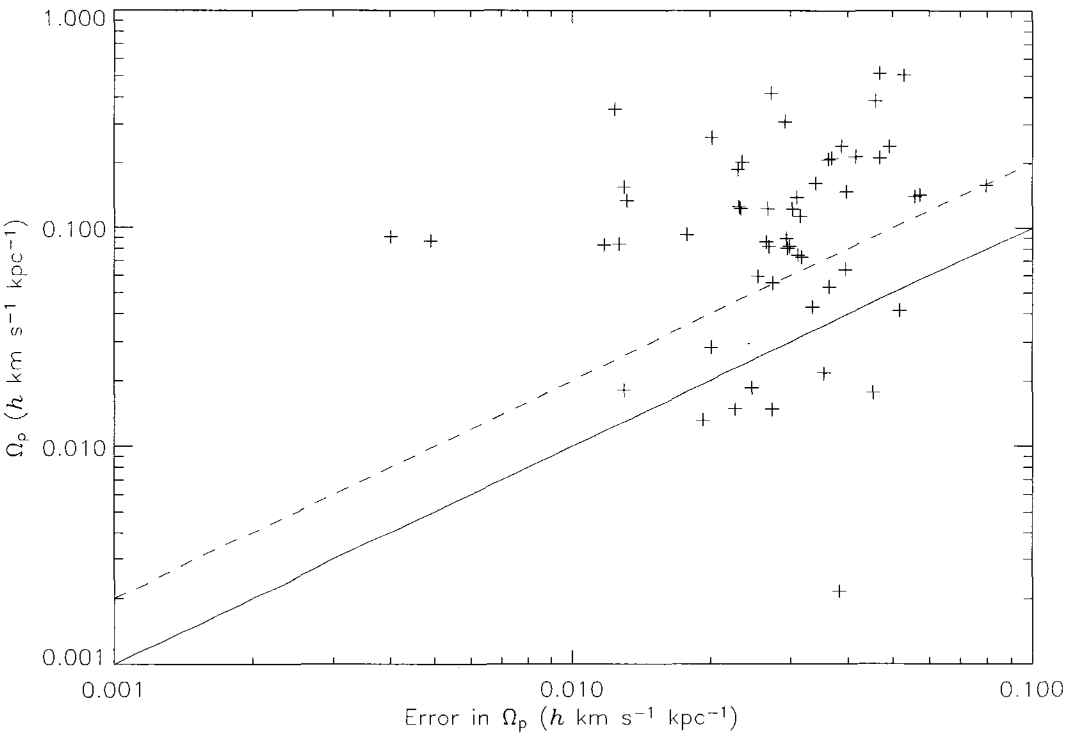

Figure 4.15: Measured pattern speed of the rotation of the minor axis, for potential major axis rotators. The x-axis is the error in the pattern speed. The solid line indicates where the measured pattern speed is equal to the estimated error, while the dashed line is the 2σ limit.

between the two determinations of the figure rotation axis lie between -0.1 and 0.5. In virtually all of these cases, the quaternion figure rotation axis is aligned with the major axis (see Figure 4.13). Some of the upper limits in Figure 4.7 may also be halos with rotation about the major axis. The quaternion figure rotation axes for most of the halos with upper limits are distributed randomly, but there is an excess of ~ 9 halos where the figure rotation axis aligns with the major axis (see Figure 4.14). We have constructed a sample of potential major axis rotators consisting of the 75 halos from Figure 4.12 where the alignment lies between -0.1 and 0.5 plus the 17 "upper limit" halos where the quaternion figure rotation axis aligns with the major axis to better than 0.8 (of which we expect statistically that ~ 9 are real; examination of the results reveals that 10 are real). We use the plane method to calculate the figure rotation of these halos, but investigate the evolution of the minor axis instead of the major axis. Out of the 92 halos, 37 have very large errors in the minor axis determination in one of the snapshots, and so are excluded. These are prolate halos where the quaternion method measures the fluctuations between the degenerate axes rather than true figure rotation. The pattern speeds for the remaining 55 halos are shown in Figure 4.15. We detect figure rotation about the major axis in 41 of the halos. Although the statistics are poorer, the range of pattern speeds for these halos is similar to the range of pattern speeds seen for the halos showing figure rotation about the minor axis. We conclude that 247 halos show minor axis rotation, 41 show major axis rotation, and 29 show no detectable figure rotation.

The rotation axis is compared to the angular momentum vector of the halo in Figure 4.16. Because the angular momentum is usually relatively well aligned with the minor axis of halos (Warren et al., 1992), it is no surprise that the rotation axis is also well aligned with the angular momentum vector. Because the

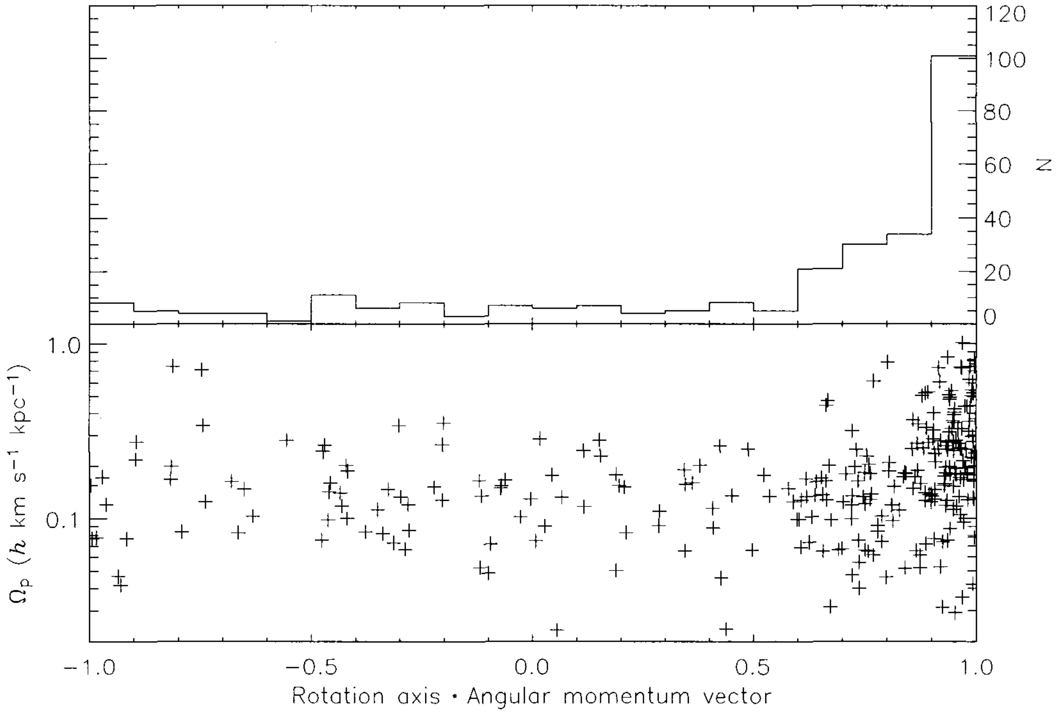

Figure 4.16: *(Top)*: Histogram of the direction cosine between the angular momentum vector of the halo and the figure rotation axis, for halos with 2σ detections of figure rotation. *(Bottom)*: Direction cosine as a function of the magnitude of figure rotation.

alignment between the minor axis and angular momentum of halos is not perfect (see § 5), some of the halos with perfectly aligned figure rotation and minor axes have less perfect alignment between the figure rotation axis and the angular momentum vector, seen as the bump in Figure 4.16 that extends down to a direction cosine of 0.6; the tail of the distribution extends all the way to anti-alignment. This indicates that the halos with retrograde figure rotation required for the polar ring and counter-rotating disk mechanism of Tremaine & Yu (2000) are rare but exist. The alignment between the rotation axis and the angular momentum vector is also plotted as a function of the pattern speed in the lower panel of Figure 4.16. There is no trend for the halos with slow figure rotation, but all but two of the halos with $\Omega_p > 0.4\ h$ km s^{-1} kpc^{-1} have figure rotation axes and angular momentum vectors that are well aligned, with a direction cosine of 0.65 or higher.

We have attempted to see if the pattern speed is correlated with other halo properties, in particular its mass and its angular momentum. Figure 4.17 shows the pattern speed of the figure rotation versus the halo mass. Error bars are 1σ errors, with 2σ upper limits plotted for halos which lie below the dashed line of Figure 4.7. There is no apparent correlation between the halo mass and its pattern speed.

Figure 4.18 shows the pattern speed versus the spin parameter λ, where

$$\lambda \equiv \frac{J\,|E|^{1/2}}{GM^{5/2}} \tag{4.15}$$

(Peebles, 1969). We use the computationally simpler λ' as an estimate for λ, where

$$\lambda' \equiv \frac{J}{\sqrt{2}MVR} \tag{4.16}$$

(Bullock et al., 2001). There is a tendency for halos with fast figure rotation to have large spin parameters; in particular, all but one of the halos with $\Omega_p > 0.4\ h$ km s^{-1} kpc^{-1} have $\lambda > 0.024$. These are the same halos which are shown to

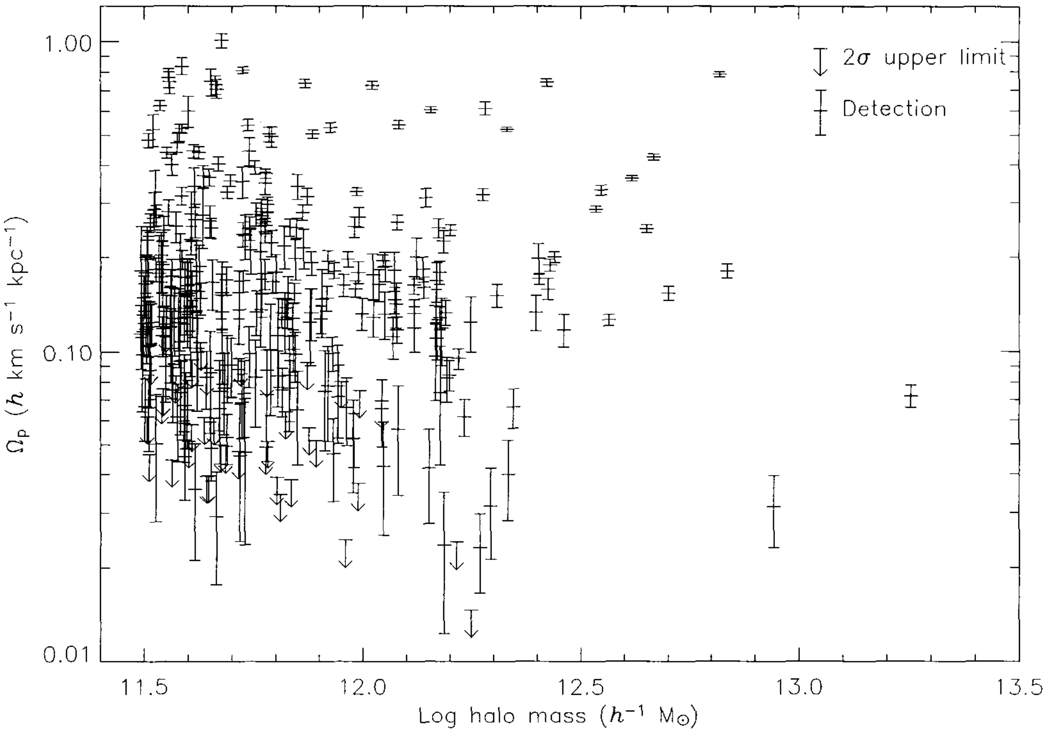

Figure 4.17: Pattern speed of the figure rotation versus the mass of the halo. Error bars are 1σ errors for halos with at least 2σ detections of figure rotation. The upper limits are the halos with a measured $\Omega_p < 2\sigma$, and are plotted at the 2σ limit.

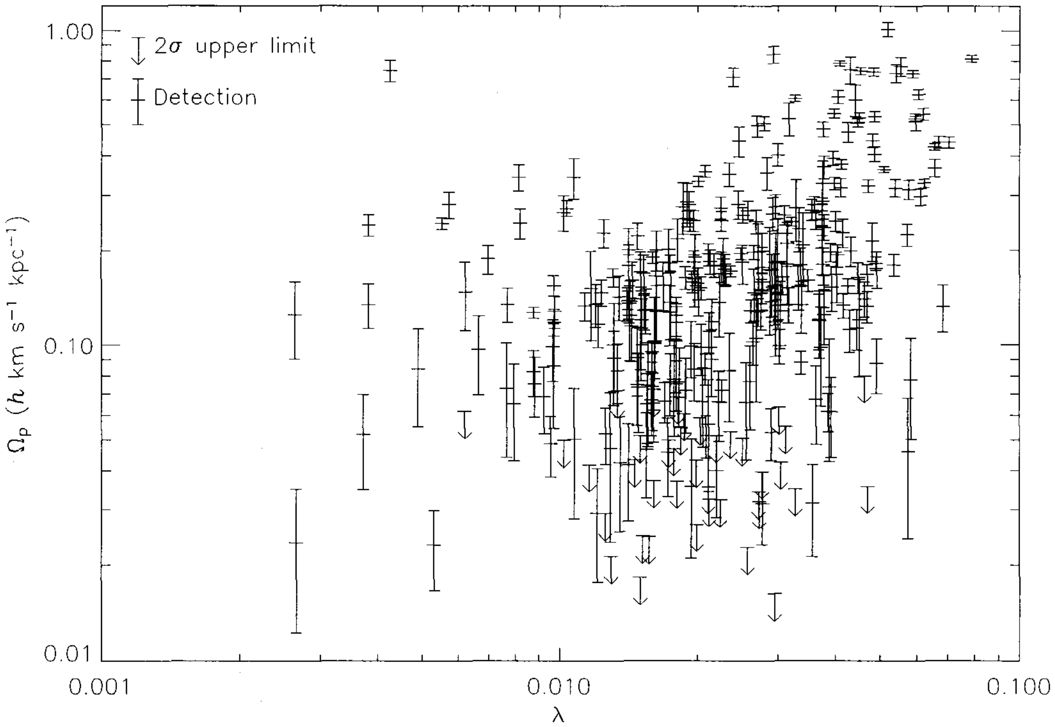

Figure 4.18: Pattern speed of the figure rotation versus the spin parameter of the halo. Error bars and upper limits are as in Figure 4.17.

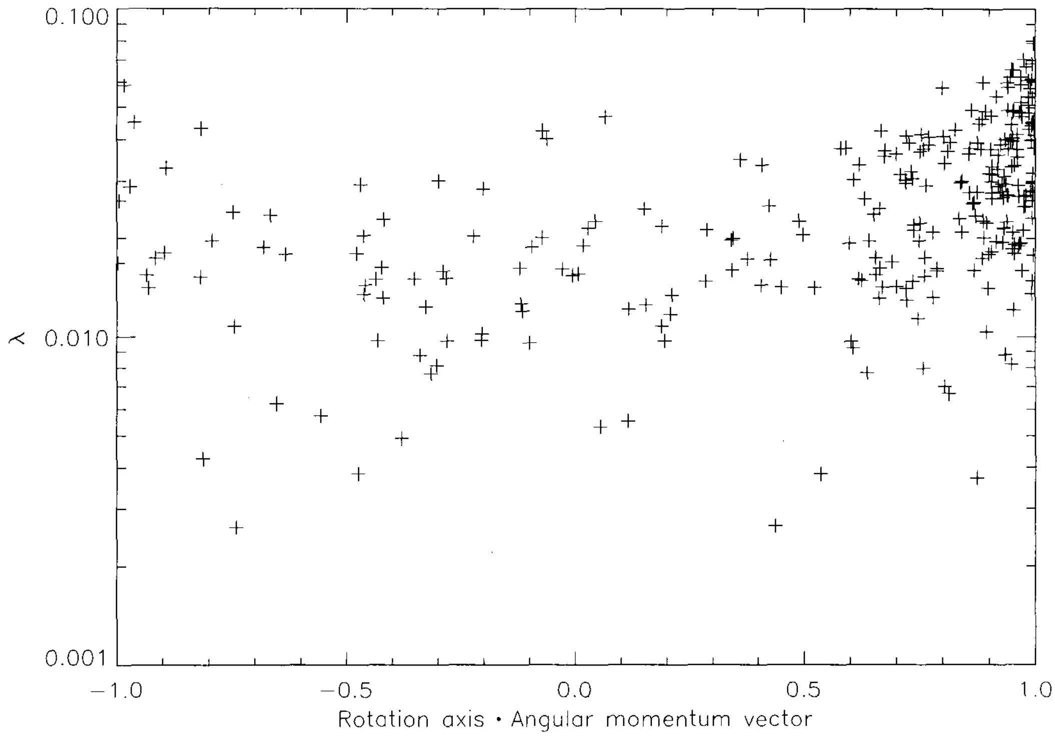

Figure 4.19: Direction cosine between the angular momentum vector of the halo and the figure rotation axis versus the spin parameter λ.

have particularly well-aligned rotation axes and angular momentum vectors in Figure 4.16. We have calculated the median value of Ω_p including the upper limits for bins of width $\Delta\lambda = 0.01$. The median rises steadily from 0.12 h km s^{-1} kpc^{-1} for $\lambda < 0.02$ to 0.44 h km s^{-1} kpc^{-1} for $\lambda > 0.06$. Note that P99 only detected figure rotation in halos with $\lambda > 0.05$ (see his Figure 5.24).

The degree of alignment between the figure rotation axis and the angular momentum vector may depend on the angular momentum content of the halo. Figure 4.19 shows how this alignment depends on the spin parameter λ. There is no trend for $\lambda < 0.05$, but the halos with $\lambda > 0.05$ show particularly good alignment. This is a natural consequence of the tendency for halos with rapid figure rotation to have well aligned figure rotation and angular momentum axes (Figure 4.16), and the correlation between figure pattern speed and spin parameter λ (Figure 4.18).

Figures 4.20 and 4.21 show how the figure rotation changes with radius. Figure 4.20 shows how the pattern speeds at different radii are related, while Figure 4.21 shows the alignment of the figure rotation axes between radii. Each panel includes only the halos that have at least 4000 particles within the inner radius, pass all of the tests of § 4.3.4.1 for both radii, and have 2σ measurements of figure rotation at both radii. Due to the smaller number of particles in spheres of smaller radii, there are progressively fewer halos with good measurements at smaller radii. The top panels show that the figure rotation in the outer regions of the halo is very coherent. To some degree, this is by construction; the test for substructure is equivalent to a cut in $\Delta\Omega_p$ between adjacent radii. However, gradual drifts of Ω_p and changes in the figure rotation axis with radius are still possible. The lower panels of Figures 4.20 and 4.21 show that this indeed happens. In particular, while the figure rotation within 0.12 $r_{\rm vir}$ and 0.6 $r_{\rm vir}$ are strongly correlated, the pattern speeds within 0.12 $r_{\rm vir}$ are slightly smaller and the alignment of the rotation axes is not quite as strong. The bottom panels show that in the innermost regions, within 0.06 $r_{\rm vir}$, the pattern speeds are significantly smaller than for the halo as a whole, particularly for those halos with high pattern speeds, and more than half of the halos show no alignment between the figure rotation axes.

We examine three possible explanations for these trends with radius. First, it may be that the halos with high pattern speeds are still affected by residual substructure in the outer regions. However, the gradual decline for all halos seen as the radius shrinks suggests that the mechanism responsible for the difference af-

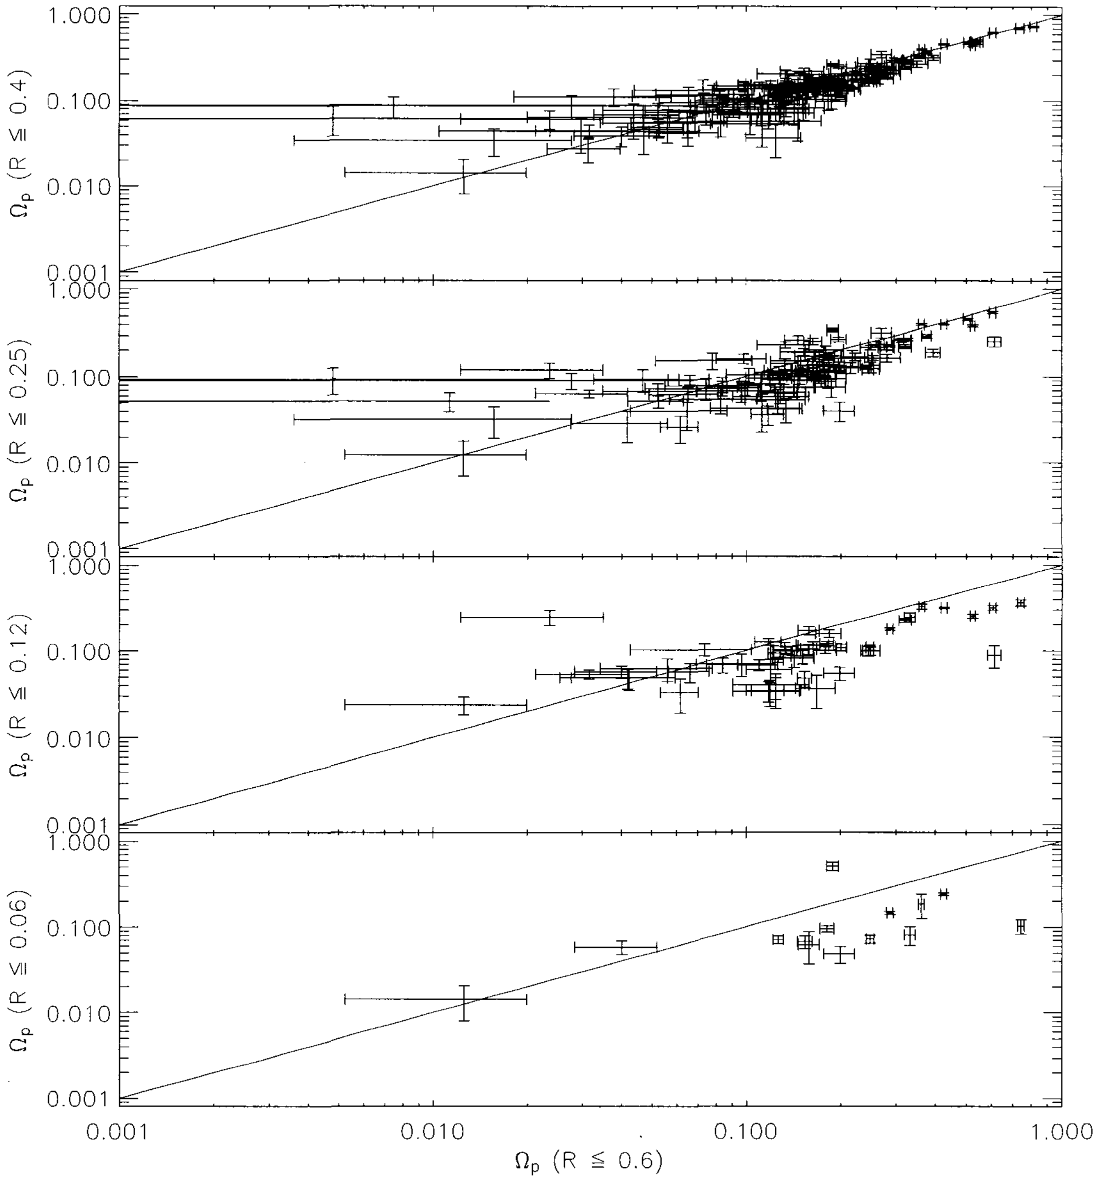

Figure 4.20: Pattern speed of figure rotation Ω_p at 0.4, 0.25, 0.12, and 0.06 of the virial radius $r_{\rm vir}$ (top to bottom) as a function of the pattern speed at 0.6 $r_{\rm vir}$. Only halos where both radii in the comparison contain at least 4000 particles, pass all of the tests of § 4.3.4.1, and have 2σ detections of figure rotation are included. All units are h km s^{-1} kpc^{-1}. The solid line corresponds to equal pattern speeds.

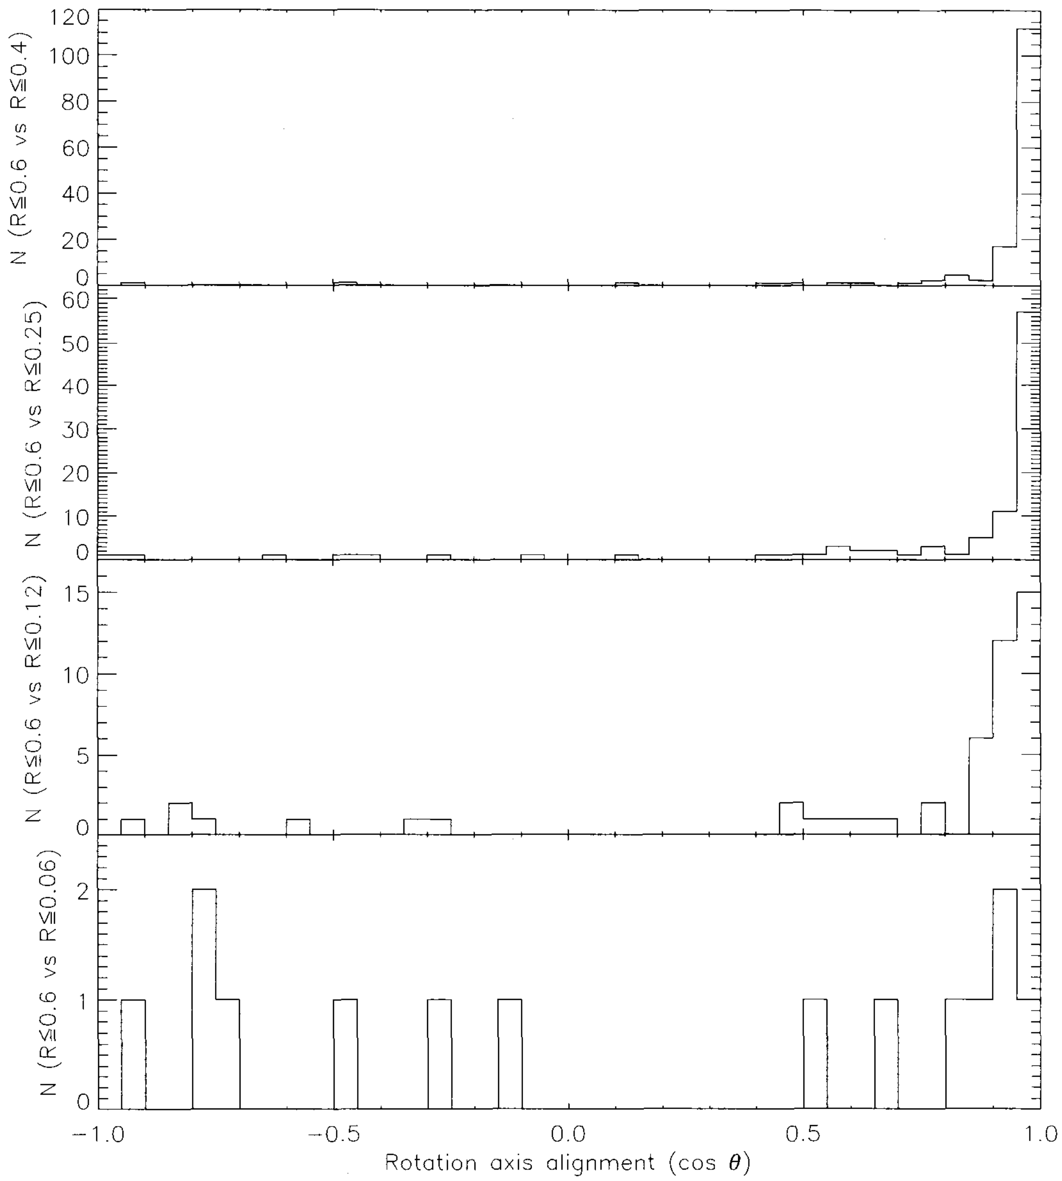

Figure 4.21: Histogram of the direction cosine of the alignment between the figure rotation axis of the halo as a whole ($r < 0.6\, r_{\rm vir}$) and of just the inner 0.4, 0.25, 0.12, and 0.06 of the virial radius $r_{\rm vir}$ (top to bottom). Only halos where both radii in the comparison contain at least 4000 particles, pass all of the tests of § 4.3.4.1, and have 2σ detections of figure rotation are included.

fects all halos equally and gradually, rather than affecting a few halos at a specific radius. Another piece of evidence that argues against this explanation is that the halos with the highest measured pattern speeds do not have preferentially high values of f_s; they have values evenly spread between 0 and the cutoff of 0.05. A second possibility is that the figure rotation, although steady on timescales of 1 Gyr, may be fundamentally a transitory feature caused by a tidal encounter or the most recent major merger. The inner region of the halo has a shorter dynamical time, and therefore the effects of such a disturbance will be erased faster in the inner regions than the outer regions. This is consistent with the gradual decrease in pattern speed with radius and the decrease in alignment. However, the halo of BFPM99 has fast figure rotation (faster than any of our halos), and yet shows steady figure rotation at all radii for 5 Gyr. We propose instead that the effects of force softening are becoming important at the smaller radii. The radius of the innermost sphere for the halos plotted in the bottom panels of Figures 4.20 and 4.21 range from 3–5 force softening lengths, where the effects of the gravitational softening can still be important (Power et al., 2003). The weaker gravitational force results in a more spherical potential, consistent with the weaker figure rotation and lack of alignment.

We have calculated the rate of change of the b/a and c/a axis ratios over the 5 snapshots using linear regression. The evolution of the axis ratios with time is linear for almost all halos. Figure 4.22 shows the fractional rate of change, $(\dot{b/a})/(b/a)$ and $(\dot{c/a})/(c/a)$, of the axis ratios as a function of the value of the axis ratio in the final snapshot. The median and standard deviation of the distribution of $(\dot{b/a})/(b/a)$ are 0.0089 h Gyr^{-1} and 0.0349 h Gyr^{-1} respectively. For c/a, they are 0.0093 h Gyr^{-1} and 0.0297 h Gyr^{-1} respectively. Therefore, there is a weak tendency for undisturbed halos to become more spherical with time. Most halos

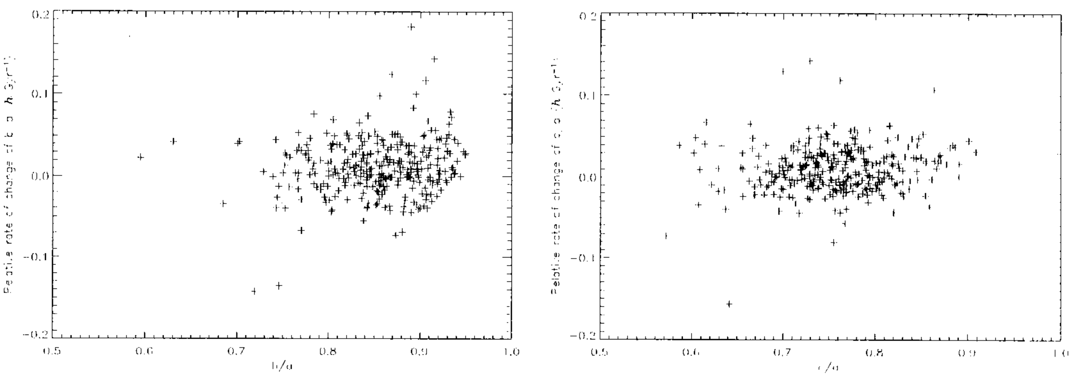

Figure 4.22: Fractional rate of change of the b/a *(left)* and c/a *(right)* axis ratio of the halos in the sample, as a function of the $z = 0$ axis ratio.

require several Gyr before their flattening changes significantly; there are, however, a few outliers with quite significant changes in their axis ratios. Figure 4.22 demonstrates that there is no trend of $(\dot{b/a})/(b/a)$ or $(\dot{c/a})/(c/a)$ with the value of the axis ratio except for the outliers with very high (low) values of $\dot{b/a}$, which could not have such high (low) rates of change if the values of b/a were not very high (low) in the final snapshot. We find no trend with any other halo property such as mass, spin parameter, pattern speed, substructure fraction, or alignment of the figure rotation axis with the angular momentum vector or minor axis.

4.5 Conclusions

We have detected rotation of the orientation the principal axes in most undisturbed halos of a ΛCDM cosmological simulation. The axis around which the figure rotates is very well aligned with the minor axis in $\sim 85\%$ of the cases, and well aligned with the major axis in the remaining halos. It is also usually well aligned with the angular momentum vector. The distribution of pattern speeds is well fit by a log normal distribution,

$$P(\Omega_p) = \frac{1}{\Omega_p \sigma \sqrt{2\pi}} \exp\left(-\frac{\ln^2(\Omega_p/\Omega_{p_0})}{2\sigma^2}\right), \tag{4.17}$$

with $\Omega_{p_0} = 0.148\, h$ km s^{-1} kpc^{-1} and $\sigma = 0.83$.

The pattern speed Ω_p is correlated with spin parameter λ. The median pattern speed rises from $0.12\, h$ km s^{-1} kpc^{-1} for halos with $\lambda < 0.02$ to $0.44\, h$ km s^{-1} kpc^{-1} for halos with $\lambda > 0.06$, with a spread of almost an order of magnitude around this median at a given value of λ. The 11% of halos in the sample with the highest pattern speeds, $\Omega_p > 0.4\, h$ km s^{-1} kpc^{-1}, not only have large spin parameters, but also show particularly strong alignment between their figure rotation axes and their angular momentum vectors. There is no obvious correlation of the figure rotation properties with mass. The pattern speed and figure rotation axis is coherent in the outer regions of the halo. Within 0.12 $r_{\rm vir}$, the pattern speed drops, particularly for those halos with fast figure rotation, and the internal alignment of the figure rotation axis deteriorates. This is probably an artifact of the numerical force softening.

BFPM99 hypothesized that the spiral structure in NGC 2915 is due to figure rotation of a triaxial halo. The required pattern speed of 7 ± 1 km s^{-1} kpc^{-1} (MB03) is much higher than the pattern speeds seen in the simulated halos, and is estimated to have a probability of 5×10^{-7}. We therefore conclude that the figure rotation of undisturbed ΛCDM halos is not able to produce this spiral structure. Halos with large values of λ tend to have more substructure (Barnes & Efstathiou, 1987), so there is a deficiency of halos with very high λ in our sample. Because Ω_p correlates with λ, we cannot exclude the possibility that there exist halos with very high λ whose figures rotate sufficiently quickly. However, halos with such high λ are themselves very rare (MB03), and if such halos fall out of our sample due to the presence of strong substructure, the effects of the substructure on the gas disk of NGC 2915 would be of more concern than the slow rotation of the halo figure, a possibility BFPM99 rule out due to the lack of any plausible companion in the vicinity.

More generally, Bekki & Freeman (2002) found very weak if any enhancement of spiral structure in disk simulations with triaxial figures rotating at 0.77 km s^{-1} kpc^{-1}, a value similar to the highest pattern speed seen in our sample. Therefore, it is unlikely that triaxial figure rotation can be detected by looking for spiral structure in extended gas disks.

The mechanism of Tremaine & Yu (2000) for creating polar rings and counter-rotating disks requires halos with retrograde figure rotation whose pattern speed slowly drops to zero or smoothly reverses direction. While our temporal sampling is not sufficient to detect slow changes in the pattern speed, we note that halos with figure rotation retrograde to their angular momentum do exist.

We have found that the axis ratios of undisturbed halos tend to become more spherical with time, with median fractional increases in the b/a and c/a axis ratios of $\approx 0.009\ h$ Gyr^{-1}. The distributions of $\dot{(b/a)}/(b/a)$ and $\dot{(c/a)}/(c/a)$ are relatively wide, with standard deviations of $\approx 0.03\ h$ Gyr^{-1}. A few outliers have axis ratios that change quite significantly over the span of 1 Gyr. The rate of change of the axis ratios is not correlated with any other halo property.

CHAPTER 5

INTERNAL AND EXTERNAL ALIGNMENT OF THE SHAPES AND ANGULAR MOMENTA OF ΛCDM HALOS

This chapter has been submitted for publication as Bailin & Steinmetz (2004b).

5.1 Introduction

The three-dimensional structure of the dark matter halos that host galaxies, groups, and clusters is an important aspect of their nature that can provide insight into their formation and affect the luminous structures within. The orientation of the halo shapes and angular momenta, both internally and with respect to surrounding halos, provide important constraints on other studies of galaxy formation and evolution.

Halos formed in cosmological simulations are generally not spherical, but have an ellipsoidal shape. There have been several studies of the shapes of halos in low resolution N-body simulations based on the standard cold dark matter (CDM) paradigm (Frenk et al., 1988; Dubinski & Carlberg, 1991; Warren et al., 1992 (hereafter W92); Cole & Lacey, 1996). These studies have found that halos are usually triaxial, with a preference for prolate figures at small radii and more oblate figures at large radii, and have minor-to-major axis ratios ranging from 0.3 to almost unity. Dubinski & Carlberg (1991), in simulations of the formation of isolated galaxies where the effects of the external tidal field were prescriptively superimposed, found that the projected two-dimensional ellipticities peak around $e = 0.5$, where

$$e \equiv 1 - q, \tag{5.1}$$

for a projected axis ratio of q. W92 found that while this is true in the inner regions, the location of the peak falls to $e = 0.25$ beyond 50 kpc. The behaviour of the axis ratios with radius is controversial; Dubinski & Carlberg (1991) and W92 find that the axis ratios increase (become more spherical) with radius, while Frenk et al. (1988) and Cole & Lacey (1996) find that they decrease with radius. W92 have also studied the internal alignment of the ellipsoid principal axes. They found that both the major and minor axes of halos are extremely well aligned out to 40 kpc.

More recently, several authors have performed large high resolution simulations using the currently-favoured ΛCDM cosmology (Bullock, 2002; Jing & Suto, 2002 (hereafter JS02); Kasun & Evrard, 2004 (hereafter KE04)). Bullock (2002) finds that the c/a axis ratios are a strong function of halo mass, and range from 0.55 at $10^{14}\ h^{-1}\ \mathrm{M}_{\odot}$ to 0.7 around $10^{12}\ h^{-1}\ \mathrm{M}_{\odot}$, with a distribution that is peaked but has a large tail to small axis ratios. He also finds that the inner 30 h^{-1} kpc of halos are more spherical than the outer regions, i.e. the axis ratios decrease with radius. On the other hand, JS02 find c/a axis ratios that increase with radius, decrease slightly with mass, and are well fit by a Gaussian centred at $c/a = 0.55$ with a width of 0.11. They find that the major axes of halos are relatively well aligned — typically $\cos\theta_{11} \sim 0.8$, where θ_{11} is the angle between the major axis at small or large radius compared to that at an intermediate radius. The alignment of the middle axes is somewhat poorer, but JS02 argue that this is due to the inclusion of nearly prolate halos whose axes are degenerate and therefore not well determined. KE04 find more spherical shapes in their very large sample of high mass halos ($M > 3 \times 10^{14}\ h^{-1}\ \mathrm{M}_{\odot}$), possibly due to the spherical outer boundary they impose, with $c/a \approx 0.65$.

The shapes of dark matter halos can have important observational consequences (for a good review, see Sackett, 1999). On galactic scales, they can affect the coherence of tidal streams. Some authors have claimed that the thinness of the tidal stream associated with the Sagittarius dwarf spheroidal indicates that the halo of of the Milky Way is nearly spherical, with $c/a \gtrsim 0.8$ (Ibata et al., 2001; Johnston et al., 2004; Martínez-Delgado et al., 2004). However, more recent studies suggest that the material that makes up the stream was stripped from the satellite too recently to have had time to undergo differential precession, which severely weakens the constraints on the halo shape (Helmi, 2004a,b). Helmi (2004b) claims that the stream is best fit by a prolate halo elongated perpendicular to the disk with $c/a \approx 0.6$. Law, Johnston, & Majewski (2004) also find that the velocities of stars in the leading stream can only be fit with a prolate halo, but that the precession of the leading stream with respect to the trailing stream can only be fit with an oblate halo. These contradictory results suggest that evolution of the satellite orbit or other effects of a live Milky Way potential (as opposed to the static potential that has been used in all of these studies) are important for determining the shape of the Milky Way halo using the Sgr stream. Other measures of the Milky Way ellipticity using the flaring of the gas disk at large radius or the anisotropy of stellar velocities suggest that the halo is oblate with a flattening of $c/a \sim 0.8$ (Olling & Merrifield, 2000). The shapes of the halos of external galaxies can be measured using the flaring of the gas layer (Olling & Merrifield, 2000), the projected shape of X-ray gas (Buote et al., 2002), or the kinematics of polar ring galaxies (Sackett et al., 1994). These methods suggest that galaxy halos have a wide range of flattenings from $c/a \sim 0.3 - 0.8$. A new method for measuring the shapes of external galaxy halos is weak gravitational lensing. By measuring the azimuthal variation of the shear with respect to the position angle of the visible lens galaxy (Natarajan & Refregier, 2000), Hoekstra et al. (2004) detected an average projected halo ellipticity of $\langle e \rangle = 0.33$ for halos with an average mass of $8 \times 10^{11}\ h^{-1}\ \mathrm{M}_{\odot}$. This detection also implies that the orientation of the visible and dark mass in galaxies must be similar. On group and cluster scales, X-ray observations and the Sunyaev-Zeldovich effect (Sunyaev & Zeldovich, 1980) can be directly used as a probe of halo ellipticities (Lee & Suto, 2003, 2004).

The orientation of the angular momentum in halos has also been studied in cosmological numerical simulations. Early low-resolution studies (Barnes & Efstathiou, 1987; Frenk et al., 1988; Cole & Lacey, 1996) gave conflicting results due to the difficulty of measuring the direction of the angular momentum with few particles. Other CDM studies (Dubinski, 1992; W92) have found that the direction of the angular momentum at different radii is usually the same, but that the distribution of alignments has a tail that stretches all the way to anti-alignment. This result is verified by recent high-resolution ΛCDM simulations (Bullock et al., 2001). CDM studies have also found that the angular momentum is most often aligned with the minor axis and perpendicular to the major axis, although there is some scatter (Dubinski, 1992; W92). This result has not yet been thoroughly tested in high-resolution ΛCDM simulations.

Internal misalignment of the angular momentum can have a number of observational consequences. It may cause galactic warps (Ostriker & Binney, 1989; Debattista & Sellwood, 1999; López-Corredoira et al., 2002a; see also § 3), or manifest itself in anisotropic distributions of the orbits of satellite galaxies (Holmberg, 1974; Zaritsky et al., 1997; Aubert et al., 2004; Knebe et al., 2004).

Going beyond individual halos, the shapes and angular momenta of nearby halos can be correlated due to initial conditions or dynamical evolution. This subject has attracted increased interest recently due to the emergence of weak gravitational lensing as a method to measure the projected mass density in front of background galaxies. Intrinsic correlations between the projected shapes of luminous galaxies act as spurious background signals in weak lensing, so predicting their magnitude is important. At a more fundamental level, the degree of correlation between structures can be tested against models, and can inform our understanding of the origin of halo shapes and angular momenta.

Measurements of halo alignments and correlations come from two sources: cluster orientations and large galaxy surveys. The study of the alignment of cluster orientations was pioneered by Binggeli (1982), who used the locations of the constituent galaxies to determine that the major axes of clusters separated by less than 15 h^{-1} Mpc tend to point toward each other. While some authors have not found any such correlation (e.g. Struble & Peebles, 1985), larger samples of both galaxies and clusters, along with improved error estimates, have confirmed this result (e.g. Flin, 1987; Rhee & Katgert, 1987; Plionis, 1994). The cluster potential is better probed by X-ray emitting gas (Sarazin, 1986; Lee & Suto, 2003, 2004). While early studies using X-ray contours found no alignment of clusters with the large scale structure (Ulmer et al., 1989), more and better data have confirmed that the orientation of both the substructure and the main cluster potential tends to point toward neighbouring clusters (West et al., 1995; Chambers et al., 2002)

While the principal axes of clusters can be determined from optical or X-ray photometry, the angular momentum direction is very difficult to determine. In disk galaxies, on the other hand, the angular momentum direction of the baryons, presumed to be perpendicular to the orientation of the disk, can be measured much more easily than the shape of the dark matter halo. Although there may be some misalignment between the angular momentum of the baryons and dark matter (van den Bosch et al., 2002; Sharma & Steinmetz, 2004), spiral galaxies still provide the best targets for looking for angular momentum alignments and correlations. Studies with small samples of spiral galaxies (less than a few hundred) have generally found no correlation between the orientation of the angular momentum and the large scale structure (Han et al., 1995; Cabanela & Dickey, 1999; see however Navarro, Abadi, & Steinmetz, 2004). With a sample of 618 lenticular and disk galaxies in the local supercluster, Kashikawa & Okamura (1992) found that while the full sample was consistent with an isotropic distribution of angular momenta, those galaxies within 2 h^{-1} Mpc of the supergalactic plane tend to have spin vectors pointing in the plane, while those above or below the plane tend to have spin vectors that point toward or away from the plane. Navarro et al. (2004) also find a clear excess of galaxies whose angular momenta lie in the supergalactic plane. Larger samples of galaxies provide further evidence of alignments between spin and the large scale structure: Pen, Lee, & Seljak (2000) have found that the spin directions of the 12,122 spiral galaxies in the Tully catalog are positively correlated at separations less than 3 h^{-1} Mpc, while Brown et al. (2002) have measured intrinsic alignments between galaxy orientations at a range of angular separations in the 2×10^6 galaxies of the SuperCOSMOS survey.

Many recent theoretical studies of halo alignment have been motivated by predicting the intrinsic alignment signal for weak lensing measurements. These studies jump straight to correlations of the ellipticities of luminous galaxies as a function of angular separation on the sky (e.g. Croft & Metzler, 2000; Heavens et al., 2000; Catelan et al., 2001; Crittenden et al., 2001). While these studies are useful for those making weak lensing observations, they do not provide any direct measurement of the intrinsic spatial alignments that can provide physical insight into the galaxy formation process.

Linear tidal torque theory (Doroshkevich, 1970; White, 1984) can be used to predict the directions of the angular momentum vectors and their correlations with the surrounding structure (Pen et al., 2000; Lee & Pen, 2000, 2001; Porciani, Dekel, & Hoffman, 2002a,b). These studies have found that the angular momenta of halos tend to lie perpendicular to the large scale structure, and that the correlation of the halo spin vectors with each other exists but is very weak. While tidal torque theory does a reasonable job of predicting the evolution of the magnitude of the angular momentum (Sugerman et al., 2000), Porciani et al. (2002a) have tested the predictions of spin directions against N-body simulations, and found that the spin axes of N-body halos show significant misalignment compared to the tidal torque predictions, with a mean misalignment of $\sim 50°$ at $z = 0$. Therefore, N-body simulations that take the full non-linear dynamics into account are necessary.

Some early numerical work at predicting the intrinsic alignments using simulations with power law or CDM power spectra found that the major axes of cluster mass halos tend to point toward other nearby clusters over scales of $\sim 15\ h^{-1}$ Mpc, and that there is a very weak tendency for the major axes to be correlated with each other over the same range of separations (Barnes & Efstathiou, 1987; West et al., 1991; van Haarlem & van de Weygaert, 1993; Splinter et al., 1997). More recent high resolution N-body simulations in a ΛCDM cosmology have been studied to search for alignments of the major axes (Onuora & Thomas, 2000; KE04), the angular momentum vectors (Hatton & Ninin, 2001), or both (Faltenbacher et al., 2002; hereafter F02). These studies have found a strong tendency for the major axes of cluster mass halos to point toward other clusters out to several tens (or even hundreds) of Mpc and to correlate with each other out to 20 h^{-1} Mpc. The situation for the angular momentum is murkier. Hatton & Ninin (2001) have found that the angular momenta of halos tend to lie parallel to the large scale structure, while F02 have found that they lie perpendicular to the large scale structure, in agreement with the prediction from the linear theory (Lee & Pen, 2001). Correlations of halo angular momentum vectors with each other are weak at best.

In this chapter, we present an extensive study of the shapes of galaxy and group-mass halos in a large high resolution ΛCDM N-body simulation. We study the internal alignments of all of the principal axes and the angular momentum. We also study the alignment of all of these quantities with the large scale structure, and how they are correlated in halos of different separations. This work improves upon earlier studies of the internal structure of halos by using large high resolution simulations in a currently-favoured ΛCDM cosmology, by studying both the alignment of the angular momentum and the shape, and by using a method that allows us to quantify our errors and therefore feel confident about the source of any measured misalignments. Previous external alignment studies have all been restricted to massive clusters; we improve upon this significantly by reaching down to galaxy mass halos, by studying the mass dependence of the correlations, by studying the alignments of both the principal axes and the angular momenta, and by including the oft-neglected intermediate and minor axes. The structure of the chapter is as follows. In § 5.2 we present the details of the simulation and describe the method used to measure the principal axes and angular momentum vectors of the halos. § 5.3 presents the overall shapes of the halos and how they change with radius. We discuss the internal alignment of the principal axes and angular momenta in § 5.4, while we explore the alignment of these quantities with external halos in § 5.5. We present a brief discussion of what these results imply for the origin of halo shapes and angular momenta in § 5.6, and summarize the results in § 5.7.

5.2 Methodology

5.2.1 The simulation

The simulation used here is the same as the one used in § 4. It consists of 512^3 N-body particles in a periodic box of length 50 h^{-1} Mpc, in a low density flat universe ($\Omega = 0.3$, $\Omega_\Lambda = 0.7$, $h = 0.7$, $\sigma_8 = 0.9$). The particle mass is 7.757×10^7 h^{-1} $\mathrm{M}_\odot$, and the force softening length is 5 h^{-1} kpc. Halos were found using the standard friends-of-friends (FOF) algorithm with a linking length of 0.2 times the mean inter-particle separation.

In order to accurately measure the direction of the principal axes, many particles are needed. The direction of an axis can generally be determined to within an angle of

$$\theta_{\mathrm{err}} = \frac{1}{2\sqrt{N}} \frac{\sqrt{r}}{1-r} \quad \text{radians,} \tag{5.2}$$

where N is the number of particles used and r is the relevant axis ratio: b/a for the major axis, c/b for the minor axis, and $\max(b/a, c/b)$ for the intermediate axis (see § 4). For the purposes of measuring internal alignments, we would like the angular errors to be less than 10°. This requires on the order of 200 particles. Since each halo is split up into 6 radial shells, and this accuracy is required in each shell, each halo should have over 1200 particles. For convenience, we adopt a cutoff of 10^{11} h^{-1} $\mathrm{M}_\odot$ for the mass of the halo, or 1289 particles.

There are 3869 halos in the sample with masses extending from 10^{11} h^{-1} $\mathrm{M}_\odot$ to 2.8×10^{14} h^{-1} $\mathrm{M}_\odot$. 451 of the halos have masses in the range $10^{12} - 10^{13}$ h^{-1} $\mathrm{M}_\odot$, while 62 of the halos have masses greater than 10^{13} h^{-1} $\mathrm{M}_\odot$.

5.2.2 Measuring the axes

A standard technique to measure halo triaxiality in simulations is to use an iterative approach, where the particles are initially chosen to lie in a sphere or spherical shell, an ellipsoid is fit to these particles, and particles are chosen for the next iteration based on the new ellipsoid (e.g. W92). While this works for simulations that have sufficiently low resolution that overmerging erases substructure, we find in agreement with JS02 that in high resolution simulations, the presence of substructure prevents this technique from converging for a large fraction of halos.

JS02 have adopted a novel approach aimed at directly measuring isodensity contours. They assign SPH-like densities to halo particles, and then measure the principal axes and axis ratios of particles with densities near the nominal density of the isodensity contour. Due to the presence of substructure, this procedure often picks out disconnected shells in addition to the particles that define the isodensity contour of the smooth distribution. JS02 use the FOF algorithm to select the largest structure that fulfils the density criterion and therefore eliminate the substructure.

While this algorithm works well for very high-resolution halos, such as those JS02 use to demonstrate the technique (all of which contain $N > 6 \times 10^5$ (Jing & Suto, 2000)), we have run into difficulties using it on more moderate resolution halos. In particular, we find that it is not possible to find an optimal FOF linking length for eliminating substructure; if the linking length is too large, many "contours" contain obvious disconnected substructures, while reducing the linking length sufficiently to eliminate this problem reduces it to the point where in many halos, the single ellipsoid corresponding to the smooth distribution is broken up by the algorithm into several disconnected pieces. JS02's algorithm also uses a relatively small number of particles to determine the shape. The error in the determination of the direction of the principal axes of a particle distribution goes as $N^{-1/2}$ (see eq. [5.2]). Measuring the internal alignment of a halo requires small well-understood errors, and therefore as many particles as possible.

We take the following approach, similar to that taken by Frenk et al. (1988). The center of mass is calculated iteratively in spheres centred on the center of mass of the sphere in the previous iteration, starting with a sphere containing all of the particles. The radius of each successive sphere is reduced by 90%, and the procedure is iterated 25 times, by which point it has converged. The particles of each halo are transformed into this center of mass frame, and the velocities are transformed into the center of velocity frame. Each halo is then divided up into six concentric spherical shells with outer radii R of 1.0, 0.6, 0.4, 0.25, 0.12, and 0.06 times the virial radius $r_{\rm vir}$. These radii are chosen to allow easy comparison to the isodensity contours of JS02. The outer radius of each shell also forms the inner radius of the next larger shell, except for the innermost "shell", which is actually a sphere extending to the halo center. One would like to use the inertia tensor to measure the principal axes of the mass distribution. However, the inertia tensor can be dominated by substructure in the outer part of the shell. Therefore, we weight particles by $1/r^2$ so that each mass unit contributes equally regardless of radius (Gerhard, 1983). Within each shell, we calculate this reduced moment of inertia tensor

$$\tilde{I}_{ij} = \sum_k \frac{r_{k,i} r_{k,j}}{r_k^2}, \qquad (5.3)$$

which we then diagonalize. The axis ratios a, b, and c are the square roots of the eigenvalues ($a \geq b \geq c$), and the eigenvectors give the directions of the principal axes. There are no particles in common between radial shells, so the measurements of the axes at different radii are completely independent. We also calculate the angular momentum of the particles in each shell.

To calculate the error in the axis and angular momentum determinations, we perform a bootstrap analysis of the particles in each radial shell (Heyl et al., 1994). If the shell contains N particles, we resample the shell by randomly selecting N particles from that set allowing for duplication and determine the axes and angular momentum from this bootstrap set. We do this 100 times for each radial shell. The dispersions of these estimates of the axis ratios, directions of the axes, magnitude of the angular momentum, and direction of the angular momentum around the measured values are taken formally as the "1σ" error of each of these quantities.

Using spherical shells, rather than the ellipsoids defined by the isodensity contours as in JS02, does not affect the orientation of the principal axes. It does, however, bias the derived axis ratios toward spherical. To calibrate the magnitude of this bias, we have constructed prolate Poisson-sampled NFW (Navarro, Frenk, & White, 1996) and singular isothermal halos with 10^5 particles each that have known c/a axis ratios ranging from 0.3 – 1. The NFW halos had concentration parameters $c_{\rm vir} = 10$. In Figure 5.1 we plot the input c/a axis ratio for each combination of radial bin and density distribution as a function of the c/a ratio measured for the halos constructed using the method described above. The points for different radii and density distributions are virtually indistinguishable; the only important parameter is the input axis ratio. The solid line is an empirical fit to these points:

$$(c/a)_{\rm true} = (c/a)^{\sqrt{3}}_{\rm measured}. \tag{5.4}$$

The relationship for b/a is identical. The corrected axis ratios $(b/a)_{\rm true}$ and $(c/a)_{\rm true}$ are used throughout the remainder of this paper.

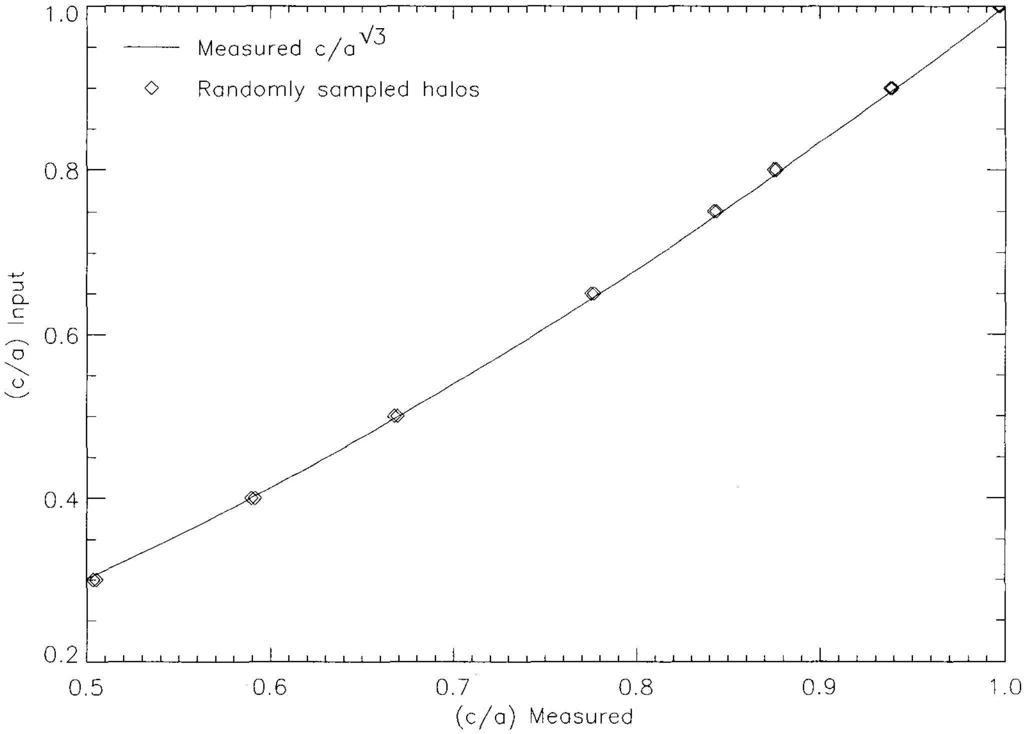

Figure 5.1: Measured versus input c/a axis ratio of randomly-sampled NFW and singular isothermal halos. At each input axis ratio, there are 12 independent points, representing the results for each of six radial bins and two different density profiles, but they are virtually indistinguishable. The solid line is the empirical fit of equation (5.4).

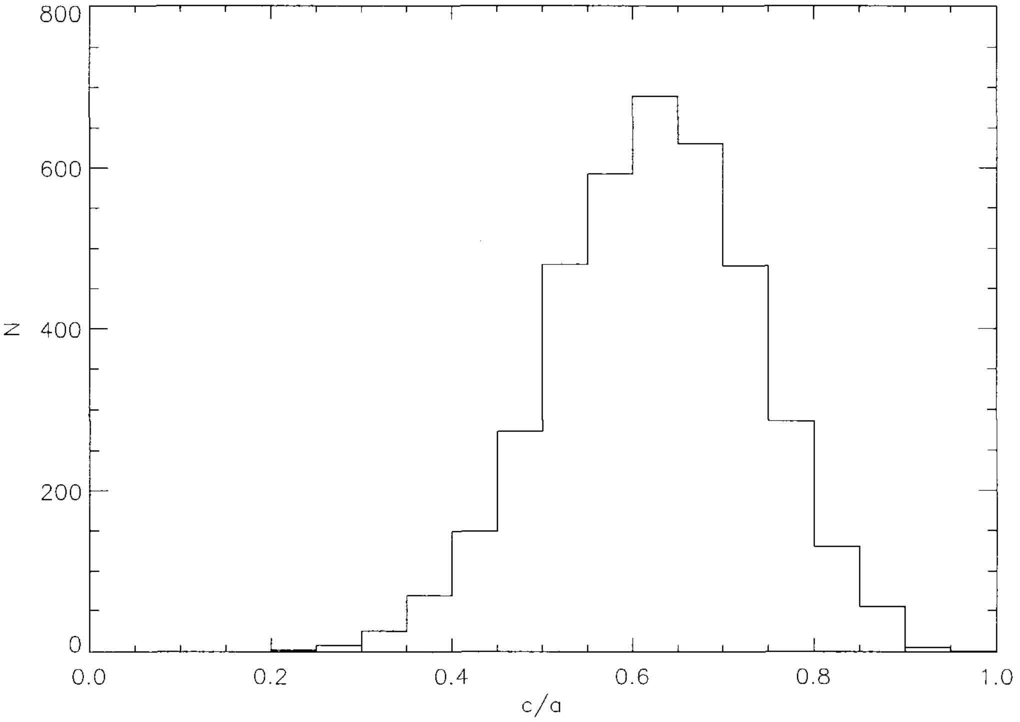

Figure 5.2: Histogram of the minor-to-major c/a axis ratio for each halo in the simulation, as measured in the $R = 0.4\,r_{\rm vir}$ shell.

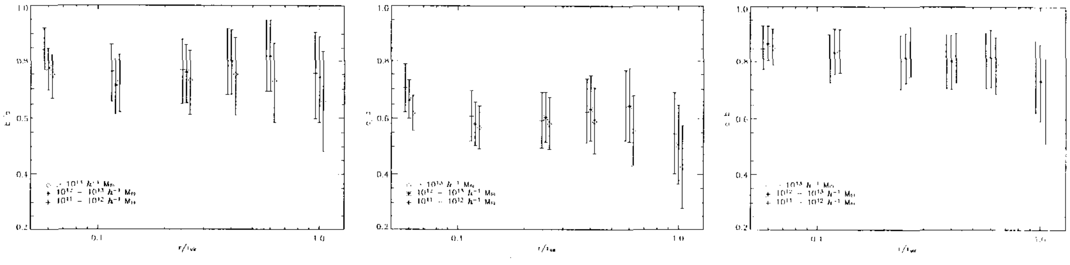

Figure 5.3: Median of the b/a *(left)*, c/a *(middle)*, and c/b *(right)* axis ratios for each radial shell. Crosses, asterisks, and diamonds represent the mass ranges 10^{11} – 10^{12} h^{-1} $M_\odot$, 10^{12} – 10^{13} h^{-1} $M_\odot$, and $10^{13} - 3 \times 10^{14}$ h^{-1} $M_\odot$ respectively. Error bars represent the 1σ width of the distribution. The error in the median is typically 0.002 for the crosses, 0.005 for the asterisks, and 0.01 for the diamonds. Crosses and diamonds are offset in radius for clarity.

5.3 Shapes

Dark matter halos are well approximated by ellipsoids (JS02) and are well described by the intermediate-to-major and minor-to-major axis ratios b/a and c/a. Figure 5.2 is a histogram of the c/a axis ratios as measured in the $R = 0.4\,r_{\rm vir}$ shell for all of the halos in our sample. Unlike JS02, we find that they are not quite Gaussian, but rather have a tail toward very flattened halos as seen by KE04, although the tail is not as extreme as that seen by Bullock (2002). The distribution of b/a and c/a values measured at other radii have a similar shape. The lack of very flattened halos in JS02 may be a result of their exclusion of halos deemed to be interacting. The axis ratios we find are intermediate between the quite flattened halos found by JS02 and the more spherical halos found by KE04.

Early studies suggested that the coldness of the Sgr stream indicates that the dark matter halo of the Milky Way has $c/a \gtrsim 0.8$ (Ibata et al., 2001; Johnston et al., 2004; Martínez-Delgado et al., 2004). Only 5% of the halos shown in Figure 5.2

have axis ratios so large. While this is not negligible, it is uncomfortably small, and forces us into the anti-Copernican situation of living in an exceptional galaxy (Copernicus, 1543). However, new models of the Sgr stream that are more careful about matching observations of the body of the dwarf find that the stars that constitute the stream were stripped from the body of the satellite too recently to have had time to undergo differential procession, thereby severely weakening the constraints on the halo ellipticity (Helmi, 2004a,b); current models are unable to simultaneously fit the velocities of stars in the leading stream and the orbital planes of the leading and trailing streams (Law et al., 2004). It should also be noted that our simulations do not take into account the effects of baryonic physics. There is some evidence that baryon cooling leads to more spherical halos (Dubinski, 1994; Kazantzidis et al., 2004). Therefore systems in which most of the baryons have cooled, such as disk galaxies like the Milky Way, may have dark matter shapes that are more spherical than those presented here. Observations of external galaxies using a variety of methods find halo flattenings that range from 0.3 to 0.8, in agreement with our results (Sackett, 1999).

The radial dependence of the axis ratios is shown in Figure 5.3. There are three distinct regions of the halo. Over most of the halo, the axis ratios increase with radius (i.e. the halos become more spherical). Near the virial radius, infalling unvirialized structure causes the axis ratios to drop. In the central 6% of the virial radius, the axis ratios rise. However, this is probably an artifact of the numerical softening, as shown in § 4. Further evidence that this is a numerical effect comes from examining the location of this increase in sphericity for halos of different mass. The increase occurs at a larger fraction of the virial radius for low mass halos, i.e. at a similar physical radius. This complicated and non-monotonic radial dependence may explain the discrepancy between studies that have found that flattening increases with radius (Frenk et al., 1988; Cole & Lacey, 1996; Bullock, 2002) and those that have found that it decreases with radius (Dubinski & Carlberg, 1991, W92, JS02). The mass dependence of the axis ratios is shown by the different symbols in Figure 5.3. The highest mass halos have smaller axis ratios at all radii than the smaller halos (W92; Bullock, 2002; JS02; KE04). Our halos extend down to masses an order of magnitude smaller than any other study; the difference between the most massive halos (the diamonds) and the galaxy mass halos is more pronounced than the difference between different masses of galaxy halos. The high mass halos show particularly strong flattening near the virial radius.

As seen in Figure 5.3, the c/b ratio falls steadily with radius, indicating a transition from prolate figures in the center to oblate figures at large radii (Dubinski & Carlberg, 1991; W92). Triaxiality can be quantified by the parameter T:

$$T = \frac{a^2 - b^2}{a^2 - c^2} \tag{5.5}$$

(Franx et al., 1991). Purely prolate halos have $T = 1$ while purely oblate halos have $T = 0$. We have measured T at three radii: $R = 1.0\ r_{\rm vir}$, where infalling material results in substantial flattening, $R = 0.6\ r_{\rm vir}$, where the halos are at their least flattened, and $R = 0.12\ r_{\rm vir}$, where the interior of the halos are at their most flattened. Histograms of T at these radii are shown in Figure 5.4. The interior regions of halos clearly tend to be prolate (solid histogram). As the flattening decreases at larger radii, many of the halos become more oblate (dashed histogram), although still more are prolate than oblate. Near the virial radius, there is a small shift back toward prolate shapes (dotted histogram). Figure 5.5 shows the full relationship between b/a and c/a for all of the halos in our sample, measured at $R = 0.4\ r_{\rm vir}$, where the values are most typical for the halo as a whole. The preponderance of prolate and triaxial halos over oblate halos is clearly seen.

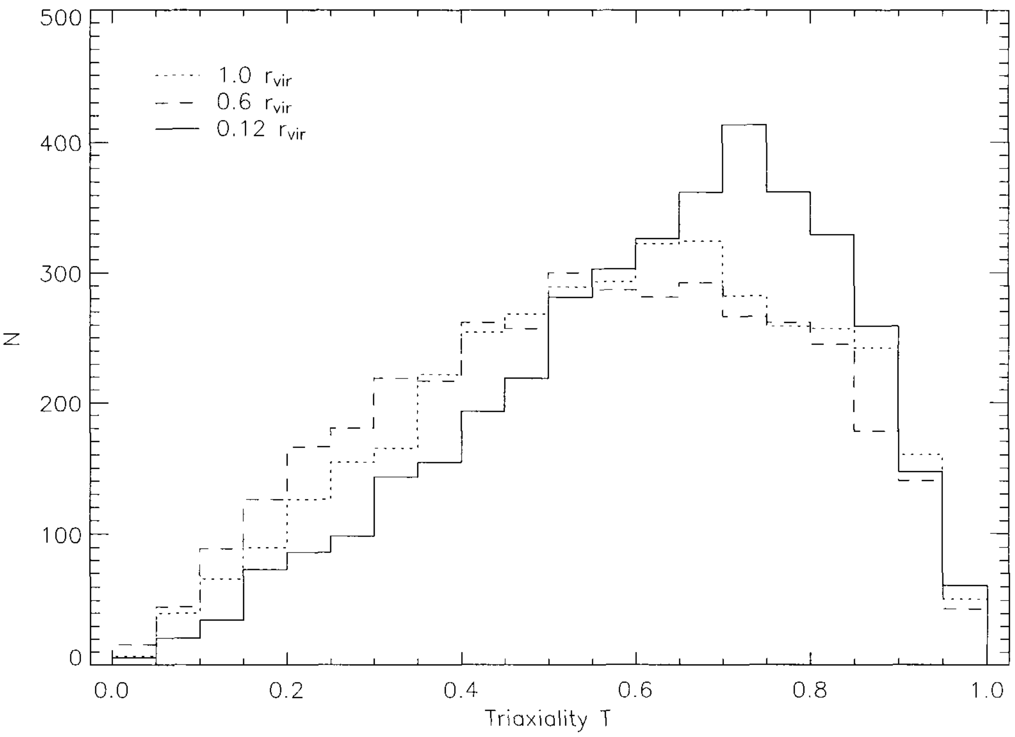

Figure 5.4: Histogram of the triaxiality T of all halos. The solid, dashed, and dotted histograms represent the halos measured at 0.12, 0.6, and 1.0 $r_{\rm vir}$ respectively.

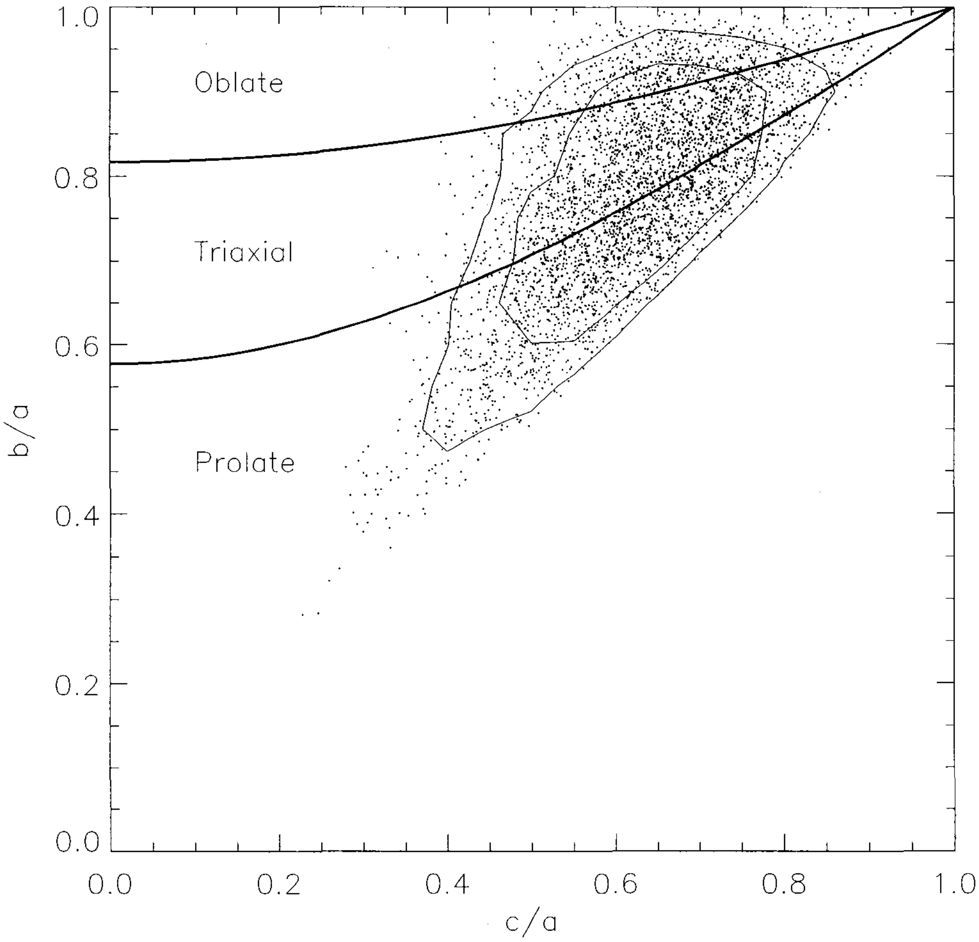

Figure 5.5: Intermediate axis ratio b/a as a function of minor axis ratio c/a for all of the halos, measured at $R = 0.4\, r_{\mathrm{vir}}$. The inner and outer contours enclose 68% and 90% of the halos respectively. The thick lines have constant values of the triaxiality parameter T. The separations between the prolate, triaxial, and oblate populations occur at $T = 1/3$ and $2/3$.

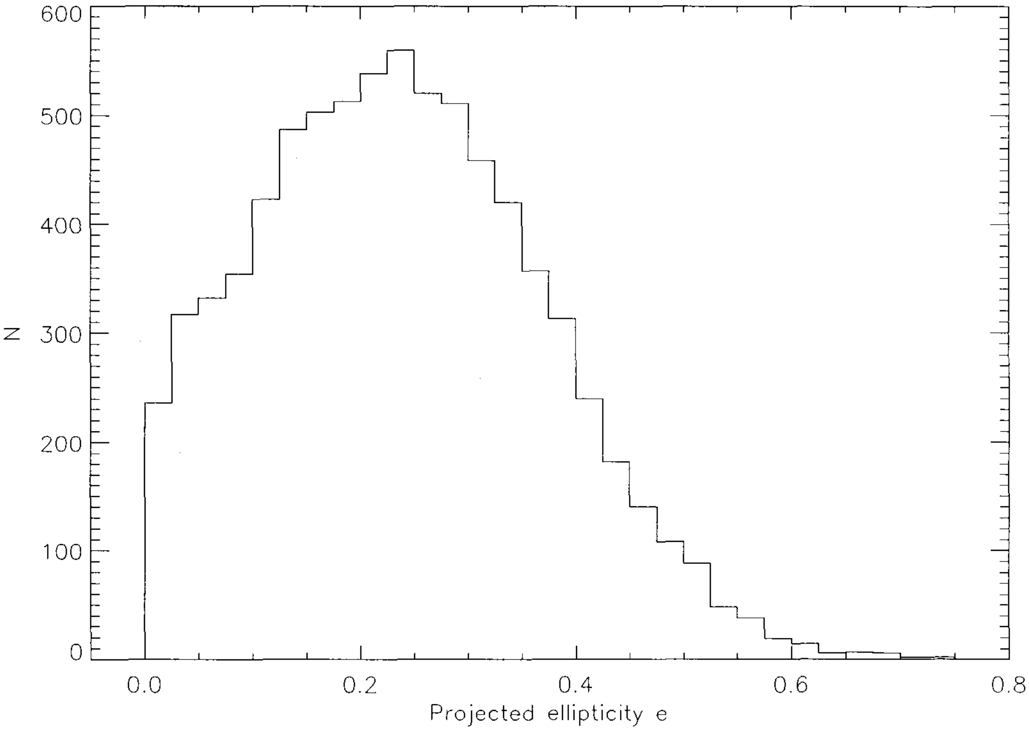

Figure 5.6: Projected ellipticity e of each halo seen from two random orientations, measured at $R = 0.4\, r_{\rm vir}$.

Weak lensing measurements have recently begun to probe the two-dimensional projected ellipticity of the lensing mass distribution (Hoekstra et al., 2004). For each halo in our sample, we have calculated the projected axis ratio q using the method of Stark (1977) for two random orientations, and calculated the ellipticity e using equation (5.1). A histogram of the results for the $R = 0.4\ r_{\rm vir}$ shell is shown in Figure 5.6. The mean and median of the distribution of ellipticities are $e_{\rm mean} = 0.24$ and $e_{\rm median} = 0.23$ respectively, with a 1σ width of 0.13. This is consistent with the lower limit of $\langle e \rangle = 0.33^{+0.07}_{-0.09}$ found by Hoekstra et al. (2004) from stacked weak lensing measurements around galaxies. It is smaller than the ellipticities found by Croft & Metzler (2000) (note that they quantify ellipticity as $e = (1 - q^2)/(1 + q^2)$, in which units our mean ellipticity is $e = 0.27$). However, most of the halos used in their study have higher masses than the galactic-mass halos studied here, and they included relatively poorly-resolved halos which may skew the results toward higher ellipticities.

5.4 Internal alignment

5.4.1 Principal axes

We compare the alignment of the principal axes within each halo to see whether the approximation of the halo as a set of concentric ellipsoids is justified (JS02). In order to determine whether the axes are aligned, the directions of the axes must be well determined, otherwise we may claim misalignments that are due to measurement error. Therefore, we restrict ourselves in this section to axes whose bootstrap error is less than 0.2 radians. The number of halos satisfying this criterion at each radius for each axis is given in Table 5.1.

To understand the effect of the error on the determination of the alignment, imagine two axes which are intrinsically perfectly aligned, but are each measured

Table 5.1. Number of halos with axes determined to within 0.2 radians

Radius ($r_{\rm vir}$)	Major axis	Intermediate axis	Minor axis
0.06	904	388	621
0.12	2370	1005	1628
0.25	2942	1762	2540
0.4	2545	1488	2421
0.6	2322	1363	2324
1.0	2973	2231	2944

with an error of 0.2 radians (note that this is the worst possible case — the median error of the sample is 0.1 radians). Due to the measurement error, these axes will appear to be misaligned by an angle $\theta_{\rm spurious}$. The component of an isotropic error in any particular plane, such as the plane containing both of the measured axes, is half of the isotropic error, so we divide the isotropic error of 0.2 radians by two and add the error of each axis in quadrature to find the typical $\theta_{\rm spurious} \approx 0.14$. Therefore, the cosine of the angle between the two axes, which is intrinsically 1.0, is measured to be $\cos\theta_{\rm spurious} = 0.99$. If the axes are intrinsically perpendicular, in which case the effect of $\theta_{\rm spurious}$ on the direction cosine is maximized, the error in the direction cosine is 0.14. Most halos have well aligned axes (see below), so the error is negligible.

Figure 5.7 shows histograms of the alignment between the major axis of the

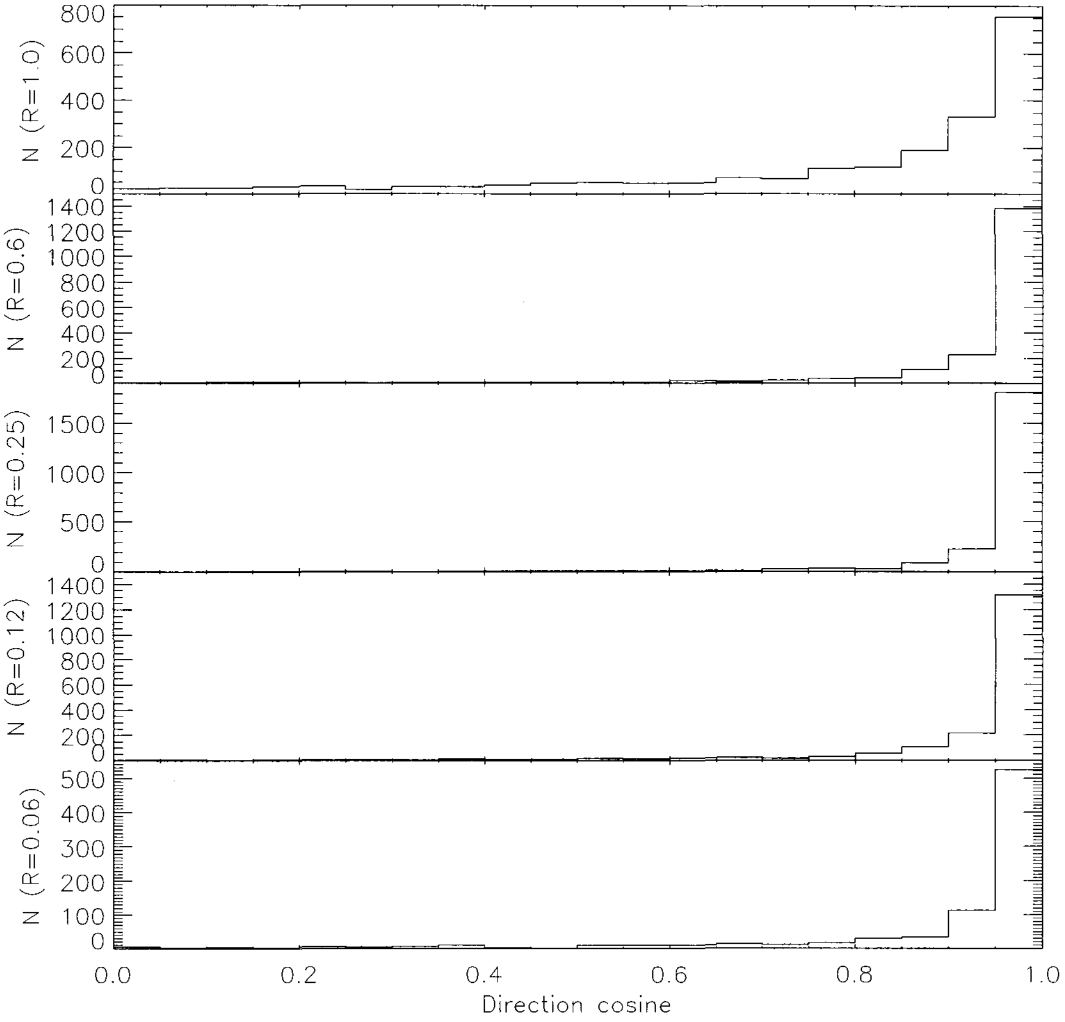

Figure 5.7: Histograms of the direction cosine between the major axis of the halo at $R = 0.4\ r_{\rm vir}$ and the major axis at $R = 1.0$, 0.6, 0.25, 0.12, and 0.06 $r_{\rm vir}$ (top to bottom). Due to the symmetry of the axes, this is always positive. Each histogram contains all halos where the major axes at both radii are determined to within 0.2 radians. If the axes were isotropic, this distribution would be uniform.

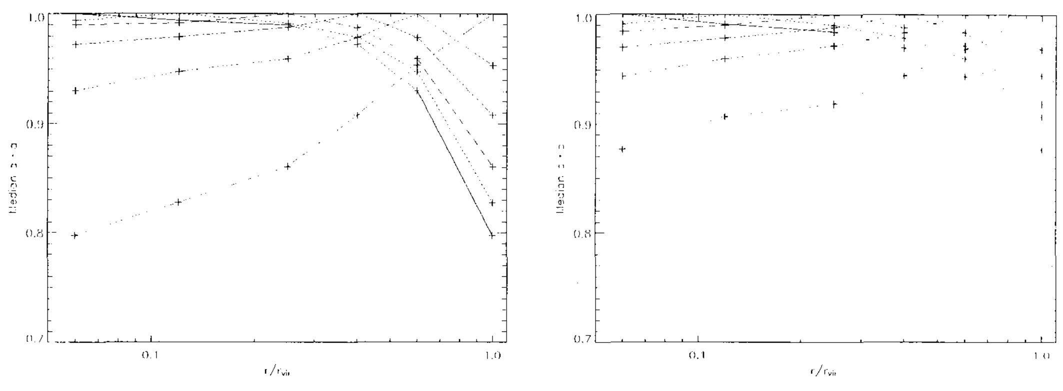

Figure 5.8: Median alignment of the major *(a-left)* and minor *(b-right)* axes at different radii. The alignment is with respect to the R = 0.06 (black/solid), R = 0.12 (red/dotted), R = 0.25 (green/short-dashed), R = 0.4 (cyan/dot-dashed), R = 0.6 (blue/dot-dot-dot-dashed), and R = 1.0 $r_{\rm vir}$ (magenta/long-dashed) shell. For each pair of radii, only halos where the direction of the axis is determined to within 0.2 radians at both radii are used.

$R = 0.4\ r_{\rm vir}$ shell and the major axis of the outer (top two panels) and inner (bottom three panels) regions of the halo. The alignment is very good at all radii. The relative alignment of the major and minor axes as a function of radius is shown in Figure 5.8. Each line shows the median alignment with respect to a different fiducial radius, recognizable as the radius where the median is exactly unity. The axes are extremely well aligned within 0.6 $r_{\rm vir}$. Near the virial radius, some halos show deviations, but there is still usually very good alignment, especially for the minor axis. The alignment is better than that seen by JS02. This confirms their suggestion that many of the halos in which they measured poor alignment were nearly prolate or oblate. Such halos have large errors in their direction determination, and so are not included in our sample.

Figure 5.9 examines the internal alignment of the major axis as a function of

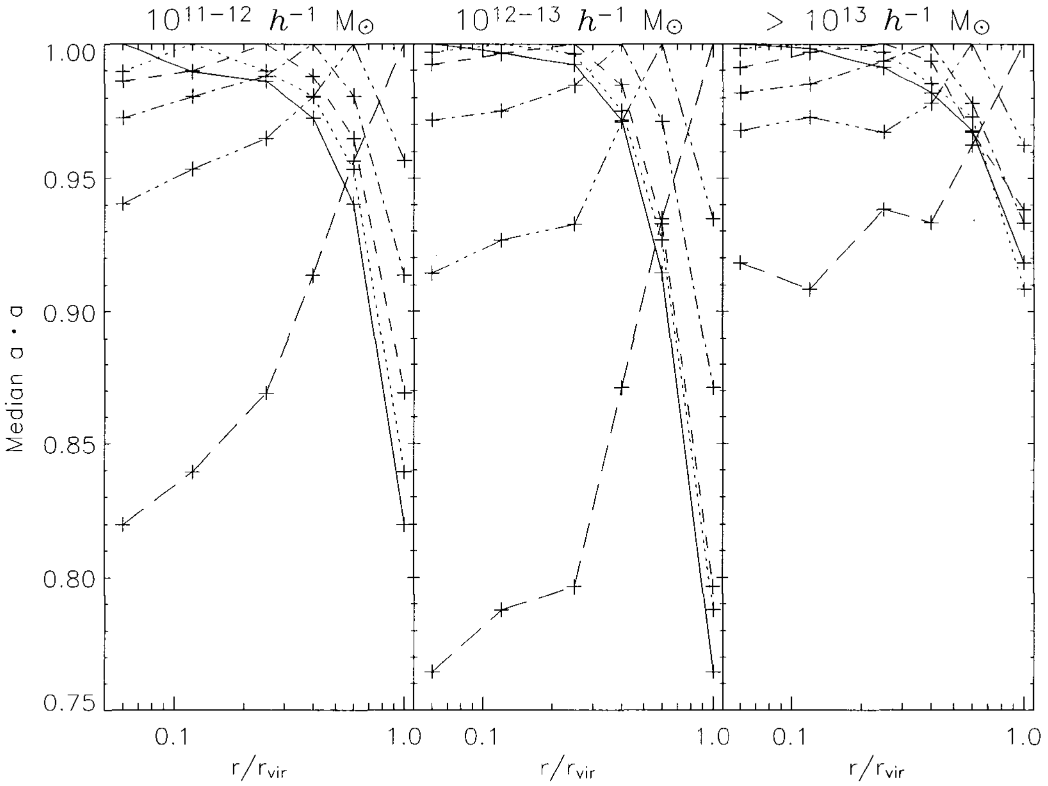

Figure 5.9: Median alignment of the major axis, as in Figure 5.8a, but only for halos with masses of 10^{11} – 10^{12} h^{-1} $\mathrm{M}_\odot$ *(left)*, 10^{12} – 10^{13} h^{-1} $\mathrm{M}_\odot$ *(middle)*, and 10^{13} – 3×10^{14} h^{-1} $\mathrm{M}_\odot$ *(right)*. The alignment is with respect to the $R = 0.06$ (black/solid), $R = 0.12$ (red/dotted), $R = 0.25$ (green/short-dashed), $R = 0.4$ (cyan/dot-dashed), $R = 0.6$ (blue/dot-dot-dot-dashed), and $R = 1.0\ r_{\mathrm{vir}}$ (magenta/long-dashed) shell. For each pair of radii, only halos where the direction of the axis is determined to within 0.2 radians at both radii are used.

Table 5.2. Number of halos with angular momentum directions determined to within 0.4 radians

Radius	Halos
0.06	702
0.12	1686
0.25	2820
0.4	3060
0.6	3094
1.0	3229

halo mass. The inner 0.4 $r_{\rm vir}$ of the halos are equally well aligned for halos of all masses. However, the outer half of the halo is better aligned with the rest of the halo in high mass halos than in low mass halos. JS02 saw a similar effect and suggested that it was because the low mass halos are intrinsically rounder and therefore have larger errors. We rule out this explanation, as halos with large errors are not included in our sample for any mass. Therefore, the better alignment within high mass halos appears to be a real physical effect.

5.4.2 Angular momentum

The direction of the angular momentum cannot be determined as precisely as the direction of the principal axes. In order to have a reasonably large sample of halos, we use angular momentum vectors whose bootstrap error is less than 0.4 radians. This is twice as large as the limit adopted for the principal axis directions. The number of halos that satisfy this criterion at each radius is given in

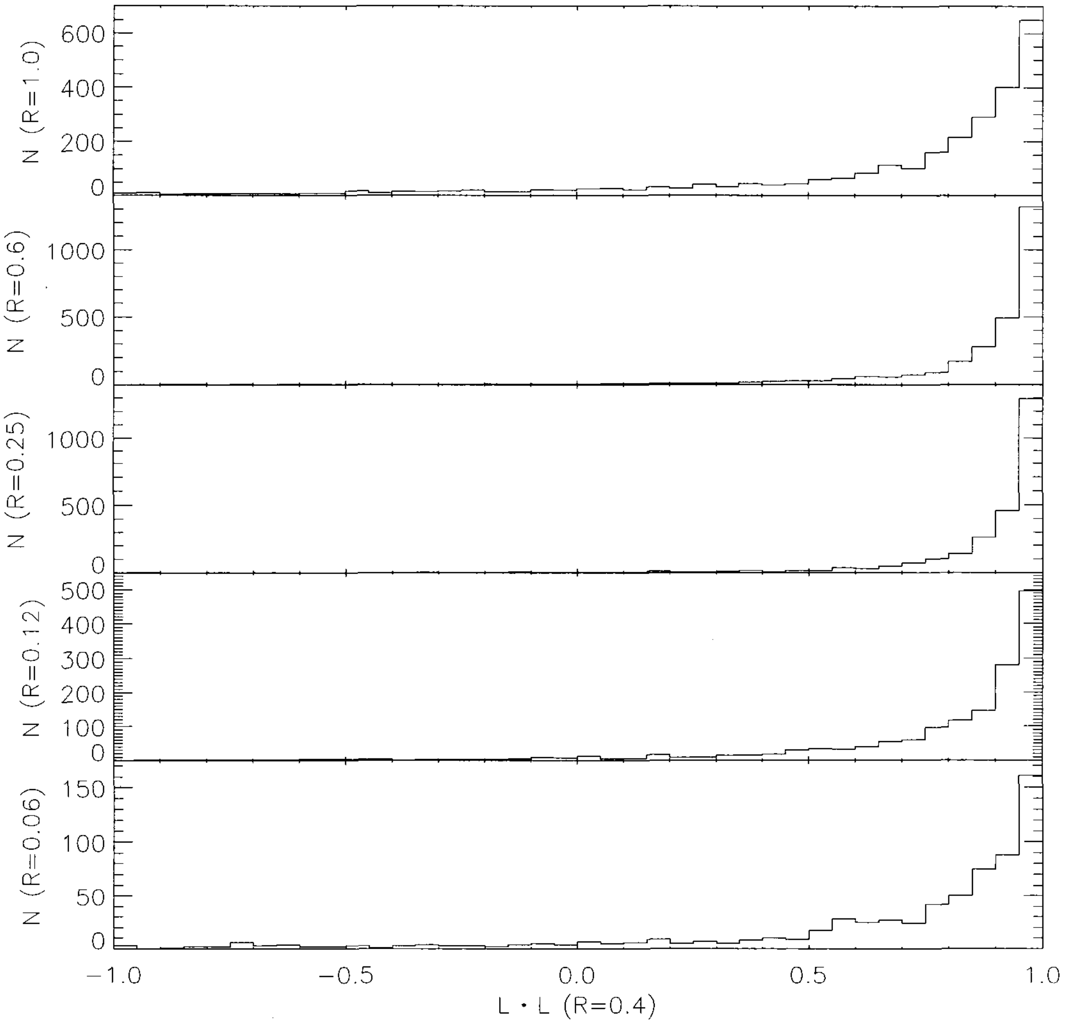

Figure 5.10: Histograms of the direction cosine between the angular momentum of the halo at $R = 0.4\ r_{\rm vir}$ and the angular momentum at $R = 1.0$, 0.6, 0.25, 0.12, and 0.06 $r_{\rm vir}$ (top to bottom). Each histogram contains halos whose angular momentum direction is determined to within 0.4 radians at both radii of the comparison. If the orientations were isotropic, this distribution would be uniform.

Table 5.2. Following the logic of § 5.4.1, the error in the direction cosine of two vectors with errors of 0.4 radians is 0.04 if they are perfectly aligned and 0.28 if they are perpendicular. The median errors of the samples are half of these worst-case scenarios. The alignment is shown to be good (see below), so the effect of the errors is negligible.

Figure 5.10 shows histograms of the relative alignment of the angular momentum of the $R = 0.4\, r_{\rm vir}$ shell with the outer regions of the halo (upper two panels) and the inner regions of the halo (lower three panels). The alignment is very good at most radii.

The relative alignment of the angular momentum as a function of radius is shown in Figure 5.11. Each line shows the median alignment of the angular momentum with respect to a different fiducial radius, recognizable as the radius where the median is exactly unity. The alignment gets progressively worse as the radii get further separated; the median cosine between the angular momenta in the innermost and outermost regions is 0.64. However, the angular momentum vector at intermediate radius, such as at 0.4 $r_{\rm vir}$, is generally representative of its direction at all radii.

Figure 5.12 examines the internal alignment of the angular momentum as a function of halo mass. The patterns seen in Figure 5.11 generally hold for all masses, although the alignment between the very innermost and outermost regions is slightly worse for the highest mass halos.

5.4.3 Alignment between the angular momentum and the halo shape

We investigate here the alignment between the angular momentum of a halo at a given radius and the principal axes of the mass distribution at that radius. If we compare an angular momentum vector whose error is 0.4 radians with a principal axis whose error is 0.2 radians, the error in the direction cosine is 0.02 if they are

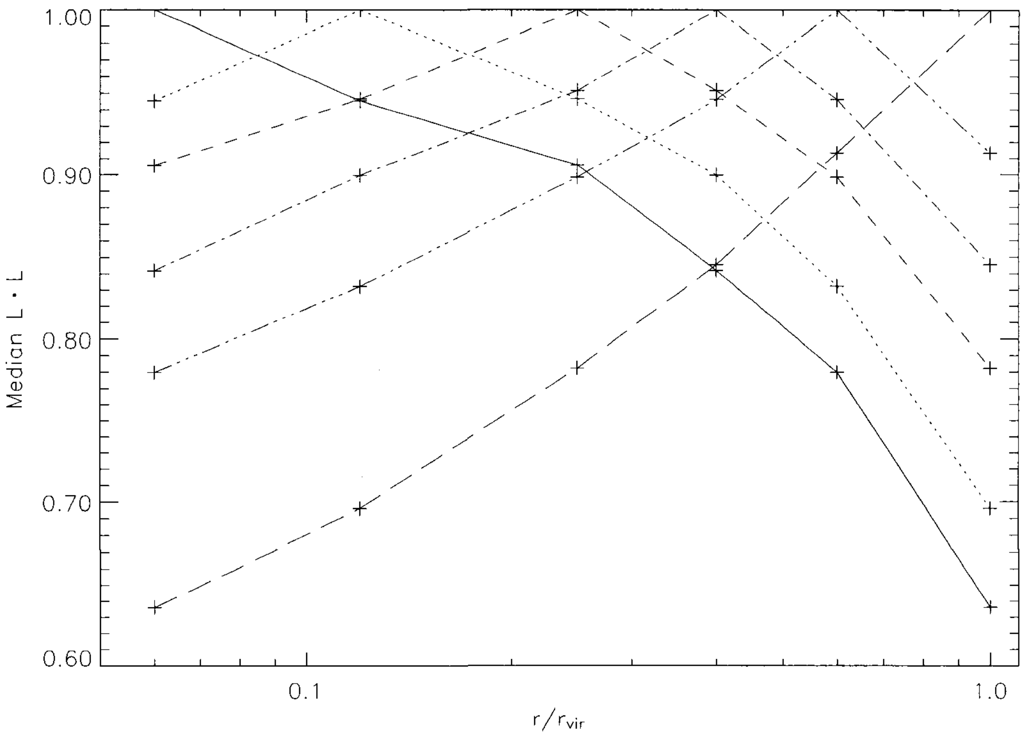

Figure 5.11: Median alignment of the angular momentum vector at different radii. The alignment is with respect to the R = 0.06 (black/solid), R = 0.12 (red/dotted), R = 0.25 (green/short-dashed), R = 0.4 (cyan/dot-dashed), R = 0.6 (blue/dot-dot-dot-dashed), and R = 1.0 $r_{\rm vir}$ (magenta/long-dashed) shell. For each pair of radii, only halos where the direction of the angular momentum vector is determined to within 0.4 radians at both radii are used.

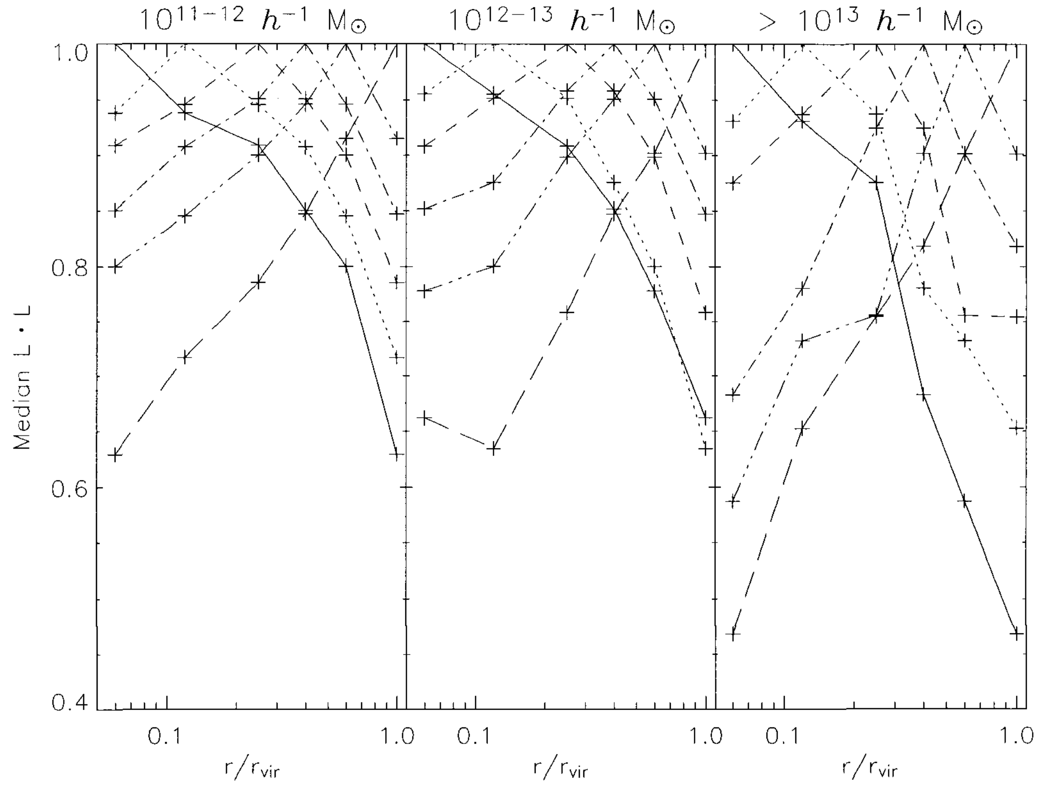

Figure 5.12: Median angular momentum alignment as in Figure 5.11, but only for halos with masses of 10^{11} – 10^{12} h^{-1} M$_\odot$ *(left)*, 10^{12} – 10^{13} h^{-1} M$_\odot$ *(middle)*, and 10^{13} – 3×10^{14} h^{-1} M$_\odot$ *(right)*. The alignment is with respect to the $R = 0.06$ (black/solid), $R = 0.12$ (red/dotted), $R = 0.25$ (green/short-dashed), $R = 0.4$ (cyan/dot-dashed), $R = 0.6$ (blue/dot-dot-dot-dashed), and $R = 1.0$ $r_{\rm vir}$ (magenta/long-dashed) shell. For each pair of radii, only halos where the direction of the angular momentum vector is determined to within 0.4 radians at both radii are used.

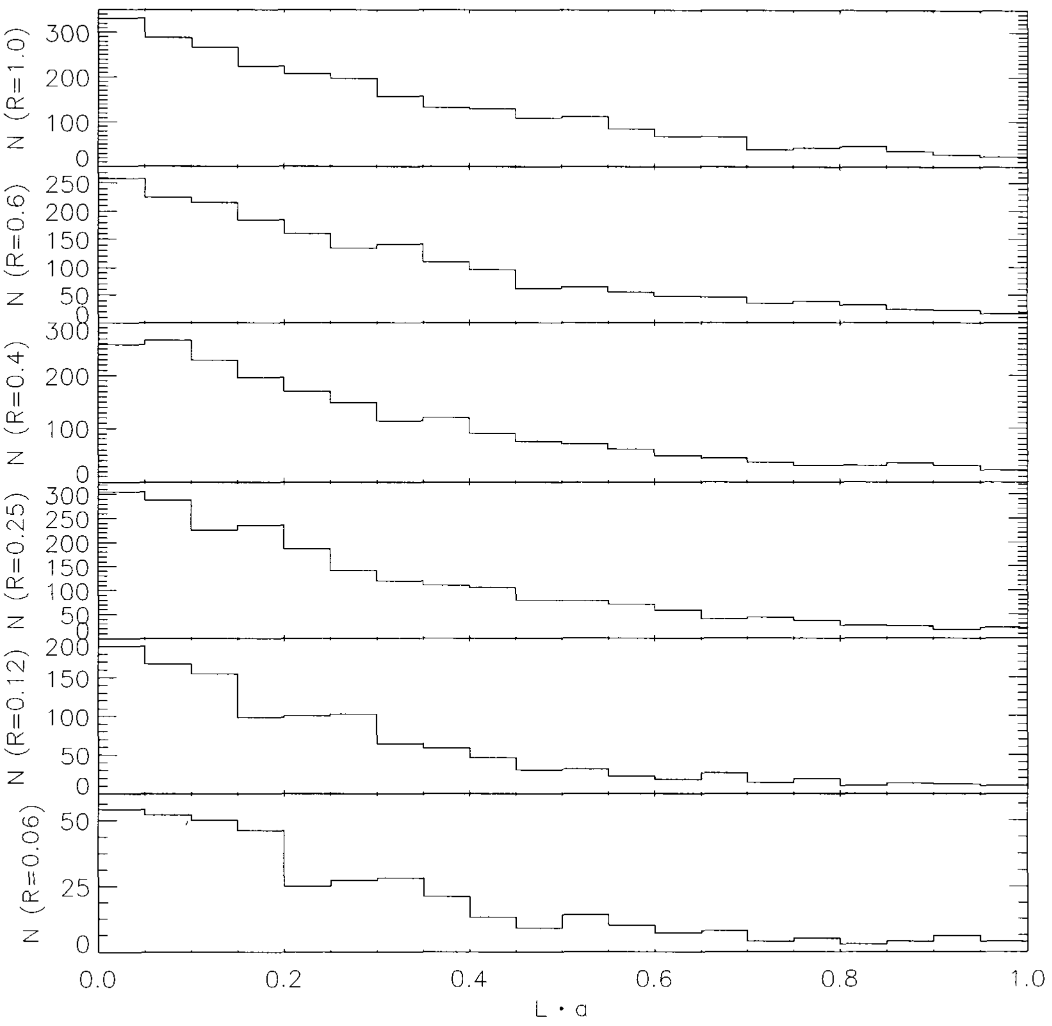

Figure 5.13: Histograms of the alignment between the angular momentum vector and the major axis of the $R = 1.0$, 0.6, 0.4, 0.25, 0.12, and 0.06 $r_{\rm vir}$ (top to bottom) shell of each halo where the error in the major axis direction is less than 0.2 radians and the error in the angular momentum direction is less than 0.4 radians at that radius. Due to the symmetry of the axes, the direction cosine is always positive. If the orientations were random, this distribution would be uniform.

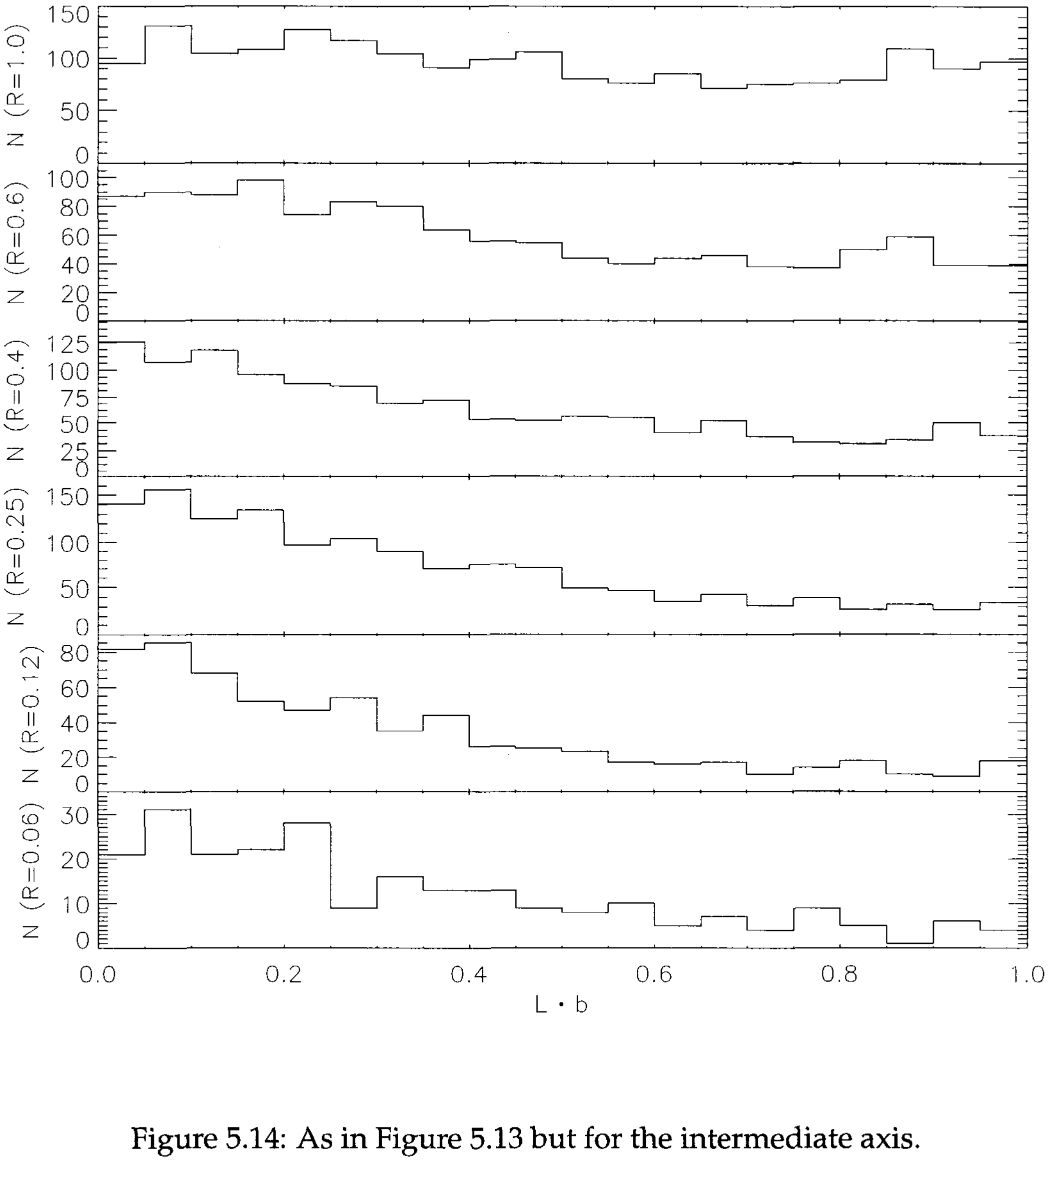

Figure 5.14: As in Figure 5.13 but for the intermediate axis.

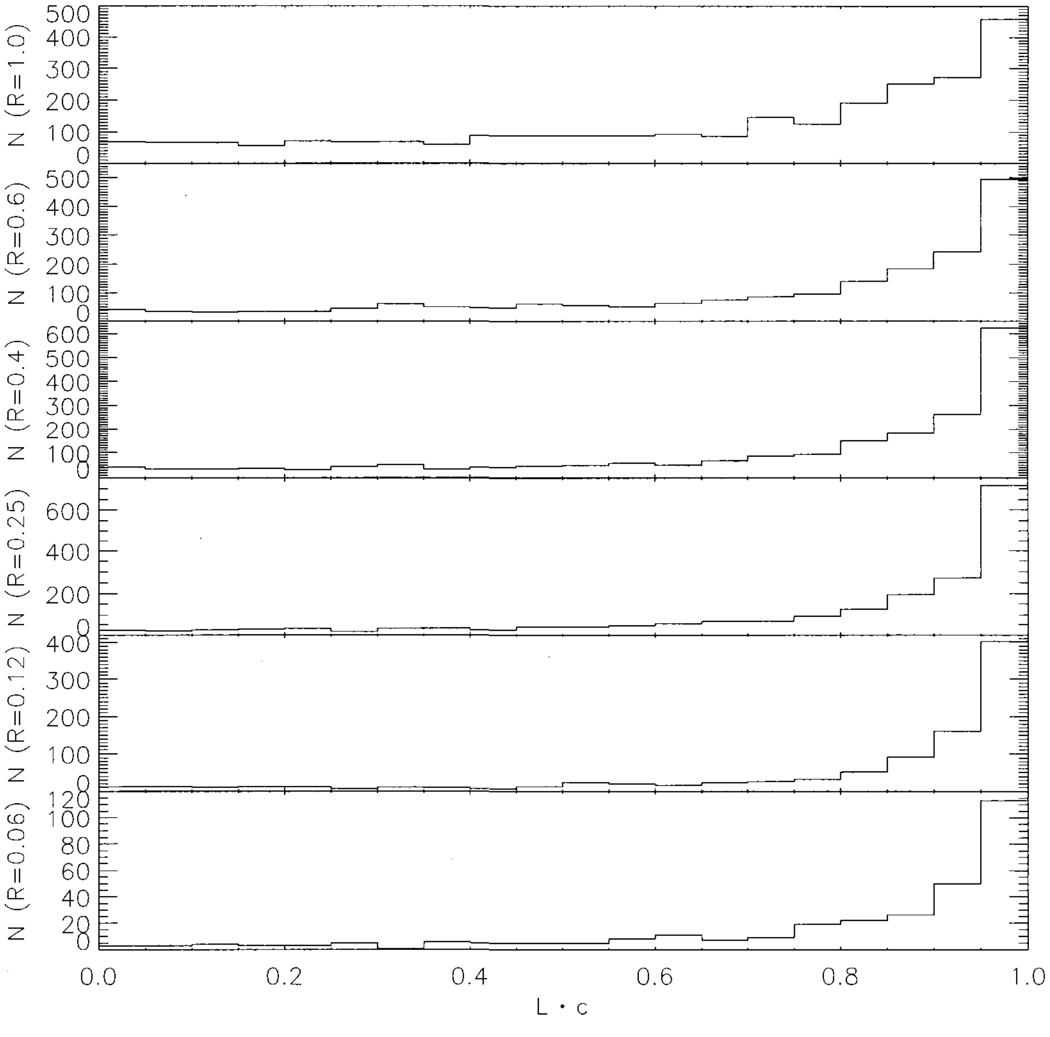

Figure 5.15: As in Figures 5.13 and 5.14 but for the minor axis.

perfectly aligned and 0.22 if they are perpendicular. The median errors are half of those values. Therefore, in the cases where the alignment is good, the effect of the errors is negligible. In the cases where the vectors are perpendicular or isotropic, we must take the errors into account when drawing conclusions.

Figures 5.13, 5.14, and 5.15 show histograms of the cosine between the direction of the angular momentum vector and the major, intermediate, and minor axis respectively. Only those halos with both angular momentum direction errors of less than 0.4 radians and axis direction errors of less than 0.2 radians are used. The angular momentum vector tends to be perpendicular to the major and intermediate axes, and parallel with the minor axis. Because of the different effects of the error on parallel and perpendicular vectors, the tendency of the angular momentum to be perpendicular to the major axis is as significant as the trend for it to be parallel to the minor axis, despite the different appearance of the histograms. The angular momentum tends to lie perpendicular to the intermediate axis, but this trend gets weaker with radius.

These results are consistent with those of Barnes & Efstathiou (1987), Dubinski (1992), and W92. Of these, only W92 quantify any change with radius; they found slightly better alignment at larger radii, in contrast to the results presented here. However, both of the radii at which they performed the comparison were well within the virial radius, well inside the radii where we see the alignment drop.

These relationships are summarized in Figure 5.16, which shows the median alignment between the angular momentum vector and each of the principal axes. The angular momentum tends to lie parallel with the minor axis and perpendicular to both the major and intermediate axes. These trends are strongest in the central 0.25 $r_{\rm vir}$ of the halo, deteriorating slightly in the outer regions. The median cosine of 0.9 corresponds to a misalignment between the angular momentum and

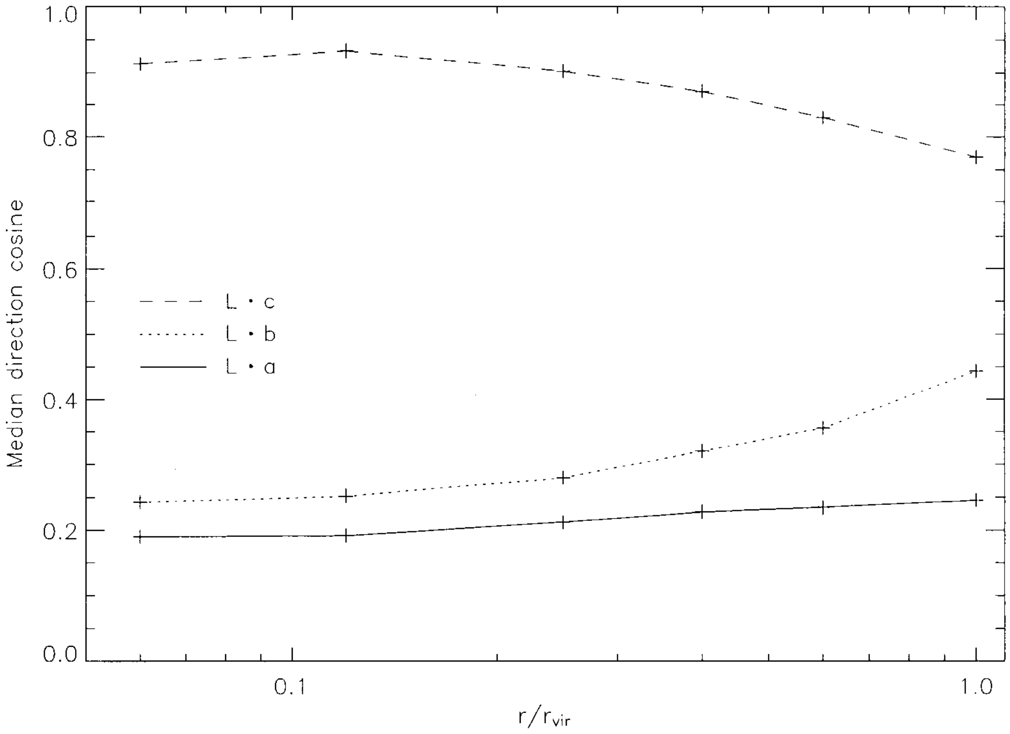

Figure 5.16: Median alignment between the angular momentum vector and the major (black/solid), intermediate (red/dotted), and minor (blue/dashed) axis of each halo as a function of radius within the halo.

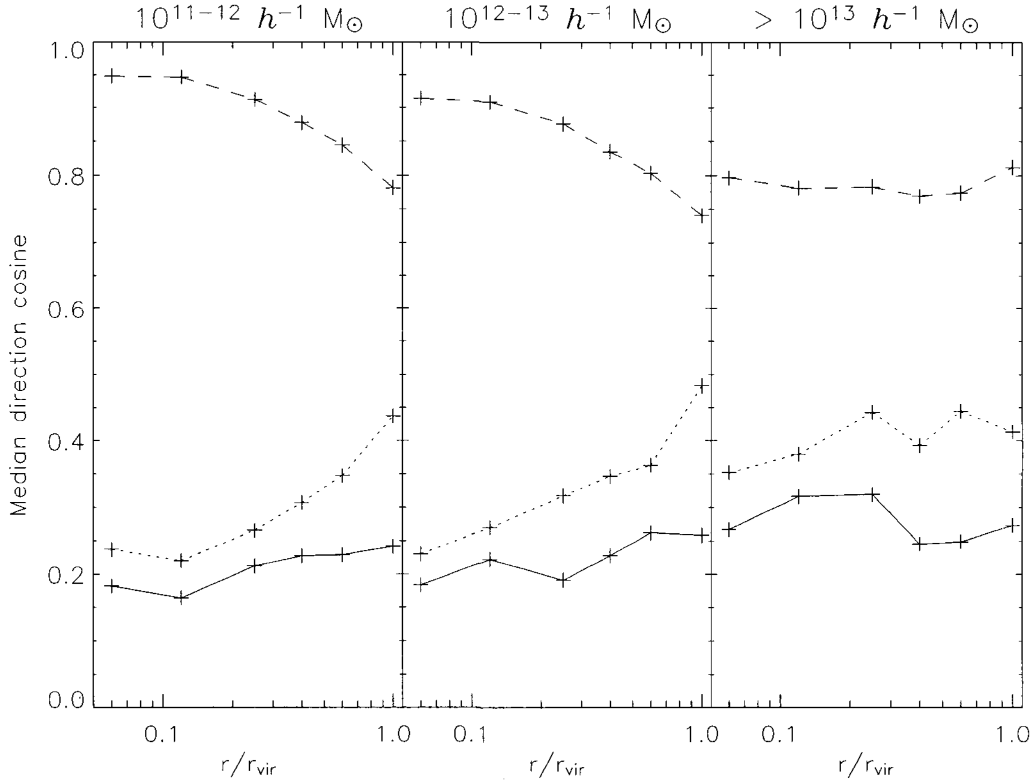

Figure 5.17: As in Figure 5.16, but only for halos with masses of $10^{11} - 10^{12}\ h^{-1}\ M_\odot$ *(left)*, $10^{12} - 10^{13}\ h^{-1}\ M_\odot$ *(middle)*, and $10^{13} - 3 \times 10^{14}\ h^{-1}\ M_\odot$ *(right)*.

minor axis of $25°$, comparable to the misalignment between the disk and halo potential we investigated in § 2.2.2. Figure 5.17 shows how these trends depend on the halo mass. For high mass halos, the alignment is slightly worse and has less of a dependence on the location within the halo.

5.5 External alignment

5.5.1 Introduction

In this section, we compare the orientation of the principal axes and angular momenta of individual halos with the location of mass around them and the orientation of those properties in surrounding halos. For each halo, the volume is split into 7 radial bins. The nearest bin spans separations from 0 to 390.625 h^{-1} kpc. The outer radii double for each subsequent bin, while the inner radius is equal to the outer radius of the interior bin. The largest bin has an outer radius of 25 h^{-1} Mpc, extending to the edge of the periodic box. The properties of the halos in the $R = 0.4\ r_{\rm vir}$ shell are characteristic of their overall properties over a large range of radii (see § 5.4). Therefore, we use the principal axes and angular momenta measured at this radius for all halos in this section.

We follow the nomenclature of Splinter et al. (1997), and use "alignment" to refer to the tendency of a vector (such as a principal axis or an angular momentum vector) to point toward or away from other halos, and "correlation" to refer to the tendency of vectors in different halos to point in the same direction.

A visual impression of how the axes and angular momenta align can be seen in Figure 5.18, which shows the axes and angular momentum vectors of all of the halos in a slab containing one quarter of the simulation. If the axis or vector lies mostly in the z direction, it appears much shorter in the xy projection shown. Figure 5.19 shows the properties of only the high mass halos, $M > 10^{13}\ h^{-1}\ {\rm M}_{\odot}$,

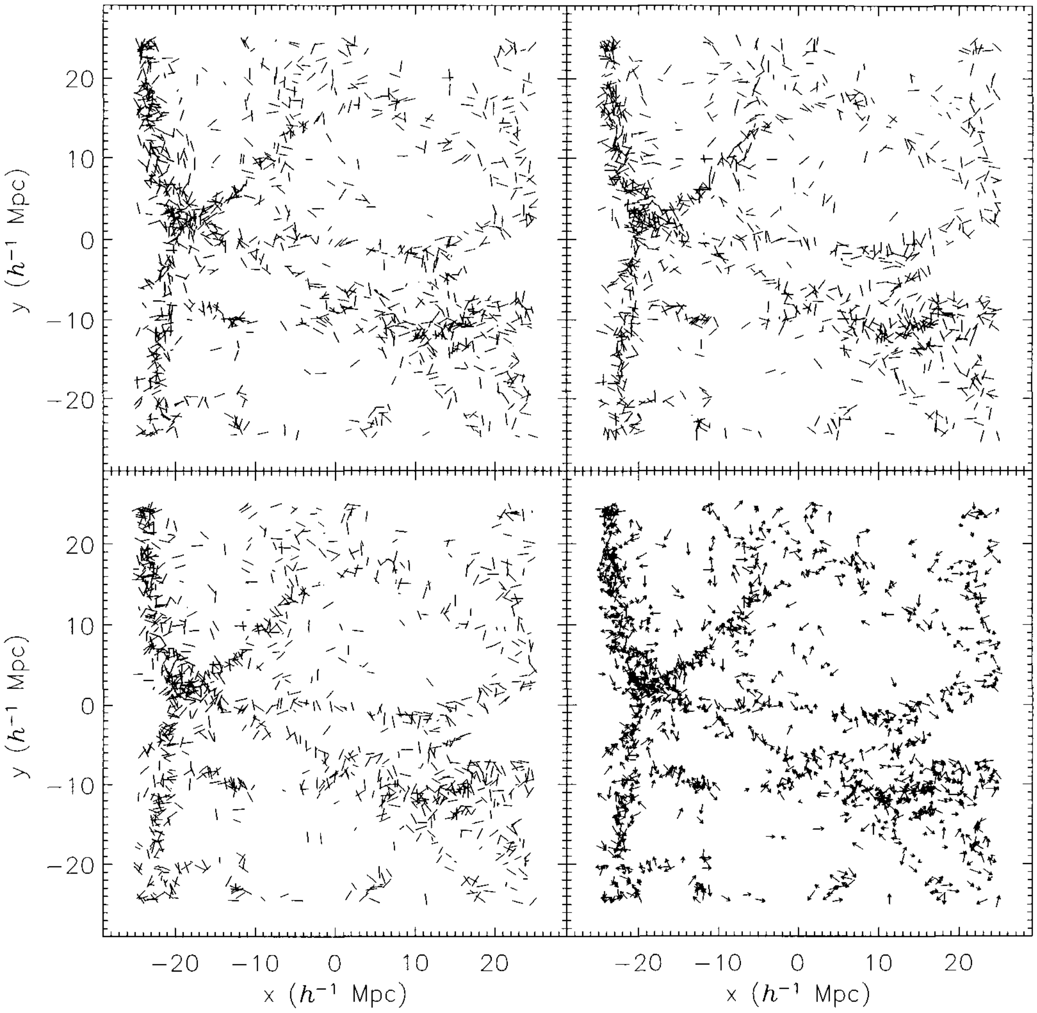

Figure 5.18: Projection in xy of the major *(top-left)*, intermediate *(top-right)*, and minor *(bottom-left)* axes, and unit angular momentum vectors *(bottom-right)* of halos in a slab of depth $\Delta z = 12.5\ h^{-1}$ Mpc (one quarter of the simulation volume). The three-dimensional length of each line or vector is 1.5 h^{-1} Mpc, including the unseen z component.

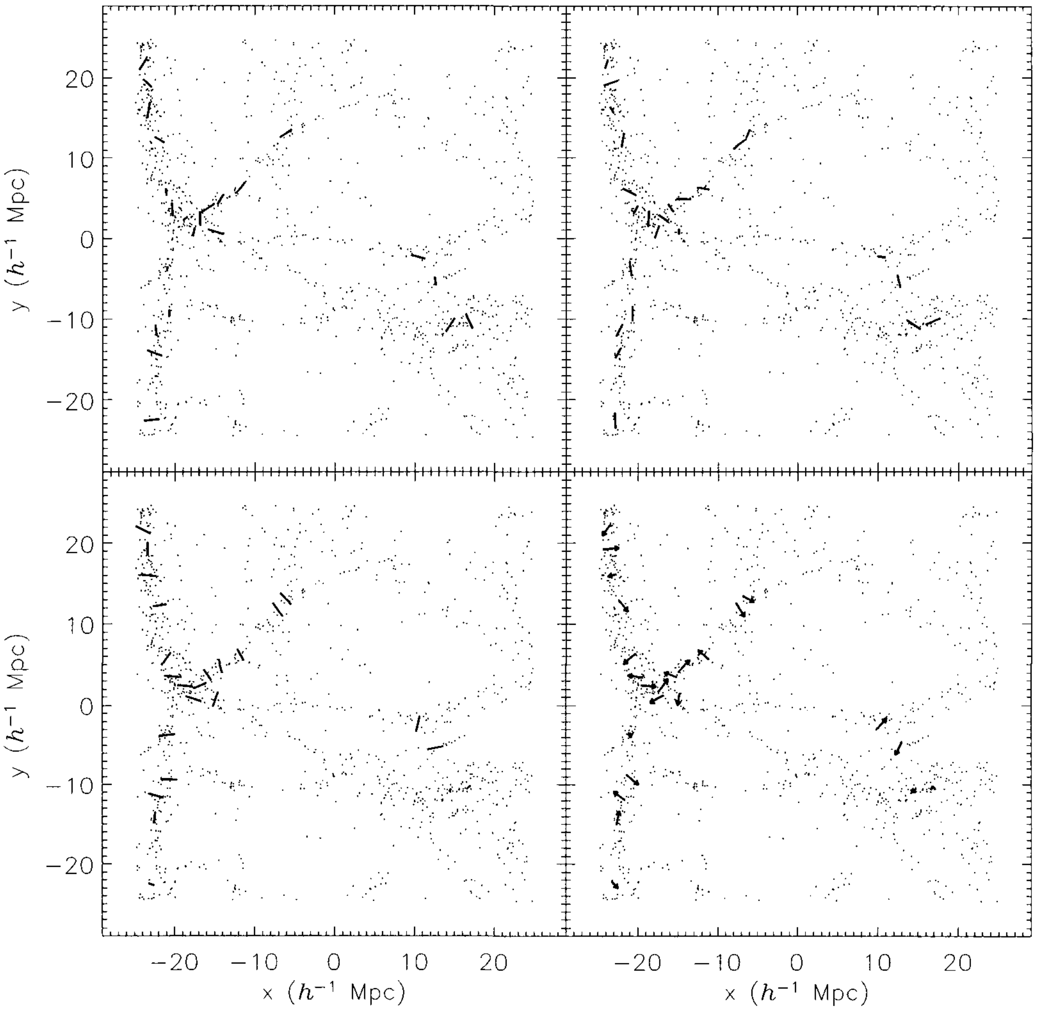

Figure 5.19: As in Figure 5.18, but for only those halos with masses greater than $10^{13}\ h^{-1}\ M_{\odot}$. The three-dimensional length of each line or vector is $2\ h^{-1}$ Mpc, including the unseen z component. Dots show the positions of all halos, regardless of mass.

with the location of the other halos shown as dots. Because of the filamentary nature of the large scale structure (e.g. Colberg et al., 1999), positive alignment indicates that a quantity tends to point along filaments.

5.5.2 Axis alignments

We compare here the alignment of the principal axes of the halos with the location of surrounding structure. Methods of measuring these alignments vary in the literature, as does the nomenclature for a given metric. We adopt an internally-consistent nomenclature ξ_{xy} and $\xi_{|xy|}$, which are defined to be the mean value of the direction cosine between direction x and direction y, and the mean of the absolute value of the direction cosine respectively. We note names other authors use for the same quantities when applicable. For example, to measure the alignment of the major axis, whose direction is defined by the unit vector $\hat{\mathbf{a}}$, with the large scale structure, we calculate the alignment $\xi_{|ar|}$:

$$\xi_{|ar|}(r) \equiv \langle |\hat{\mathbf{a}} \cdot \hat{\mathbf{r}}| \rangle \equiv \frac{1}{N} \sum_{i,j} |\hat{\mathbf{a}}_i \cdot \hat{\mathbf{r}}_{ij}|, \tag{5.6}$$

where the sum over i is over all halos in the primary sample, the sum over j is over all halos in the secondary sample, $\hat{\mathbf{r}}_{ij}$ is a unit vector in the direction of the displacement from halo i to halo j, and N is the number of terms in the double sum. The primary sample consists of all halos whose major axis is determined to within 0.2 radians, while the secondary sample consists of all halos. We define the alignments $\xi_{|br|}$ and $\xi_{|cr|}$ for the intermediate axis $\hat{\mathbf{b}}$ and the minor axis $\hat{\mathbf{c}}$ similarly. Note that the primary samples used to define $\xi_{|ar|}$, $\xi_{|br|}$, and $\xi_{|cr|}$ are not identical, as the set of halos with good major axis determinations is deficient in very oblate halos, the set of halos with good minor axis determinations is deficient in very prolate halos, and the set of halos with good intermediate axis determinations is deficient in both very prolate and very oblate halos.

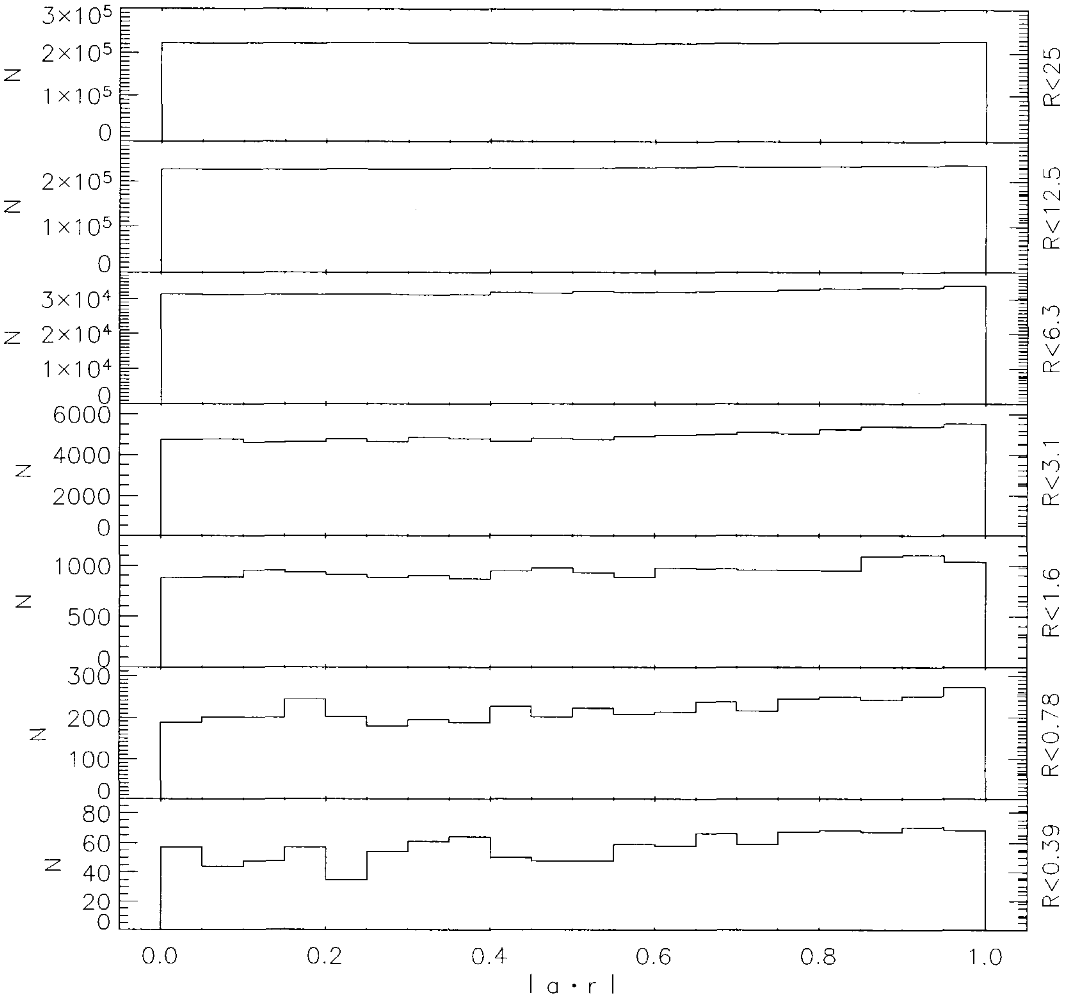

Figure 5.20: Histograms of the direction cosine $|\hat{a} \cdot \hat{r}|$ between the major axes of halos in the primary sample and the displacement from each primary halo to all surrounding halos, binned by radial separation between the centers of the halos. The direction cosine is always positive due to the symmetry of the axes. If the axes were randomly oriented, the distributions would be uniform. The radial bins consist of halos separated by (top to bottom): 12.5 – 25, 6.25 – 12.5, 3.125 – 6.25, 1.5625 – 3.125, 0.78125 – 1.5625, 0.390625 – 0.78125, and 0 – 0.390625 h^{-1} Mpc respectively. The primary sample consists of halos whose major axes are determined to within 0.2 radians.

Figure 5.20 shows histograms of the distribution of direction cosines $|\hat{\mathbf{a}} \cdot \hat{\mathbf{r}}|$ for halos at a variety of separations. If the axes were randomly oriented, the distributions would be uniform. If $|\hat{\mathbf{a}} \cdot \hat{\mathbf{r}}| > 0.5$ then the axes tend to be parallel to the lines connecting halos, while if $|\hat{\mathbf{a}} \cdot \hat{\mathbf{r}}| < 0.5$ then the axes tend to point perpendicular to the lines connecting halos. The distribution is mostly isotropic, but there is an excess of halos with $|\hat{\mathbf{a}} \cdot \hat{\mathbf{r}}| > 0.5$. We quantify this by calculating the mean, $\xi_{|ar|}$ (note that this is similar but not identical to the quantity $w(r)$ used by KE04).

Figure 5.21 shows $\xi_{|ar|}$ for the halos. Each bin is plotted at the geometric mean between its inner and outer radius, except for the innermost bin which is plotted at $1/\sqrt{2}$ times its outer radius. The inner bin has no formal inner radius, but in practice is limited by twice the radial extent of the typical halo, or 250 h^{-1} kpc. Halos whose centers of mass are closer to each other than this are merging, and are detected as a single object by the group finder. We find that the alignment is well fit by a power law over a wide range of separations. The fit is shown as the solid line in Figure 5.21, and is given by

$$\xi_{|ar|}(r) = \frac{1}{2} + m_1 r^{\alpha}, \tag{5.7}$$

where r is the separation in units of h^{-1} Mpc, the alignment at 1 h^{-1} Mpc is $m_1 = 0.015$, and the slope is $\alpha = -0.64$. F02 find even stronger alignment for their cluster-mass halos, as do KE04 (except at the very smallest separations, which is at the spatial limit of their simulations). To see whether the halo mass is important, we recalculate $\xi_{|ar|}$, restricting the primary sample to those halos with masses greater than 10^{13} h^{-1} $M_{\odot}$. This is the tendency for the major axes of group or cluster mass halos to point along the filaments. The results are shown as the diamonds in Figure 5.21. The group and cluster mass halos are much more strongly aligned than the full sample (which is dominated by galaxy mass ha-

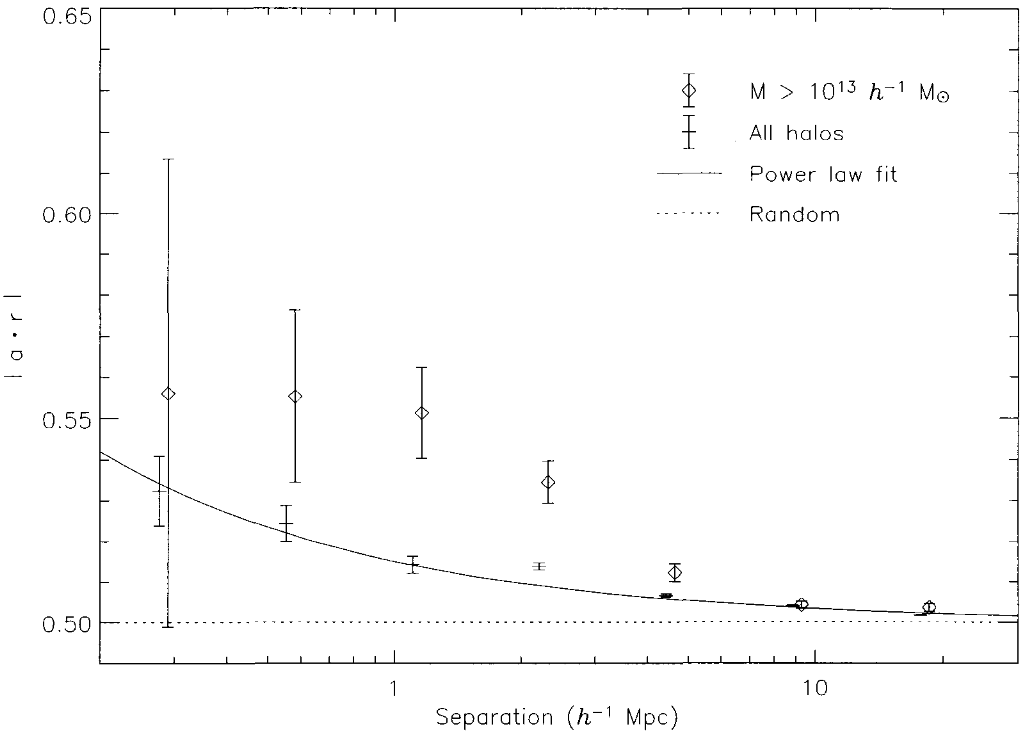

Figure 5.21: Mean alignment $\xi_{|ar|}$ between the major axes of halos in the primary sample and the displacement from the primary halos to all surrounding halos. This is the tendency for the major axes to point along filaments. Only halos where the major axis is determined to within 0.2 radians are used for the primary sample. The different symbols are for primary samples consisting of all such halos (crosses) or of only those with masses greater than 10^{13} h^{-1} $M_\odot$ (diamonds). The diamonds are shifted slightly to the right for clarity. The values are binned by radial separation between the centers of the halos. The ordinate for each bin is the outer radius of the bin divided by $\sqrt{2}$, which is the geometric mean radius of the bin for all but the central bin. Error bars represent the 1σ Poisson sampling error in the mean. Values greater than 0.5 indicate that the axes lie parallel to the filaments, while values less than 0.5 indicate that the axes lie perpendicular to the filaments. The solid line is the power law fit for the full sample given by equation (5.7). The dotted line is the value expected for random orientations.

los). The alignment $\xi_{|ar|}$ is constant out to 3 h^{-1} Mpc, after which it drops until it agrees with the results of the full sample by 6 h^{-1} Mpc. The orientation of clusters is influenced in large part by the accretion of substructure (van Haarlem & van de Weygaert, 1993). The substructure is accreted along filaments, and therefore groups and clusters tend to be strongly aligned with surrounding halos to larger radii. KE04 find stronger alignments out to even larger separations. Their simulations are performed in a much larger (Hubble volume) box, which contains power on longer wavelengths than exist in our smaller simulation volume (and allows them to measure alignments at separations far exceeding our entire box length). Alternatively, it may be a consequence of the higher mass of their halos; while there is virtually no overlap in the range of halo masses between this work and KE04, we find that the alignment is a strong function of halo mass. As KE04 do not present independent results for their lowest mass halos (which are the closest match to our "high mass halo" sample), we cannot tell which of these factors dominates the difference between the results.

The intermediate axes are also aligned with the filaments. The mean alignment as a function of radius is shown in Figure 5.22. While the innermost bins show no alignment, there is clear evidence of alignment at separations of 1 h^{-1} Mpc and beyond. The solid line is a power law fit to the outer 5 points, i.e. for $r > 781\ h^{-1}$ kpc, and is given by

$$\xi_{|br|} = \frac{1}{2} + m_1 r^{\alpha}, \tag{5.8}$$

where $m_1 = 0.008$ and $\alpha = -0.7$. The high mass halos again show constant alignment out to several Mpc.

As both the major and intermediate axes tend to point along filaments, the minor axis must tend to lie perpendicular to the filaments. This trend can be clearly seen in the bottom-left panels of Figures 5.18 and 5.19 as sequences of

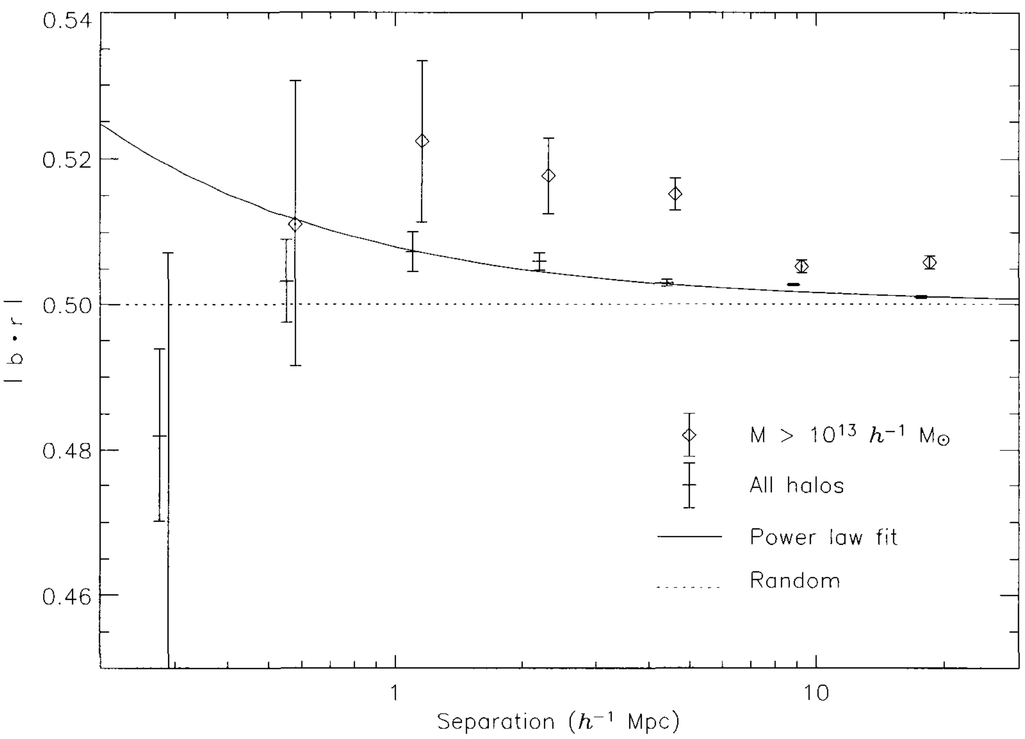

Figure 5.22: As in Figure 5.21, but for the intermediate axis, $\xi_{|br|}$. The solid line is the power law fit for the outer 5 points of the full sample, i.e. for $r > 781$ h^{-1} kpc, and is given by equation (5.8).

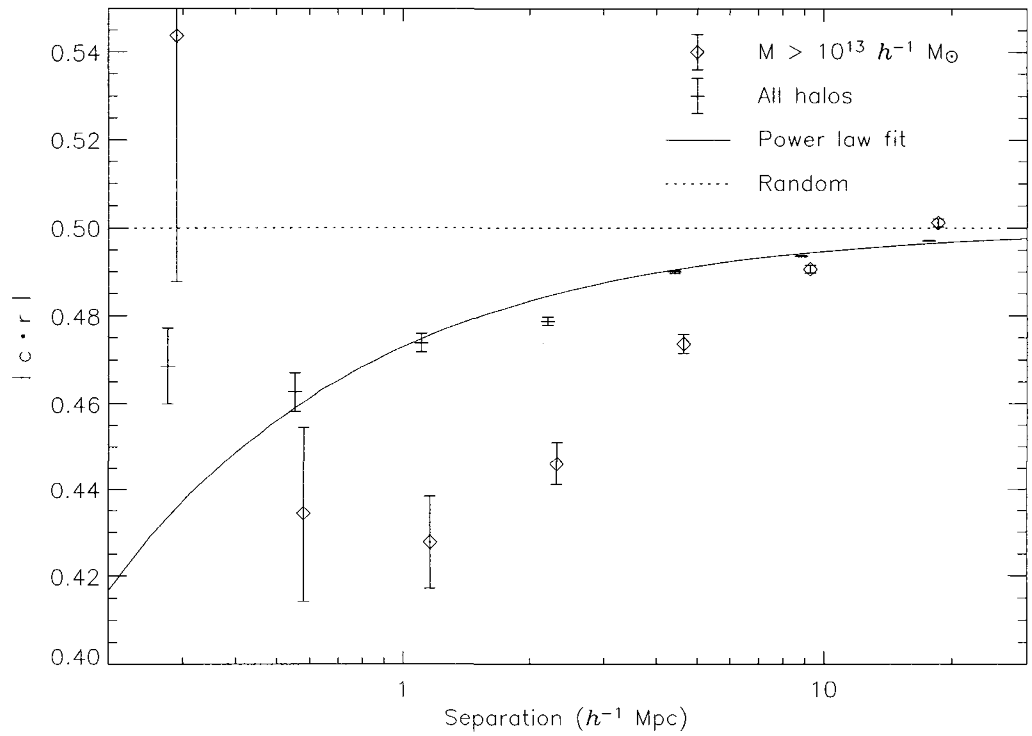

Figure 5.23: As in Figures 5.21 and 5.22, but for the minor axis, $\xi_{|cr|}$. The solid line is the power law fit for the full sample, excluding the innermost bin, i.e. for $r > 391$ h^{-1} kpc, and is given by equation (5.9).

parallel lines running along the filaments. Figure 5.23 shows $\xi_{|cr|}$. The solid line is a power law fit for the full sample excluding the innermost bin, i.e. for $r > 391\ h^{-1}$ kpc, and is given by

$$\xi_{|cr|}(r) = \frac{1}{2} - m_1 r^{\alpha}, \tag{5.9}$$

where the alignment at 1 h^{-1} Mpc is $m_1 = 0.027$ and $\alpha = -0.7$. The minor axes of group and cluster mass halos show even stronger alignment, as seen by the diamonds in Figure 5.23.

The relative strength of the alignment for the different axes is well described by m_1, the value of the power law fit at a separation of 1 h^{-1} Mpc. This is 0.015, 0.008, and 0.027 for the major, intermediate, and minor axes respectively. Therefore it is the minor axis, not the major axis, that is most influenced by the presence of surrounding material. This also explains how both the major and intermediate axes can be positively aligned with the filament: if the minor axis of a halo lies perpendicular to the filament, then both the major and intermediate axes are constrained to lie within a plane that contains the filament. Therefore, they are both more likely to point along the filament then a randomly-oriented three-dimensional axis, and therefore both show positive alignment (for example, if all of the minor axes were perfectly perpendicular to the filament and there were no difference between the major and intermediate axes, then $\xi_{|cr|}$ would vanish and both $\xi_{|ar|}$ and $\xi_{|br|}$ would equal $2/\pi \approx 0.64$). The geometry of the environment is not strictly linear, especially at larger scales. However, despite these complications, the relationships we have found provide useful quantitative predictions that can be tested when full three-dimensional observations of halo shapes are available.

5.5.3 Axis correlations

We compare here the tendency for the principal axes of neighbouring halos to point in the same direction. The procedure we use is completely analogous to that used to calculate the alignments in § 5.5.2. The correlation between the major axes is defined as

$$\xi_{|aa|}(r) \equiv \langle |\hat{\mathbf{a}} \cdot \hat{\mathbf{a}}| \rangle \equiv \frac{1}{N} \sum_{i,j} |\hat{\mathbf{a}}_i \cdot \hat{\mathbf{a}}_j|, \tag{5.10}$$

where the sum is over all unique pairs (i, j) because the measurement is symmetric with respect to each pair of halos. If the major axes lie parallel to each other, $|\hat{\mathbf{a}} \cdot \hat{\mathbf{a}}| > 0.5$, while if they lie perpendicular to each other, $|\hat{\mathbf{a}} \cdot \hat{\mathbf{a}}| < 0.5$. Only halos whose major axes are determined to within 0.2 radians are used for both the primary and secondary sample.

Figure 5.24 plots the mean correlation $\xi_{|aa|}$ of the major axes as a function of the halo separation (this is identical to the quantity defined as $u(r)$ in KE04). Although almost all bins are individually consistent with isotropy, they all lie above 0.5, and taken together are evidence that the directions of the major axes are correlated. We fit a power law and find

$$\xi_{|aa|}(r) = \frac{1}{2} + m_1 r^{\alpha}, \tag{5.11}$$

where r is the separation in units of h^{-1} Mpc, the correlation at 1 h^{-1} Mpc is $m_1 = 0.004$, and the slope is $\alpha = -1.2$, although the large errors introduce considerable uncertainties in these values. The correlations found by both F02 and KE04 are considerably stronger than those found here. We have recalculated $\xi_{|aa|}$ using only halos with masses greater than 10^{13} h^{-1} $\mathrm{M}_{\odot}$ to see if the behaviour of high mass halos differs from those of lower mass, but due to the small number of halos in our sample in this mass range, the errors are too large draw any conclusions.

In Figures 5.25 and 5.26, we plot the mean correlations of the intermediate and

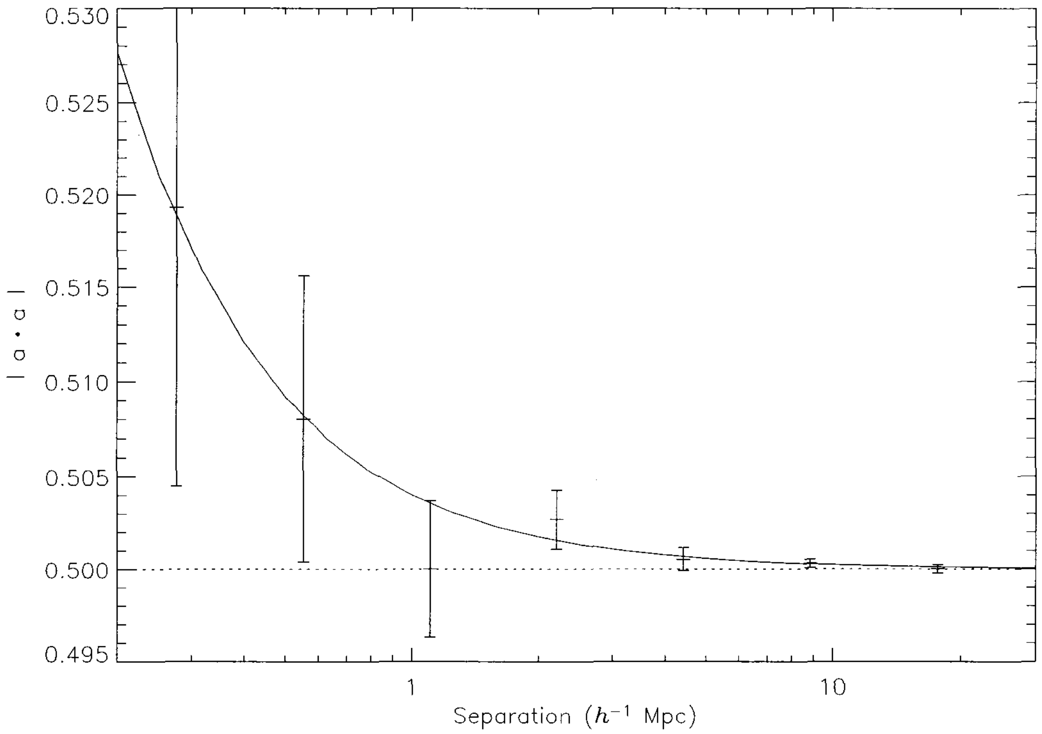

Figure 5.24: Mean correlation $\xi_{|aa|}$ of the major axes of halos for which the direction of the major axes is determined to within 0.2 radians for both halos. The solid line is the power law fit, and is given by equation (5.11). The dotted line is the expected value for random orientations. Error bars represent the 1σ Poisson sampling error in the mean. Values greater than 0.5 indicate that the major axes are parallel to each other, while values less than 0.5 indicate that the axes are perpendicular to each other.

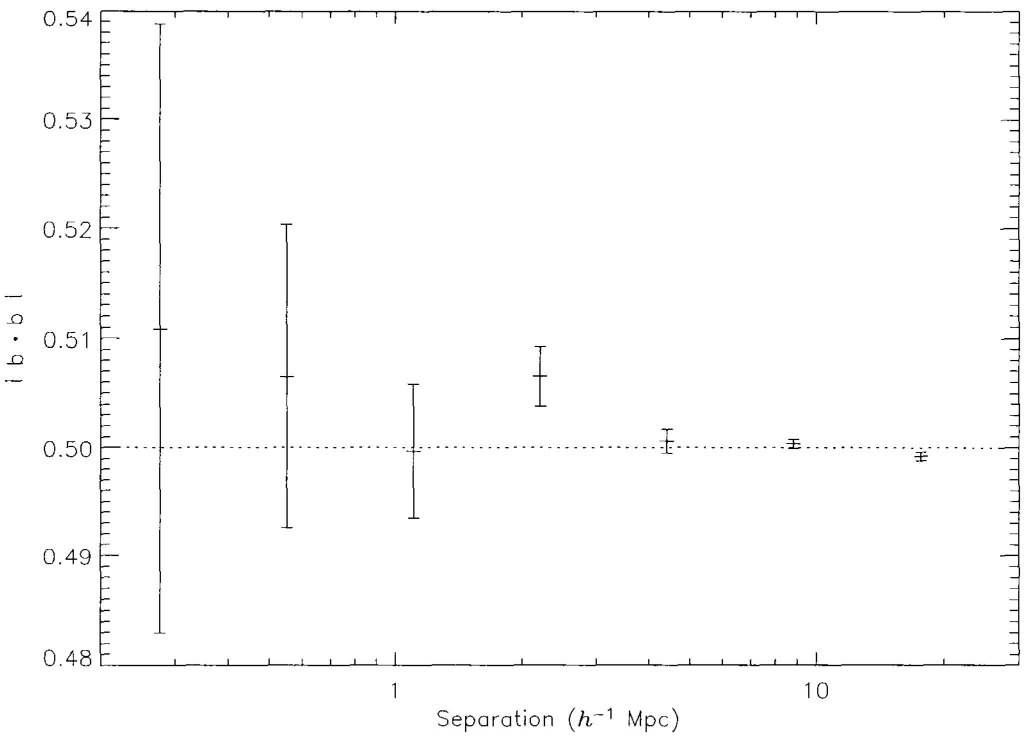

Figure 5.25: As in Figure 5.24, but for the intermediate axis, $\xi_{|bb|}$.

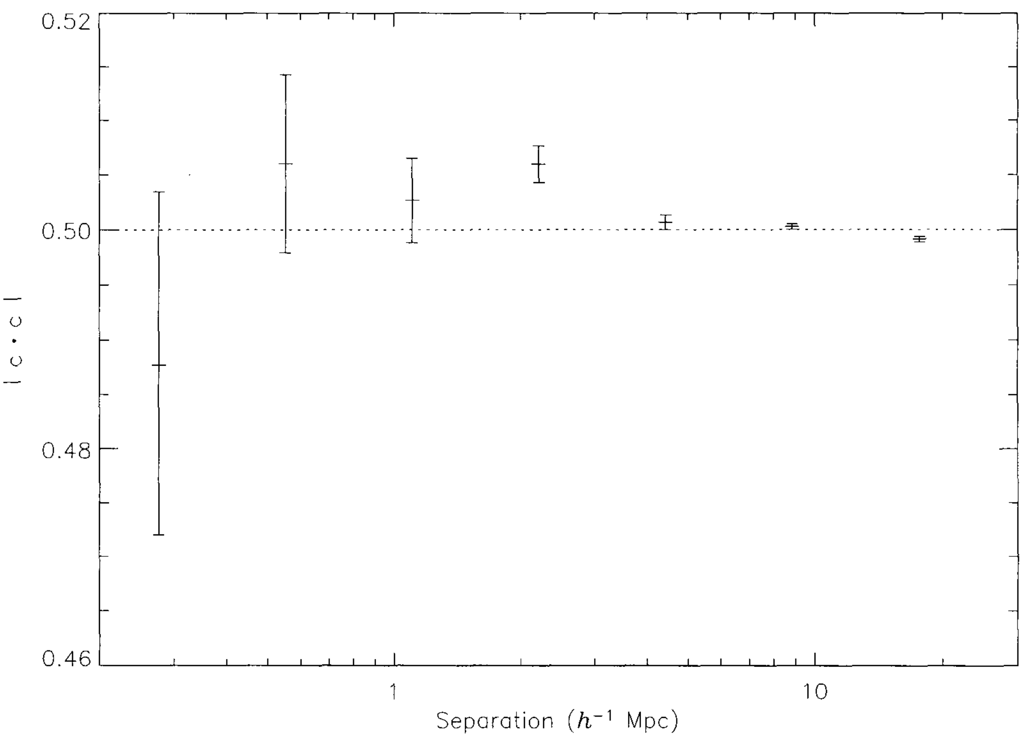

Figure 5.26: As in Figures 5.24 and 5.25, but for the minor axis, $\xi_{|cc|}$.

minor axes of the halos, $\xi_{|bb|}$ and $\xi_{|cc|}$. The error bars are too large to robustly detect any correlation, though the preponderance of bins in Figure 5.26 with $\xi_{|cc|} > 0.5$ is suggestive. A larger sample of halos simulated at equally high spatial resolution is required to see if this is real.

5.5.4 Angular momentum

We investigate the tendency for the angular momentum to point toward or away from nearby halos, for it to lie parallel or perpendicular to the filamentary structure, and for the angular momentum vectors of neighbouring halos to point in the same direction. To measure the tendency of halo angular momenta to point toward or away from other halos, we calculate the alignment

$$\xi_{Lr}(r) \equiv \left\langle \hat{\mathbf{L}} \cdot \hat{\mathbf{r}} \right\rangle \equiv \frac{1}{N} \sum_{i,j} \hat{\mathbf{L}}_i \cdot \hat{\mathbf{r}}_{ij}, \tag{5.12}$$

where $\hat{\mathbf{L}}_i$ is a unit vector in the direction of the angular momentum vector for halo i, and $\hat{\mathbf{r}}_{ij}$ is a unit vector in the direction of the displacement from halo i in the primary sample to halo j in the secondary sample. The primary sample consists of all halos where the direction of the angular momentum is determined to within 0.4 radians, while the secondary sample consists of all halos. Figure 5.27 shows the results. There may be a weak tendency for the angular momentum to point toward local density enhancements on scales of 0.5 – 3 h^{-1} Mpc, but the size of the error bars makes such a result tentative. The high mass halos show no alignment.

Kashikawa & Okamura (1992) and Navarro et al. (2004) have found that galaxies within the local supercluster have their spins pointing within the supercluster plane, while those at least 2 h^{-1} Mpc from the plane have spins that point toward or away from the plane. To see if the behaviour of simulated halos in low and high density regions differs, we have split the sample into those halos that have

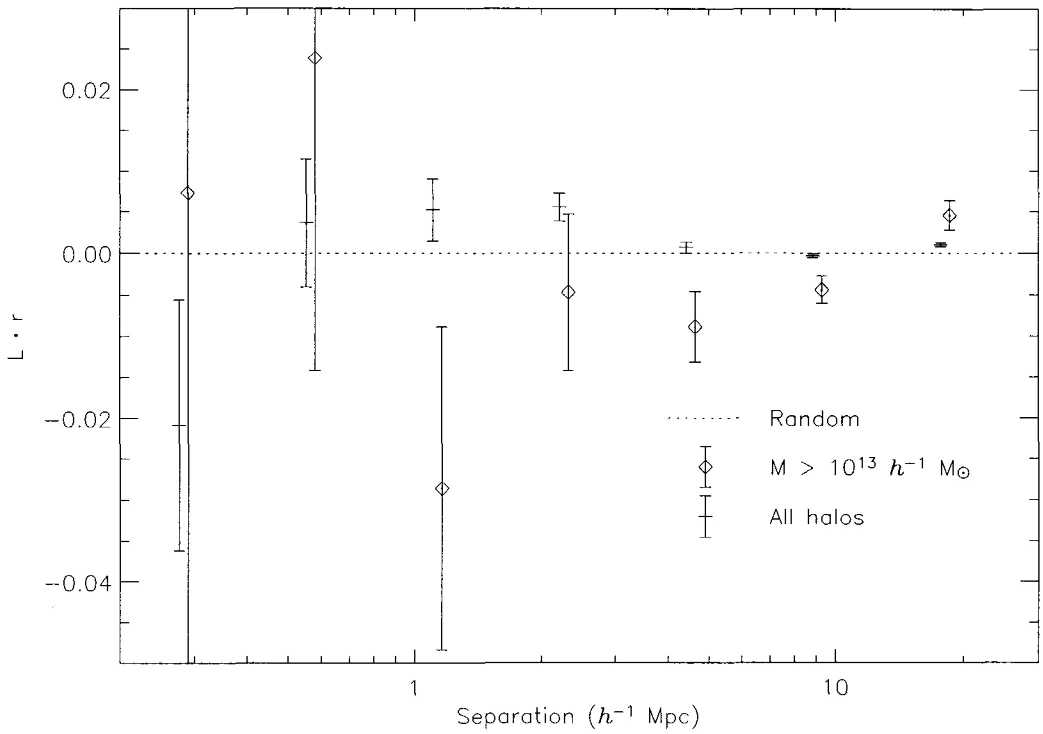

Figure 5.27: Mean alignment ξ_{Lr} between the angular momentum vector of a halo and the location of surrounding halos, as a function of radial separation. Only halos where the direction of the angular momentum is determined to within 0.4 radians are used for the primary samples. The different symbols are for primary samples consisting of all such halos (crosses), or of only those with masses greater than 10^{13} h^{-1} $M_\odot$ (diamonds). The diamonds are shifted slightly to the right for clarity. Error bars represent the 1σ Poisson sampling error in the mean. Positive values indicate that the angular momentum tends to point toward local density enhancements, while negative values indicate that the angular momentum points away from them. The dotted line is the expected value for random orientations.

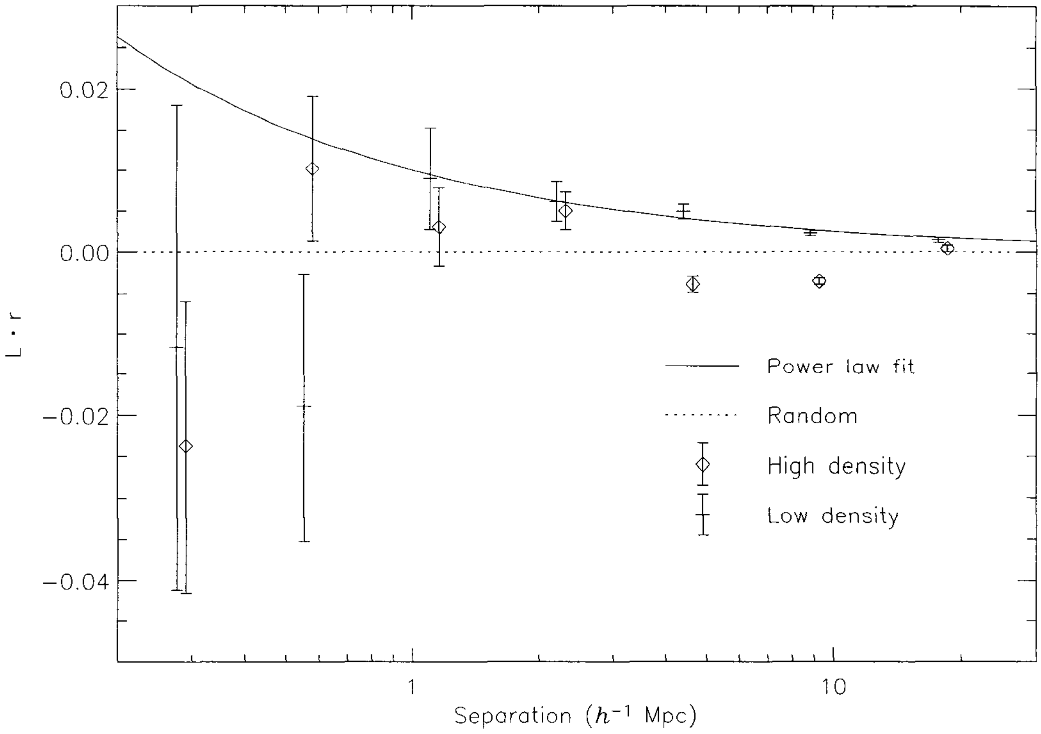

Figure 5.28: Mean alignment ξ_{Lr} as in Figure 5.27, with primary samples selected by local density. Crosses designate halos that have 3 or fewer neighbouring halos within 2 h^{-1} Mpc, while diamonds designate halos with 4 or more neighbouring halos within 2 h^{-1} Mpc. The solid line is the power law fit for the low density sample at $r > 781\ h^{-1}$ kpc, and is given by equation (5.13).

3 or fewer neighbouring halos within 2 h^{-1} Mpc (the low density sample), and those that have 4 or more neighbours within 2 h^{-1} Mpc (the high density sample). There are 2155 and 1714 halos in the low and high density samples respectively. In Figure 5.28, we have plotted the alignment ξ_{Lr} against separation for the two samples. By construction, there are very few pairs (and therefore large error bars) at small separations in the low density sample, but beyond 1 h^{-1} Mpc there is a clear detection of positive alignment in this sample. The alignment is well fit by a power law of the form

$$\xi_{Lr} = m_1 r^{\alpha}, \tag{5.13}$$

where r is the separation in units of h^{-1} Mpc, the correlation at 1 h^{-1} Mpc is $m_1 = 0.01$, and the slope is $\alpha = -0.6$. The high density sample shows no coherent tendency for the angular momentum vectors to point toward or away from density enhancements. The low and high density samples do not show significant deviations from the full sample for any of the other statistics studied.

The tendency for the angular momentum vectors to lie parallel versus perpendicular to the filaments is measured by

$$\xi_{|Lr|}(r) \equiv \left\langle |\hat{\mathbf{L}} \cdot \hat{\mathbf{r}}| \right\rangle \equiv \frac{1}{N} \sum_{i,j} |\hat{\mathbf{L}}_i \cdot \hat{\mathbf{r}}_{ij}|, \tag{5.14}$$

defined as $L_{||}$ in Hatton & Ninin (2001), or

$$\omega(r) \equiv \left\langle |\hat{\mathbf{L}} \cdot \hat{\mathbf{r}}|^2 \right\rangle - \frac{1}{3} \equiv \frac{1}{N} \sum_{i,j} |\hat{\mathbf{L}}_i \cdot \hat{\mathbf{r}}_{ij}|^2 - \frac{1}{3}, \tag{5.15}$$

as used by Lee & Pen (2001). If $\xi_{|Lr|} > 0.5$ or $\omega(r) > 0$ then the angular momentum vectors lie parallel to the filaments, while if $\xi_{|Lr|} < 0.5$ or $\omega(r) < 0$ then the angular momentum vectors point perpendicular to the filaments. In Figures 5.29 and 5.30, we plot $\xi_{|Lr|}$ and $\omega(r)$ respectively. The solid line in Figure 5.30 shows the prediction of Lee & Pen (2001) from linear tidal torque theory. On scales less

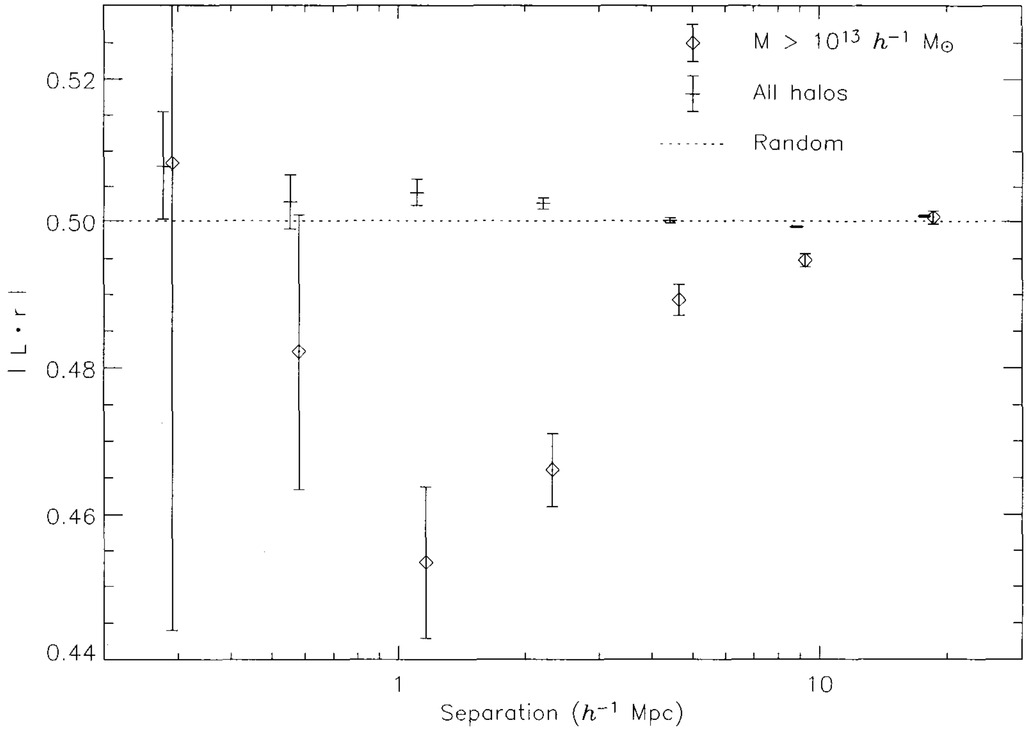

Figure 5.29: Mean alignment $\xi_{|Lr|}$ between the angular momentum vector of primary halos and the location of all surrounding halos. Only halos where the direction of the angular momentum is determined to within 0.4 radians are used for the primary samples. The different symbols are for primary samples consisting of all such halos (crosses), or of only those with masses greater than 10^{13} h^{-1} M$_\odot$ (diamonds). The diamonds are shifted slightly to the right for clarity. Error bars represent the 1σ Poisson sampling error in the mean. Values greater than 0.5 indicate that the angular momentum vectors point parallel to the filaments, while values less than 0.5 indicate that the angular momentum vectors point perpendicular to the filaments. The dotted line is the expected value for random orientations.

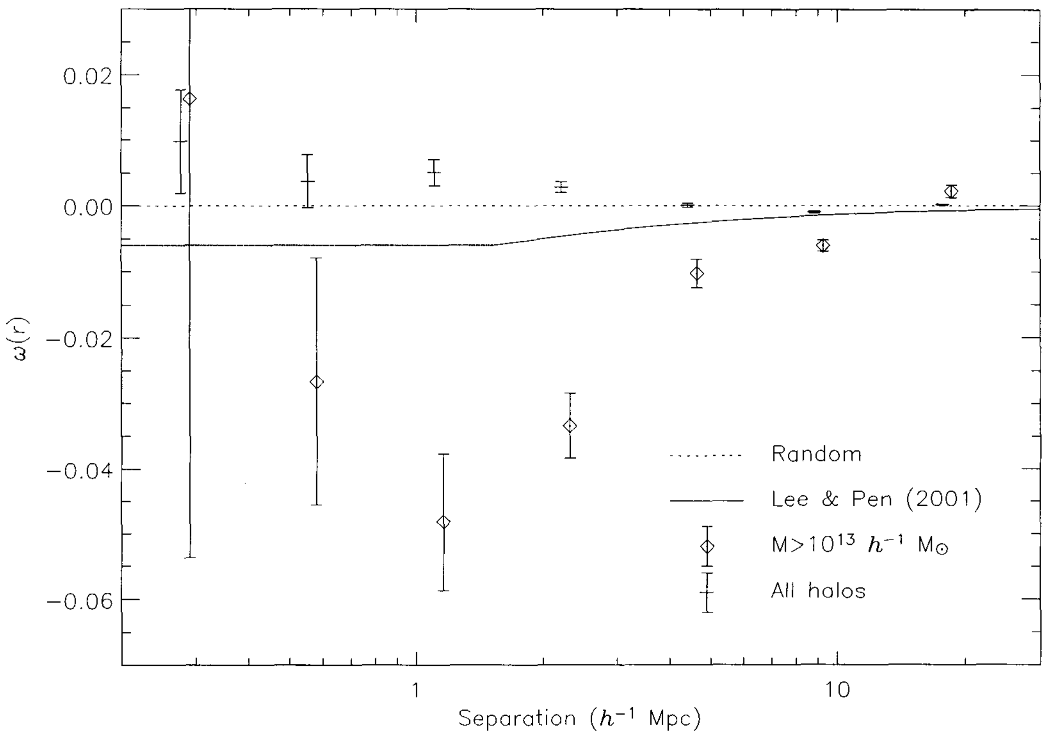

Figure 5.30: Alignment $\omega(r)$ between the angular momentum vector of a primary halo and the location of all surrounding halos. Only halos where the direction of the angular momentum is determined to within 0.4 radians are used for the primary samples. The different symbols are for primary samples consisting of all such halos (crosses), or of only those with masses greater than 10^{13} h^{-1} M$_\odot$ (diamonds). The diamonds are shifted slightly to the right for clarity. Error bars represent the 1σ Poisson sampling error. Positive values indicate that the angular momentum vectors tend to point parallel to the filaments, while negative values indicate that the angular momentum vectors point perpendicular to the filaments. The dotted line is the expected value for random orientations, while the solid line is the prediction from linear tidal torque theory (Lee & Pen, 2001).

than 3 h^{-1} Mpc, the angular momentum tends to lie parallel to the filaments. The values are consistent with those found by Hatton & Ninin (2001). However, both Lee & Pen (2001) and F02 find that the angular momenta of halos tend to lie perpendicular to the filaments. This discrepancy may lie in the different mass ranges probed. The halos of Hatton & Ninin (2001) cover a very similar mass range to those in this work, while the sample of F02 consists entirely of cluster mass halos (their *smallest* halo has a mass of 1.4×10^{14} h^{-1} $\mathrm{M}_{\odot}$, nearly the mass of our *largest* halo). We have recalculated $\xi_{|Lr|}$ and $\omega(r)$ using only halos with masses greater than 10^{13} h^{-1} $\mathrm{M}_{\odot}$ in the primary sample and plotted them as diamonds in Figures 5.29 and 5.30. The behaviour is radically different; the angular momenta of groups and clusters tend to point perpendicular to the filaments.

The correlations between the angular momentum vectors of halos are calculated analogously. In particular, the tendency for halo angular momenta to point in the same direction versus in the opposite direction as the angular momenta of other halos is defined as

$$\xi_{LL} \equiv \left\langle \hat{\mathbf{L}} \cdot \hat{\mathbf{L}} \right\rangle \equiv \frac{1}{N} \sum_{i,j} \hat{\mathbf{L}}_i \cdot \hat{\mathbf{L}}_j, \tag{5.16}$$

where the sum is over all unique pairs (i, j). This is equivalent to the quantity defined is $\eta(r)$ in Porciani et al. (2002a). As seen in Figure 5.31, we detect no deviations from random. We also measure the tendency for halo angular momentum vectors to be either parallel or perpendicular to each other. We define two quantities to measure this,

$$\xi_{|LL|} \equiv \left\langle |\hat{\mathbf{L}} \cdot \hat{\mathbf{L}}| \right\rangle \equiv \frac{1}{N} \sum_{i,j} |\hat{\mathbf{L}}_i \cdot \hat{\mathbf{L}}_j|, \tag{5.17}$$

which is the same as the quantity defined as $\mu(r)$ in Hatton & Ninin (2001), and

$$\eta(r) \equiv \left\langle |\hat{\mathbf{L}} \cdot \hat{\mathbf{L}}|^2 \right\rangle - \frac{1}{3} \equiv \frac{1}{N} \sum_{i,j} |\hat{\mathbf{L}}_i \cdot \hat{\mathbf{L}}_j|^2 - \frac{1}{3}, \tag{5.18}$$

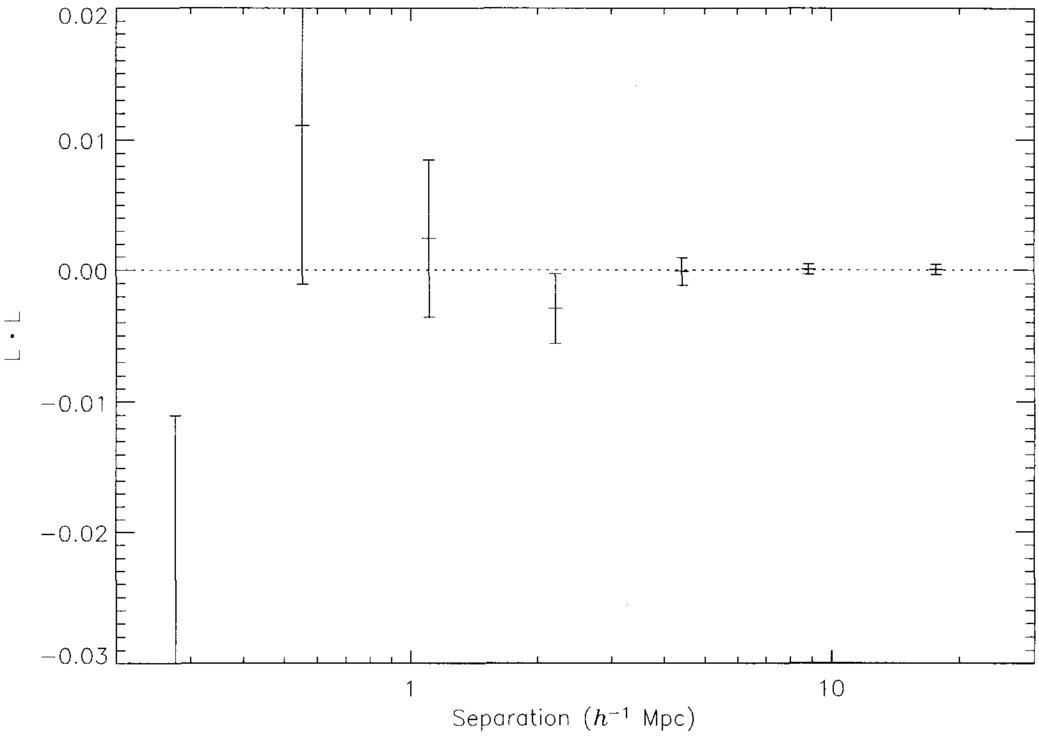

Figure 5.31: Mean correlation ξ_{LL} of the angular momentum vectors of halos as a function of their separation. All halos where the direction of the angular momentum is determined to within 0.4 radians are used. Error bars represent the 1σ Poisson sampling error in the mean. Positive values indicate that the angular momenta tend to point in the same direction, while negative values indicate that the angular momenta tend to point in opposite directions. The dotted line is the expected value for random orientations.

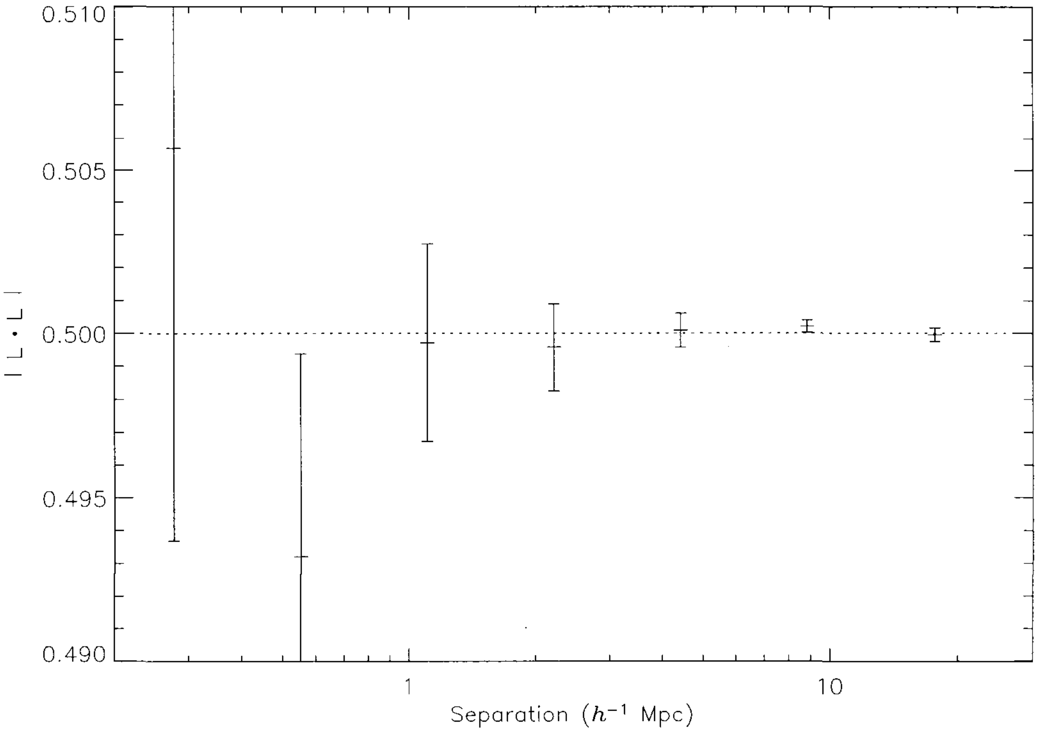

Figure 5.32: Mean correlation $\xi_{|LL|}$ between the angular momentum vectors of halos as a function of their separation. All halos where the direction of the angular momentum is determined to within 0.4 radians are used. Error bars represent the 1σ Poisson sampling error in the mean. Positive values indicate that the angular momentum vectors tend to lie parallel to each other, while negative values indicate that the angular momentum vectors tend to lie perpendicular. The dotted line is the expected value for random orientations.

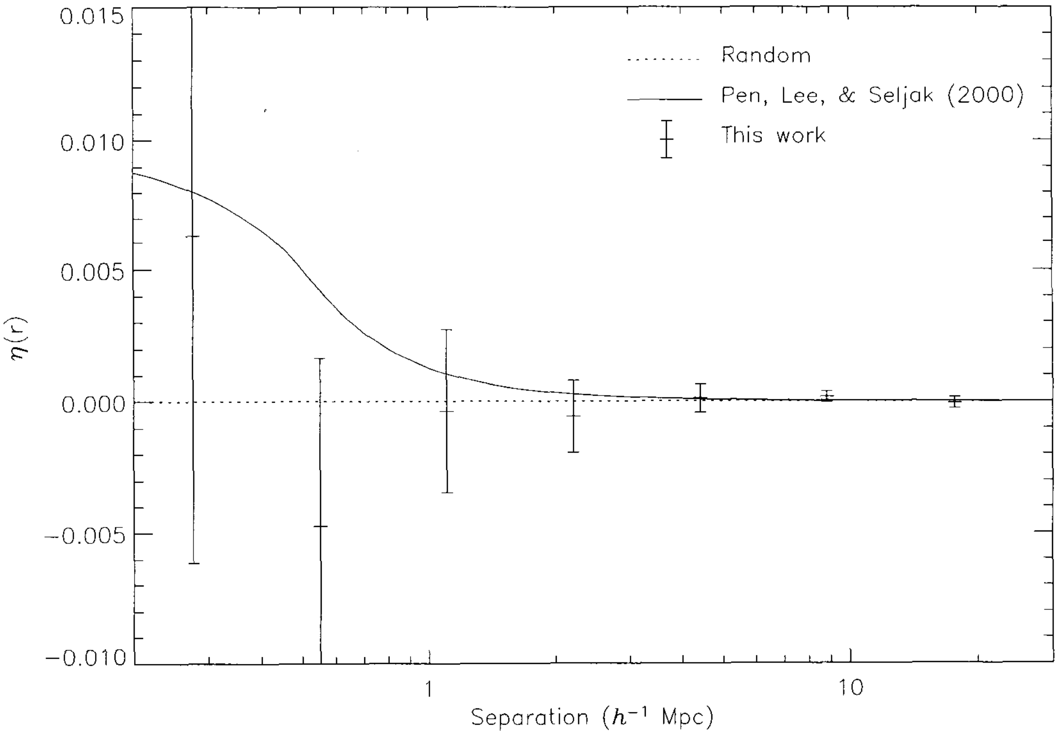

Figure 5.33: Correlation $\eta(r)$ between the angular momentum vectors of halos as a function of their separation. All halos where the direction of the angular momentum is determined to within 0.4 radians are used. Error bars represent the 1σ Poisson sampling error. Positive values indicate that the angular momentum vectors tend to lie parallel to each other, while negative values indicate that the angular momentum vectors tend to lie perpendicular. The dotted line is the expected value for random orientations. The solid line is the linear tidal torque prediction (Pen et al., 2000).

which is the same as Pen et al. (2000)'s $\eta(r)$ and Porciani et al. (2002a)'s $\eta_2(r)$. The quantities $\xi_{|LL|}(r)$ and $\eta(r)$ are shown in Figures 5.32 and 5.33 respectively. We detect no deviations from isotropy. Due to the size of the error bars, these results are consistent with the non-detections and weak tendencies toward alignment found by other authors (Pen et al., 2000; Hatton & Ninin, 2001; Porciani et al., 2002a; F02). We do not find any difference if we restrict the sample to halos of any particular mass range.

5.6 Discussion

The differences between the properties of galaxy mass halos compared to those of group and cluster mass halos suggest that different factors determine their orientations. For example, the orientation of the major axis of a cluster is strongly affected by the direction of the most recently accreted subhalo, as seen in simulations (van Haarlem & van de Weygaert, 1993; Tormen, 1997) and observations (Ebeling et al., 2004). If these are accreted from filaments, then there should be a strong alignment of cluster principal axes with the filaments, as we have found. We find that the tendency of the minor axes to lie perpendicular to the filaments is much stronger than the tendency of the major axes to lie parallel to the filaments. While previous authors have neglected the minor axis, it appears that the effect of these mergers on the minor axis is stronger than their effect on the major axis. Another possibility is that the initial alignment of the major axis is stronger, but that figure rotation scrambles this alignment. In § 4, we found that most halos show slow figure rotation about the minor axis, with a median pattern speed of $\approx 0.15\ h\ \mathrm{km\ s^{-1}\ kpc^{-1}}$, suggesting that the major axis of a typical halo can change by 90° in a Hubble time. The much weaker alignment for galaxy mass halos suggests that either the orientation of the axes is less affected by recent mergers (although we find no relationship between the pattern speed and halo mass, the lower mass halos are on average dynamically older, and therefore figure rotation may have had a longer time to modify their orientation), or that the direction of the accretion is more isotropic (Vitvitska et al., 2002). Aubert et al. (2004) find that the accretion onto halos at masses down to $5 \times 10^{12}\ M_{\odot}$ is quite anisotropic; however, 97% of our "galaxy" mass halos lie below this limit.

The differences between the angular momentum alignments for high and low mass halos are also intriguing. Angular momentum is usually thought to either arise from the tidal torquing of an asymmetric protohalo (White, 1984), or by the accretion of substructure on non-radial orbits (Tormen, 1997; Vitvitska et al., 2002). In a sense, these are not distinct scenarios; accreted subhalos *are* protogalactic material that has been tidally torqued. However, the relation between the direction of the angular momentum and surrounding matter may differ depending on the clumpiness of the accretion and other non-linear effects (Porciani et al., 2002a).

It is interesting to compare our results for collisionless dark matter halos with the results of Navarro et al. (2004) for simulations that include baryonic physics. The main results of Navarro et al. (2004) are that the angular momenta of baryonic galactic disks in gasdynamical simulations tend to align with the intermediate axis of the local baryon distribution on scales of $\approx 2\ h^{-1}$ Mpc, and that observed edge-on disk galaxies in the local supercluster have their spin axes lying within the supergalactic plane, as expected in such a scenario. We find that on the scale of an individual halo, the angular momentum of the dark matter aligns with its minor axis. However, we find that the angular momentum of the dark matter in galaxy mass halos is parallel to the large scale structure, as expected if the spin vector lies within local sheet-like structures as suggested by Navarro et al. (2004). We also find that in low density regions, halo spins point toward nearby filaments, and are therefore aligned with the intermediate axis of the local density field. Therefore, despite the common presence of misalignments between the angular momenta of the baryons and the dark matter in individual galaxies (Sharma & Steinmetz, 2004), each retains some memory of the large scale structure which provided the initial torques.

In the group and cluster mass halos, on the other hand, the angular momentum shows a strong tendency to point perpendicular to the large scale structure. These objects tend to occur at the intersections of large filaments, and have been built up by the recent accretion of smaller halos. These halos are accreted along the filaments, and therefore add angular momentum perpendicular to the filaments (note that this confirms the suggestion of Maller et al. (2002) that accretion along a preferred axis is required to explain the distribution of spin parameters seen in simulated halos). This process appears to dominate the direction of the angular momentum for halos of mass greater than $10^{13}\,h^{-1}\,\mathrm{M}_{\odot}$. Unlike for galaxy mass halos, for which we can confirm that the angular momenta of the dark matter and baryons share similar relationships to the large scale structure in simulations, we do not have high resolution gasdynamical simulations of these more massive objects with which to compare. It may be possible to determine the rotation axis of baryons in these systems observationally. In theory, the spin vector of a relaxed cluster can be deduced from the presence of a redshift gradient of the galaxies in the cluster; however, confusion due to structure along the line of sight, the small magnitude of the rotation compared to the intrinsic velocity dispersion, ambiguity in the orientation of the ellipsoidal shape, intrinsic distance gradients, the lack of a large sample of substructureless relaxed clusters, and the uncertain relationship between the angular momentum of the cluster galaxies and that of the smooth X-ray emitting gas that dominates the baryonic mass make this measurement difficult. The rotation of the X-ray gas itself, however, may be measured using the kinematic Sunyaev-Zeldovich effect in future Cosmic Microwave Background (CMB) surveys (Chluba & Mannheim, 2002; Cooray & Chen, 2002).

5.7 Summary

We have studied the internal shapes and angular momenta of galaxy and group mass dark matter halos formed in a ΛCDM N-body simulation, and studied how they are correlated with the large scale structure and the properties of neighbouring halos.

Internally, halos are triaxial with b/a and c/a ratios of 0.75 ± 0.15 and 0.6 ± 0.1 respectively. The distribution of axis ratios has a tail to low values. The two-dimensional projected ellipticities cover a broad range of values from 0 to 0.5, with a mean of 0.24, consistent with the weak lensing results of Hoekstra et al. (2004). The axis ratios rise between 0.12 and 0.6 $r_{\rm vir}$, beyond which they drop. Within 0.12 $r_{\rm vir}$, the measurement is probably compromised by the force softening in the simulations. Halos are most often prolate in the inner regions, but tend to a more even mix of prolate and oblate at large radii.

The internal alignment of the halos within 0.6 $r_{\rm vir}$ is very good, particularly for the minor axis, with a slight decrease in alignment in the outermost regions. High mass halos have particularly well-aligned axes. The orientation of the angular momentum is also relatively constant, though it changes more noticeably as a function of radius than do the axes. At any given radius, the angular momentum vector tends to be aligned with the minor axis and be perpendicular to the major axis. It also shows a weaker tendency to be perpendicular to the intermediate axis. The alignment between the angular momentum and the principal axes gets weaker at larger radii. The properties of the halo at 0.4 $r_{\rm vir}$ are quite characteristic of their values at most other radii.

The minor axes of halos show a strong tendency to lie perpendicular to the filaments. As a consequence, the major (and, to a lesser degree, intermediate) axes tend to point along filaments. Figure rotation about the minor axis may be responsible for the smaller degree of major axis alignment. These alignments fall off with distance as a power law. In all cases, the alignment for group and cluster mass halos extends to much larger separations than for galaxy mass halos. The major axes show a weak correlation with those of other nearby halos. There is no robust detection of a correlation for the other axes.

The angular momenta of halos in low density environments tend to point toward local density enhancements, in agreement with the results of Kashikawa & Okamura (1992) for galaxies 2 h^{-1} Mpc or more away from the local supergalactic plane. The angular momenta of galaxy mass halos show a weak tendency to point along filaments on scales up to 3 h^{-1} Mpc, but those of group and cluster mass halos show a very strong tendency to point perpendicular to the filaments. We detect no correlations of the angular momentum directions of nearby halos with each other, but due to the size of the error bars, this is consistent with previous linear and N-body studies that predict weak correlations. Comparisons with recent gasdynamical simulations and observations of edge-on disk galaxies in the local supercluster suggest that both the baryons and dark matter in galaxies share a memory of the orientation of the large scale structure that provided the initial torque. Groups and clusters, on the other hand, appear to acquire most of their angular momentum from the accretion of subhalos along filaments. This may be tested with large samples of galaxy redshifts within relaxed clusters, or by kinematic Sunyaev-Zeldovich studies in future CMB experiments.

CHAPTER 6

CONCLUSIONS

6.1 Summary

We have used numerical N-body simulations to study several aspects of the alignment between the angular momentum of the disk, halo, and surrounding structure in galaxy formation.

The angular momenta of the inner and outer disks in warped galaxies are misaligned with each other. Angular momentum is transferred by torques, so the cause of this misalignment is related to how a massive galactic disk reacts to torques. We have calculated the torques a misaligned dark matter halo imparts on a disk, and found that the tilting timescale for the disk is quite short, significantly shorter than a Hubble time. We have performed N-body simulations of galactic disks subject to torques of this strength. The inner disk reacts more quickly to the torque, tilting as a solid body while the outer disk follows behind, resulting in a trailing warp. Galactic disks are marginally unstable to warping when the local surface density drops below $70\ M_{\odot}\ \mathrm{pc}^{-2}$; this is the radius at which the warp first develops, regardless of the disk mass. Vertical bending waves transmit the warp to the end of the disk on a timescale of 500 Myr. The warped angular momentum in these simulated disk is similar to the warped angular momentum observed in the Milky Way.

Satellite galaxies may provide another source of misaligned angular momentum that can trigger warps. We have compared the misaligned angular momentum responsible for the Milky Way warp to that of the Sagittarius dSph galaxy and 6 other nearby satellites. There is a striking coincidence between the magnitude and direction of the warp and Sagittarius: they both have a magnitude of between $2-8\times10^{12}\ M_{\odot}$ kpc km s^{-1} (depending on the assumed mass model), and are almost perfectly anti-aligned. We estimate the probability that this is a chance configuration to be 1.4%. This suggests that they are coupled, and that Sgr is responsible for the Milky Way warp. We have performed a series of N-body simulations of satellite-disk interactions in order to test this hypothesis, and found that the key parameters that determine whether a satellite can efficiently excite a warp are the orientation of the orbit (inclined orbits excite warps more easily than polar orbits), the satellite mass, and the number of perigalactic passages (for eccentric orbits). The warps excited by satellite galaxies are asymmetric. The orientation of the warp with respect to the satellite orbit is quickly scrambled by the processes of differential precession and the vertical bending waves which dissipate the warp. Even our most disturbed model disks have only a fraction of the warp angular momentum that we calculate for the Milky Way, which may suggest that the mass models are mismatched. However, even when we normalize by the total angular momentum, which should reduce any discrepancy due to mismatched mass models, we find that the warp of the Milky Way is several times stronger than in any of the simulated disk-satellite systems. The typical warp angular momentum excited by a satellite galaxy is $\sim 1\%$ of the satellite's orbital angular momentum. Therefore, Sgr with its current mass on its current orbit cannot provide the torque necessary to explain the Milky Way warp, despite the promising coincidence of its angular momentum.

The torques that warp galaxies may come from misalignments between the angular momentum of the disk and the angular momentum or shape of its dark matter halo. Therefore, we have studied the properties of galaxy and group-mass dark matter halos formed in cosmological simulations. We have studied the triaxial shapes of the halos, and found that the minor-to-major axis ratios are typically $c/a = 0.6$, and are more spherical at larger radii throughout most of the halo, but become more flattened again near the virial radius. We have detected rotation of the triaxial figure in most undisturbed halos, with the rotation axis coinciding with the minor axis in 85% of the halos and the major axis in 15% of the halos. The pattern speeds follow a log normal distribution centred on $\Omega_p = 0.148\ h\ \mathrm{km\ s^{-1}\ kpc^{-1}}$, which is not sufficiently fast to account for the spiral structure in NGC 2915 as had been hypothesized.

The shapes of halos are internally very well aligned, particularly for high mass halos. The angular momentum is relatively well aligned within the halo, and tends to align with the minor axis. However, the median misalignment between the angular momentum and minor axis is still $\approx 25^\circ$; misalignment of this magnitude between the angular momentum of a galactic disk and the minor axis of its halo can produce realistic-looking warps. The alignment between the angular momentum of the halo and its shape becomes poorer at larger radii. The minor axes of halos show a strong tendency to lie perpendicular to large scale filaments, while the major axes show a weaker tendency to lie along the filaments. These alignments are much stronger for group and cluster mass halos than for galaxy mass halos. The angular momenta of halos in low density environments tend to point toward local density enhancements. The angular momenta of galaxy mass halos tend to point parallel to the large scale structure, while those of group and cluster mass halos point perpendicular to the surrounding mass distribution.

The Milky Way satellites do not appear to be sufficiently massive to account for the warp of the disk. On the other hand, torques from misaligned dark matter halos are sufficiently strong to warp the disk by the observed amount. Misalignments of the required degree between the angular momentum and shape of cosmological dark matter halos are seen in cosmological simulations. Although such misalignments are expected to quickly dissipate if isolated, the figures of dark matter halos rotate; this figure rotation may be responsible for sustaining the misalignment between the shape of the disk and halo, and therefore for the prevalence of warped galaxies.

6.2 Outlook

There are still many unanswered questions regarding the alignment of angular momentum within galaxies and the effects misalignment may have on observed galactic disks. In this section, we pose several questions prompted by our results, and outline what future work can be done to begin to answer them.

- *How are warps produced in cosmological galactic disks?*

  We have argued that tidal torques in a cosmological context, in the form of halos whose shape or angular momentum is misaligned with that of the galactic disk, provide a plausible mechanism for exciting galactic warps. We have found that the necessary misalignments between the angular momentum and the shape of dark matter halos exist in N-body simulations, and are of the appropriate magnitude. However, galactic disks are not formed in purely gravitational simulations. The ideal way to check this scenario is with extremely high-resolution cosmological simulations that include hydrodynamics and other relevant non-gravitational physics, where the disks of individual galaxies are well enough resolved for warps to be detected. The current resolution of cosmological simulations with large enough box sizes to sample a representative volume of the universe is not high enough; even in simulations that focus on single galaxies, warps are just barely detectable[1], but simulations of only a handful of such galaxies exist. Using such multimass resimulation techniques, it is currently feasible to assemble a representative sample of high resolution disks from which warp statistics can be drawn and the responsible environmental effects can be directly found. Soon, computing power will reach a sufficient level where resimulations will not be necessary, and we will be able to draw these statistics directly from completely self-consistent cosmological simulations. In the meantime, we recommend going to the beach (Gottbrath et al., 1999).

- *What is the intrinsic distribution of satellite orbits around the Milky Way and other galaxies, and what effects might they have on the ability of satellites to warp galactic disks?*

  The initial conditions of our simulations in Chapter 3 are motivated by the orbits of observed Milky Way satellites, especially Sgr. The Milky Way satellites tend to be on polar orbits, the so-called "Holmberg effect", which has also been seen in samples of external galaxies by several authors (Zaritsky et al., 1997; Sales & Lambas, 2004). If galactic disks tend to lie perpendicular to large scale sheets (Navarro et al., 2004), this may explain such an overabundance of satellites in the polar regions of galactic disks. However, cosmological simulations suggest that infalling substructures are preferentially accreted near the equatorial plane, with their angular momenta aligned reasonably well with the halo angular momentum (Tormen, 1997; Aubert et al., 2004; Benson, 2004; Knebe et al., 2004), and an analysis of SDSS galaxies also finds an overabundance of satellites in the plane of disk galaxies (Brainerd, 2004). Therefore, the true intrinsic distribution of satellite orbits is still an open question, both theoretically and observationally. It may be that observations of the Milky Way environment have preferentially missed satellites on equatorial orbits (the recently-discovered Canis Major structure may be one such satellite; Martin et al., 2004), dynamical processes may preferentially destroy satellites on equatorial orbits (e.g. Peñarrubia et al., 2002), the samples of external galaxies may be compromised by projection effects, or the cosmological simulations may be incorrect, Until the intrinsic distribution of satellite orbits is understood, simulations like those presented in Chapter 3 must either use the observed distribution of Milky Way satellite orbits for the initial conditions (as we have done), or the distribution of substructure orbits found in cosmological simulations. On the observational side, this problem can be attacked by conducting a thorough census of the Galactic environment using all sky surveys such as SDSS[2], 2MASS[3], RAVE[4], and GASS[5]. Better proper motions of the known satellites, which will improve as the time baseline grows longer and as astrometric missions such as GAIA[6] fly, will also reduce the uncertainties in the orbital parameters considerably. On the theoretical side, an improved treatment of star formation, feedback, the intergalactic radiation field, and of course better resolution, may help resolve the discrepancy.

---

[1] see, for example, `http://www.aip.de/People/MSteinmetz/Movies/AVI/spiral_gas.avi`, in particular during and after satellite passages at $z < 0.5$.

- *What additional physics may affect the results of satellite-disk simulations?*

  The disk simulations in Chapters 2 and 3 only include the effects of gravity. Observed galactic disks can contain large amounts of gas which experience hydrodynamic forces in addition to the gravitational forces we have modelled. Although our proposed warp mechanisms are entirely gravitational, and therefore the torques imposed on the disk are independent of the gas content of the disk, the response of a gaseous disk to these torques may differ somewhat from that of an N-body disk (Ibata & Razoumov, 1998). Furthermore, the satellites have been modelled as single particles without any internal degrees of freedom. Simulations with self-consistent satellites will allow us to model the tidal disruption of the satellite and more confidently follow the simulations for a longer period of time.

---

[2] `http://www.sdss.org/`

[3] `http://www.ipac.caltech.edu/2mass/`

[4] `http://www.aip.de/RAVE`

[5] a recently-commenced all sky survey at Parkes for overlooked H I clouds in the environment of the Galaxy and the Local Group (B. Gibson, private communication)

[6] `http://astro.estec.esa.nl/GAIA/`

- *What further effects might halo figure rotation have on the embedded galaxy?*
  The magnitude of the figure rotation determined in Chapter 4 is not sufficient to explain the spiral structure of NGC 2915. However, it may be partly responsible for the misalignment between baryonic disks and the shape of their dark matter halos, and therefore indirectly responsible for galactic warps. Simulations such as those performed by Bekki & Freeman (2002) and MB03 should be performed for the range of pattern speeds we have found to discover whether figure rotation of this magnitude has any directly observable effects. In addition, calculations of the figure rotation with much finer time resolution would allow us to discover if changes in the figure rotation rate, such as those required by the Tremaine & Yu (2000) polar ring and counter-rotating disk mechanism, occur.

- *What determines the figure rotation?*
  We have demonstrated that the magnitude of the figure rotation is correlated with the spin parameter of the halo, while the figure rotation axis is usually aligned with the minor axis but occasionally aligned with the major

axis. However, the factors in each halo's history that determine the figure rotation are still not understood. Clues to this may come from a detailed look at each halo's formation history, to see if there are correlations between the figure rotation and the parameters of recent major or minor mergers. There may also be environmental factors such as the strength and direction of the local tidal field. These studies can be complemented by detailed simulations of the effects of mergers and the tidal field on initially static triaxial halos.

I look forward to tackling some of these issues in the future.

REFERENCES

Aarseth, S. J. 1963, MNRAS, 126, 223

Athanassoula, E. 2002, ApJ, 569, L83

Athanassoula, E., Bosma, A., & Papaioannou, S. 1987, A&A, 179, 23

Aubert, D., Pichon, C., & Colombi, S. 2004, MNRAS, 352, 376

Bailin, J. 2003, ApJ, 583, L79

Bailin, J., & Steinmetz, M. 2004a, ApJ, in press, (astro-ph/0405442)

—. 2004b, ApJ, submitted, (astro-ph/0408163)

Bardeen, J. M., & Petterson, J. A. 1975, ApJ, 195, L65

Barnes, J., & Efstathiou, G. 1987, ApJ, 319, 575

Battaner, E., Florido, E., & Sanchez-Saavedra, M. L. 1990, A&A, 236, 1

Bekki, K., & Freeman, K. C. 2002, ApJ, 574, L21

Bell, E. F., & de Jong, R. S. 2001, ApJ, 550, 212

Benson, A. J. 2004, MNRAS, submitted, (astro-ph/0407428)

Binggeli, B. 1982, A&A, 107, 338

Binney, J. 1981, MNRAS, 196, 455

—. 1992, ARA&A, 30, 51

Binney, J., & Dehnen, W. 1997, MNRAS, 287, L5

Binney, J., Jiang, I., & Dutta, S. 1998, MNRAS, 297, 1237

Binney, J., & Merrifield, M. 1998, Galactic astronomy (Princeton, NJ: Princeton University Press)

Binney, J., & Tremaine, S. 1987, Galactic dynamics (Princeton, NJ: Princeton University Press)

Bolatto, A. D., Simon, J. D., Leroy, A., & Blitz, L. 2002, ApJ, 565, 238

Bontekoe, T. R., & van Albada, T. S. 1987, MNRAS, 224, 349

Bosma, A. 1981, AJ, 86, 1791

Brainerd, T. G. 2004, ApJL, submitted, (astro-ph/0408559)

Briggs, F. H. 1990, ApJ, 352, 15

Broeils, A. H., & Courteau, S. 1997, in ASP Conf. Ser. 117: Dark and Visible Matter in Galaxies and Cosmological Implications, 74

Brown, M. L., Taylor, A. N., Hambly, N. C., & Dye, S. 2002, MNRAS, 333, 501

Bullock, J. S. 2002, in "The shapes of galaxies and their dark halos", Proceedings of the Yale Cosmology Workshop, New Haven, Connecticut, USA, 28-30 May 2001., ed. Priyamvada Natarajan (Singapore: World Scientific), 109

Bullock, J. S., Dekel, A., Kolatt, T. S., Kravtsov, A. V., Klypin, A. A., Porciani, C., & Primack, J. R. 2001, ApJ, 555, 240

Buote, D. A., Jeltema, T. E., Canizares, C. R., & Garmire, G. P. 2002, ApJ, 577, 183

Bureau, M., Freeman, K. C., Pfitzner, D. W., & Meurer, G. R. 1999, AJ, 118, 2158, (BFPM99)

Byrd, G. G., Saarinen, S., & Valtonen, M. J. 1986, MNRAS, 220, 619

Cabanela, J. E., & Dickey, J. M. 1999, AJ, 118, 46

Catelan, P., Kamionkowski, M., & Blandford, R. D. 2001, MNRAS, 320, L7

Chambers, S. W., Melott, A. L., & Miller, C. J. 2002, ApJ, 565, 849

Chandrasekhar, S. 1943, ApJ, 97, 255

Chluba, J., & Mannheim, K. 2002, A&A, 396, 419

Christodoulou, D. M., Tohline, J. E., & Steiman-Cameron, T. Y. 1993, ApJ, 416, 74

Clowe, D., Gonzalez, A., & Markevitch, M. 2004, ApJ, 604, 596

Colberg, J. M., White, S. D. M., Jenkins, A., & Pearce, F. R. 1999, MNRAS, 308, 593

Cole, S., & Lacey, C. 1996, MNRAS, 281, 716

Colpi, M., Mayer, L., & Governato, F. 1999, ApJ, 525, 720

Cooray, A., & Chen, X. 2002, ApJ, 573, 43

Copernicus, N. 1543, De Revolutionibus Orbium Cœlestium (Nuremberg: Ioh. Peterium)

Cora, S. A., Muzzio, J. C., & Vergne, M. M. 1997, MNRAS, 289, 253

Courteau, S., & Rix, H. 1999, ApJ, 513, 561

Creze, M., Chereul, E., Bienayme, O., & Pichon, C. 1998, A&A, 329, 920

Crittenden, R. G., Natarajan, P., Pen, U., & Theuns, T. 2001, ApJ, 559, 552

Croft, R. A. C., & Metzler, C. A. 2000, ApJ, 545, 561

Davis, M., Efstathiou, G., Frenk, C. S., & White, S. D. M. 1985, ApJ, 292, 371

de Blok, W. J. G., & McGaugh, S. S. 1997, MNRAS, 290, 533

De Lucia, G., Kauffmann, G., Springel, V., White, S. D. M., Lanzoni, B., Stoehr, F., Tormen, G., & Yoshida, N. 2004, MNRAS, 348, 333

Debattista, V. P., & Sellwood, J. A. 1999, ApJ, 513, L107

Dehnen, W., & Binney, J. 1998, MNRAS, 294, 429, (DB98)

Dekel, A., & Shlosman, I. 1983, in IAU Symp. 100: Internal Kinematics and Dynamics of Galaxies, 187

Dinescu, D. I., Keeney, B. A., Majewski, S. R., & Girard, T. M. 2004, AJ, 128, 687

Diplas, A., & Savage, B. D. 1991, ApJ, 377, 126

Doroshkevich, A. G. 1970, Astrofizika, 6, 581

Drimmel, R., Smart, R. L., & Lattanzi, M. G. 2000, A&A, 354, 67

Dubinski, J. 1992, ApJ, 401, 441

—. 1994, ApJ, 431, 617

Dubinski, J., & Carlberg, R. G. 1991, ApJ, 378, 496

Ebeling, H., Barrett, E., & Donovan, D. 2004, ApJ, 609, L49

Eisenstein, D. J., & Loeb, A. 1995, ApJ, 439, 520

Eke, V. R., Navarro, J. F., & Frenk, C. S. 1998, ApJ, 503, 569

Faltenbacher, A., Gottlöber, S., Kerscher, M., & Müller, V. 2002, A&A, 395, 1, (F02)

Flin, P. 1987, MNRAS, 228, 941

Franx, M., Illingworth, G., & de Zeeuw, T. 1991, ApJ, 383, 112

Freeman, K. C. 1966, MNRAS, 134, 1

Frenk, C. S., White, S. D. M., Davis, M., & Efstathiou, G. 1988, ApJ, 327, 507

Frenk, C. S., White, S. D. M., Efstathiou, G., & Davis, M. 1985, Nature, 317, 595

Gao, L., White, S. D. M., Jenkins, A., Stoehr, F., & Springel, V. 2004, MNRAS, submitted, (astro-ph/0404589v2)

García-Ruiz, I., Kuijken, K., & Dubinski, J. 2002, MNRAS, 337, 459

Gerhard, O. E. 1983, MNRAS, 202, 1159

Gerhard, O. E., & Binney, J. 1985, MNRAS, 216, 467

Ghigna, S., Moore, B., Governato, F., Lake, G., Quinn, T., & Stadel, J. 2000, ApJ, 544, 616

Gill, S. P. D., Knebe, A., & Gibson, B. K. 2004, MNRAS, 351, 399

Gottbrath, C., Bailin, J., Meakin, C., Thompson, T., & Charfman, J. J. 1999, preprint (astro-ph/9912202)

Grimm, J., & Grimm, H. 1857, Kinder- und Hausmärchen (Berlin), no. 53

Han, C., Gould, A., & Sackett, P. D. 1995, ApJ, 445, 46

Hartwick, F. D. A. 2000, AJ, 119, 2248

Hatton, S., & Ninin, S. 2001, MNRAS, 322, 576

Heavens, A., Refregier, A., & Heymans, C. 2000, MNRAS, 319, 649

Helmi, A. 2004a, MNRAS, 351, 643

—. 2004b, ApJ, 610, L97

Helmi, A., & White, S. D. M. 2001, MNRAS, 323, 529

Hernquist, L. 1993, ApJS, 86, 389

Heyl, J. S., Hernquist, L., & Spergel, D. N. 1994, ApJ, 427, 165

Hoekstra, H., Yee, H. K. C., & Gladders, M. D. 2004, ApJ, 606, 67

Hofner, P., & Sparke, L. S. 1994, ApJ, 428, 466

Holmberg, E. 1974, Arkiv for Astronomi, 5, 305

Hoyle, F. 1949, in Problems of Cosmical Aero-dynamics, Proceedings of the Symposium on the Motion of Gaseous Masses of Cosmical Dimensions held at Paris, August 16-19, 1949 (Ohio: Central Air Documents Office), 195–197

Huang, S., & Carlberg, R. G. 1997, ApJ, 480, 503

Hunter, C., & Toomre, A. 1969, ApJ, 155, 747

Ibata, R., Lewis, G. F., Irwin, M., Totten, E., & Quinn, T. 2001, ApJ, 551, 294

Ibata, R. A., Gilmore, G., & Irwin, M. J. 1994, Nature, 370, 194

Ibata, R. A., & Lewis, G. F. 1998, ApJ, 500, 575

Ibata, R. A., & Razoumov, A. O. 1998, A&A, 336, 130

Ibata, R. A., Wyse, R. F. G., Gilmore, G., Irwin, M. J., & Suntzeff, N. B. 1997, AJ, 113, 634

Ideta, M., Hozumi, S., Tsuchiya, T., & Takizawa, M. 2000, MNRAS, 311, 733

Jiang, I., & Binney, J. 2000, MNRAS, 314, 468

Jing, Y. P., & Suto, Y. 2000, ApJ, 529, L69

—. 2002, ApJ, 574, 538, (JS02)

Johnston, K. V., Law, D. R., & Majewski, S. R. 2004, ApJ, submitted, (astro-ph/0407565)

Johnston, K. V., Zhao, H., Spergel, D. N., & Hernquist, L. 1999, ApJ, 512, L109

Kahn, F. D., & Woltjer, L. 1959, ApJ, 130, 705

Kashikawa, N., & Okamura, S. 1992, PASJ, 44, 493

Kasun, S. F., & Evrard, A. E. 2004, ApJ, submitted, (astro-ph/0408056) (KE04)

Katz, N. 1991, ApJ, 368, 325

Kazantzidis, S., Kravtsov, A. V., Zentner, A. R., Allgood, B., Nagai, D., & Moore, B. 2004, ApJ, 611, L73

Klypin, A., Kravtsov, A. V., Valenzuela, O., & Prada, F. 1999, ApJ, 522, 82

Knebe, A., Gill, S. P. D., Gibson, B. K., Lewis, G. F., Ibata, R. A., & Dopita, M. A. 2004, ApJ, 603, 7

Kroupa, P., & Bastian, U. 1997, New Astronomy, 2, 77

Kuijken, K. 1991, ApJ, 376, 467

López-Corredoira, M., Betancort-Rijo, J., & Beckman, J. E. 2002a, A&A, 386, 169

López-Corredoira, M., Cabrera-Lavers, A., Garzón, F., & Hammersley, P. L. 2002b, A&A, 394, 883

Law, D. R., Johnston, K. V., & Majewski, S. R. 2004, ApJ, submitted, (astro-ph/0407566)

Lee, J., & Pen, U. 2000, ApJ, 532, L5

—. 2001, ApJ, 555, 106

Lee, J., & Suto, Y. 2003, ApJ, 585, 151

—. 2004, ApJ, 601, 599

Lin, D. N. C. 1996, in Gravitational dynamics, Proceedings of the 36th Herstmonceux conference (Cambridge: Cambridge Univeristy Press), 15

Lynden-Bell, D. 1965, MNRAS, 129, 299

—. 1976, MNRAS, 174, 695

Lynden-Bell, D., & Lynden-Bell, R. M. 1995, MNRAS, 275, 429

Maller, A. H., Dekel, A., & Somerville, R. 2002, MNRAS, 329, 423

Martin, N. F., Ibata, R. A., Bellazzini, M., Irwin, M. J., Lewis, G. F., & Dehnen, W. 2004, MNRAS, 348, 12

Martínez-Delgado, D., Gómez-Flechoso, M. Á., Aparicio, A., & Carrera, R. 2004, ApJ, 601, 242

Masset, F. S., & Bureau, M. 2003, ApJ, 586, 152, (MB03)

Mateo, M. L. 1998, ARA&A, 36, 435

Merrifield, M. R. 1992, AJ, 103, 1552

Meurer, G. R., Carignan, C., Beaulieu, S. F., & Freeman, K. C. 1996, AJ, 111, 1551

Milgrom, M. 1983, ApJ, 270, 365

Moore, B., Ghigna, S., Governato, F., Lake, G., Quinn, T., Stadel, J., & Tozzi, P. 1999, ApJ, 524, L19

Natarajan, P., & Refregier, A. 2000, ApJ, 538, L113

Navarro, J. F., Abadi, M. G., & Steinmetz, M. 2004, ApJ, 613, L41

Navarro, J. F., Frenk, C. S., & White, S. D. M. 1996, ApJ, 462, 563

—. 1997, ApJ, 490, 493

Nelson, R. W., & Tremaine, S. 1995, MNRAS, 275, 897

Newton, I. 1687, Philosophiae naturalis principia mathematica (London: The Royal Society and Joseph Streater)

Nyquist, H. 1928, AIEE Transactions, 47, 617

Okumura, S. K., Makino, J., Ebisuzaki, T., Fukushige, T., Ito, T., Sugimoto, D., Hashimoto, E., Tomida, K., & Miyakawa, N. 1993, PASJ, 45, 329

Olling, R. P., & Merrifield, M. R. 2000, MNRAS, 311, 361

Onuora, L. I., & Thomas, P. A. 2000, MNRAS, 319, 614

Ostriker, E. C., & Binney, J. J. 1989, MNRAS, 237, 785

Palma, C., Majewski, S. R., & Johnston, K. V. 2002, ApJ, 564, 736

Peñarrubia, J., Kroupa, P., & Boily, C. M. 2002, MNRAS, 333, 779

Peebles, P. J. E. 1969, ApJ, 155, 393

Pen, U., Lee, J., & Seljak, U. 2000, ApJ, 543, L107

Pfenniger, D., Combes, F., & Martinet, L. 1994, A&A, 285, 79

Pfitzner, D. W. 1999, PhD thesis, Australian National University, (P99)

Piatek, S., Pryor, C., Olszewski, E. W., Harris, H. C., Mateo, M., Minniti, D., Monet, D. G., Morrison, H., & Tinney, C. G. 2002, AJ, 124, 3198

Piatek, S., Pryor, C., Olszewski, E. W., Harris, H. C., Mateo, M., Minniti, D., & Tinney, C. G. 2003, AJ, 126, 2346

Plionis, M. 1994, ApJS, 95, 401

Plummer, H. C. 1911, MNRAS, 71, 460

Porciani, C., Dekel, A., & Hoffman, Y. 2002a, MNRAS, 332, 325

—. 2002b, MNRAS, 332, 339

Power, C., Navarro, J. F., Jenkins, A., Frenk, C. S., White, S. D. M., Springel, V., Stadel, J., & Quinn, T. 2003, MNRAS, 338, 14

Quinn, T., & Binney, J. 1992, MNRAS, 255, 729

Reed, B. C. 1996, AJ, 111, 804

Reshetnikov, V., & Combes, F. 1998, A&A, 337, 9

Rhee, G. F. R. N., & Katgert, P. 1987, A&A, 183, 217

Sackett, P. D. 1997, ApJ, 483, 103

Sackett, P. D. 1999, in ASP Conf. Ser. 182: Galaxy Dynamics - A Rutgers Symposium, ed. D. Merritt, J. A. Sellwood, & M. Valluri, 393

Sackett, P. D., Rix, H., Jarvis, B. J., & Freeman, K. C. 1994, ApJ, 436, 629

Sales, L., & Lambas, D. G. 2004, MNRAS, 348, 1236

Sanders, R. H., & McGaugh, S. S. 2002, ARA&A, 40, 263

Sarazin, C. L. 1986, Rev. Mod. Phys., 58, 1

Schwarzkopf, U., & Dettmar, R.-J. 2001, A&A, 373, 402

Schwarzschild, M. 1982, ApJ, 263, 599

Schweitzer, A. E. 1996, PhD thesis, University of Wisconsin-Madison

Schweitzer, A. E., Cudworth, K. M., Majewski, S. R., & Suntzeff, N. B. 1995, AJ, 110, 2747

Sharma, S., & Steinmetz, M. 2004, ApJ, submitted, (astro-ph/0406533)

Simon, J. D., Bolatto, A. D., Leroy, A., & Blitz, L. 2003, ApJ, 596, 957

Sparke, L. S., & Casertano, S. 1988, MNRAS, 234, 873

Spergel, D. N., Verde, L., Peiris, H. V., Komatsu, E., Nolta, M. R., Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Page, L., Tucker, G. S., Weiland, J. L., Wollack, E., & Wright, E. L. 2003, ApJS, 148, 175

Splinter, R. J., Melott, A. L., Linn, A. M., Buck, C., & Tinker, J. 1997, ApJ, 479, 632

Springel, V., Yoshida, N., & White, S. D. M. 2001, New Astronomy, 6, 79

Stark, A. A. 1977, ApJ, 213, 368

Steinmetz, M. 1996, MNRAS, 278, 1005

Strömberg, G. 1934, ApJ, 79, 460

Struble, M. F., & Peebles, P. J. E. 1985, AJ, 90, 582

Subramanian, K. 1988, MNRAS, 234, 459

Sugerman, B., Summers, F. J., & Kamionkowski, M. 2000, MNRAS, 311, 762

Sunyaev, R. A., & Zeldovich, I. B. 1980, ARA&A, 18, 537

Taylor, J. E., & Babul, A. 2001, ApJ, 559, 716

—. 2004, MNRAS, 348, 811

Toomre, A. 1964, ApJ, 139, 1217

Toomre, A. 1983, in IAU Symp. 100: Internal Kinematics and Dynamics of Galaxies, 177

Tormen, G. 1997, MNRAS, 290, 411

Tremaine, S., & Yu, Q. 2000, MNRAS, 319, 1

Tsuchiya, T. 2002, New Astronomy, 7, 293

Tully, R. B., & Fisher, J. R. 1977, A&A, 54, 661

Ulmer, M. P., McMillan, S. L. W., & Kowalski, M. P. 1989, ApJ, 338, 711

Valenzuela, O., & Klypin, A. 2003, MNRAS, 345, 406

van den Bosch, F. C., Abel, T., Croft, R. A. C., Hernquist, L., & White, S. D. M. 2002, ApJ, 576, 21

van Haarlem, M., & van de Weygaert, R. 1993, ApJ, 418, 544

Vitvitska, M., Klypin, A. A., Kravtsov, A. V., Wechsler, R. H., Primack, J. R., & Bullock, J. S. 2002, ApJ, 581, 799

Warren, M. S., Quinn, P. J., Salmon, J. K., & Zurek, W. H. 1992, ApJ, 399, 405, (W92)

Weinberg, M. D. 1998, MNRAS, 299, 499

West, M. J., Jones, C., & Forman, W. 1995, ApJ, 451, L5

West, M. J., Villumsen, J. V., & Dekel, A. 1991, ApJ, 369, 287

White, S. D. M. 1984, ApJ, 286, 38

White, S. D. M., & Rees, M. J. 1978, MNRAS, 183, 341

Zaritsky, D., Smith, R., Frenk, C. S., & White, S. D. M. 1997, ApJ, 478, L53

Zavala, J., Avila-Reese, V., Hernández-Toledo, H., & Firmani, C. 2003, A&A, 412, 633

Zeldovich, Y. B. 1970, A&A, 5, 84