From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

Doctoral Dissertations, 2011-2026 Astrophysical Sciences

2014

Constraints on Cosmology and the Physics of the Intracluster Medium from Secondary Anisotropies in the Cosmic Microwave Background

Hill, James Colin

Princeton, NJ : Princeton University

http://arks.princeton.edu/ark:/88435/dsp016t053j213

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Constraints on Cosmology and the Physics of the Intracluster Medium from Secondary Anisotropies in the Cosmic Microwave Background

James Colin Hill

A Dissertation
Presented to the Faculty
of Princeton University
in Candidacy for the Degree
of Doctor of Philosophy

Recommended for Acceptance
by the Department of
Astrophysical Sciences
Adviser: Prof. David N. Spergel

September 2014

© Copyright by James Colin Hill, 2014.

All rights reserved.

Abstract

This thesis presents a number of new statistical approaches to constrain cosmology and the physics of the intracluster medium (ICM) using measurements of secondary anisotropies in the cosmic microwave background (CMB), the remnant thermal radiation left over from the Big Bang that still permeates the universe today. In particular, I consider the thermal Sunyaev-Zel'dovich (tSZ) effect, in which the hot, ionized gas in massive galaxy clusters casts a "shadow" against the CMB, and gravitational lensing of the CMB, in which the gravity of intervening structures bends the path of CMB photons as they travel to us from the early universe. I use data from the Planck satellite to obtain the first detection of the tSZ – CMB lensing cross-power spectrum, a signal which is sensitive to a previously unstudied population of high-redshift, low-mass groups and clusters. I combine this measurement with Planck measurements of the tSZ power spectrum to obtain cosmological and astrophysical constraints, and also investigate the ability of future tSZ power spectrum measurements to constrain primordial non-Gaussianity and the sum of the neutrino masses. In addition to these power spectrum measurements, I also present new techniques to obtain cosmological and astrophysical constraints from non-Gaussian statistics of the tSZ signal, starting with the tSZ skewness and proceeding to consider the entire one-point PDF of the tSZ field. I apply these methods to data from the Atacama Cosmology Telescope, yielding some of the tightest constraints yet achieved on the amplitude of cosmic matter density fluctuations. Moreover, these methods allow the long-standing degeneracy between cosmology and ICM physics to be directly broken with tSZ data alone for the first time.

Acknowledgements

I am deeply indebted to a great number of people without whom this thesis would not exist. First, I am grateful to my advisor, David Spergel, for his guidance and direction over the past five years. He has provided me with countless interesting ideas to work on, and more importantly has helped me to develop the ability and expertise necessary to formulate such ideas independently. His perspectives on cosmology, physics, and science as a whole have greatly shaped my own. I feel extremely lucky to have had the opportunity to collaborate with him for the past five years.

I also feel extremely fortunate to have been able to learn cosmology in Princeton, surrounded by scientists who founded the field and are amongst its greatest practitioners. I am very grateful to those who have shared their time and expertise, including Matias Zaldarriaga, Kendrick Smith, Lyman Page, Tobias Marriage, Rachel Mandelbaum, Enrico Pajer, Jerry Ostriker, and Michael Strauss. I am likewise thankful to the members of the Atacama Cosmology Telescope collaboration, including David, Lyman, Suzanne Staggs, and many postdocs and graduate students, who provided me with the opportunity to work at the forefront of cosmic microwave background research.

In addition, I am very grateful to those who encouraged my interest in physics and mathematics earlier in my career, and allowed me to start doing serious research. This includes Max Tegmark, Kenneth Rines, and Scott Hughes from my undergraduate years at MIT, as well as John Nord and a number of deeply committed math and science teachers from my earlier years in Spokane. Without their support and instruction, I never would have made it to this point.

Beyond these mentors, I am very thankful for the friends who have not only kept me sane during graduate school, but have also been responsible for some of the most fun years of my life. Special amongst these people is David Reinecke, who is responsible for most of what I now know about the social sciences, guitars, and the intersection of the two. Playing music with David, Keith McKnight, Noah Buckley-Farlee, Mike Fleder, and Manish Nag has been

the best part of my non-physics life over the past few years. I am also thankful for the great friends I have made within the physics world, including Kendrick, Nick Battaglia, Marilena Loverde, Laura Newburgh, Alex Dahlen, Elisa Chisari, David McGady, Tim Brandt, and many others, who make the world of astrophysics a fun place to be. A special mention is due to Blake Sherwin, who spans all of these categories and has been both a great friend and a fantastic collaborator for several years. Talking with all of these people about physics, code, math, and solving problems in general has been as crucial to my success as reading any textbook.

Finally, and most importantly, I am deeply thankful to my family, who have encouraged and supported me in everything I have done. In some ways my entire career has simply been a series of dominoes falling that began with opening up a few popular physics books my parents gave me when I was young. I am thankful for the environment my parents and my brothers created that always pushed me to strive to do the best I possibly could. Their support over the years is responsible for all I have achieved, and I dedicate this thesis to them.

For my parents.

Relation to Published Work

This thesis consists of five chapters. The first chapter is primarily introductory and provides background on the cosmic microwave background (CMB) and the Sunyaev-Zel'dovich (SZ) effect. However, it also includes derivations of useful statistics of the thermal SZ effect, which will mostly be used in Chapter 2, but provide the general theoretical framework underlying much of the thesis. These results exist elsewhere in the literature, but I am not aware of completely equivalent derivations in any other work, and thus some of this chapter is original material.

Chapters two through five comprise the main part of this thesis. Chapters two through four have been published in peer-reviewed journals (as has the appendix). Chapter five will soon be submitted for publication as well; it is currently under review within the Atacama Cosmology Telescope collaboration. In this section, I will describe the publication or submission status of each chapter and — as each chapter is the result of collaborative research — will detail my own contribution to each of the chapters.

Chapter 2: Cosmology from the Thermal Sunyaev-Zel'dovich Power Spectrum: Primordial non-Gaussianity and Massive Neutrinos

Hill, J. C., & Pajer, E. 2013, Physical Review D, 88, 063526

Nearly all of the results and text in this chapter are my own work, with the sole exception of Section 2.5, which was primarily written by Enrico Pajer, though with significant input from me. I wrote the underlying code for all of the thermal SZ calculations and worked

out all of the relevant theoretical derivations presented in the chapter (with additional details in Chapter 1). Enrico and I both independently performed the Fisher calculations in order to cross-check each other's results. The project originated in conversations with Enrico, David Spergel, Marilena LoVerde, and Kendrick Smith, but it was Enrico's interest in forecasting results for the PIXIE experiment that led to the chapter's eventual completion.

Chapter 3: Detection of Thermal SZ – CMB Lensing Cross-Correlation in Planck Nominal Mission Data

Hill, J. C., & Spergel, D. N. 2014, Journal of Cosmology and Astroparticle Physics, 02, 030

The work described in this chapter was performed almost entirely by myself, with guidance and assistance from David Spergel. I wrote the theoretical code for thermal SZ and CMB lensing calculations, the code to analyze Planck data, the code for computing cosmological and astrophysical constraints, and the text. David provided many suggestions and directions throughout the project, in particular with regards to the data analysis methods. He also edited and contributed to the text. The idea for the project was jointly developed between David and myself.

Chapter 4: Cosmological Constraints from Moments of the Thermal Sunyaev-Zel'dovich Effect

Hill, J. C., & Sherwin, B. D. 2013, Physical Review D, 87, 023527

This chapter is essentially entirely my own work. I wrote all of the code and performed all of the calculations, with the exception of the simulation analysis, which was performed by Blake Sherwin. I wrote all of the text, with some editing and contributions from Blake. The underlying idea grew out of conversations with Blake and David Spergel.

Chapter 5: The Atacama Cosmology Telescope: A Measurement of the Thermal Sunyaev-Zel'dovich One-Point PDF

Hill, J. C., Sherwin, B. D., Smith, K. M., et al. 2014, to be submitted to the Astrophysical Journal

This chapter is primarily my own work, with additional contributions from Blake Sherwin, Kendrick Smith, and David Spergel. The theoretical background was developed in collaboration with Kendrick; I then implemented these results and wrote the necessary code for the theoretical calculations. The measurement using Atacama Cosmology Telescope data directly followed from the results presented in Appendix A and was done in collaboration with Blake. The simulation pipeline was a collaborative effort combining tools developed by both myself and Blake, as well as Sudeep Das. Finally, I wrote the likelihood code for the cosmological and astrophysical interpretation. The text is almost entirely my own, with editing and contributions from Blake, Kendrick, and David. The idea for the project grew out of the results presented in Appendix A, and was mostly developed on my own and in discussions with Kendrick and David.

Appendix A: The Atacama Cosmology Telescope: A Measurement of the Thermal Sunyaev-Zel'dovich Effect Using the Skewness of the CMB Temperature Distribution

Wilson, M. J., Sherwin, B. D., Hill, J. C., et al. 2012, Physical Review D, 86, 122005

This appendix was a joint effort between Michael Wilson (an Oxford master's student who visited Princeton in summer 2011), Blake Sherwin, and myself, under the guidance of David Spergel. Michael and Blake wrote the data analysis code and performed the measurement. I wrote the thermal SZ code to perform the relevant theoretical calculations and worked with Blake on the interpretation of the measurement. I wrote the section describing the theoretical work and contributed to the section describing the interpretation, while Blake and Michael wrote most of the remainder, with contributions and editing from myself. Other members of the Atacama Cosmology Telescope collaboration contributed comments and suggestions.

Contents

Abstract iii

Acknowledgements iv

Relation to Published Work vii

List of Tables xv

List of Figures xvii

1 Introduction: The Sunyaev-Zel'dovich Effect **1**

1.1 Introduction 1

1.2 The Thermal Sunyaev-Zel'dovich Effect 5

1.2.1 Background 5

1.2.2 Basic Definitions 6

1.2.3 Halo Model Derivation of tSZ Statistics 7

1.2.4 The One-Halo Term 10

1.2.5 The Two-Halo Term 14

1.2.6 The Covariance Matrix 18

2 Cosmology from the Thermal Sunyaev-Zel'dovich Power Spectrum: Primordial non-Gaussianity and Massive Neutrinos **23**

2.1 Abstract 23

2.2 Introduction 24

2.3 Modeling Large-Scale Structure 32

2.3.1 Halo Mass Function . . . 33
2.3.2 Halo Bias . . . 39
2.4 Thermal SZ Power Spectrum . . . 41
2.4.1 Halo Model Formalism . . . 42
2.4.2 Modeling the ICM . . . 45
2.4.3 Parameter Dependences . . . 50
2.5 Experimental Considerations . . . 56
2.5.1 Multifrequency Subtraction . . . 59
2.5.2 Foregrounds . . . 63
2.5.3 Noise After Multifrequency Subtraction . . . 66
2.6 Covariance Matrix of the tSZ Power Spectrum . . . 68
2.7 Forecasted Constraints . . . 80
2.7.1 $f_{\rm NL}$ . . . 83
2.7.2 M_ν . . . 85
2.7.3 Other Parameters . . . 89
2.7.4 Forecasted SNR . . . 90
2.8 Discussion and Outlook . . . 93
2.9 Comparison with Planck Results . . . 96
2.10 Acknowledgments . . . 97

3 Detection of Thermal SZ – CMB Lensing Cross-Correlation in Planck Nominal Mission Data 108

3.1 Abstract . . . 108
3.2 Introduction . . . 109
3.3 Theory . . . 114
3.3.1 Thermal SZ Effect . . . 114
3.3.2 CMB Lensing . . . 119
3.3.3 Power Spectra . . . 120

3.4 Thermal SZ Reconstruction . . . 130
3.4.1 Data and Cuts . . . 130
3.4.2 ILC . . . 133
3.4.3 Thermal SZ Auto-Power Spectrum . . . 140
3.5 Thermal SZ – CMB Lensing Cross-Power Spectrum . . . 142
3.5.1 CMB Lensing Potential Map . . . 142
3.5.2 Measurement . . . 145
3.5.3 CIB Contamination Correction . . . 147
3.5.4 Null Tests and Systematic Errors . . . 152
3.6 Interpretation . . . 160
3.6.1 ICM Constraints . . . 160
3.6.2 Cosmological Constraints . . . 165
3.6.3 Simultaneous Constraints on Cosmology and the ICM . . . 173
3.7 Discussion and Outlook . . . 176
3.8 Acknowledgments . . . 177

4 Cosmological Constraints from Moments of the Thermal Sunyaev-Zel'dovich Effect 185

4.1 Abstract . . . 185
4.2 Introduction . . . 186
4.3 Calculating tSZ Moments . . . 188
4.3.1 Background . . . 188
4.3.2 Results . . . 196
4.4 Isolating the Dependence on ICM Astrophysics . . . 201
4.5 Application to ACT and SPT Data . . . 204
4.6 Future Cosmological Constraints . . . 208
4.6.1 Extension to other parameters . . . 208
4.6.2 Estimating the constraints from *Planck* . . . 213

4.7 Acknowledgments . . . 214
4.8 Addendum . . . 214
4.8.1 Directly Canceling the Gas Physics . . . 214
4.8.2 Higher-Order tSZ Moments . . . 220

5 The Atacama Cosmology Telescope: A Measurement of the Thermal Sunyaev-Zel'dovich One-Point PDF 229

5.1 Abstract . . . 229
5.2 Introduction . . . 230
5.3 Data . . . 236
5.4 Theory . . . 239
5.4.1 Thermal SZ Effect . . . 239
5.4.2 Noiseless tSZ-Only PDF . . . 241
5.4.3 Noisy PDF . . . 245
5.4.4 Covariance Matrix . . . 256
5.5 Simulations . . . 259
5.6 Interpretation . . . 263
5.6.1 Likelihood Function . . . 263
5.6.2 Constraints . . . 269
5.7 Discussion and Outlook . . . 274
5.8 Acknowledgments . . . 276

A The Atacama Cosmology Telescope: A Measurement of the Thermal Sunyaev-Zel'dovich Effect Using the Skewness of the CMB Temperature Distribution 283

A.1 Abstract . . . 283
A.2 Introduction . . . 284
A.3 Skewness of the tSZ Effect . . . 287

A.4 Map Processing 290
A.4.1 Filtering the Maps 290
A.4.2 Removing Point Sources 291
A.5 Results 294
A.5.1 Evaluating the Skewness 294
A.5.2 The Origin of the Signal 296
A.5.3 Testing for Systematic Infrared Source Contamination 298
A.6 Cosmological Interpretation 300
A.7 Conclusions 303
A.8 Acknowledgments 304

List of Tables

2.1 For the nine frequency bands for Planck we report the central frequency (in GHz), the Full Width at Half Maximum (FWHM, in arcminutes, to be converted into radians in the noise computation) of each pixel, and the 1σ sensitivity to temperature per square pixel [23]. 63

3.1 Measured tSZ – CMB lensing cross-spectrum bandpowers. All values are dimensionless. These bandpowers have been corrected for contamination from CIB emission following the procedure described in Section 3.5.3; they correspond to the blue squares shown in Fig. 3.10. Uncertainties due to the CIB subtraction have been propagated into the final errors provided here. Note that the bins have been chosen to match those in [48]. 153

3.2 Channel map weights computed by our ILC pipeline for three different sky cuts (see Eqs. (3.22) and (3.26)). All values are in units of $\mathrm{K}_{\mathrm{CMB}}^{-1}$. The fiducial case used throughout this chapter is $f_{\mathrm{sky}} = 0.3$. Note that these f_{sky} values are computed by simply thresholding the 857 GHz map (see Section 3.4.1); the final sky fractions used in each case are smaller than these values due to the inclusion of the point source mask and S1 and S2 season map masks. The weights change slightly as f_{sky} is increased as the ILC adjusts to remove the increased Galactic dust emission. 157

4.1 Amplitudes and power-law scalings with σ_8 for the tSZ variance and skewness, as defined in Eq. (4.5). The first column lists the pressure profile used in the calculation (note that all calculations use the Tinker mass function). The amplitudes are specified at $\sigma_8 = 0.817$, the WMAP5 maximum-likelihood value. All results are computed at $\nu = 150$ GHz. 193

4.2 Amplitudes and power-law scalings with σ_8 for the tSZ kurtosis. The first column lists the pressure profile used in the calculation (note that all calculations use the Tinker mass function). The amplitudes are specified at $\sigma_8 = 0.817$, the WMAP5 maximum-likelihood value. All results are computed at $\nu = 150$ GHz. 222

A.1 Constraints on σ_8 derived from our skewness measurement using two different simulations and three different scalings of the skewness and its variance with σ_8. The top row lists the simulations used to calculate the expected skewness for $\sigma_8 = 0.8$ [35, 33]; the left column lists the pressure profiles used to calculate the scaling of the skewness and its variance with σ_8 [16, 17, 25]. The errors on σ_8 shown are the 68% and 95% confidence levels. 303

List of Figures

2.1 This plot shows the unmasked tSZ power spectrum for our fiducial model (black curves), as specified in Section 2.4.3, as well as variations with $f_{\rm NL} = 100$ (blue curves) and $f_{\rm NL} = -100$ (red curves). $f_{\rm NL}$ values of this magnitude are highly disfavored by current constraints, but we plot them here to clarify the influence of primordial non-Gaussianity on the tSZ power spectrum. The effect of $f_{\rm NL}$ on the one-halo term is simply an overall amplitude shift due to the corresponding increase or decrease in the number of massive clusters in the universe, as described in Section 2.3.1. The effect of $f_{\rm NL}$ on the two-halo term includes not only an amplitude shift due to the change in the mass function, but also a steep upturn at low-ℓ due to the influence of the scale-dependent halo bias, as described in Section 2.3.2. Note that for $f_{\rm NL} < 0$ the two effects cancel for $\ell \approx 4 - 5$. The relative smallness of the two-halo term (compared to the one-halo term) makes the scale-dependent bias signature subdominant for all ℓ values except $\ell \lesssim 7 - 8$. However, masking of nearby massive clusters suppresses the low-ℓ one-halo term in the power spectrum (in addition to decreasing the cosmic variance, as discussed in Section 2.6), which increases the relative importance of the two-halo term and thus the dependence of the total signal on $f_{\rm NL}$ at low-ℓ. 46

2.2 This plot shows the unmasked tSZ power spectrum for our fiducial model (black curves), as specified in Section 2.4.3, as well as variations with $M_\nu = 0.05$ eV (blue curves) and $M_\nu = 0.10$ eV (red curves). M_ν values of this magnitude are at the lower bound allowed by neutrino oscillation measurements, and thus an effect of at least this magnitude is expected in our universe. The effect of M_ν on both the one- and two-halo terms is effectively an overall amplitude shift, although the effect tapers off very slightly at high-ℓ. It may be puzzling at first to see an increase in the tSZ power when $M_\nu > 0$, but the key fact is that we are holding σ_8 constant when varying M_ν (indeed, we hold all of the other parameters constant). In order to keep σ_8 fixed despite the late-time suppression of structure growth due to $M_\nu > 0$, we must increase the primordial amplitude of scalar perturbations, A_s. Thus, the net effect is an increase in the tSZ power spectrum amplitude, counterintuitive though it may be. If our ΛCDM parameter set included A_s rather than σ_8, and we held A_s constant while $M_\nu > 0$, we would indeed find a corresponding decrease in the tSZ power (and in σ_8, of course). 47

2.3 The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $f_{\rm NL} = \pm 100$. 57

2.4 The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $M_\nu = 0.05$ eV and 0.10 eV. 57

2.5 The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $\Omega_b h^2 = 0.02262$ and 0.02218. 57

2.6 The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $\Omega_c h^2 = 0.11575$ and 0.11345. 57

2.7 The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $\Omega_\Lambda = 0.7252$ and 0.7108. 58

2.8 The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $\sigma_8 = 0.825$ and 0.809. . 58
2.9 The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $n_s = 0.97425$ and 0.95495. 58
2.10 The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $C_{P_0} = 1.01$ and 0.99. . . 58
2.11 The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $C_\beta = 1.01$ and 0.99. . . 59
2.12 The two plots show the various contributions to the total noise per ℓ and m after multifrequency subtraction for Planck (top panel) and PIXIE (bottom panel). Because of the many frequency channels both experiments can subtract the various foregrounds and the total noise is not significantly different from the instrumental noise alone. 60
2.13 The plots show the weights appearing in Eq. (2.31) for PIXIE (left) and Planck (right) as a function of ν, for $\ell = 30$. 67
2.14 The plot compares the noise $N_\ell[f_{sky}(2\ell + 1)]^{-1/2}$ for Planck (black dashed line) and PIXIE (black dotted line) with the expected tSZ power spectrum $C_\ell^{SZ} \equiv C_\ell^y$ (continuous orange line) at 150 GHz with $f_{sky} = 0.7$. 68
2.15 This plot shows the effect of different masking scenarios on the tSZ power spectrum calculated for our fiducial model. The amplitude of the one-halo term is suppressed relative to the two-halo term, leading to an increase in the multipole where their contributions are equal. In addition there is an overall decrease in the total signal at $\ell \lesssim 200$, as one would expect after masking the massive, nearby clusters that dominate the one-halo term in this regime. . . 71

2.16 This plot shows the fractional difference between the tSZ power spectrum of our fiducial model and three different parameter variations, as labeled in the figure. In addition, we show the 1σ fractional errors on the tSZ power spectrum for each of the three experiments we consider, computed from the square root of the diagonal of the covariance matrix. In this unmasked calculation, the overwhelming influence of cosmic variance due to the tSZ trispectrum is clearly seen at low-ℓ. 73

2.17 This plot is identical to Fig. 2.16 except that it has been computed for the ROSAT-masked scenario (see Section 2.6). Compared to Fig. 2.16, it is clear that the errors at low-ℓ for PIXIE and the CV-limited case have been dramatically reduced by the masking procedure. The errors from Planck are less affected because it is dominated by instrumental noise over most of its multipole range. 74

2.18 This plot is identical to Fig. 2.16 except that it has been computed for the eROSITA-masked scenario (see Section 2.6). Compared to Fig. 2.16, it is clear that the errors have been significantly reduced at low-ℓ, though again the effect is minimal for Planck. The comparison to the ROSAT-masked case in Fig. 2.17 is more complicated: it is clear that the errors are further reduced for the CV-limited case, but the PIXIE fractional errors actually increase at low-ℓ as a result of the reduction in the amplitude of the signal there due to the masking. However, the PIXIE fractional errors at higher multipoles are in fact slightly reduced compared to the ROSAT-masked case, although it is hard to see by eye in the plot. 75

2.19 This plot shows the square root of the diagonal elements of the covariance matrix of the tSZ power spectrum, as well as the contributions from the Gaussian and non-Gaussian terms in Eq. (2.56). Results are shown for the total in each experiment, as well as the different Gaussian terms in each experiment and the trispectrum contribution that is identical for all three. See the text for discussion. 78

2.20 This plot is identical to Fig. 2.19 except that it has been computed for the ROSAT-masked scenario (see Section 2.6). Results are shown for the total in each experiment, as well as the different Gaussian terms in each experiment and the trispectrum contribution that is identical for all three. See the text for discussion. 79

2.21 This plot is identical to Fig. 2.19 except that it has been computed for the eROSITA-masked scenario (see Section 2.6). Results are shown for the total in each experiment, as well as the different Gaussian terms in each experiment and the trispectrum contribution that is identical for all three. See the text for discussion. 80

2.22 The tables show the estimated unmarginalized and marginalized 1σ error bars for the indicated parameters and for a total of nine different experimental specifications (Planck, PIXIE, and CV-limited) and cluster masking choices (no masking, ROSAT masking, and eROSITA masking). The unmarginalized errors (top table) are derived using only the tSZ power spectrum. The marginalized errors (two central tables) are derived by adding as an external prior the forecasted Planck constraints from the CMB TT power spectrum plus a prior on H_0 from [134]. Since adding the tSZ power spectrum leads to effectively no improvement in the errors on the ΛCDM parameters with respect to the Planck TT priors, we do not show those numbers in the two central tables (marginalized tSZ plus Planck CMB+H_0 priors). They are equal to the numbers in the bottom table, which show the marginalized Planck CMB+H_0 priors, to $\approx 1\%$ precision. Note that all constraints on M_ν are in units of eV, while the other parameters are dimensionless. 84

2.23 The same as in the right central table of Fig. 2.22, but for a fiducial $f_{\rm NL} = 37$, the central value of WMAP9 [37]. A CV-limited experiment could lead to a 3σ detection. 85

2.24 The same as in the left central table of Fig. 2.22, but for a fiducial $M_\nu = 0.1$ eV, similar to the minimum allowed value in the inverted neutrino hierarchy [39]. Note that all constraints on M_ν are in units of eV. 87

2.25 These plots show the effect of strengthening the prior on the ICM gas physics parameters C_{P_0} and C_β on our forecasted 1σ error on the sum of the neutrino masses (vertical axis, in units of eV). The left panel assumes a fiducial neutrino mass of $M_\nu = 0.1$ eV, while the right panel assumes a fiducial value of $M_\nu = 0$ eV. The solid red curves show the results for Planck, assuming (from top to bottom) no masking, ROSAT masking, or eROSITA masking. The blue dashed curves show the results for PIXIE, with the masking scenarios in the same order. The dotted black curves show the results for a CV-limited experiment, again with the masking scenarios in the same order. The horizontal axis in both plots is the value of the 1σ Gaussian prior $\Delta C_{P_0} = \Delta C_\beta \equiv \Delta_{gas}$ placed on the gas physics parameters (our fiducial value is 0.2). Note that we vary the priors simultaneously, keeping them fixed to the same value. The plots demonstrate that even modest improvements in the external priors on C_{P_0} and C_β could lead to significant decreases in the expected error on M_ν from tSZ power spectrum measurements. 88

2.26 This plot shows the cumulative SNR achievable on the tSZ power spectrum for each of nine different experimental and making scenarios. The solid curves display results for a CV-limited experiment, the short dashed curves show results for Planck, and the long dashed curves show results for PIXIE. The different colors correspond to different masking options, as noted in the figure. Note that PIXIE is close to the CV limit over its signal-dominated multipole range. The total SNR using the imminent Planck data is ≈ 35, essentially independent of the masking option used (note that masking can increase the cumulative SNR up to lower multipoles, however, as compared to the unmasked case). 92

3.1 The tSZ – CMB lensing cross-power spectrum computed for our fiducial model (WMAP9 cosmological parameters and the "AGN feedback" ICM pressure profile fit from [29]). The total signal is the green solid curve, while the one-halo and two-halo contributions are the blue and red solid curves, respectively. The plot also shows the total signal computed for variations around our fiducial model: $\pm 5\%$ variations in σ_8 are the dashed green curves, while $\pm 5\%$ variations in Ω_m are the dotted green curves. In both cases, an increase (decrease) in the parameter's value yield an increase (decrease) in the amplitude of the cross-spectrum. 125

3.2 The left panel shows the power spectra of the tSZ and CMB lensing potential fields for our fiducial model, as labeled in the figure. The right panel shows the normalized cross-correlation coefficient of the tSZ and CMB lensing potential signals as a function of multipole (see Eq. (3.15)). The typical normalized correlation is 30–40% over the ℓ range where we measure $C_\ell^{y\phi}$ ($100 < \ell < 1600$). The strength of the correlation increases at smaller angular scales, where both signals are dominated by the one-halo term. 126

3.3 Mass and redshift contributions to the tSZ – CMB lensing cross-spectrum (left panel), tSZ auto-spectrum (middle panel), and CMB lensing potential auto-spectrum (right panel), computed at $\ell = 100$, where the two-halo term dominates the cross-spectrum. Each curve includes contributions from progressively higher mass scales, as labelled in the figures, while the vertical axis encodes the differential contribution from each redshift. At this multipole, the tSZ auto-spectrum is heavily dominated by massive, low-redshift halos, while the CMB lensing auto-spectrum receives contributions from a much wider range of halo masses and redshifts. The cross-spectrum, as expected, lies between these extremes. Note that the sharp features in the cross-spectrum curves are an artifact of combining the various fitting functions used in our calculation. See the text for further discussion. 127

3.4 Identical to Fig. 3.3, but computed at $\ell = 1000$, where the one-halo term dominates the cross-spectrum. The tSZ auto-spectrum contributions here extend to lower masses and higher redshifts than at $\ell = 100$, leading to a stronger expected cross-correlation with the CMB lensing signal. See the text for further discussion. 128

3.5 Cumulative contributions to the tSZ – CMB lensing cross-spectrum (black), CMB lensing auto-spectrum (blue), and tSZ auto-spectrum (red) as a function of halo mass, at $\ell = 100$ (left panel) and $\ell = 1000$ (right panel). These results are integrated over all redshifts. Note that the cross-spectrum receives significant contributions from much lower mass scales than the tSZ auto-spectrum. See the text for further discussion. 129

3.6 The Compton-y map reconstructed by our ILC pipeline for the fiducial $f_{\rm sky} = 0.30$ case, plotted in Galactic coordinates. Note that the mean of the map has been removed before plotting for visual clarity. 137

3.7 Histogram of Compton-y values in our fiducial ILC y-map. The non-Gaussian tail extending to positive y-values provides some evidence that the map does indeed contain tSZ signal. The negative tail is likely due to residual Galactic dust in the map. See [59, 60, 61] for related work on interpreting the moments of Compton-y histograms. 138

3.8 Sub-map of a 5°-by-5° region centered on the Coma cluster in the ILC y-map (after removing the map mean). The masked areas are the locations of point sources in this field. 139

3.9 The tSZ power spectrum estimated from our fiducial S1 and S2 Compton-y maps (red circles), with statistical errors only. The Planck results from [44] are shown as blue circles and green squares — the green squares have been corrected to subtract residual power from leakage of CIB and IR and radio point sources into the y-map constructed in [44], while the blue circles are the uncorrected results (with statistical errors only). We have not attempted to subtract the residual power due to these foregrounds; the red circles are the raw cross-spectrum of our S1 and S2 ILC y-maps, which are in reasonable agreement with the uncorrected Planck points (blue circles) over the region in which our ILC algorithm minimizes residual contamination ($300 < \ell < 1000$, as delineated on the plot; see Section 3.4.2). The thick black error bars are the statistical errors alone on the foreground-corrected Planck points, while the thin green error bars are the combined statistical and foreground-subtraction uncertainties (from Table 3 of [44]). The solid magenta curve shows the tSZ power spectrum predicted by our fiducial model. The cyan and yellow points are the most recent ACT [4] and SPT [57] constraints on the tSZ power spectrum amplitude at $\ell = 3000$, respectively (the SPT point is slightly offset for visual clarity). Note that the ACT and SPT constraints are derived in a completely different data analysis approach, in which the tSZ signal is constrained through its indirect contribution to the total CMB power measured at $\ell = 3000$ by these experiments; in other words, a y-map is not constructed. We use the corrected Planck points with their full errors when deriving cosmological and astrophysical constraints in Section 3.6. 143

3.10 The tSZ – CMB lensing cross-power spectrum estimated from our fiducial co-added Compton-y map and the publicly released Planck lensing potential map is shown in the red circles, with statistical errors computed via Eq. (3.30). The blue squares show the cross-power spectrum after a correction for CIB leakage into the y-map has been applied (see Section 3.5.3 for a description of the subtraction procedure). The errors on the blue squares include statistical uncertainty and additional systematic uncertainty arising from the CIB subtraction. The final significance of the CIB-corrected measurement of $C_\ell^{y\phi}$ is 6.2σ. The solid green curve shows the theoretical prediction of our fiducial model for the tSZ – CMB lensing cross-spectrum (see Section 3.3) — we emphasize that this curve is not a fit to the data, but rather an independent prediction. It agrees well with our measurement: $\chi^2 = 14.6$ for 12 degrees of freedom. The dashed blue and dotted red curves show the one- and two-halo contributions to the fiducial theoretical cross-spectrum, respectively. Our measurement shows clear evidence of contributions from both terms. 146

3.11 The left panel shows the auto-power spectrum of the Planck 857 GHz map (circles with error bars), computed for three sky cuts — note that the fiducial case is $f_{\rm sky} = 0.3$ (red circles). The right panel shows the cross-power spectrum between our fiducial co-added y-map and the Planck 857 GHz map (for $f_{\rm sky} = 0.3$ only). These measurements provide a method with which to assess the level of residual contamination in the y-map due to dust emission from the CIB and from the Galaxy. We model the 857 GHz map as a combination of CIB emission (computed from the best-fit model in [81], shown as a dashed green curve in the left panel) and Galactic dust emission (computed by subtracting the CIB model from the measured 857 GHz power spectrum, shown as squares in the left panel). We then model the y-map – 857 GHz cross-power spectrum as a linear combination of these emission sources using two free amplitude parameters (see the text for further information). We fit these parameters using the measured cross-power spectrum; the results of the fit are shown as dotted magenta (Galactic) and dashed green (CIB curves), with the sum of the two in solid blue. The model provides a reasonable fit to the cross-power spectrum ($\chi^2 = 16.9$ for $12 - 2 = 10$ degrees of freedom) and allows us to effectively constrain the level of CIB leakage into the y-map. 154

3.12 Cross-power spectrum between the null Compton y-map and the CMB lensing potential map (black points). The null y-map is constructed by subtracting (rather than co-adding) the S1 and S2 season y-maps output by our ILC pipeline in Section 3.4.2. The error bars, shown in black, are computed by applying Eq. (3.30) to the null y-map and CMB lensing potential map. The result is consistent with a null signal: $\chi^2 = 13.2$ for 12 degrees of freedom. The cyan error bars show the 1σ errors on our fiducial tSZ – CMB lensing cross-spectrum (i.e., the error bars on the blue squares in Fig. 3.10), which allow a quick assessment of the possible importance of any systematic effect in our measurement. 155

3.13 Difference between the tSZ – CMB lensing cross-power spectrum measured using our fiducial $f_{\rm sky} = 0.3$ Galactic mask and using an $f_{\rm sky} = 0.2$ mask (black squares in left panel) or $f_{\rm sky} = 0.4$ mask (black squares in right panel). The error bars on the black points are computed by applying Eq. (3.30) to the $f_{\rm sky} = 0.2$ or 0.4 case, i.e., using the same procedure as for the fiducial $f_{\rm sky} = 0.3$ analysis. In both cases, the difference between the fiducial cross-power spectrum and the $f_{\rm sky}$ variation cross-power spectrum is consistent with null: $\chi^2 = 9.7$ (12 degrees of freedom) for the $f_{\rm sky} = 0.2$ case, and $\chi^2 = 2.8$ (12 degrees of freedom) for the $f_{\rm sky} = 0.4$ case. The errors on the fiducial signal are shown in cyan, as in Fig. 3.12. Note that the entire pipeline described in Sections 3.4 and 3.5, including the CIB correction, is re-run to obtain the $f_{\rm sky} = 0.2$ and 0.4 results. The ILC weights for each sky cut are given in Table 3.2. 156

3.14 Difference between the tSZ – CMB lensing cross-power spectrum corrected for CIB contamination using our fiducial approach and corrected using an alternative approach based on the model fitting results presented in [48] (black squares). In the alternative approach, we simply weight the $C_\ell^{\mathrm{CIB}\times\phi}$ results at each HFI frequency by the relevant ILC weight computed by our Compton-y reconstruction pipeline (given in Table 3.2), and then sum the results to get the total leakage into the $C_\ell^{y\phi}$ measurement. The difference between our fiducial CIB-corrected cross-spectrum and the cross-spectrum corrected using this method is consistent with null: $\chi^2 = 9.4$ for 12 degrees of freedom. The slight disagreement at low-ℓ is likely due to discrepancies between the model fit from [48] and the measured CIB – CMB lensing cross-spectrum at 100, 143, and 217 GHz — see the text for discussion. The errors on our fiducial tSZ – CMB lensing cross-spectrum measurement are shown in cyan, as in Figs. 3.12 and 3.13. 158

3.15 ICM model comparison results. In both panels, the blue squares show a re-binned version of the tSZ – CMB lensing cross-power spectrum results presented in Fig. 3.10. The re-binning was performed using inverse-variance weighting, and is only done for visual purposes; all model fitting results are computed using the twelve bins shown in Fig. 3.10. The left panel shows a comparison between the predictions of several ICM physics models, all computed using a WMAP9 background cosmology. The fiducial "AGN feedback" model [29] is the best fit, with $\chi^2 = 14.6$ for 12 degrees of freedom (see the text for comparisons to the other models shown here). The right panel shows the result of a fit for the amplitude of the pressure–mass relation P_0 with respect to its standard value in our fiducial model (solid black curve), as well as the result of a fit for the HSE mass bias in the UPP of Arnaud et al. [62] (dashed black curve). Our measurements are consistent with the fiducial value of P_0 ($P_0/P_{0,\rm fid} = 1.10 \pm 0.22$) in our standard "AGN feedback" model, and similarly with standard values of the HSE mass bias for the UPP, $(1-b) = 1.06^{+0.11}_{-0.14}$ (all results at 68% C.L.). 162

3.16 Constraints on σ_8 and Ω_m from the tSZ – CMB lensing cross-power spectrum measurement presented in this work (blue shaded), the tSZ auto-power spectrum measurement from [44] re-analyzed in this work (red shaded), the Planck+WMAP polarization CMB analysis [3] (green dashed), WMAP9 CMB analysis [1] (black dashed), number counts of Planck tSZ clusters [12] (cyan dashed), and a recent cross-correlation of tSZ signal from Planck and WMAP with an X-ray "δ"-map based on ROSAT data [67] (magenta shaded). All contours shown are 68% confidence intervals only (for visual clarity). See the text for further discussion. 169

3.17 Normalized one-dimensional likelihoods for the degenerate parameter combination $\sigma_8(\Omega_m/0.282)^{0.26}$ that is best constrained by the tSZ auto- and cross-power spectra measurements. The probes shown are identical to those in Fig. 3.16, except for an additional result from a recent re-analysis of the Planck CMB + WMAP polarization data [93]. See the text for further discussion. . 170

3.18 Constraints on $\sigma_8(\Omega_m/0.282)^{0.26}$ and P_0 (the amplitude of the pressure–mass relation normalized to its fiducial value in our ICM model [29]) from the tSZ – CMB lensing cross-power spectrum measurement presented in this work (blue), the tSZ auto-power spectrum measurement from [44] re-analyzed in this work (red), and the combination of the two probes (black dashed). The different dependences of the two probes on these parameters allow the degeneracy between the ICM and cosmology to be broken (albeit weakly in this initial application) using tSZ measurements alone. The contours show the 68% and 95% confidence intervals for the tSZ probes; in addition, the vertical lines show the 68% confidence interval for the WMAP9 constraint on this cosmological parameter combination, which is obviously independent of P_0. Standard values for the pressure–mass normalization ($P_0 \approx 0.8$–1.0) are compatible with all of the probes. 174

4.1 The tSZ variance versus σ_8 obtained from five different pressure profiles (using the Tinker mass function) and one direct simulation measurement. The scalings with σ_8 are similar for all the models: $\langle T^2 \rangle \propto \sigma_8^{6.6-7.9}$. It is evident that $\langle T^2 \rangle$ is very sensitive to σ_8, but the scatter due to uncertainties in the ICM astrophysics (which is greater than twice the signal for large σ_8) makes precise constraints from this quantity difficult. 194

4.2 Similar to Fig. 1, but now showing the tSZ skewness. The scalings with σ_8 are again similar for all the models: $\langle T^3 \rangle \propto \sigma_8^{9.7-11.5}$. As for the variance, the sensitivity to σ_8 is pronounced, but degraded due to uncertainties in the ICM astrophysics (as represented by the different choices of pressure profile). . . . 195

4.3 Fraction of the tSZ variance contributed by clusters with virial mass $M < M_{\rm max}$. 197

4.4 Fraction of the tSZ skewness contributed by clusters with virial mass $M < M_{\rm max}$. Comparison with Fig. 3 indicates that the skewness signal arises from higher-mass clusters than those that comprise the variance. 199

4.5 The rescaled skewness $\tilde{S}_{3,\beta}$ for $\beta = 1.4$ plotted against σ_8. It is evident that this statistic is nearly independent of σ_8, as expected based on the scalings in Table I. However, it is still dependent on the ICM astrophysics, as represented by the different pressure profiles. Thus, this statistic can be used for determining the correct gas physics model. 200

4.6 Fraction of the rescaled skewness $\tilde{S}_{3,\beta}$ for $\beta = 1.4$ contributed by clusters with virial mass $M < M_{\rm max}$. A significant fraction of the signal comes from low-mass objects — even more so than the variance (Fig. 3). One can thus interpret this statistic as a measure of the gas fraction in low-mass clusters: pressure profiles that yield $f_{gas} \approx \Omega_b/\Omega_m$ in low-mass objects lead to a small value for this statistic (since they give much larger values of $\langle T^2 \rangle$), while pressure profiles that include significant feedback effects — thus lowering the gas fraction in low-mass objects — lead to larger values for this statistic. . . 202

4.7 Similar to Fig. 5, but calculated in terms of observed tSZ statistics, namely, C_{3000} (the amplitude of the tSZ power spectrum at $\ell = 3000$) and $\langle \tilde{T}^3 \rangle$ (the filtered skewness as defined in [38]). The green point shows the constraint using the SPT measurement of C_{3000} [10] and the ACT measurement of $\langle \tilde{T}^3 \rangle$ [38]. The magenta point shows an updated constraint using the SPT measurement of the tSZ bispectrum (converted to $\langle \tilde{T}^3 \rangle$) [52] instead of the ACT skewness measurement. Note that this figure has been updated from the published version. 205

4.8 The tSZ variance versus the sum of the neutrino masses Σm_ν, with $\Delta^2_{\mathcal{R}} = 2.46 \times 10^{-9}$ (its WMAP5 value). The dependence is not precisely captured by a simple scaling as for σ_8 (note that the axes are log-linear), but the curves are well-fit by quadratic polynomials. 207

4.9 The tSZ skewness versus the sum of the neutrino masses Σm_ν, with $\Delta^2_{\mathcal{R}} = 2.46 \times 10^{-9}$ (its WMAP5 value). The dependence is not precisely captured by a simple scaling as for σ_8 (note that the axes are log-linear), but the curves are well-fit by cubic polynomials. 209

4.10 The tSZ variance versus f_{NL}, with $\Delta^2_{\mathcal{R}} = 2.46 \times 10^{-9}$ (its WMAP5 value). The dependence is not precisely captured by a simple scaling as for σ_8 (note that the axes are log-linear), but the curves are well-fit by quadratic polynomials. 210

4.11 The tSZ skewness versus f_{NL}, with $\Delta^2_{\mathcal{R}} = 2.46 \times 10^{-9}$ (its WMAP5 value). The dependence is not precisely captured by a simple scaling as for σ_8 (note that the axes are log-linear), but the curves are well-fit by cubic polynomials. 212

4.12 The rescaled skewness $\tilde{S}_{3,\beta}$ for $\beta = 0.7$ plotted against σ_8. Clearly there is significantly less scatter for this statistic between the different pressure profiles, but it still scales as σ_8^{5-6}. Thus, this value of β provides a method for constraining cosmological parameters that is fairly independent of the details of the gas physics. The simulation point is somewhat low, but the disagreement is not statistically significant. This point was also not used in determining the optimal value of β. 215

4.13 Contributions to the rescaled skewness with $\beta = 0.7$ from clusters of various masses, for each of the different pressure profiles. The contributions to this statistic are very similar to those found for the skewness in Fig. 4, as one would expect based on the definition of $\tilde{S}_{3,\beta=0.7}$. 216

4.14 Similar to Fig. 12, but now showing the rescaled skewness with $\beta = 0.7$ plotted against the sum of the neutrino masses Σm_ν. We assume $\Delta^2_{\mathcal{R}} = 2.46\times10^{-9}$ (its WMAP5 value). There is significantly reduced scatter between the pressure profiles for this statistic, and it retains a leading-order quadratic dependence on Σm_ν. 219

4.15 The rescaled kurtosis $\langle T^4\rangle/\langle T^2\rangle^{1.9}$ plotted against σ_8. It is evident that this statistic is nearly independent of σ_8, as expected based on the scalings in Tables I and II. However, it is still dependent on the ICM astrophysics, as represented by the different pressure profiles. Thus, this statistic — like the rescaled skewness — can be used for determining the correct gas physics model. 221

4.16 Fraction of the rescaled kurtosis $\langle T^4\rangle/\langle T^2\rangle^{1.9}$ contributed by clusters with virial mass $M < M_{\rm max}$. Although a non-trivial fraction of the signal comes from low-mass objects, comparison with Fig. 6 indicates that the rescaled kurtosis is mostly sourced by more massive halos than those responsible for the rescaled skewness. 223

5.1 PDF of pixel temperature values in the ACT 148 GHz Equatorial CMB map after filtering and point source masking. The solid red vertical line denotes $\Delta\tilde{T} = 0$. The tSZ effect is responsible for the significant non-Gaussian tail on the negative side of the PDF ($\Delta\tilde{T} < 0$). This figure is effectively a re-binned version of Fig. 2 from [19], although the positive side of the PDF is slightly different as we do not apply the 218 GHz CIB mask constructed in that work to the 148 GHz map here (see text for details). The mask is unnecessary because we do not consider the positive side of the PDF ($\Delta\tilde{T} > 0$) in this analysis, since it contains essentially no tSZ signal. 238

5.2 Comparison of the noiseless halo model computation of the tSZ PDF using Eq. (5.3) to the PDF directly measured from the cosmological hydrodynamics simulations of [33]. The error bars are computed from the scatter amongst the 390 maps extracted from the simulation. Note that no filtering or noise convolution has been performed for these calculations, and thus the ΔT values cannot be directly compared to those shown in the other plots in the paper. The comparison demonstrates that the halo model approach works extremely well, except for discrepancies at low-$|\Delta T|$ values arising from the breakdown of the halo model assumptions (see the text for discussion). 246

5.3 The noisy tSZ PDF computed for our fiducial cosmology and pressure profile model using the semi-analytic halo model framework described in Section 5.4.3. The dashed red curve shows the tSZ contribution given by the second term in Eq. (5.7), the dotted magenta curve shows the noise contribution given by the first term in Eq. (5.7), and the solid blue line shows the total signal. 252

5.4 The noisy tSZ PDF computed for our fiducial cosmology and pressure profile model using the semi-analytic halo model framework described in Section 5.4.3, but with the lower mass cutoff in the integrals set to $1.5\times10^{14}M_{\odot}/h$ (left panel) and $3\times10^{14}M_{\odot}/h$ (right panel), slightly below and above the fiducial cutoff of $2\times10^{14}M_{\odot}/h$ (Fig. 5.3). Comparing the two panels and Fig. 5.3, it is evident that the total predicted tSZ PDF remains unchanged as the lower mass cutoff is varied — the contributions from low-mass clusters are simply shifted between the tSZ and noise terms in Eq. (5.7), as seen in the dashed red and dotted magenta curves, respectively. The non-Gaussian, tSZ-dominated tail of the PDF is dominated by clusters well above these mass cutoffs, and is thus insensitive to the exact cutoff choice as well. 253

5.5 Dependence of the tSZ PDF on cosmological and astrophysical parameters. The plot shows the ratio with respect to the fiducial tSZ PDF (see Fig. 5.3) for $\pm5\%$ variations for each parameter in the model $\{\sigma_8, \Omega_m, P_0\}$. The sensitivity to σ_8 is quite pronounced in the tSZ-dominated tail, with the values in these bins scaling as $\sim \sigma_8^{10-16}$. As expected, the bins in the noise-dominated regime are almost completely insensitive to variations in the cosmological or astrophysical parameters. 254

5.6 Mass and redshift contributions to the noisy tSZ PDF. The curves show the fraction of the total signal in each bin that is sourced by clusters at a mass or redshift below the specified value. In the low-$|\Delta\tilde{T}|$ bins, the total signal is dominated by noise (see Fig. 5.3), and thus there is essentially no information about the mass or redshift contributions. Below $\Delta\tilde{T} \approx -40\,\mu$K, the tSZ signal dominates over the noise, and robust inferences can be made about the mass and redshift contributions, as shown. Note that these results are specific to the noise levels and angular resolution of the filtered ACT Equatorial map, and must be recomputed for different experimental scenarios. See the text for further discussion. 256

5.7 Comparison of the semi-analytic halo model computation of the noisy tSZ PDF using Eq. (5.7) to the PDF directly measured from the simulated maps described in Section 5.5. The error bars are computed from the scatter amongst the 476 simulated maps. The two curves are nearly indistinguishable. This validates the halo model approximation that clusters do not overlap along the LOS (see Section 5.4.2). 262

5.8 Covariance matrix of the noisy tSZ PDF computed for our fiducial cosmology and pressure profile model using the semi-analytic halo model framework described in Section 5.4.4 with corrections computed from the Monte Carlo simulations described in Section 5.5. The axes are labeled by bin number, with the top left corner corresponding to the lowest $|\Delta\tilde{T}|$ bin shown in Fig. 5.3 (i.e. the noise-dominated regime) and the bottom right corner corresponding to the highest $|\Delta\tilde{T}|$ bin shown in Fig. 5.3 (i.e. the tSZ-dominated regime). The specific quantity plotted is $\log_{10}(|\mathrm{Cov}|)$, with values labeled by the colors (the absolute value is necessary because some of the off-diagonal bins representing correlations between the noise-dominated region and the tSZ-dominated region are negative). Note the significant amount of bin-to-bin correlation, especially amongst the tSZ-dominated bins in the tail of the PDF. 264

5.9 Histogram of ML σ_8 values recovered from 476 Monte Carlo simulations processed through the likelihood in Eq. (5.12) with P_0 fixed to unity and the nuisance parameter σ^2_{nuis} marginalized over. The input value, $\sigma_8 = 0.817$ (indicated by the red vertical line), is recovered in an unbiased fashion: the mean recovered value is $\langle\sigma_8^{\mathrm{ML}}\rangle = 0.818 \pm 0.018$ (indicated by the orange vertical lines). 266

5.10 Scatter plot of ML recovered values of σ_8 and P_0 using 476 Monte Carlo simulations processed through the likelihood in Eq. (5.12) with both parameters allowed to vary (the nuisance parameter σ^2_{nuis} can also vary and is marginalized over). The approximate degeneracy between σ_8 and P_0 seen here can be inferred from the scalings shown in Fig. 5.5. The input values, $\sigma_8 = 0.817$ and $P_0 = 1$, are represented by the red square; the mean recovered ML result is represented by the orange diamond. A bias can clearly be seen; this is further investigated in the 1D marginalized results in Fig. 5.11. 267

5.11 Histogram of ML σ_8 values (left panel) and P_0 values (right panel) recovered from 476 Monte Carlo simulations processed through the likelihood in Eq. (5.12) after marginalizing over all other parameters. The input values, $\sigma_8 = 0.817$ and $P_0 = 1$, are indicated by the red vertical lines in each panel. There is evidence of a bias in the likelihood: the mean recovered values are $\langle \sigma_8^{\rm ML} \rangle = 0.806 \pm 0.030$ and $\langle P_0^{\rm ML} \rangle = 1.05 \pm 0.10$ (indicated by the orange lines in each panel). The cause of this bias is presently unclear (see text for discussion), and we treat it as a systematic which must be corrected for in the ACT data analysis. 269

5.12 Likelihood of σ_8 computed from the ACT 148 GHz PDF using the likelihood function described in Section 5.6.1 with P_0 fixed to unity (i.e., only σ_8 and $\sigma^2_{\rm nuis}$ are free to vary, and the latter is marginalized over). 271

5.13 Likelihood contours for σ_8 and P_0 computed from the ACT 148 GHz PDF using the likelihood function described in Section 5.6.1 with both parameters (and $\sigma^2_{\rm nuis}$) allowed to vary. The blue, red, and green curves represent the 68%, 95%, and 99% confidence level contours, respectively. The orange diamond represents the ML point in the 2D parameter space. These results are not corrected for the biases found in Section 5.6.1. Note that the PDF measurement is indeed able to weakly break the degeneracy between σ_8 and P_0, using tSZ data alone. 272

5.14 One-dimensional likelihood profiles for σ_8 (left panel) and P_0 (right panel) after marginalizing the 2D likelihood shown in Fig. 5.13 over all other parameters. These results are not corrected for the biases found in Section 5.6.1. 273

5.15 Comparison of the measured ACT PDF to the ML PDF model for the σ_8-only likelihood (left panel) and (σ_8, P_0) likelihood (right panel). 274

5.16 Difference between the measured ACT PDF and the ML model (for both likelihood scenarios presented in Fig. 5.15) divided by the square root of the diagonal elements of the covariance matrix of the ML model. Note that this plot obscures the significant bin-to-bin correlations in the PDF (see Fig. 5.8), and thus extreme caution must be used in any interpretation. However, the plot clearly indicates that the empty outermost bins in the tail of the ACT PDF are responsible for pulling the best-fit amplitude of the ML models down from the value that would likely be preferred by the populated bins. 275

A.1 The Wiener filter applied to the ACT temperature maps before calculating the unnormalized skewness. This filter upweights scales on which the tSZ signal is large compared to other sources of anisotropy. 292

A.2 Histogram of the pixel temperature values in the filtered, masked ACT CMB temperature maps. A Gaussian curve is overlaid in red. 293

A.3 Likelihood of the skewness measurement described in the text (with Gaussian statistics assumed). 294

A.4 Plot of the skewness signal as a function of the minimum S/N of the clusters that are masked (this indicates how many known clusters are left in the data, unmasked). The blue line is calculated using the full cluster candidate catalog obtained via matched filtering, while the green line uses a catalog containing only optically-confirmed clusters [38]. Both lines have identical errors, but we only plot them for the green line for clarity. Confirmed clusters source approximately two-thirds of the signal, which provides strong evidence that it is due to the tSZ effect. Note that one expects a positive bias of $\approx 4\ \mu\mathrm{K}^3$ for the S/N = 4 point of the blue line due to impurities in the full candidate catalog masking the tail of the Gaussian distribution. 296

A.5 A test for IR source contamination: similar to the blue line in Fig. 4, but with a range of values of the cutoff used to construct an IR source mask in the 218 GHz band. Any cutoff below $\approx 3.2\sigma$ gives similarly negative results and thus appears sufficient for point source removal, where $\sigma = 10.3$ μK is the standard deviation of the 148 GHz maps. For comparison, the standard deviation of the 218 GHz maps is ≈ 2.2 times larger. The percentages of the map which are removed for the masking levels shown, from the least to the most strict cut, are 0.7, 2.5, 8.4, 14.5, 23.7, and 36.6%. 297

Chapter 1

Introduction: The Sunyaev-Zel'dovich Effect

1.1 Introduction

Since its discovery nearly five decades ago, the cosmic microwave background radiation (CMB) has proven to be the most valuable source of cosmological information known to science. Its exquisite blackbody spectrum confirms that our universe was in an extremely hot, dense phase in its earliest moments — the Big Bang. The near-perfect homogeneity of the CMB temperature over the sky provides strong support for the isotropy and homogeneity of the universe, assumptions that lie at the foundation of modern cosmology.

However, the study of fluctuations in the CMB temperature between different points on the sky has arguably been an even richer line of scientific inquiry. These one-part-per-hundred-thousand fluctuations are the earliest known signposts of the cosmic overdensities and underdensities that later grew into the large-scale structure of the universe. Under the influence of gravity, these initially minute overdensities collapsed during the subsequent ≈ 13.7 billion years to form the galaxies and clusters of galaxies seen in the present-day universe. The seeds of structure observed in the CMB are of particular importance be-

cause they provide an essentially unspoiled view of the initial perturbations laid down in the earliest moments of the universe, currently believed to have occurred during a phase of super-exponential "inflation" that preceded the conventional hot Big Bang. The CMB fluctuations present an opportunity to determine with high precision the fundamental parameters governing our universe, including those specifying its shape, age, and matter-energy content. This opportunity has been seized by a number of highly successful experiments in the past few decades, culminating in the Wilkinson Microwave Anisotropy Probe (WMAP) and Planck satellite missions [1, 2].

WMAP and Planck have provided exquisite measurements of the anisotropies imprinted in the CMB at the "surface of last scattering" $\approx 380,000$ years after the Big Bang. However, in addition to these "primary" anisotropies, which provide a window onto the early universe, the CMB also contains fluctuations imprinted by processes occurring much later in the history of the universe — "secondary" anisotropies. On large angular scales, these fluctuations are generally subdominant to the primary anisotropies, but on smaller angular scales they can dominate, as the primary fluctuations are damped by photon diffusion in the primordial plasma (Silk damping). Thus, high-resolution, high-sensitivity experiments developed in recent years, including the Atacama Cosmology Telescope (ACT) and the South Pole Telescope (SPT), have allowed detailed mapping of these secondary anisotropies for the first time [3, 4].

These secondary anisotropies result from the interaction of CMB photons with intervening matter along the line-of-sight (LOS) as they pass to our detectors from the surface of last scattering. This thesis is primarily concerned with the Sunyaev-Zel'dovich (SZ) effect, which arises due to the inverse Compton scattering of CMB photons off of free electrons along the LOS [5]. The SZ effect is typically further subdivided into two components: the thermal SZ (tSZ) effect, which is sourced by the thermal energy of collections of hot electrons, and the kinematic SZ (kSZ) effect, which is sourced by the bulk velocity of electrons with respect

to the CMB rest frame. This thesis is almost entirely focused on the former, although the latter appears occasionally throughout the chapters, usually as a contaminant.

In addition to the SZ effect, secondary anisotropies are imprinted in the CMB through the lensing of CMB photons by the gravitational influence of intervening structures along the LOS. Unlike the SZ effect, both dark matter and ordinary (baryonic) matter contribute to the gravitational lensing effect. This signal thus represents a tool with which to map out the distribution of all matter in the universe back to the surface of last scattering.

Unraveling the properties of these secondary anisotropies holds the promise of a vast amount of cosmological information beyond that contained in the primary CMB fluctuations. These anisotropies are a sensitive probe of the late universe, allowing measurements of cosmic distances and the growth of cosmic structure, both of which contain information about dark matter, dark energy, inflationary physics, and the physics of the neutrino sector. For example, the tSZ effect in principle allows for the mapping of all galaxy clusters in the observable universe, as the hot electrons sourcing the signal are predominantly found in these objects. These clusters are the rarest, most massive gravitationally bound objects to have formed in the current universe, and they are thus a sensitive probe of many underlying cosmological parameters.

However, the cosmological potential of these signals comes with significant additional complications due to theoretical modeling uncertainties, especially for the tSZ signal. The tSZ effect has long been recognized as a particularly precise probe of the amplitude of cosmic density fluctuations (e.g., [6]); this point is emphasized in further detail in Chapter 2 and throughout much of the thesis. Unfortunately, constraints derived from tSZ measurements still remain limited by systematic errors arising from poorly understood gas physics in the intracluster medium (ICM) [7, 8, 9]. Much of the work in this thesis is devoted to making progress toward solving this problem, while simultaneously maximizing the cosmological constraining power of the tSZ signal. With the unprecedented statistical power of upcoming datasets, including those from Planck, ACTPol [10], and SPTPol [11], it is imperative to de-

velop optimal frameworks with which to understand and interpret the secondary anisotropies in the CMB, a goal which this thesis aims to fulfill.

The remainder of this thesis is organized as follows. The rest of the this chapter provides a basic introduction to the tSZ effect and presents complete derivations of the tSZ two-point statistics in the framework of the halo model. This framework will underlie much of the theoretical material throughout the other chapters. Chapter 2 presents a comprehensive study of the power spectrum of temperature anisotropies sourced by the tSZ effect, focusing on their ability to constrain a number of cosmological and astrophysical parameters in present and forthcoming data sets. Chapter 3 describes the first detection of the cross-correlation of the tSZ and CMB lensing fields using data from the Planck satellite, including detailed theoretical background, a novel approach to constructing a tSZ map from multifrequency CMB data, and constraints on cosmological and astrophysical parameters derived from the cross-correlation measurements. Moving beyond the two-point level, Chapters 4 and 5 consider higher N-point statistics of the tSZ signal, which probe its non-Gaussian structure. Chapter 4 presents theoretical calculations of real-space moments of the tSZ signal, and describes methods with which to use measurements of multiple tSZ moments to break degeneracies between cosmological and ICM parameters in tSZ-derived constraints. We also demonstrate applications of these methods to data from ACT and SPT. Chapter 5 extends this approach to consider the entire one-point probability distribution function (PDF) of the tSZ signal. We use analytic and simulation-based methods to compute the tSZ PDF and its covariance matrix, and then apply these results to the PDF measured in ACT data, yielding tight constraints on cosmology and the ICM. The Appendix presents an earlier measurement of the tSZ three-point function (skewness) using ACT data, a result which motivated much of the work described in Chapters 4 and 5.

1.2 The Thermal Sunyaev-Zel'dovich Effect

1.2.1 Background

The thermal Sunyaev-Zel'dovich (tSZ) effect is a spectral distortion of the cosmic microwave background (CMB) that arises due to the inverse Compton scattering of CMB photons off hot electrons that lie between our vantage point and the surface of last scattering [5]. The vast majority of these hot electrons are located in the intracluster medium (ICM) of galaxy clusters, and thus the tSZ signal is dominated by contributions from these massive objects. The tSZ effect has been used for many years to study individual clusters in pointed observations (e.g., [12, 13, 14, 15]) and in recent years has been used as a method with which to find and characterize massive clusters in blind millimeter-wave surveys [16, 17, 18, 8, 7]. Moreover, recent years have brought the first detections of the angular power spectrum of the tSZ effect through its contribution to the power spectrum in arcminute-resolution maps of the microwave sky made by the Atacama Cosmology Telescope (ACT)[1] and the South Pole Telescope (SPT)[2] [19, 20, 21, 22]. More recently, the Planck collaboration has recently released a measurement of the tSZ power spectrum over a much wider range of angular scales using a tSZ map constructed from the multifrequency Planck data [23]. In addition, three-point statistics of the tSZ signal have been detected recently for the first time: first, the real-space skewness was detected in ACT data [24] using methods first anticipated by [25], and second, the Fourier-space bispectrum was very recently detected in SPT data [26]. The amplitudes of these measurements were shown to be consistent in the SPT analysis, despite observing different regions of sky and using different analysis methods. Note that the tSZ signal is highly non-Gaussian since it is dominated by contributions from massive collapsed objects in the late-time density field; thus, higher-order tSZ statistics contain significant information beyond that found in the power spectrum. Furthermore, the combination of multiple different N-point tSZ statistics provides an avenue to extract tighter

[1] `http://www.princeton.edu/act/`
[2] `http://pole.uchicago.edu/`

constraints on cosmological parameters and the astrophysics of the ICM than the use of the power spectrum alone, through the breaking of degeneracies between ICM and cosmological parameters [27, 28]. The recent SPT bispectrum detection used such methods in order to reduce the error bar on the tSZ power spectrum amplitude by a factor of two [26]. In Chapter 5, we present a generalization of these methods to consider the full one-point PDF of the tSZ effect.

1.2.2 Basic Definitions

The tSZ effect produces a frequency-dependent shift in the CMB temperature observed in the direction of a galaxy group or cluster. The temperature shift ΔT at angular position $\vec{\theta}$ with respect to the center of a cluster of mass M at redshift z is given by [5]

$$\begin{aligned} \frac{\Delta T(\vec{\theta}, M, z)}{T_{\rm CMB}} &= g_\nu y(\vec{\theta}, M, z) \\ &= g_\nu \frac{\sigma_T}{m_e c^2} \int_{\rm LOS} P_e\left(\sqrt{l^2 + d_A^2 |\vec{\theta}|^2}, M, z\right) dl\,, \end{aligned} \tag{1.1}$$

where $g_\nu = x\coth(x/2) - 4$ is the tSZ spectral function with $x \equiv h\nu/k_B T_{\rm CMB}$, y is the Compton-y parameter, σ_T is the Thomson scattering cross-section, m_e is the electron mass, and $P_e(\vec{r})$ is the ICM electron pressure at location $\vec{r}$ with respect to the cluster center. We have neglected relativistic corrections in Eq. (1.1) (e.g., [29]), as these effects are relevant only for the most massive clusters in the universe ($\gtrsim 10^{15}\ M_\odot/h$). For the power spectrum statistics considered in Chapters 2 and 3, these clusters do not dominate the signal, and thus we do not include relativistic corrections in our calculations.

Note that we only consider spherically symmetric pressure profiles in this work, i.e. $P_e(\vec{r}) = P_e(r)$ in Eq. (1.1). The integral in Eq. (1.1) is computed along the LOS such that $r^2 = l^2 + d_A(z)^2\theta^2$, where $d_A(z)$ is the angular diameter distance to redshift z and $\theta \equiv |\vec{\theta}|$ is the angular distance between $\vec{\theta}$ and the cluster center in the plane of the sky (note that this formalism assumes the flat-sky approximation is valid; we provide exact

full-sky results for the tSZ power spectrum in the following section). In the flat-sky limit, a spherically symmetric pressure profile implies that the temperature decrement (or Compton-y) profile is azimuthally symmetric in the plane of the sky, i.e., $\Delta T(\vec{\theta}, M, z) = \Delta T(\theta, M, z)$. Finally, note that the electron pressure $P_e(\vec{r})$ is related to the thermal gas pressure via $P_{th} = P_e(5X_H + 3)/2(X_H + 1) = 1.932P_e$, where $X_H = 0.76$ is the primordial hydrogen mass fraction. We calculate all tSZ power spectra in Chapter 2 at $\nu = 150$ GHz, where the tSZ effect is observed as a decrement in the CMB temperature ($g_{150\,\mathrm{GHz}} = -0.9537$). We make this choice simply because recent tSZ measurements have been performed at this frequency using ACT and SPT (e.g., [24, 26, 22, 20]), and thus the temperature values in this regime are perhaps more familiar and intuitive. All of our calculations can be phrased in a frequency-independent manner in terms of the Compton-y parameter, and we will often use "y" as a label for tSZ quantities, although they are calculated numerically in Chapter 2 at $\nu = 150$ GHz. In Chapter 3, we will work predominantly in terms of the y-parameter, while in Chapters 4 and 5 we will primarily compute quantities at $\nu = 150$ GHz.

1.2.3 Halo Model Derivation of tSZ Statistics

In the halo model, it is assumed that all matter in the universe is bound in halos. Each halo of virial mass M is assumed to have a density profile, $\rho(\vec{x}; M)$, and (for our purposes) an electron pressure profile, $P_e(\vec{x}; M)$. The mass density field at position $\vec{x}$ is then given by the sum of the contributions from all halos in the universe:

$$\rho(\vec{x}) = \sum_{i \in \mathrm{halos}} \rho(\vec{x} - \vec{x}_i; M_i)\,. \tag{1.2}$$

Similarly, the electron pressure field at position $\vec{x}$ is given by:

$$P_e(\vec{x}) = \sum_{i \in \mathrm{halos}} P_e(\vec{x} - \vec{x}_i; M_i)\,. \tag{1.3}$$

For calculations involving the tSZ effect, it is convenient to define a "3D Compton-y" field that is simply a re-scaling of the electron pressure field:

$$y_{3D}(\vec{x}) = \frac{\sigma_T}{m_e c^2} P_e(\vec{x})\,. \tag{1.4}$$

Note that the 3D Compton-y field has dimensions of inverse length (it is thus important to be careful about comoving versus physical units — in our calculations using the Battaglia pressure profile, the pressure is given in physical units, and thus so is y_{3D}). The usual (2D) Compton-y field is then given by the LOS projection of $y_{3D}(\vec{x})$:

$$\begin{aligned} y(\hat{n}) &= \int c\,dt\, y_{3D}(\vec{x}(\chi(t), \hat{n})) \\ &= \int d\chi\, a(\chi) y_{3D}(\vec{x}(\chi, \hat{n}))\,, \end{aligned} \tag{1.5}$$

where t is the age of the universe at a given epoch, $\chi(t)$ is the comoving distance to that epoch, $a(\chi)$ is the scale factor at that epoch, and $\hat{n}$ is a unit vector on the sky. We have used $dt/da = 1/(aH)$ and $d\chi/da = -c/(Ha^2)$ in going from the first line to the second line, where $H(a)$ is the Hubble parameter. Defining the projection kernel $W^y(\chi)$ for the Compton-y field via $y(\vec{\hat{n}}) = \int d\chi W^y(\chi) y_{3D}(\vec{x}(\chi, \hat{n}))$, we thus have:

$$W^y(\chi) = a(\chi)\,. \tag{1.6}$$

Starting from the 3D Compton-y field defined in Eq. (1.4), we derive the angular power spectrum of the 2D Compton-y field. First, we calculate the relevant 3D power spectrum by means of the halo model (N.B. in this expression and many others in the following, the redshift dependence will be suppressed for notational simplicity):

$$P_{y_{3D}}(\vec{k}) = P^{1h}_{y_{3D}}(\vec{k}) + P^{2h}_{y_{3D}}(\vec{k})\,, \tag{1.7}$$

where the one-halo term is

$$P^{1h}_{y_{3D}}(\vec{k}) = \int dM \frac{dn}{dM} \left| \tilde{y}_{3D}(\vec{k}; M) \right|^2 \tag{1.8}$$

and the two-halo term is

$$P^{2h}_{y_{3D}}(\vec{k}) = \int dM_1 \frac{dn}{dM_1} b(M_1) \tilde{y}_{3D}(\vec{k}; M_1) \int dM_2 \frac{dn}{dM_2} b(M_2) \tilde{y}_{3D}(\vec{k}; M_2) P_{\rm lin}(\vec{k}) \,. \tag{1.9}$$

In these equations, $\tilde{y}_{3D}(\vec{k}; M)$ is the Fourier transform of the 3D Compton-y profile for a halo of virial mass M:

$$\begin{aligned} \tilde{y}_{3D}(\vec{k}; M) &= \int d^3r \, e^{-i\vec{k}\cdot\vec{r}} y_{3D}(\vec{r}; M) \\ &= \int dr \, 4\pi r^2 \frac{\sin(kr)}{kr} y_{3D}(r; M) \,, \end{aligned} \tag{1.10}$$

where $r = |\vec{r}|$, $k = |\vec{k}|$, and we have assumed that $y_{3D}(\vec{r}; M)$ is spherically symmetric to obtain the second expression. Also, in Eqs. (1.8) and (1.9), $dn(M, z)/dM$ is the comoving number density of halos of mass M at redshift z, $b(M, z)$ is the bias of halos of mass M at redshift z (which we will later consider to be scale-dependent) and $P_{\rm lin}(\vec{k})$ is the linear matter power spectrum, as defined in Section 2.3.1. At this point, it is worth emphasizing that expressions analogous to Eqs. (1.7)–(1.9) can be written for any field defined at all points in the universe after its profile for each halo of mass M is specified. These expressions are generic consequences of the halo model. The primary assumption made is that the halo-halo power spectrum for halos of mass M_1 and M_2 is given by the linear matter power spectrum multiplied by the relevant bias parameters:

$$P_{hh}(\vec{k}; M_1, M_2) = b(M_1) b(M_2) P_{\rm lin}(\vec{k}) \,. \tag{1.11}$$

1.2.4 The One-Halo Term

We now compute the contribution of the one-halo term to the tSZ power spectrum. We do the exact calculation first, and then consider the flat-sky limit. Note that for the one-halo term, the notion of the "Limber approximation" is not well-defined — there are no LOS cancellations to consider, since one halo is by definition fixed at a single redshift. Thus, in going from the exact calculation to the small-angle (high-ℓ) limit, we only need to consider the projection of the electron pressure profile from 3D to 2D. Note that this generalization will only affect very massive, low-redshift clusters, which subtend a significant solid angle on the sky.

Consider a cluster of virial mass M at comoving separation $\vec{\chi}$ with respect to our location. The 3D Compton-y field due to this cluster at comoving separation $\vec{r}$ with respect to the cluster center is given by:

$$y_{3D}(\vec{r}; M) = \int \frac{d^3k}{(2\pi)^3} \tilde{y}_{3D}(\vec{k}; M) e^{i\vec{k}\cdot(\vec{r}-\vec{\chi})} \,. \tag{1.12}$$

This expression is simply the inverse transform of Eq. (1.10). Projecting along the LOS as in Eq. (1.5) and using the Rayleigh plane wave expansion, we obtain:

$$\begin{aligned} y(\hat{n}; M) &= \int d\chi' W^y(\chi') y_{3D}(\vec{r}(\chi', \hat{n}); M) \\ &= \int d\chi' a(\chi') \int \frac{d^3k}{(2\pi)^3} \tilde{y}_{3D}(\vec{k}; M) \times \\ &\quad \left[\sum_{\ell m} 4\pi i^\ell Y^*_{\ell m}(\hat{k}) Y_{\ell m}(\hat{n}) j_\ell(k\chi') \right] e^{-i\vec{k}\cdot\vec{\chi}} \,, \end{aligned} \tag{1.13}$$

where $\hat{n}$ is a unit vector on the sky and $\hat{k}$ is the direction of $\vec{k}$. Defining the expansion coefficients $y_{\ell m}(M)$ via

$$y(\hat{n}; M) = \sum_{\ell m} y_{\ell m}(M) Y_{\ell m}(\hat{n}) \,, \tag{1.14}$$

we can read them off from Eq. (1.13):

$$\begin{aligned} y_{\ell m}(M) &= \int d\chi' a(\chi') \int \frac{d^3k}{(2\pi)^3} \tilde{y}_{3D}(\vec{k}; M) 4\pi i^\ell Y^*_{\ell m}(\hat{k}) j_\ell(k\chi') \times \\ &\quad \left[\sum_{\ell' m'} 4\pi (-i)^{\ell'} Y_{\ell' m'}(\hat{k}) Y^*_{\ell' m'}(\hat{\chi}) j_{\ell'}(k\chi) \right] \\ &= \int d\chi' a(\chi') \int \frac{2}{\pi} k^2 \, dk \, \tilde{y}_{3D}(k; M) j_\ell(k\chi') j_\ell(k\chi) Y^*_{\ell m}(\hat{\chi}) \,, \end{aligned} \tag{1.15}$$

where we have again used the Rayleigh expansion and have used the orthonormality of the spherical harmonics to do the integral over $\hat{k}$ in going from the first line to the second. From this expression, we can read off the exact result for the 2D Fourier transform of the projected y-profile due to a cluster of mass M at redshift z:

$$\begin{aligned} \tilde{y}_{2D}(\ell; M, z) &= \int d\chi' a(\chi') \int \frac{2}{\pi} k^2 \, dk \, j_\ell(k\chi') j_\ell(k\chi(z)) \, \tilde{y}_{3D}(k; M, z) \\ &= \frac{1}{\sqrt{\chi(z)}} \int \frac{d\chi'}{\sqrt{\chi'}} a(\chi') \int k \, dk \, J_{\ell+1/2}(k\chi') J_{\ell+1/2}(k\chi) \times \\ &\quad \tilde{y}_{3D}(k; M, z) \,, \end{aligned} \tag{1.16}$$

where we have rewritten the spherical Bessel functions in terms of Bessel functions of the first kind using $j_\nu(x) = \sqrt{\frac{\pi}{2x}} J_{\nu+1/2}(x)$ and we have explicitly included a possible dependence of the y_{3D} profile on redshift (in addition to mass).

The total one-halo term in the tSZ power spectrum is then given by the sum of the individual contributions from every cluster in the universe:

$$C_\ell^{y,1h} = \int dz \frac{d^2V}{dz d\Omega} \int dM \frac{dn(M,z)}{dM} \left| \tilde{y}_{2D}(\ell; M, z) \right|^2 \,, \tag{1.17}$$

where $d^2V/dzd\Omega = c\chi^2(z)/H(z)$ is the comoving volume element per steradian. Substituting Eq. (1.16) into this expression, converting the χ' integral to a redshift integral, and

rearranging the order of the integrals then yields the final result for the exact one-halo term:

$$C_\ell^{y,1h} = \int \frac{dz}{\chi(z)} \frac{d^2V}{dz d\Omega} \int dM \frac{dn}{dM} \times \left| \int k\, dk\, J_{\ell+1/2}(k\chi(z)) \tilde{y}_{3D}(k; M, z) \int \frac{c\, dz'}{H(z')(1+z')\sqrt{\chi(z')}} J_{\ell+1/2}(k\chi(z')) \right|^2 \quad (1.18)$$

Note that this expression is exact: no flat-sky approximation (or any other) has been used in deriving Eq. (1.18).

In order to recover the flat-sky (i.e., small-angle) limit of Eq. (1.18), we use the following $\ell \to \infty$ limit for the spherical Bessel functions:

$$j_\ell(x) \to \sqrt{\frac{\pi}{2\ell+1}} \delta_D(\ell + 1/2 - x)\,. \quad (1.19)$$

Applying this limit to the first line of Eq. (1.16) yields

$$\tilde{y}_{2D}(\ell \gg 1; M, z) \approx \frac{a(z)}{\chi^2(z)} \tilde{y}_{3D}\left(\frac{\ell+1/2}{\chi(z)}; M, z\right)\,. \quad (1.20)$$

Using Eq. (1.10), we can simplify this expression into a familiar form:

$$\begin{aligned} \tilde{y}_{2D}(\ell \gg 1; M, z) &\approx \frac{a(z)}{\chi^2(z)} \int dr\, 4\pi r^2 \frac{\sin((\ell+1/2)r/\chi)}{(\ell+1/2)r/\chi} y_{3D}(r; M) \\ &= \frac{4\pi r_s}{\ell_s^2} \int dx\, x^2 \frac{\sin((\ell+1/2)x/\ell_s)}{(\ell+1/2)x/\ell_s} y_{3D}(x; M)\,, \end{aligned} \quad (1.21)$$

where we have performed the following change of variables in the integral over y_{3D}: $x \equiv a(z)r/r_s$, where r_s is a characteristic scale radius of the y_{3D} profile. Finally, $\ell_s = a(z)\chi(z)/r_s = d_A(z)/r_s$ is the characteristic multipole moment associated with the scale radius, with $d_A(z)$ the angular diameter distance. Note that the change of variables involved the scale factor because we transformed from comoving coordinates to physical coordinates. Eq. (1.21) is identical to the quantity $\tilde{y}_\ell(M, z)$ defined in (for example) Eq. (2) of [6], although we have explicitly used $\ell+1/2$ rather than ℓ. This is both technically

correct and reduces the error in the approximation from $\mathcal{O}(\ell^{-1})$ to $\mathcal{O}(\ell^{-2})$ [30]. Eq. (1.21) is simply the flat-sky limit of Eq. (1.16). Following the long-standing convention, we will use the same definition as that established in [6]:

$$\begin{aligned} \tilde{y}_\ell(M,z) &\equiv \tilde{y}_{2D}(\ell \gg 1; M, z) \\ &\approx \frac{4\pi r_s}{\ell_s^2} \int dx\, x^2 \frac{\sin((\ell+1/2)x/\ell_s)}{(\ell+1/2)x/\ell_s} y_{3D}(x; M, z)\,. \end{aligned} \tag{1.22}$$

The flat-sky limit of the one-halo term given in Eq. (1.18) is thus given by

$$C_{\ell\gg1}^{y,1h} \approx \int dz \frac{d^2V}{dz d\Omega} \int dM \frac{dn(M,z)}{dM} \left|\tilde{y}_\ell(M,z)\right|^2\,, \tag{1.23}$$

as written down in (for example) Eq. (1) of [6].

Evaluating Eq. (1.18) numerically is somewhat computationally expensive, as it contains five nested integrals (including the Fourier transform to obtain $\tilde{y}_{3D}$), two of which involve highly oscillatory Bessel functions. However, the flat-sky limit in Eq. (1.23) contains only three nested integrals, and involves no oscillatory functions. Furthermore, for a WMAP9 cosmology, we find that the flat-sky result in Eq. (1.23) only overestimates the exact result in Eq. (1.18) by $\approx 13\%$, 5%, and 3% at $\ell = 2$, 10, and 20, respectively. By $\ell = 60$, the two results are identical within our numerical precision. In addition, at $\ell = 2$ where the correction is largest, the one-halo term is only $\approx 67\%$ as large as the two-halo term, and thus the total tSZ power spectrum is only overestimated by $\approx 5\%$. Note that for non-Gaussian cosmologies this overestimate is far smaller, because the two-halo term dominates by a much larger amount at low-ℓ than in a Gaussian cosmology (for example, the two-halo term at $\ell = 2$ is 2.2 times as large as the one-halo term for $f_{\rm NL} = 50$). Given the small size of this correction and the significant computational expense required to evaluate the exact expression, we thus use the flat-sky result in Eqs. (1.22) and (1.23) to compute the one-halo contribution to the tSZ power spectrum in Chapter 2, and for the tSZ – CMB lensing cross-power spectrum in Chapter 3.

1.2.5 The Two-Halo Term

We now compute the contribution of the two-halo term to the tSZ power spectrum. We do the exact calculation first, and then consider the Limber-approximated (small-angle) limit. The exact result is necessary for studying the signature of the scale-dependent halo bias induced by primordial non-Gaussianity on the tSZ power spectrum, since the effect is only significant at very low ℓ. To calculate the two-halo contribution to the angular power spectrum of the (2D) Compton-y field, we project Eq. (1.9) along the LOS using the projection kernel in Eq. (1.6), which gives:

$$
\begin{aligned}
C_\ell^{y,2h} &= \int d\chi_1 W^y(\chi_1) \int d\chi_2 W^y(\chi_2) \int \frac{2k^2 dk}{\pi} j_\ell(k\chi_1) j_\ell(k\chi_2) P^{2h}_{y_{3D}}(k) \\
&= \int d\chi_1 \frac{W^y(\chi_1)}{\sqrt{\chi_1}} \int d\chi_2 \frac{W^y(\chi_2)}{\sqrt{\chi_2}} \int dk\, k\, J_{\ell+1/2}(k\chi_1) J_{\ell+1/2}(k\chi_2) P^{2h}_{y_{3D}}(k) \\
&= \int d\chi_1 \frac{a(\chi_1)}{\sqrt{\chi_1}} \int d\chi_2 \frac{a(\chi_2)}{\sqrt{\chi_2}} \int dk\, k\, J_{\ell+1/2}(k\chi_1) J_{\ell+1/2}(k\chi_2) P^{2h}_{y_{3D}}(k) \\
&= \int dz_1 \frac{c}{H(z_1)} \frac{a(z_1)}{\sqrt{\chi(z_1)}} \int dz_2 \frac{c}{H(z_2)} \frac{a(z_2)}{\sqrt{\chi(z_2)}} \int dk\, k\, J_{\ell+1/2}(k\chi(z_1)) \times \\
&\quad J_{\ell+1/2}(k\chi(z_2)) \int dM_1 \frac{dn}{dM_1} b(k, M_1, z_1) \tilde{y}_{3D}(k; M_1, z_1) \int dM_2 \frac{dn}{dM_2} \times \\
&\quad b(k, M_2, z_2) \tilde{y}_{3D}(k; M_2, z_2) P_{\rm lin}(k; z_1, z_2) \\
&= \int \frac{dz_1}{\sqrt{\chi(z_1)}} \frac{d^2V}{dz_1 d\Omega} \int \frac{dz_2}{\sqrt{\chi(z_2)}} \frac{d^2V}{dz_2 d\Omega} \int dk\, k\, J_{\ell+1/2}(k\chi(z_1)) \times \\
&\quad J_{\ell+1/2}(k\chi(z_2)) P_{\rm lin}(k; z_1, z_2) \int dM_1 \frac{dn}{dM_1} b(k, M_1, z_1) \tilde{y}_{k\chi_1}(M_1, z_1) \times \\
&\quad \int dM_2 \frac{dn}{dM_2} b(k, M_2, z_2) \tilde{y}_{k\chi_2}(M_2, z_2) \\
&= \int dk\, k\, \frac{P_{\rm lin}(k; z_{in})}{D^2(z_{in})} \times \\
&\quad \left[\int \frac{dz}{\sqrt{\chi(z)}} \frac{d^2V}{dz d\Omega} J_{\ell+1/2}(k\chi(z)) D(z) \int dM \frac{dn}{dM} b(k, M, z) \tilde{y}_{k\chi(z)}(M, z) \right]^2 \quad (1.24)
\end{aligned}
$$

where we have again used $j_\nu(x) = \sqrt{\frac{\pi}{2x}} J_{\nu+1/2}(x)$ and the notation $P_{\rm lin}(k; z_1, z_2)$ refers to the re-scaling of the linear matter power spectrum by the growth factor $D(z)$:

$$P_{\rm lin}(k; z_1, z_2) = \frac{D(z_1)D(z_2)}{D^2(z_{in})} P_{\rm lin}(k; z_{in}) \,, \tag{1.25}$$

where z_{in} is a reasonable input redshift for the linear theory matter power spectrum (e.g., our choice is $z_{in} = 30$). Also, note that we have explicitly included the possible scale-dependence of the bias, $b(k, M, z)$, as arises in cosmologies with local primordial non-Gaussianity. Finally, the notation $\tilde{y}_{k\chi}(M, z)$ in Eq. (1.24) refers to the expression for $\tilde{y}_\ell(M, z)$ given in Eq. (1.22) evaluated with $\ell+1/2 = k\chi$. This notation is simply a mathematical convenience; no flat-sky or Limber approximation was used in deriving Eq. (1.24), and no ℓ appears in $\tilde{y}_{k\chi}(M, z)$. Note that this expression only requires the evaluation of four nested integrals (whereas the exact one-halo term required five), although the redshift integrand is highly oscillatory due to the Bessel function.

In order to recover the Limber-approximated (i.e., small-angle) limit of Eq. (1.24), we again use Eq. (1.19) given above. This step is most easily accomplished starting from the

first line of the derivation that led to Eq. (1.24), which yields:

$$
\begin{aligned}
C_{\ell\gg1}^{y,2h} &\approx \int d\chi_1 W^y(\chi_1) \int d\chi_2 W^y(\chi_2) \times \\
&\quad \int \frac{k^2 dk}{\ell+1/2} \delta_D(\ell+1/2-k\chi_1)\delta_D(\ell+1/2-k\chi_2) P_{y_{3D}}^{2h}(k) \\
&= \int \frac{d\chi_1}{\chi_1} a(\chi_1) \int d\chi_2 a(\chi_2) \times \\
&\quad \int \frac{k^2 dk}{\ell+1/2} \delta_D\left(k-\frac{\ell+1/2}{\chi_1}\right) \delta_D(\ell+1/2-k\chi_2) P_{y_{3D}}^{2h}(k) \\
&= \int \frac{d\chi_1}{\chi_1^3} a(\chi_1) \int d\chi_2 a(\chi_2)(\ell+1/2)\delta_D\left(\ell+1/2-\frac{\ell+1/2}{\chi_1}\chi_2\right) P_{y_{3D}}^{2h}\left(\frac{\ell+1/2}{\chi_1}\right) \\
&= \int \frac{d\chi_1}{\chi_1^2} a(\chi_1) \int d\chi_2 a(\chi_2)\delta_D\left(\chi_2-\chi_1\right) P_{y_{3D}}^{2h}\left(\frac{\ell+1/2}{\chi_1}\right) \\
&= \int d\chi \left(\frac{a(\chi)}{\chi}\right)^2 P_{y_{3D}}^{2h}\left(\frac{\ell+1/2}{\chi}\right) \\
&= \int d\chi \left(\frac{a(\chi)}{\chi}\right)^2 \left[\int dM \frac{dn}{dM} b(M) \tilde{y}_{3D}\left(\frac{\ell+1/2}{\chi}; M\right)\right]^2 P_{\rm lin}\left(\frac{\ell+1/2}{\chi}\right) \\
&= \int dz \frac{c}{H(z)} \frac{a^2}{\chi^2} \left[\int dM \frac{dn}{dM} b(M) \int dr\, 4\pi r^2 \frac{\sin((\ell+1/2)r/\chi)}{(\ell+1/2)r/\chi} y_{3D}(r;M)\right]^2 \times \\
&\quad P_{\rm lin}\left(\frac{\ell+1/2}{\chi}\right) \\
&= \int dz \frac{d^2V}{dz d\Omega} \frac{a^2}{\chi^4} \left[\int dM \frac{dn}{dM} b(M) \int dr\, 4\pi r^2 \frac{\sin((\ell+1/2)r/\chi)}{(\ell+1/2)r/\chi} y_{3D}(r;M)\right]^2 \times \\
&\quad P_{\rm lin}\left(\frac{\ell+1/2}{\chi}\right) \\
&= \int dz \frac{d^2V}{dz d\Omega} \left[\int dM \frac{dn}{dM} b(M) \frac{4\pi r_s}{\ell_s^2} \int dx\, x^2 \frac{\sin((\ell+1/2)x/\ell_s)}{(\ell+1/2)x/\ell_s} y_{3D}(x;M)\right]^2 \times \\
&\quad P_{\rm lin}\left(\frac{\ell+1/2}{\chi}\right) \\
&= \int dz \frac{d^2V}{dz d\Omega} \left[\int dM \frac{dn(M,z)}{dM} b(k,M,z) \tilde{y}_\ell(M,z)\right]^2 P_{\rm lin}\left(\frac{\ell+1/2}{\chi(z)}; z\right) , \qquad (1.26)
\end{aligned}
$$

where we have restored all of the mass, redshift, and scale dependences in the final expression, and $\tilde{y}_\ell(M,z)$ is given by Eq. (1.22). Eq. (1.26) precisely matches the result written down for the Limber-approximated two-halo term in [31], although again we have explicitly used

$\ell + 1/2$ in the Limber approximation (rather than ℓ), as this choice increases the accuracy of the calculation (and is formally correct).

As noted above, the exact expression for the two-halo term in Eq. (1.24) requires the evaluation of four nested integrals; the Limber-approximated result in Eq. (1.26) requires three. Thus, the computational expense is not vastly different, although the Limber case is roughly an order of magnitude faster. For a WMAP9 cosmology, we find that the Limber result in Eq. (1.26) overestimates the exact result in Eq. (1.24) by $\approx 7\%$, 2%, and 1% at $\ell = 2$, 4, and 20, respectively. By $\ell = 30$, the two results are identical within our numerical precision. Note that although the fractional difference between the exact and flat-sky results at low-ℓ is smaller for the two-halo term than for the one-halo term, the two-halo term dominates in this regime, and thus greater precision is required in its computation in order to predict the total C_ℓ^y precisely. Note that using the exact result for the two-halo term is more important for $f_{\rm NL} \neq 0$ cosmologies, for which the Limber approximation has been found to be less accurate [32]. For a cosmology with $f_{\rm NL} = 100$, we find that the Limber result in Eq. (1.26) overestimates the exact result in Eq. (1.24) by $\approx 18\%$, 5%, and 1% at $\ell = 2$, 4, and 20, respectively. To be conservative, we thus use the exact result for the two-halo term for all calculations at $\ell < 50$ (in Chapter 2), while we use the Limber-approximated result at higher multipoles. We note that the fairly small size of the correction to the Limber approximation, even at $\ell = 2$, can be explained using arguments from [30] regarding the width of the tSZ projection kernel, which is very broad (see Eq. (1.6)). In particular, their results imply that the Limber approximation is reliable when $\ell + 1/2 \gtrsim \bar{r}/\sigma_r$, where $\bar{r}$ is the distance at which the projection kernel peaks and σ_r is the width of the projection kernel, which are effectively comparable for the tSZ signal. Thus the Limber approximation is reliable for $\ell + 1/2 \gtrsim 1$, which our numerical calculations verify.

1.2.6 The Covariance Matrix

In order to obtain a complete expression for the covariance matrix of the tSZ power spectrum, we need to compute the tSZ angular trispectrum. Trispectrum configurations are quadrilaterals in ℓ-space, characterized by four sides and one diagonal. The configurations that contribute to the power spectrum covariance matrix are of a "collapsed" shape characterized by two lines of length ℓ and ℓ' with zero diagonal [6]. Analogous derivations to those that led to Eqs. (1.18) and (1.23) lead to the exact and flat sky-approximated expressions for the one-halo contribution to these configurations of the tSZ trispectrum:

$$\begin{aligned} T^{y,1h}_{\ell\ell'} &= \int \frac{dz}{\chi^2(z)} \frac{d^2V}{dz d\Omega} \int dM \frac{dn}{dM} \times \\ &\left| \int k\, dk\, J_{\ell+1/2}(k\chi(z)) \tilde{y}_{3D}(k;M,z) \int \frac{c\, dz'}{H(z')(1+z')\sqrt{\chi(z')}} J_{\ell+1/2}(k\chi(z')) \right|^2 \times \\ &\left| \int k'\, dk'\, J_{\ell'+1/2}(k'\chi(z)) \tilde{y}_{3D}(k';M,z) \int \frac{c\, dz''}{H(z'')(1+z'')\sqrt{\chi(z'')}} J_{\ell'+1/2}(k'\chi(z'')) \right|^2 \\ &\text{(exact)} \end{aligned} \tag{1.27}$$

$$T^{y,1h}_{\ell\ell'\gg 1} \approx \int dz \frac{d^2V}{dz d\Omega} \int dM \frac{dn}{dM} \left|\tilde{y}_\ell(M,z)\right|^2 \left|\tilde{y}_{\ell'}(M,z)\right|^2 \text{ (flat sky)}\,. \tag{1.28}$$

For computational efficiency, we choose to implement the flat-sky result at all ℓ values in our calculations. Based on the errors discussed earlier for the flat-sky version of the one-halo contribution to the power spectrum compared to the exact result, we estimate that the error in the trispectrum due to this approximation may be $\sim 25-30\%$ at $\ell = 2$ (where the discrepancy would be maximal). However, the only parameter forecast in Chapter 2 that would likely be affected is the $f_{\rm NL}$ constraint (due to the necessity of measuring the influence of the scale-dependent bias in order to constrain this parameter), for which we do not find competitive results compared to other probes. If our forecasts for $f_{\rm NL}$ were in need of percent-level precision, we would certainly want to use the exact trispectrum; however, this is clearly not the case, and thus we neglect this small error in our results (the constraints

on all other parameters are insensitive to moderate changes in the errors at the lowest few ℓ values). Moreover, in the masked cases (which present the greatest promise for cosmological constraints), the trispectrum contribution is heavily suppressed at low-ℓ (see Figs. 2.19, 2.20, and 2.21), and the total errors are dominated by the Gaussian term. Thus, for the masked cases, the exact vs. flat-sky correction should be vanishingly small.

Note that we neglect the two-halo, three-halo and four-halo contributions to the trispectrum, as it is dominated even more heavily by the Poisson term than the power spectrum is [33]. The two-halo term will contribute to some extent at low-ℓ, but it is unlikely that higher-order terms will be significant even in this regime. For the masked calculations in Chapter 2, the two-halo term may be somewhat important, though likely only at very low-ℓ, where, as we have argued above, it seems we do not need percent-level accuracy on the errors (since the forecasts for $f_{\rm NL}$ are not particularly promising, and it is the only parameter very sensitive to this region of the power spectrum).

The full covariance matrix of the tSZ power spectrum, $M^y_{\ell\ell'}$ is then given by [6]:

$$\begin{aligned} M^y_{\ell\ell'} &\equiv \langle (C^{y,\rm obs}_\ell - C^y_\ell)(C^{y,\rm obs}_{\ell'} - C^y_{\ell'})\rangle \\ &= \frac{1}{4\pi f_{\rm sky}} \left(\frac{4\pi (C^y_\ell + N_\ell)^2}{\ell + 1/2} \delta_{\ell\ell'} + T^y_{\ell\ell'} \right) , \end{aligned} \tag{1.29}$$

where the angular brackets denote an ensemble average, $f_{\rm sky}$ is the sky fraction covered by a given experiment (we assume $f_{\rm sky} = 0.7$ throughout Chapter 2), N_ℓ is the power spectrum due to instrumental noise after multifrequency subtraction (computed for Planck and PIXIE in Section 2.5), and we approximate $T^y_{\ell\ell'} \approx T^{y,1h}_{\ell\ell' \gg 1}$. Note that Eq. (1.29) does not include the so-called "halo sample variance" term, as discussed in Section 2.6, as this term is negligible for a (nearly) full-sky survey.

We can then compute the covariance matrix $\mathrm{Cov}(p_i, p_j)$ for the cosmological and astrophysical parameters of interest $p_i = \{\Omega_b h^2, \Omega_c h^2, \Omega_\Lambda, \sigma_8, n_s, f_{\rm NL}, M_\nu, C_{P_0}, C_\beta\}$:

$$\mathrm{Cov}(p_i, p_j) = \left[\frac{\partial C_\ell^y}{\partial p_i} \left(M_{\ell\ell'}^y\right)^{-1} \frac{\partial C_{\ell'}^y}{\partial p_j} \right]^{-1} , \tag{1.30}$$

where summation over the repeated indices is implied. The Fisher matrix for these parameters is then simply given by the inverse of their covariance matrix:

$$F_{ij} = \mathrm{Cov}^{-1}(p_i, p_j) \, . \tag{1.31}$$

The Fisher matrix encodes the constraining power of the tSZ power spectrum on the cosmological and astrophysical parameters.

Bibliography

[1] Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, ApJS, 208, 20

[2] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5062

[3] Swetz, D. S., Ade, P. A. R., Amiri, M., et al. 2011, ApJS, 194, 41

[4] Carlstrom, J. E., Ade, P. A. R., Aird, K. A., et al. 2011, PASP, 123, 568

[5] Sunyaev, R. A., & Zeldovich, Y. B. 1970, Ap&SS, 7, 3

[6] Komatsu, E., & Seljak, U. 2002, Mon. Not. R. Astron. Soc, 336, 1256

[7] Reichardt, C. L., Stalder, B., Bleem, L. E., et al. 2013, ApJ, 763, 127

[8] Hasselfield, M., Hilton, M., Marriage, T. A., et al. 2013, J. Cosmology Astropart. Phys., 7, 8

[9] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5080

[10] Niemack, M. D., Ade, P. A. R., Aguirre, J., et al. 2010, Proc. SPIE, 7741

[11] Austermann, J. E., Aird, K. A., Beall, J. A., et al. 2012, Proc. SPIE, 8452

[12] Reese, E. D., Mroczkowski, T., Menanteau, F., et al. 2011, arXiv:1108.3343

[13] Plagge, T. J., Marrone, D. P., Abdulla, Z., et al. 2012, arXiv:1203.2175

[14] Lancaster, K., Birkinshaw, M., Gawroński, M. P., et al. 2011, Mon. Not. R. Astron. Soc, 418, 1441

[15] AMI Consortium, Shimwell, T. W., Rodriguez-Gonzalvez, C., et al. 2011, arXiv:1101.5590

[16] Marriage, T. A., Acquaviva, V., Ade, P. A. R., et al. 2011, ApJ, 737, 61

[17] Williamson, R., Benson, B. A., High, F. W., et al. 2011, ApJ, 738, 139

[18] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2011, A&A, 536, A8

[19] Dunkley, J., Hlozek, R., Sievers, J., et al. 2011, ApJ, 739, 52

[20] Sievers, J. L., Hlozek, R. A., Nolta, M. R., et al. 2013, arXiv:1301.0824

[21] Reichardt, C. L., Shaw, L., Zahn, O., et al. 2012, ApJ, 755, 70

[22] Story, K. T., Reichardt, C. L., Hou, Z., et al. 2012, arXiv:1210.7231

[23] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5081

[24] Wilson, M. J., Sherwin, B. D., Hill, J. C., et al. 2012, Phys. Rev. D, 86, 122005

[25] Rubiño-Martín, J. A., & Sunyaev, R. A. 2003, Mon. Not. R. Astron. Soc, 344, 1155

[26] Crawford, T. M., Schaffer, K. K., Bhattacharya, S., et al. 2013, arXiv:1303.3535

[27] Hill, J. C., & Sherwin, B. D. 2013, Phys. Rev. D, 87, 023527

[28] Bhattacharya, S., Nagai, D., Shaw, L., Crawford, T., & Holder, G. P. 2012, ApJ, 760, 5

[29] Nozawa, S., Itoh, N., Suda, Y., & Ohhata, Y. 2006, Nuovo Cimento B Serie, 121, 487

[30] Loverde, M., & Afshordi, N. 2008, Phys. Rev. D, 78, 123506

[31] Komatsu, E., & Kitayama, T. 1999, ApJ, 526, L1

[32] Pillepich, A., Porciani, C., & Reiprich, T. H. 2012, Mon. Not. R. Astron. Soc, 422, 44

[33] Cooray, A. 2001, Phys. Rev. D, 64, 063514

Chapter 2

Cosmology from the Thermal Sunyaev-Zel'dovich Power Spectrum: Primordial non-Gaussianity and Massive Neutrinos

2.1 Abstract

We carry out a comprehensive analysis of the possible constraints on cosmological and astrophysical parameters achievable with measurements of the thermal Sunyaev-Zel'dovich (tSZ) power spectrum from upcoming full-sky CMB observations, with a particular focus on one-parameter extensions to the ΛCDM standard model involving local primordial non-Gaussianity (described by $f_{\rm NL}$) and massive neutrinos (described by M_ν). We include all of the relevant physical effects due to these additional parameters, including the change to the halo mass function and the scale-dependent halo bias induced by local primordial non-Gaussianity. We model the pressure profile of the intracluster medium (ICM) using a parametrized fit that agrees well with existing observations, and include uncertainty in

the ICM modeling by including the overall normalization and outer logarithmic slope of the profile as free parameters. We compute forecasts for Planck, PIXIE, and a cosmic variance (CV)-limited experiment, using multifrequency subtraction to remove foregrounds and implementing two masking criteria based on the ROSAT and eROSITA cluster catalogs to reduce the significant CV errors at low multipoles. We find that Planck can detect the tSZ power spectrum with $> 30\sigma$ significance, regardless of the masking scenario. However, neither Planck or PIXIE is likely to provide competitive constraints on $f_{\rm NL}$ from the tSZ power spectrum due to CV noise at low-ℓ overwhelming the unique signature of the scale-dependent bias. A future CV-limited experiment could provide a 3σ detection of $f_{\rm NL} \simeq 37$, which is the WMAP9 maximum-likelihood result. The outlook for neutrino masses is more optimistic: Planck can reach levels comparable to the current upper bounds $\lesssim 0.3$ eV with conservative assumptions about the ICM; stronger ICM priors could allow Planck to provide $1-2\sigma$ evidence for massive neutrinos from the tSZ power spectrum, depending on the true value of the sum of the neutrino masses. We also forecast a $< 10\%$ constraint on the outer slope of the ICM pressure profile using the unmasked Planck tSZ power spectrum.

2.2 Introduction

The thermal Sunyaev-Zel'dovich (tSZ) effect is a spectral distortion of the cosmic microwave background (CMB) arising from the inverse Compton scattering of CMB photons off hot electrons that lie along the line-of-sight (LOS) [1]. The tSZ signal is dominated by contributions from electrons located in the intracluster medium (ICM) of massive galaxy clusters. However, contributions from lower-mass, higher-redshift objects are important in statistical measures of the unresolved, diffuse tSZ signal, e.g., the tSZ power spectrum. First detections of the amplitude of the angular power spectrum of the tSZ effect have been made by the Atacama Cosmology Telescope (ACT)[1] and the South Pole Telescope (SPT)[2] [11, 12, 13, 14].

[1] `http://www.princeton.edu/act/`
[2] `http://pole.uchicago.edu/`

More recently, the Planck collaboration has reported a measurement of the tSZ power spectrum over a wide range of angular scales [137].

At the time of writing (March 2013), tSZ power spectrum detections have been limited to measurements or constraints on the power at a single multipole (typically $\ell = 3000$) because ACT and SPT do not have sufficient frequency coverage to fully separate the tSZ signal from other components in the microwave sky using its unique spectral signature. However, this situation will shortly change with the imminent release of full-sky temperature maps from the Planck satellite[3]. Planck has nine frequency channels that span the spectral region around the tSZ null frequency ≈ 218 GHz. Thus, it should be possible to separate the tSZ signal from other components in the sky maps to high accuracy, allowing for a measurement of the tSZ power spectrum over a wide range of multipoles, possibly $\sim 100 \lesssim \ell \lesssim 1500$, as we demonstrate in this chapter. The proposed Primordial Inflation Explorer (PIXIE) experiment [24] will also be able to detect the tSZ power spectrum at high significance, as its wide spectral coverage and high spectral resolution will allow for very accurate extraction of the tSZ signal. However, its angular resolution is much lower than Planck's, and thus the tSZ power spectrum will be measured over a much smaller range of multipoles ($\ell \lesssim 200$). However, the PIXIE data (after masking using X-ray cluster catalogs — see below) will permit tSZ measurements on large angular scales ($\ell \lesssim 100$) that are essentially inaccessible to Planck due to its noise levels; these multipoles are precisely where one would expect the signature of scale-dependent bias induced by primordial non-Gaussianity to arise in the tSZ power spectrum. Assessing the amplitude and detectability of this signature is a primary motivation for this chapter.

The tSZ power spectrum has been suggested as a potential cosmological probe by a number of authors over the past few decades (e.g., [18, 20, 19]). Nearly all studies in the last decade have focused on the small-scale tSZ power spectrum ($\ell \gtrsim 1000$) due to its role as a foreground in high angular resolution CMB measurements, and likely because without

[3] `http://planck.esa.int`

multi-frequency information, the tSZ signal has only been able to be isolated by looking for its effects on small scales (e.g., using an ℓ-space filter to upweight tSZ-dominated small angular scales in CMB maps). Much of this work was driven by the realization that the tSZ power spectrum is a very sensitive probe of the amplitude of matter density fluctuations, σ_8 [19]. The advent of multi-frequency data promises measurements of the large-scale tSZ power spectrum very shortly, and thus we believe it is timely to reassess its value as a cosmological probe, including parameters beyond σ_8 and including a realistic treatment of the uncertainties due to modeling of the ICM. We build on the work of [20] to compute the full angular power spectrum of the tSZ effect, including both the one- and two-halo terms, and moving beyond the Limber/flat-sky approximations where necessary.

Our primary interest is in assessing constraints from the tSZ power spectrum on currently unknown parameters beyond the ΛCDM standard model: the amplitude of local primordial non-Gaussianity, $f_{\rm NL}$, and the sum of the neutrino masses, $M_\nu \equiv \sum m_\nu$. The values of these parameters are currently unknown, and determining their values is a key goal of modern cosmology.

Primordial non-Gaussianity is one of the few known probes of the physics of inflation. Models of single-field, minimally-coupled slow-roll inflation predict negligibly small deviations from Gaussianity in the initial curvature perturbations [25, 26]. In particular, a detection of a non-zero bispectrum amplitude in the so-called "squeezed" limit ($k_1 \ll k_2, k_3$) would falsify essentially all single-field models of inflation [26, 27]. This type of non-Gaussianity can be parametrized using the "local" model, in which $f_{\rm NL}$ describes the lowest-order deviation from Gaussianity [28, 29, 30]:

$$\Phi(\vec{x}) = \Phi_G(\vec{x}) + f_{\rm NL}\left(\Phi_G^2(\vec{x}) - \langle\Phi_G^2\rangle\right) + \cdots , \tag{2.1}$$

where Φ is the primordial potential and Φ_G is a Gaussian field. Note that $\Phi = \frac{3}{5}\zeta$, where ζ is the initial adiabatic curvature perturbation. Local-type non-Gaussianity can be generated

in multi-field inflationary scenarios, such as the curvaton model [31, 32, 33], or by non-inflationary models for the generation of perturbations, such as the new ekpyrotic/cyclic scenario [34, 35, 36]. It is perhaps most interesting when viewed as a method with which to rule out single-field inflation, however. Current constraints on $f_{\rm NL}$ are consistent with zero [37, 38], but the errors will shrink significantly very soon with the imminent CMB results from Planck. We review the effects of $f_{\rm NL} \neq 0$ on the large-scale structure of the universe in Section 2.3.

In contrast to primordial non-Gaussianity, massive neutrinos are certain to exist at a level that will be detectable within the next decade or so; neutrino oscillation experiments have precisely measured the differences between the squared masses of the three known species, leading to a lower bound of ≈ 0.05 eV on the total summed mass [39]. The remaining questions surround their detailed properties, especially their absolute mass scale. The contribution of massive neutrinos to the total energy density of the universe today can be expressed as

$$\Omega_\nu \approx \frac{M_\nu}{93.14\,h^2\ \mathrm{eV}} \approx 0.0078 \frac{M_\nu}{0.1\ \mathrm{eV}}\,, \tag{2.2}$$

where M_ν is the sum of the masses of the three known neutrino species. Although this contribution appears to be small, massive neutrinos can have a significant influence on the small-scale matter power spectrum. Due to their large thermal velocities, neutrinos free-stream out of gravitational potential wells on scales below their free-streaming scale [41, 40]. This suppresses power on scales below the free-streaming scale. Current upper bounds from various cosmological probes assuming a flat ΛCDM+M_ν model are in the range $M_\nu \lesssim 0.3$ eV [12, 101, 42, 43], although a 3σ detection near this mass scale was recently claimed in [44]. Should the true total mass turn out to be near 0.3 eV, its effect on the tSZ power spectrum may be marginally detectable using the Planck data even for fairly conservative assumptions about the ICM physics, as we show in this chapter. With stronger ICM priors,

Planck could achieve a $\sim 2-3\sigma$ detection for masses at this scale, using only the primordial CMB temperature power spectrum and the tSZ power spectrum. We review the effects of massive neutrinos on the large-scale structure of the universe in Section 2.3.

In addition to the effects of both known and currently unknown cosmological parameters, we also model the effects of the physics of the ICM on the tSZ power spectrum. This subject has attracted intense scrutiny in recent years after the early measurements of tSZ power from ACT [11] and SPT [45] were significantly lower than the values predicted from existing ICM pressure profile models (e.g., [122]) in combination with WMAP5 cosmological parameters. Subsequent ICM modeling efforts have ranged from fully analytic approaches (e.g., [123]) to cosmological hydrodynamics simulations (e.g., [117, 116]), with other authors adopting semi-analytic approaches between these extremes (e.g., [46, 118]), in which dark matter-only N-body simulations are post-processed to include baryonic physics according to various prescriptions. In addition, recent SZ and X-ray observations have continued to further constrain the ICM pressure profile from data, although these results are generally limited to fairly massive, nearby systems (e.g., [119, 120]). We choose to adopt a parametrized form of the ICM pressure profile known as the Generalized NFW (GNFW) profile, with our fiducial parameter values chosen to match the constrained pressure profile fit from hydrodynamical simulations in [116]. In order to account for uncertainty in the ICM physics, we free two of the parameters in the pressure profile (the overall normalization and the outer logarithmic slope) and treat them as additional parameters in our model. This approach is discussed in detail in Section 2.4.2. We use this profile to compute the tSZ power spectrum following the halo model approach.

In addition to a model for the tSZ signal, we must compute the tSZ power spectrum covariance matrix in order to forecast parameter constraints. There are two important issues that must be considered in computing the expected errors or signal-to-noise ratio (SNR) for a measurement of the tSZ power spectrum. First, we must assess how well the tSZ signal can actually be separated from the other components in maps of the microwave sky, including the

primordial CMB, thermal dust, point sources, and so on. Following [47] and [48], we choose to implement a multi-frequency subtraction technique that takes advantage of the unique spectral signature of the tSZ effect, and also takes advantage of the current state of knowledge about the frequency- and multipole-dependence of the foregrounds. Although there are other approaches to this problem, such as using internal linear combination techniques to construct a Compton-y map from the individual frequency maps in a given experiment (e.g., [50, 51, 52]), we find this method to be fairly simple and robust. We describe these calculations in detail in Section 2.5.

Second, we must account for the extreme cosmic variance induced in the large-angle tSZ power spectrum by massive clusters at low redshifts. The one-halo term from these objects dominates the angular trispectrum of the tSZ signal, even down to very low multipoles [59]. The trispectrum represents a large non-Gaussian contribution to the covariance matrix of the tSZ power spectrum [19], which is especially problematic at low multipoles. However, the trispectrum can be greatly suppressed by masking massive, low-redshift clusters using existing X-ray, optical, or SZ catalogs [20]. This procedure can greatly increase the SNR for the tSZ power spectrum at low multipoles. For constraints on $f_{\rm NL}$ it also has the advantage of enhancing the relative importance of the two-halo term compared to the one-halo, thus showing greater sensitivity to the scale-dependent bias at low-ℓ. Moreover, even in a Gaussian cosmology, the inclusion of the two-halo term slightly changes the shape of the tSZ power spectrum, which likely helps break degeneracies amongst the several parameters which effectively only change the overall amplitude of the one-halo term; the relative enhancement of the two-halo term due to masking should help further in this regard. We consider two masking scenarios motivated by the flux limits of the cluster catalogs from all-sky surveys performed with the ROSAT[4] X-ray telescope and the upcoming eROSITA[5] X-ray telescope. These scenarios are detailed in Section 2.6; by default all calculations and figures are computed for the unmasked scenario unless they are labelled otherwise.

[4] `http://www.dlr.de/en/rosat`

[5] Extended ROentgen Survey with an Imaging Telescope Array, `http://www.mpe.mpg.de/erosita/`

Earlier studies have investigated the consequences of primordial non-Gaussianity for the tSZ power spectrum [53, 54], though we are not aware of any calculations including the two-halo term (and hence the scale-dependent bias) or detailed parameter constraint forecasts. We are also not aware of any previous work investigating constraints on massive neutrinos from the tSZ power spectrum, although previous authors have computed their signature [55]. Other studies have investigated detailed constraints on the primary ΛCDM parameters from the combination of CMB and tSZ power spectrum measurements [57]. Many authors have investigated constraints on $f_{\rm NL}$ and M_ν from cluster counts, though the results depend somewhat on the cluster selection technique and mass estimation method. Considering SZ cluster count studies only, [55] and [56] investigated constraints on M_ν from a Planck-derived catalog of SZ clusters (in combination with CMB temperature power spectrum data). The earlier paper found a 1σ uncertainty of $\Delta M_\nu \approx 0.28$ eV while the later paper found $\Delta M_\nu \approx 0.06 - 0.12$ eV; the authors state that the use of highly degenerate nuisance parameters degraded the results in the former study. In either case, the result is highly sensitive to uncertainties in the halo mass function, as the clusters included are deep in the exponential tail of the mass function. We expect that our results using the tSZ power spectrum should be less sensitive to uncertainties in the tail of the mass function, as the power spectrum is dominated at most angular scales by somewhat less massive objects ($10^{13} - 10^{14} M_\odot/h$) [19]. Finally, a very recent independent study [58] found $\Delta M_\nu \approx 0.3 - 0.4$ eV for Planck SZ cluster counts (with CMB temperature power spectrum information added), although they estimated that this bound could be improved to $\Delta M_\nu \approx 0.08$ eV with the inclusion of stronger priors on the ICM physics.

Our primary findings are as follows:

- The tSZ power spectrum can be detected with a total SNR > 30 using the imminent Planck data up to $\ell = 3000$, regardless of masking;

- The tSZ power spectrum can be detected with a total SNR between ≈ 6 and 22 using the future PIXIE data up to $\ell = 300$, with the result being sensitive to the level of masking applied to remove massive, nearby clusters;

- Adding the tSZ power spectrum information to the forecasted constraints from the Planck CMB temperature power spectrum and existing H_0 data is unlikely to significantly improve constraints on the primary cosmological parameters, but may give interesting constraints on the extensions we consider:

  - If the true value of $f_{\rm NL}$ is near the WMAP9 ML value of ≈ 37, a future CV-limited experiment combined with eROSITA-masking could provide a 3σ detection, completely independent of the primordial CMB temperature bispectrum; alternatively, PIXIE could give $1-2\sigma$ evidence for such a value of $f_{\rm NL}$ with this level of masking;

  - If the true value of M_ν is near 0.1 eV, the Planck tSZ power spectrum with eROSITA masking can provide upper limits competitive with the current upper bounds on M_ν ; with stronger external constraints on the ICM physics, Planck with eROSITA masking could provide $1-2\sigma$ evidence for massive neutrinos from the tSZ power spectrum, depending on the true neutrino mass;

- Regardless of the cosmological constraints, Planck will allow for a very tight constraint on the logarithmic slope of the ICM pressure profile in the outskirts of galaxy clusters, and may also provide some information on the overall normalization of the pressure profile (which sets the zero point of the $Y-M$ relation).

The remainder of this chapter is organized as follows. In Section 2.3, we describe our models for the halo mass function and halo bias, as well as the effects of primordial non-Gaussianity and massive neutrinos on large-scale structure. In Section 2.4, we describe our halo model-based calculation of the tSZ power spectrum, including the relevant ICM physics. We also demonstrate the different effects of each parameter in our model on the tSZ power

spectrum. In Section 2.5, we consider the extraction of the tSZ power spectrum from the other components in microwave sky maps via multifrequency subtraction techniques. Having determined the experimental noise levels, in Section 2.6 we detail our calculation of the covariance matrix of the tSZ power spectrum, and discuss the role of masking massive nearby clusters in reducing the low-ℓ cosmic variance. In Section 2.7, we use our tSZ results to forecast constraints on cosmological and astrophysical parameters from a variety of experimental set-ups and masking choices. We also compute the expected SNR of the tSZ power spectrum detection for each possible scenario. We discuss our results and conclude in Section 2.8. Finally, in Section 2.9, we provide a brief comparison between our forecasts and the Planck tSZ power spectrum results that were publicly released while this manuscript was under review [137].

The WMAP9+eCMB+BAO+H_0 maximum-likelihood parameters [69] define our fiducial model (see Section 2.4.3 for details). All masses are quoted in units of $M_\odot/h$, where $h \equiv H_0/(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}})$ and H_0 is the Hubble parameter today. All distances and wavenumbers are in comoving units of Mpc/h. All tSZ observables are computed at $\nu = 150$ GHz, since ACT and SPT have observed the tSZ signal at (or very near) this frequency, where the tSZ effect leads to a temperature decrement in the CMB along the LOS to a galaxy cluster.

2.3 Modeling Large-Scale Structure

In order to compute statistics of the tSZ signal, we need to model the comoving number density of halos as a function of mass and redshift (the halo mass function) and the bias of halos with respect to the underlying matter density field as a function of mass and redshift. Moreover, in order to extract constraints on $f_{\rm NL}$ and M_ν from the tSZ power spectrum, we must include the effects of these parameters on large-scale structure. We describe our approach to these computations in the following.

2.3.1 Halo Mass Function

We define the mass of a dark matter halo by the spherical overdensity (SO) criterion: $M_{\delta,c}$ ($M_{\delta,d}$) is the mass enclosed within a sphere of radius $r_{\delta,c}$ ($r_{\delta,d}$) such that the enclosed density is δ times the critical (mean matter) density at redshift z. To be clear, c subscripts refer to masses referenced to the critical density at redshift z, $\rho_{cr}(z) = 3H^2(z)/8\pi G$ with $H(z)$ the Hubble parameter at redshift z, whereas d subscripts refer to masses referenced to the mean matter density at redshift z, $\bar{\rho}_m(z) \equiv \bar{\rho}_m$ (this quantity is constant in comoving units).

We will generally work in terms of a particular SO mass, the virial mass, which we denote as M. The virial mass is the mass enclosed within a radius r_{vir} [70]:

$$r_{vir} = \left(\frac{3M}{4\pi \Delta_{cr}(z) \rho_{cr}(z)} \right)^{1/3}, \tag{2.3}$$

where $\Delta_{cr}(z) = 18\pi^2 + 82(\Omega(z)-1) - 39(\Omega(z)-1)^2$ and $\Omega(z) = \Omega_m(1+z)^3/(\Omega_m(1+z)^3+\Omega_\Lambda)$. For many calculations, we need to convert between M and various other SO masses (e.g., M_{200c} or M_{200d}). We use the NFW density profile [66] and the concentration-mass relation from [67] in order to do these conversions, which require solving the following non-linear equation for $r_{\delta,c}$ (or $r_{\delta,d}$):

$$\int_0^{r_{\delta,c}} 4\pi r'^2 \rho_{\rm NFW}(r', M, c_{vir}) dr' = \frac{4}{3}\pi r_{\delta,c}^3 \rho_{cr}(z) \delta \tag{2.4}$$

where $c_{vir} \equiv r_{vir}/r_{NFW}$ is the concentration parameter (r_{NFW} is the NFW scale radius) and we replace the critical density $\rho_{cr}(z)$ with the mean matter density $\bar{\rho}_m$ in this equation in order to obtain $r_{\delta,d}$ instead of $r_{\delta,c}$. After solving Eq. (2.4) to find $r_{\delta,c}$, we calculate $M_{\delta,c}$ via $M_{\delta,c} = \frac{4}{3}\pi r_{\delta,c}^3 \rho_{cr}(z)\delta$.

The halo mass function, $dn(M,z)/dM$ describes the comoving number density of halos per unit mass as a function of redshift. We employ the approach developed from early work by Press and Schechter [71] and subsequently refined by many other authors (e.g.,

[72, 73, 74, 75]):

$$\begin{aligned}\frac{dn(M,z)}{dM} &= \frac{\bar{\rho}_m}{M}\frac{d\ln(\sigma^{-1}(M,z))}{dM}f(\sigma(M,z)) \\ &= -\frac{\bar{\rho}_m}{2M^2}\frac{R(M)}{3\sigma^2(M,z)}\frac{d\sigma^2(M,z)}{dR(M)}f(\sigma(M,z))\,, \end{aligned} \tag{2.5}$$

where $\sigma^2(M,z)$ is the variance of the linear matter density field smoothed with a (real space) top-hat filter on a scale $R(M) = \left(\frac{3M}{4\pi\bar{\rho}_m}\right)^{1/3}$ at redshift z:

$$\sigma^2(M,z) = \frac{1}{2\pi^2}\int k^3\, P_{\rm lin}(k,z)\, W^2(k,R(M))\, d\ln k\,, \tag{2.6}$$

where $P_{\rm lin}(k,z)$ is the linear theory matter power spectrum at wavenumber k and redshift z. Note that the window function $W(k,R)$ is a top-hat filter in real space, which in Fourier space is given by

$$W(k,R) = \frac{3}{x^2}\left(\frac{\sin x}{x} - \cos x\right)\,, \tag{2.7}$$

where $x \equiv kR$. In Eq. (2.5), the function $f(\sigma(M,z))$ is known as the halo multiplicity function. It has been measured to increasingly high precision from large N-body simulations over the past decade [74, 76, 75, 77]. However, many of these calibrated mass functions are specified in terms of the friends-of-friends (FOF) mass rather than the SO mass, hindering their use in analytic calculations such as ours. For this reason, we use the parametrization and calibration from [75], where computations are performed in terms of the SO mass with respect to the mean matter density, $M_{\delta,d}$, for a variety of overdensities. The halo multiplicity function in this model is parametrized by

$$f(\sigma(M,z)) = A\left[\left(\frac{\sigma}{b}\right)^{-a} + 1\right]e^{-c/\sigma^2} \tag{2.8}$$

where $\{A, a, b, c\}$ are (redshift- and overdensity-dependent) parameters fit from simulations. We use the values of these parameters appropriate for the $M_{200,d}$ halo mass function from [75] with the redshift-dependent parameters given in their Eqs. (5)–(8); we will hereafter refer to this as the Tinker mass function. Note that the authors of that study caution against extrapolating their parameters beyond the highest redshift measured in their simulations ($z = 2.5$) and recommend setting the parameters equal to their $z = 2.5$ values at higher redshifts; we adopt this recommendation in our calculations. Also, note that our tSZ power spectrum calculations in Section 2.4 are phrased in terms of the virial mass M, and thus we compute the Jacobian $dM_{200,d}/dM$ using the procedure described in Eq. (2.4) in order to convert the Tinker mass function $dn/dM_{200,d}$ to a virial mass function $\frac{dn}{dM} = \frac{dn}{dM_{200,d}}\frac{dM_{200,d}}{dM}$.

We compute the smoothed matter density field in Eq. (2.6) by first obtaining the linear theory matter power spectrum from CAMB[6] at $z_{in} = 30$ and subsequently rescaling $\sigma^2(M, z)$ by $D^2(z)$, where $D(z)$ is the linear growth factor. We normalize $D(z)$ by requiring that $D(z) \to 1/(1 + z)$ deep in the matter-dominated era (e.g., at z_{in}). The resulting $\sigma^2(M, z)$ is then used to compute the mass function in Eq. (2.5).

Note that we assume the mass function to be known to high enough precision that the parameters describing it can be fixed; in other words, we do not consider $\{A, a, b, c\}$ to be free parameters in our model. These parameters are certainly better constrained at present than those describing the ICM pressure profile (see Section 2.4.2), and thus this assumption seems reasonable for now. However, precision cosmological constraints based on the mass function should in principle consider variations in the mass function parameters in order to obtain robust results, as has been done in some recent X-ray cluster cosmology analyses [78]. However, we leave the implications of these uncertainties for tSZ statistics as a topic for future work.

[6] `http://camb.info/`

Effect of Primordial non-Gaussianity

The influence of primordial non-Gaussianity on the halo mass function has been studied by many authors over the past two decades using a variety of approaches (e.g., [79, 80, 81, 82, 83, 84, 85, 86, 87]). The physical consequences of the model specified in Eq. (2.1) are fairly simple to understand for the halo mass function, especially in the exponential tail of the mass function where massive clusters are found. Intuitively, the number of clusters provides information about the tail of the probability distribution function of the primordial density field, since these are the rarest objects in the universe, which have only collapsed recently. For positive skewness in the primordial density field ($f_{\rm NL} > 0$), one obtains an increased number of massive clusters at late times relative to the $f_{\rm NL} = 0$ case, because more regions of the smoothed density field have $\delta > \delta_c$, the collapse threshold (≈ 1.686 in the spherical collapse model). Conversely, for negative skewness in the primordial density field ($f_{\rm NL} < 0$), one obtains fewer massive clusters at late times relative to the $f_{\rm NL} = 0$ case, because fewer regions of the smoothed density field are above the collapse threshold. As illustrated in recent analytic calculations and simulation measurements [88, 90, 89, 91, 111], these changes can be quite significant for the number of extremely massive halos ($\sim 10^{15}\, M_\odot/h$) in the late-time universe; for example, the $z = 0$ abundance of such halos for $f_{\rm NL} \approx 250$ can be $\approx 1.5 - 2$ times larger than in a Gaussian cosmology. These results have been used as a basis for recent studies constraining $f_{\rm NL}$ by looking for extremely massive outliers in the cluster distribution (e.g., [92, 93, 94, 95, 96, 97]).

We model the effect of $f_{\rm NL}$ on the halo mass function by multiplying the Tinker mass function by a non-Gaussian correction factor:

$$\left(\frac{dn}{dM}\right)_{NG} = \frac{dn}{dM} R_{NG}(M, z, f_{\rm NL})\,, \tag{2.9}$$

where dn/dM is given by Eq. (2.5). We use the model for $R_{NG}(M, z, f_{\rm NL})$ given by Eq. (35) in [88] (the "log-Edgeworth" mass function). In this approach, the density field is approxi-

mated via an Edgeworth expansion, which captures small deviations from Gaussianity. The Press-Schechter approach is then applied to the Edgeworth-expanded density field to obtain an expression for the halo mass function in terms of cumulants of the non-Gaussian density field. The results of [88] include numerical fitting functions for these cumulants obtained from N-body simulations. We use both the expression for $R_{NG}(M, z, f_{\rm NL})$ and the cumulant fitting functions from [88] to compute the non-Gaussian correction to the mass function. This prescription was shown to accurately reproduce the non-Gaussian halo mass function correction factor measured directly from N-body simulations in [88], and in particular improves upon the similar prescription derived in [86] (the "Edgeworth" mass function).

Note that we apply the non-Gaussian correction factor $R_{NG}(M, z, f_{\rm NL})$ to the Tinker mass function in Eq. (2.9), which is an SO mass function, as mentioned above. The prescription for computing $R_{NG}(M, z, f_{\rm NL})$ makes no assumption about whether M is an FOF or SO mass, so there is no logical flaw in this procedure. However, the comparisons to N-body results in [88] were performed using FOF halos. Thus, without having tested the results of Eq. (2.9) on SO mass functions from simulations, our calculation assumes that the change in the mass function due to non-Gaussianity is quasi-universal, even if the underlying Gaussian mass function itself is not. This assumption was tested in [91] for the non-Gaussian correction factor from [86] (see Fig. 9 in [91]) and found to be valid; thus, we choose to adopt it here. We will refer to the non-Gaussian mass function computed via Eq. (2.9) using the prescription from [88] as the LVS mass function.

Effect of Massive Neutrinos

It has long been known that massive neutrinos suppress the amplitude of the matter power spectrum on scales below their free-streaming scale, k_{fs} [41]:

$$k_{fs} \approx 0.082 \frac{H(z)}{H_0(1+z)^2} \left(\frac{M_\nu}{0.1 \text{ eV}}\right) \ h/\text{Mpc}\,. \tag{2.10}$$

Neutrinos do not cluster on scales much smaller than this scale (i.e., $k > k_{fs}$), as they are able to free-stream out of small-scale gravitational potential wells. This effect leads to a characteristic decrease in the small-scale matter power spectrum of order $\Delta P/P \approx -8\Omega_\nu/\Omega_m$ in linear perturbation theory [41, 40]. Nonlinear corrections increase this suppression to $\Delta P/P \approx -10\Omega_\nu/\Omega_m$ for modes with wavenumbers $k \sim 0.5 - 1$ Mpc$/h$ [40].

The neutrino-induced suppression of the small-scale matter power spectrum leads one to expect that the number of massive halos in the low-redshift universe should also be decreased. Several papers in recent years have attempted to precisely model this change in the halo mass function using both N-body simulations and analytic theory [98, 99, 100]. In [98], N-body simulations are used to show that massive neutrinos do indeed suppress the halo mass function, especially for the largest, latest-forming halos (i.e., galaxy clusters). Moreover, the suppression is found to arise primarily from the suppression of the initial transfer function in the linear regime, and not due to neutrino clustering effects in the N-body simulations. This finding suggests that an analytic approach similar to the Press-Schecter theory should work for massive neutrino cosmologies as well, and the authors subsequently show that a modified Sheth-Tormen formalism [72] gives a good fit to their simulation results. Similar N-body simulations are examined in [99], who find generally similar results to those in [98], but also point out that the effect of $M_\nu > 0$ on the mass function cannot be adequately represented by simply rescaling σ_8 to a lower value in an analytic calculation without massive neutrinos. Finally, [100] study the effect of massive neutrinos on the mass function using analytic calculations with the spherical collapse model. Their results suggest that an accurate approximation is to simply input the M_ν-suppressed linear theory (cold+baryonic-only) matter power spectrum computed at z_{in} to a ΛCDM-calibrated mass function fit (note that a similar procedure was used in some recent X-ray cluster-based constraints on M_ν [101]). The net result of this suppression can be quite significant at the high-mass end of the mass function; for example, $M_\nu = 0.1$ eV leads to a factor of ~ 2 decrease in the abundance of $10^{15}\ M_\odot/h$ halos at $z = 1$ as compared to a massless-neutrino cosmology [100]. We follow

the procedure used in [100] in our work, although we input the suppressed linear theory matter power spectrum to the Tinker mass function rather than that of [77], as was done in [100]. We will refer to the M_ν -suppressed mass function computed with this prescription as the IT mass function.

2.3.2 Halo Bias

Dark matter halos are known to cluster more strongly than the underlying matter density field; they are thus biased tracers. This bias can depend on scale, mass, and redshift (e.g., [102, 103, 104]). We define the halo bias $b(k, M, z)$ by

$$b(k, M, z) = \sqrt{\frac{P_{hh}(k, M, z)}{P(k, z)}}\,, \tag{2.11}$$

where $P_{hh}(k, M, z)$ is the power spectrum of the halo density field and $P(k, z)$ is the power spectrum of the matter density field. Knowledge of the halo bias is necessary to model and extract cosmological information from the clustering of galaxies and galaxy clusters. For our purposes, it will be needed to compute the two-halo term in the tSZ power spectrum, which requires knowledge of $P_{hh}(k, M, z)$.

In a Gaussian cosmology, the halo bias depends on mass and redshift but is independent of scale for $k \lesssim 0.05$ Mpc$/h$, i.e. on large scales (e.g., [105]). We compute this linear Gaussian bias, $b_G(M, z)$, using the fitting function in Eq. (6) of [105] with the parameters appropriate for $M_{200,d}$ SO masses (see Table 2 in [105]). This fit was determined from the results of many large-volume N-body simulations with a variety of cosmological parameters and found to be quite accurate. We will refer to this prescription as the Tinker bias model.

Although the bias becomes scale-dependent on small scales even in a Gaussian cosmology, it becomes scale-dependent on large scales in the presence of local primordial non-Gaussianity, as first shown in [106]. The scale-dependence arises due to the coupling of long- and short-wavelength density fluctuations induced by local $f_{\rm NL} \neq 0$. We model this effect as

a correction to the Gaussian bias described in the preceding paragraph:

$$b(k, M, z) = b_G(M, z) + \Delta b_{NG}(k, M, z)\,, \tag{2.12}$$

where the non-Gaussian correction is given by [106]

$$\Delta b_{NG}(k, M, z) = 2\delta_c \left(b_G(M, z) - 1\right) \frac{f_{\rm NL}}{\alpha(k, z)}\,. \tag{2.13}$$

Here, $\delta_c = 1.686$ (the spherical collapse threshold) and

$$\alpha(k, z) = \frac{2k^2 T(k) D(z) c^2}{3\Omega_m H_0^2} \tag{2.14}$$

relates the linear density field to the primordial potential via $\delta(k, z) = \alpha(k, z)\Phi(k)$. Note that $T(k)$ is the linear matter transfer function, which we compute using CAMB. Since the original derivation in [106], the results in Eqs. (2.13) and (2.14) have subsequently been confirmed by other authors [107, 108, 109] and tested extensively on N-body simulations (e.g., [110, 106, 90, 112]). The overall effect is a steep increase in the large-scale bias of massive halos, which is even larger for highly biased tracers like galaxy clusters. We will refer to this effect simply as the scale-dependent halo bias.

The influence of massive neutrinos on the halo bias has been studied far less thoroughly than that of primordial non-Gaussianity. Recent N-body simulations analyzed in [99] indicate that massive neutrinos lead to a nearly scale-independent increase in the large-scale halo bias. This effect arises because of the mass function suppression discussed in Section 2.3.1: halos of a given mass are rarer in an $M_\nu > 0$ cosmology than in a massless neutrino cosmology (for fixed A_s), and thus they are more highly biased relative to the matter density field. However, the amplitude of this change is far smaller than that induced by $f_{\rm NL} \neq 0$, especially on very large scales. For example, the results of [99] indicate an overall increase of $\sim 10\%$ in the mean bias of massive halos at $z = 1$ for $M_\nu = 0.3$ eV as compared to $M_\nu = 0$. Our

implementation of the scale-dependent bias due to local $f_{\rm NL}$ yields a factor of $\sim 100-1000$ increase in the large-scale ($k \sim 10^{-4}$ h/Mpc) bias of objects in the same mass range at $z = 1$ for $f_{\rm NL} = 50$. Clearly, the effect of $f_{\rm NL}$ is much larger than that of massive neutrinos, simply because it is so strongly scale-dependent, while M_ν only leads to a small scale-independent change (at least on large scales; the small-scale behavior may be more complicated). Moreover, the change in bias due to M_ν is larger at higher redshifts ($z \gtrsim 1$), whereas most of the tSZ signal originates at lower redshifts. Lastly, due to the smallness of the two-halo term in the tSZ power spectrum compared to the one-halo term (see Section 2.4), small variations in the Gaussian bias cause essentially no change in the total signal. For all of these reasons, we choose to neglect the effect of massive neutrinos on the halo bias in our calculations.

2.4 Thermal SZ Power Spectrum

Basic tSZ definitions and results are presented in Chapter 1. We briefly collect them here for reference. The temperature change ΔT at angular position $\vec{\theta}$ with respect to the center of a cluster of mass M at redshift z is given by [1]

$$\begin{aligned} \frac{\Delta T(\vec{\theta}, M, z)}{T_{\rm CMB}} &= g_\nu y(\vec{\theta}, M, z) \\ &= g_\nu \frac{\sigma_T}{m_e c^2} \int_{\rm LOS} P_e\left(\sqrt{l^2 + d_A^2 |\vec{\theta}|^2}, M, z\right) dl\,, \end{aligned} \tag{2.15}$$

where $g_\nu = x\coth(x/2) - 4$ is the tSZ spectral function with $x \equiv h\nu/k_B T_{\rm CMB}$, y is the Compton-y parameter, σ_T is the Thomson scattering cross-section, m_e is the electron mass, and $P_e(\vec{r})$ is the ICM electron pressure at location $\vec{r}$ with respect to the cluster center. We have neglected relativistic corrections in Eq. (2.15) (e.g., [113]), as these effects are relevant only for the most massive clusters in the universe ($\gtrsim 10^{15}$ $M_\odot/h$). Such clusters contribute non-negligibly to the tSZ power spectrum at low-ℓ, and thus our results in unmasked calculations may be slightly inaccurate; however, the optimal forecasts for cosmological constraints

arise from calculations in which such nearby, massive clusters are masked (see Section 2.7), and thus these corrections will not be relevant. Therefore, we do not include them in our calculations.

We only consider spherically symmetric pressure profiles in this work, i.e. $P_e(\vec{r}) = P_e(r)$ in Eq. (2.15). The integral in Eq. (2.15) is computed along the LOS: $r^2 = l^2 + d_A(z)^2\theta^2$, with $d_A(z)$ the angular diameter distance to redshift z and $\theta \equiv |\vec{\theta}|$ the angular distance on the sky between $\vec{\theta}$ and the cluster center. The thermal pressure P_{th} and electron pressure P_e are related via $P_{th} = P_e(5X_H + 3)/2(X_H + 1) = 1.932P_e$, where $X_H = 0.76$ is the primordial hydrogen mass fraction. We calculate all tSZ power spectra in this chapter at $\nu = 150$ GHz ($g_{150\,\mathrm{GHz}} = -0.9537$), simply because recent tSZ measurements have been performed at this frequency using ACT and SPT (e.g., [15, 17, 14, 12]). We will often use "y" as a label for tSZ quantities, although they are calculated numerically at $\nu = 150$ GHz.

In the remainder of this section, we outline the halo model-based calculations used to compute the tSZ power spectrum, discuss our model for the gas physics of the ICM, and explain the physical effects of each cosmological and astrophysical parameter on the tSZ power spectrum.

2.4.1 Halo Model Formalism

We compute the tSZ power spectrum using the halo model approach (see [114] for a review). We provide complete derivations of all the relevant expressions in Chapter 1, first obtaining completely general full-sky results and then specializing to the flat-sky/Limber-approximated case. Here, we simply quote the necessary results and refer the interested reader to Section 1.2 for the derivations. Note that we will work in terms of the Compton-y parameter; the results can easily be multiplied by the necessary g_ν factors to obtain results at any frequency.

The tSZ power spectrum, C_ℓ^y, is given by the sum of the one-halo and two-halo terms:

$$C_\ell^y = C_\ell^{y,1h} + C_\ell^{y,2h}\,. \tag{2.16}$$

The exact expression for the one-halo term is given by Eq. (1.18):

$$C_\ell^{y,1h} = \int \frac{dz}{\chi(z)} \frac{d^2V}{dz d\Omega} \int dM \frac{dn}{dM} \times \left| \int k\, dk\, J_{\ell+1/2}(k\chi(z)) \tilde{y}_{3D}(k; M, z) \int \frac{c\, dz'}{H(z')(1+z')\sqrt{\chi(z')}} J_{\ell+1/2}(k\chi(z')) \right|^2 \quad (2.17)$$

where $\chi(z)$ is the comoving distance to redshift z, $d^2V/dzd\Omega$ is the comoving volume element per steradian, dn/dM is the halo mass function discussed in Section 2.3.1, $\tilde{y}_{3D}(k; M, z)$ is given in Eq. (1.10), and $J_{\ell+1/2}(x)$ is a Bessel function of the first kind. In the flat-sky limit, the one-halo term simplifies to the following widely-used expression (given in e.g. Eq. (1) of [19]), which we derive in Eq. (1.23):

$$C_{\ell \gg 1}^{y,1h} \approx \int dz \frac{d^2V}{dz d\Omega} \int dM \frac{dn(M,z)}{dM} \left| \tilde{y}_\ell(M,z) \right|^2 , \quad (2.18)$$

where

$$\tilde{y}_\ell(M,z) \approx \frac{4\pi r_s}{\ell_s^2} \int dx\, x^2 \frac{\sin((\ell+1/2)x/\ell_s)}{(\ell+1/2)x/\ell_s} y_{3D}(x; M, z) . \quad (2.19)$$

Here, r_s is a characteristic scale radius (not the NFW scale radius) of the y_{3D} profile given by $y_{3D}(\vec{r}) = \frac{\sigma_T}{m_e c^2} P_e(\vec{r})$ and $\ell_s = a(z)\chi(z)/r_s = d_A(z)/r_s$ is the multipole moment associated with the scale radius. For the pressure profile from [116] used in our calculations, the natural scale radius is $r_{200,c}$. In our calculations, we choose to implement the flat-sky result for the one-halo term at all ℓ — see Chapter 1 for a justification of this decision and an assessment of the associated error at low-ℓ (the only regime where this correction would be relevant).

The exact expression for the two-halo term is given by Eq. (1.24):

$$C_\ell^{y,2h} = \int dk\, k \frac{P_{\rm lin}(k; z_{in})}{D^2(z_{in})} \times \left[\int \frac{dz}{\sqrt{\chi(z)}} \frac{d^2V}{dz d\Omega} J_{\ell+1/2}(k\chi(z)) D(z) \int dM \frac{dn}{dM} b(k, M, z) \tilde{y}_{k\chi(z)}(M, z) \right]^2 \quad (2.20)$$

where $P_{\rm lin}(k, z_{in})$ is the linear theory matter power spectrum at z_{in} (which we choose to set equal to 30), $b(k, M, z)$ is the halo bias discussed in Section 2.3.2, and $\tilde{y}_{k\chi(z)}(M, z)$ refers to the expression for $\tilde{y}_\ell(M, z)$ given in Eq. (2.19) evaluated with $\ell + 1/2 = k\chi$. This notation is simply a mathematical convenience; no flat-sky or Limber approximation was used in deriving Eq. (1.24), and no ℓ appears in $\tilde{y}_{k\chi}(M, z)$. In the Limber approximation [115], the two-halo term simplifies to the result given in [20], which we derive in Eq. (1.26):

$$C_{\ell\gg1}^{y,2h} \approx \int dz \frac{d^2V}{dz d\Omega} \left[\int dM \frac{dn(M,z)}{dM} b(k, M, z)\tilde{y}_\ell(M, z)\right]^2 P_{\rm lin}\left(\frac{\ell + 1/2}{\chi(z)}; z\right) . \qquad (2.21)$$

We investigate the validity of the Limber approximation in detail in Chapter 1. We find that it is necessary to compute the exact expression in Eq. (2.20) in order to obtain sufficiently accurate results at low-ℓ, where the signature of the scale-dependent bias induced by $f_{\rm NL}$ is present (looking for this signature is our primary motivation for computing the two-halo term to begin with). In particular, we compute the exact expression in Eq. (2.20) for $\ell < 50$, while we use the Limber-approximated result in Eq. (2.21) at higher multipoles.

The fiducial integration limits in our calculations are $0.005 < z < 4$ for all redshift integrals, $5 \times 10^{11} M_\odot/h < M < 5 \times 10^{15} M_\odot/h$ for all mass integrals, and 10^{-4} h/Mpc $< k <$ 3 h/Mpc for all wavenumber integrals. We check that extending the wavenumber upper limit further into the nonlinear regime does not affect our results. Note that the upper limit in the mass integral becomes redshift-dependent in the masked calculations that we discuss below, in which the most massive clusters at low redshifts are removed from the computation.

We use the halo mass functions discussed in Section 2.3.1 (Tinker, LVS, and IT) and the bias models discussed in Section 2.3.2 (Tinker and scale-dependent bias) in Eqs. (2.18), (2.20), and (2.21). The only remaining ingredient needed to complete the tSZ power spectrum calculation is a prescription for the ICM electron pressure profile as a function of mass and redshift. Note that this approach to the tSZ power spectrum calculation separates the cosmology-dependent component (the mass function and bias) from the ICM-dependent

component (the pressure profile). This separation arises from the fact that the small-scale baryonic physics that determines the structure of the ICM pressure profile effectively decouples from the large-scale physics described by the background cosmology and linear perturbation theory. Thus, it is a standard procedure to constrain the ICM pressure profile from cosmological hydrodynamics simulations (e.g., [116, 118]) or actual observations of galaxy clusters (e.g., [119, 120], which are obtained for a fixed cosmology in either case (at present, it is prohibitively computationally expensive to run many large hydrodynamical simulations with varying cosmological parameters). Of course, it is also possible to model the ICM analytically and obtain a pressure profile (e.g., [122, 123]. Regardless of its origin (observations/simulations/theory), the derived ICM pressure profile can then be applied to different background cosmologies by using the halo mass function and bias model appropriate for that cosmology in the tSZ power spectrum calculations. We follow this approach.

Note that because the tSZ signal is heavily dominated by contributions from collapsed objects, the halo model approximation gives very accurate results when compared to direct LOS integrations of numerical simulation boxes (see Figs. 7 and 8 in [116] for direct comparisons). In particular, the halo model agrees very well with the simulation results for $\ell \lesssim 1000$, which is predominantly the regime we are interested in for this chapter (on smaller angular scales effects due to asphericity and substructure become important, which are not captured in the halo model approach). These results imply that contributions from the intergalactic medium, filaments, and other diffuse structures are unlikely to be large enough to significantly impact the calculations and forecasts in the remainder of the chapter. Contamination from the Galaxy is a separate issue, which we assume can be minimized to a sufficient level through sky cuts and foreground subtraction (see Section 2.5).

2.4.2 Modeling the ICM

We adopt the parametrized ICM pressure profile fit from [116] as our fiducial model. This profile is derived from cosmological hydrodynamics simulations described in [117]. These

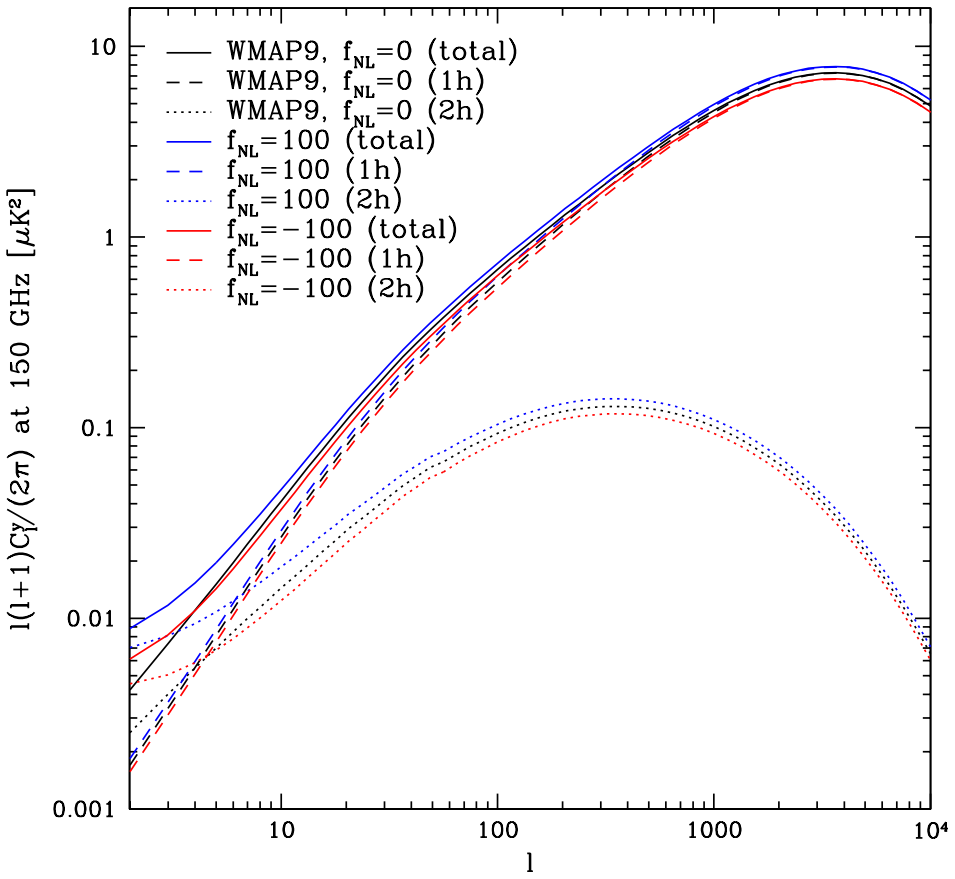

Figure 2.1: This plot shows the unmasked tSZ power spectrum for our fiducial model (black curves), as specified in Section 2.4.3, as well as variations with $f_{\rm NL} = 100$ (blue curves) and $f_{\rm NL} = -100$ (red curves). $f_{\rm NL}$ values of this magnitude are highly disfavored by current constraints, but we plot them here to clarify the influence of primordial non-Gaussianity on the tSZ power spectrum. The effect of $f_{\rm NL}$ on the one-halo term is simply an overall amplitude shift due to the corresponding increase or decrease in the number of massive clusters in the universe, as described in Section 2.3.1. The effect of $f_{\rm NL}$ on the two-halo term includes not only an amplitude shift due to the change in the mass function, but also a steep upturn at low-ℓ due to the influence of the scale-dependent halo bias, as described in Section 2.3.2. Note that for $f_{\rm NL} < 0$ the two effects cancel for $\ell \approx 4 - 5$. The relative smallness of the two-halo term (compared to the one-halo term) makes the scale-dependent bias signature subdominant for all ℓ values except $\ell \lesssim 7 - 8$. However, masking of nearby massive clusters suppresses the low-ℓ one-halo term in the power spectrum (in addition to decreasing the cosmic variance, as discussed in Section 2.6), which increases the relative importance of the two-halo term and thus the dependence of the total signal on $f_{\rm NL}$ at low-ℓ.

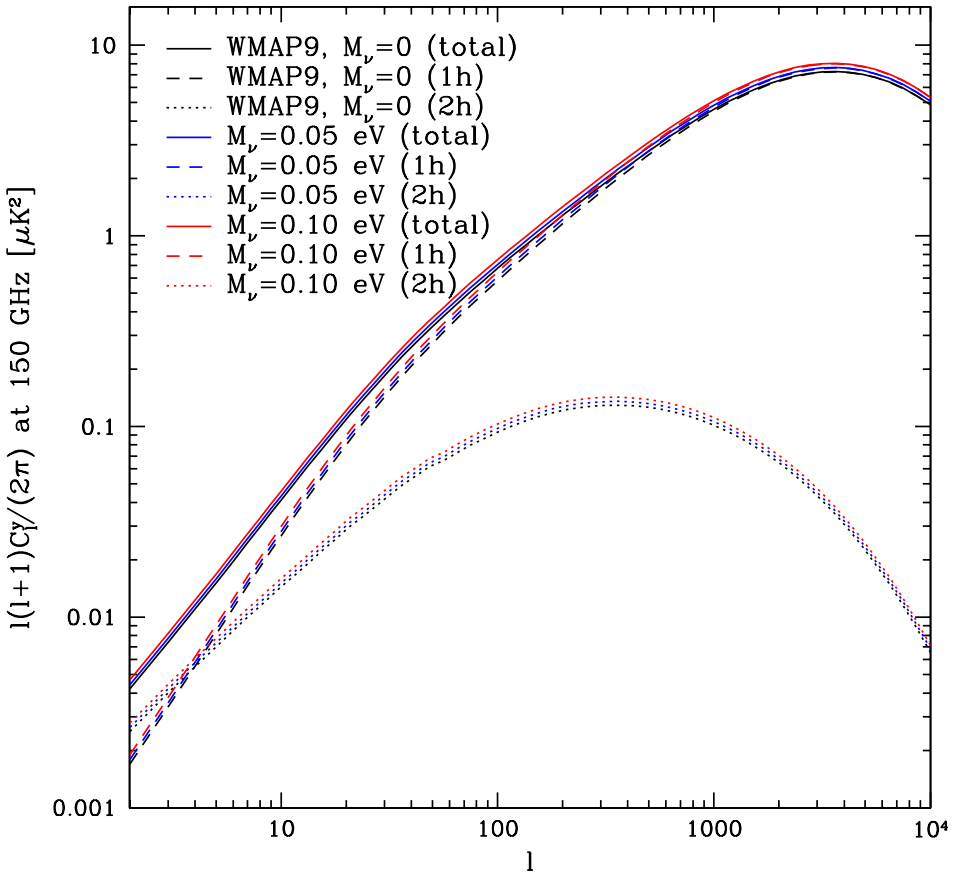

Figure 2.2: This plot shows the unmasked tSZ power spectrum for our fiducial model (black curves), as specified in Section 2.4.3, as well as variations with $M_\nu = 0.05$ eV (blue curves) and $M_\nu = 0.10$ eV (red curves). M_ν values of this magnitude are at the lower bound allowed by neutrino oscillation measurements, and thus an effect of at least this magnitude is expected in our universe. The effect of M_ν on both the one- and two-halo terms is effectively an overall amplitude shift, although the effect tapers off very slightly at high-ℓ. It may be puzzling at first to see an increase in the tSZ power when $M_\nu > 0$, but the key fact is that we are holding σ_8 constant when varying M_ν (indeed, we hold all of the other parameters constant). In order to keep σ_8 fixed despite the late-time suppression of structure growth due to $M_\nu > 0$, we must increase the primordial amplitude of scalar perturbations, A_s. Thus, the net effect is an increase in the tSZ power spectrum amplitude, counterintuitive though it may be. If our ΛCDM parameter set included A_s rather than σ_8, and we held A_s constant while $M_\nu > 0$, we would indeed find a corresponding decrease in the tSZ power (and in σ_8, of course).

simulations include (sub-grid) prescriptions for radiative cooling, star formation, supernova feedback, and feedback from active galactic nuclei (AGN). Taken together, these feedback processes typically decrease the gas fraction in low-mass groups and clusters, as the injection of energy into the ICM blows gas out of the cluster potential. In addition, the smoothed particle hydrodynamics used in these simulations naturally captures the effects of non-thermal pressure support due to bulk motions and turbulence, which must be modeled in order to accurately characterize the cluster pressure profile in the outskirts.

The ICM thermal pressure profile in this model is parametrized by a dimensionless GNFW form, which has been found to be a useful parametrization by many observational and numerical studies (e.g., [125, 119, 3, 120]):

$$\frac{P_{th}(x)}{P_{200,c}} = \frac{P_0\,(x/x_c)^{\gamma}}{[1+(x/x_c)^{\alpha}]^{\beta}} \quad , x \equiv r/r_{200,c}\,, \tag{2.22}$$

where $P_{th}(x) = 1.932 P_e(x)$ is the thermal pressure profile, x is the dimensionless distance from the cluster center, x_c is a core scale length, P_0 is a dimensionless amplitude, α, β, and γ describe the logarithmic slope of the profile at intermediate ($x \sim x_c$), large ($x \gg x_c$), and small ($x \ll x_c$) radii, respectively, and $P_{200,c}$ is the self-similar amplitude for pressure at $r_{200,c}$ given by [126, 127]:

$$P_{200,c} = \frac{200\,G M_{200,c}\rho_{cr}(z)\Omega_b}{2\,\Omega_m r_{200,c}}\,. \tag{2.23}$$

In [116] this parametrization is fit to the stacked pressure profiles of clusters extracted from the simulations described above. Note that due to degeneracies the parameters α and γ are not varied in the fit; they are fixed to $\alpha = 1.0$ and $\gamma = -0.3$, which agree with many other studies (e.g., [125, 119, 3, 120]. In addition to constraining the amplitude of the remaining

parameters, [116] also fit power-law mass and redshift dependences, with the following results:

$$P_0(M_{200,c}, z) = 18.1 \left(\frac{M_{200,c}}{10^{14}\ M_\odot} \right)^{0.154} (1+z)^{-0.758} \quad (2.24)$$

$$x_c(M_{200,c}, z) = 0.497 \left(\frac{M_{200,c}}{10^{14}\ M_\odot} \right)^{-0.00865} (1+z)^{0.731} \quad (2.25)$$

$$\beta(M_{200,c}, z) = 4.35 \left(\frac{M_{200,c}}{10^{14}\ M_\odot} \right)^{0.0393} (1+z)^{0.415}\ . \quad (2.26)$$

Note that the denominator of the mass-dependent factor has units of $M_\odot$ rather than $M_\odot/h$ as used elsewhere in this chapter. The mass and redshift dependence of these parameters captures deviations from simple self-similar cluster pressure profiles. These deviations arise from non-gravitational energy injections due to AGN and supernova feedback, star formation in the ICM, and non-thermal processes such as turbulence and bulk motions [116, 121]. Eqs. (2.22)–(2.26) completely specify the ICM electron pressure profile as a function of mass and redshift, and provide the remaining ingredient needed for the halo model calculations of the tSZ power spectrum described in Section 2.4.1, in addition to the halo mass function and halo bias. We will refer to this model of the ICM pressure profile as the Battaglia model.

Although it is derived solely from numerical simulations, we note that the Battaglia pressure profile is in good agreement with a number of observations of cluster pressure profiles, including those based on the REXCESS X-ray sample of massive, $z < 0.3$ clusters [119], independent studies of low-mass groups at $z < 0.12$ with Chandra [128], and early Planck measurements of the stacked pressure profile of $z < 0.5$ clusters [120].

We allow for a realistic degree of uncertainty in the ICM pressure profile by freeing the amplitude of the parameters that describe the overall normalization (P_0) and the outer logarithmic slope (β). To be clear, we do not free the mass and redshift dependences for these parameters given in Eqs. (2.24) and (2.26), only the overall amplitudes in those expressions. The outer slope β is known to be highly degenerate with the scale radius x_c (e.g., [116, 3]), and thus it is only feasible to free one of these parameters. The other slope parameters in Eq. (2.22) are fixed to their Battaglia values, which match the standard values in the

literature. We parametrize the freedom in P_0 and β by introducing new parameters C_{P_0} and C_β defined by:

$$P_0(M_{200,c}, z) = C_{P_0} \times 18.1 \left(\frac{M_{200,c}}{10^{14}\ M_\odot} \right)^{0.154} (1+z)^{-0.758} \tag{2.27}$$

$$\beta(M_{200,c}, z) = C_\beta \times 4.35 \left(\frac{M_{200,c}}{10^{14}\ M_\odot} \right)^{0.0393} (1+z)^{0.415} \ . \tag{2.28}$$

These parameters thus describe multiplicative overall changes to the amplitudes of the P_0 and β parameters. The fiducial Battaglia profile corresponds to $\{C_{P_0}, C_\beta\} = \{1, 1\}$. We discuss our priors for these parameters in Section 2.7.

2.4.3 Parameter Dependences

Including both cosmological and astrophysical parameters, our model is specified by the following quantities:

$$\left\{ \Omega_b h^2, \Omega_c h^2, \Omega_\Lambda, \sigma_8, n_s, C_{P_0}, C_\beta, (f_{\rm NL}, M_\nu) \right\} \ , \tag{2.29}$$

which take the following values in our (WMAP9+BAO+H_0 [69]) fiducial model:

$$\{0.02240, 0.1146, 0.7181, 0.817, 0.9646, 1.0, 1.0, (0.0, 0.0)\} \ . \tag{2.30}$$

As a reminder, the ΛCDM parameters are (in order of their appearance in Eq. (2.29)) the physical baryon density, the physical cold dark matter density, the vacuum energy density, the rms matter density fluctuation on comoving scales of 8 Mpc/h at $z = 0$, and the scalar spectral index. The ICM physics parameters C_{P_0} and C_β are defined in Eqs. (2.27) and (2.28), respectively, $f_{\rm NL}$ is defined by Eq. (2.1), and M_ν is the sum of the neutrino masses in units of eV. We have placed $f_{\rm NL}$ and M_ν in parentheses in Eq. (2.29) in order to make it clear that we only consider scenarios in which these parameters are varied separately: for all

cosmologies that we consider with $f_{\rm NL} \neq 0$, we set $M_\nu = 0$, and for all cosmologies that we consider with $M_\nu > 0$, we set $f_{\rm NL} = 0$. In other words, we only investigate one-parameter extensions of the ΛCDM concordance model.

For the primary ΛCDM cosmological parameters, we use the parametrization adopted by the WMAP team (e.g., [69]), as the primordial CMB data best constrain this set. The only exception to this convention is our use of σ_8, which stands in place of the primordial amplitude of scalar perturbations, A_s. We use σ_8 both because it is conventional in the tSZ power spectrum literature and because it is a direct measure of the low-redshift amplitude of matter density perturbations, which is physically related more closely to the tSZ signal than A_s. However, this choice leads to slightly counterintuitive results when considering cosmologies with $M_\nu > 0$, because in order to keep σ_8 fixed for such scenarios we must increase A_s (to compensate for the suppression induced by M_ν in the matter power spectrum).

For the fiducial model specified by the values in Eq. (2.30), we find that the tSZ power spectrum amplitude at $\ell = 3000$ is $\ell(\ell+1)C_\ell^y/2\pi = 7.21\ \mu\mathrm{K}^2$ at $\nu = 150$ GHz. This corresponds to $7.59\ \mu\mathrm{K}^2$ at $\nu = 148$ GHz (the relevant ACT frequency) and $6.66\ \mu\mathrm{K}^2$ at $\nu = 152.9$ GHz (the relevant SPT frequency). The most recent measurements from ACT and SPT find corresponding constraints at these frequencies of $3.4 \pm 1.4\ \mu\mathrm{K}^2$ [12] and $3.09 \pm 0.60\ \mu\mathrm{K}^2$ [17] (using their more conservative error estimate). Note that the SPT constraint includes information from the tSZ bispectrum, which reduces the error by a factor of ∼2. Although it appears that our fiducial model predicts a level of tSZ power too high to be consistent with these observations, the results are highly dependent on the true value of σ_8, due to the steep dependence of the tSZ power spectrum on this parameter. For example, recomputing our model predictions for $\sigma_8 = 0.79$ gives $5.52\ \mu\mathrm{K}^2$ at $\nu = 148$ GHz and $4.84\ \mu\mathrm{K}^2$ at $\nu = 152.9$ GHz, which are consistent at 3σ with the corresponding ACT and SPT constraints. Given that $\sigma_8 = 0.79$ is within the 2σ error bar for WMAP9 [69], it is difficult to assess the extent to which our fiducial model may be discrepant with the ACT and SPT results. The difference can easily be explained by small changes in σ_8 and is also

sensitive to variations in the ICM physics, which we have kept fixed in these calculations. We conclude that our model is not in significant tension with current tSZ measurements (or other cosmological parameter constraints), and is thus a reasonable fiducial case around which to consider variations.

Figs. 2.1 and 2.2 show the tSZ power spectra for our fiducial model and several variations around it, including the individual contributions of the one- and two-halo terms. In the fiducial case, the two-halo term is essentially negligible for $\ell \gtrsim 300$, as found by earlier studies [20], and it only overtakes the one-halo term at very low-ℓ ($\ell \lesssim 4$). However, for $f_{\rm NL} \neq 0$, the influence of the two-halo term is greatly enhanced due to the scale-dependent bias described in Section 2.3.2, which leads to a characteristic upturn in the tSZ power spectrum at low-ℓ. In addition, $f_{\rm NL}$ induces an overall amplitude change in both the one- and two-halo terms due to its effect on the halo mass function described in Section 2.3.1. While this amplitude change is degenerate with the effects of other parameters on the tSZ power spectrum (e.g., σ_8), the low-ℓ upturn caused by the scale-dependent bias is a unique signature of primordial non-Gaussianity, which motivates our assessment of forecasts on $f_{\rm NL}$ using this observable later in the chapter.

Fig. 2.2 shows the results of similar calculations for $M_\nu > 0$. In this case, the effect is simply an overall amplitude shift in the one- and two-halo terms, and hence the total tSZ power spectrum. The amplitude shift is caused by the change in the halo mass function described in Section 2.3.1. Note that the sign of the amplitude change is somewhat counterintuitive, but arises due to our choice of σ_8 as a fundamental parameter instead of A_s, as mentioned above. In order to keep σ_8 fixed while increasing M_ν , we must increase A_s, which leads to an increase in the tSZ power spectrum amplitude. Although this effect is degenerate with that of σ_8 and other parameters, the change in the tSZ power spectrum amplitude is rather large even for small neutrino masses ($\approx 12\%$ for $M_\nu = 0.1$ eV, which is larger than the amplitude change caused by $f_{\rm NL} = 100$). This sensitivity suggests that the tSZ power spectrum may be a useful observable for constraints on the neutrino mass sum.

We demonstrate the physical effects of each parameter in our model on the tSZ power spectrum in Figs. 2.3–2.11, including the effects on both the one- and two-halo terms individually. Note that the limits on the vertical axis in each plot differ, so care must be taken in assessing the amplitude of the change caused by each parameter. Except for $f_{\rm NL}$ and M_ν , the figures show $\pm 1\%$ variations in each of the parameters, which facilitates easier comparisons between their relative influences on the tSZ power spectrum. On large angular scales ($\ell \lesssim 100$), the most important parameters (neglecting $f_{\rm NL}$ and M_ν) are σ_8, Ω_Λ, and C_β. On very large angular scales ($\ell < 10$), the effect of $f_{\rm NL}$ is highly significant, but its relative importance is difficult to assess, since the true value of $f_{\rm NL}$ may be unmeasurably small. Note, however, that M_ν is important over the entire ℓ range we consider, even if its true value is as small as 0.1 eV. Comparison of Figs. 2.4 and 2.8 indicates that the amplitude change induced by $M_\nu = 0.1$ eV (for fixed σ_8) is actually slightly larger than that caused by a 1% change in σ_8 around its fiducial value.

We now provide physical interpretations of the effects shown in Figs. 2.3–2.11:

- $f_{\rm NL}$ (2.3): The change to the halo mass function discussed in Section 2.3.1 leads to an overall increase (decrease) in the amplitude of the one-halo term for $f_{\rm NL} > 0$ (< 0). This increase or decrease is essentially ℓ-independent, is also seen at $\ell > 100$ in the two-halo term, and is $\simeq \pm 10\%$ for $f_{\rm NL} = \pm 100$. More significantly, the influence of the scale-dependent halo bias induced by $f_{\rm NL} \neq 0$ is clearly seen in the dramatic increase of the two-halo term at low-ℓ. This increase is significant enough to be seen in the total power spectrum despite the typical smallness of the two-halo term relative to the one-halo term for a Gaussian cosmology.

- M_ν (2.4): The presence of massive neutrinos leads to a decrease in the number of galaxy clusters at late times, as discussed in Section 2.3.1. This decrease would lead one to expect a corresponding decrease in the amplitude of the tSZ signal, but Fig. 2.4 shows an increase. This increase is a result of our choice of parameters — we hold σ_8 constant while increasing M_ν , which means that we must simultaneously increase A_s,

the initial amplitude of scalar fluctuations. This increase in A_s (for fixed σ_8) leads to the increase in the tSZ power spectrum amplitude seen in Fig. 2.4. The effect appears to be essentially ℓ-independent, although it tapers off slightly at very high-ℓ.

- $\Omega_b h^2$ (2.5): Increasing (decreasing) the amount of baryons in the universe leads to a corresponding increase (decrease) in the amount of gas in galaxy clusters, and thus a straightforward overall amplitude shift in the tSZ power spectrum (which goes like f_{gas}^2).

- $\Omega_c h^2$ (2.6): In principle, one would expect that changing $\Omega_c h^2$ should change the tSZ power spectrum, but it turns out to have very little effect, as pointed out in [19], who argue that the effect of increasing (decreasing) $\Omega_c h^2$ on the halo mass function is cancelled in the tSZ power spectrum by the associated decrease (increase) in the comoving volume to a given redshift. We suspect that the small increase (decrease) seen in Fig. (2.6) when decreasing (increasing) $\Omega_c h^2$ is due to the fact that we hold Ω_Λ constant when varying $\Omega_c h^2$. Thus, $\Omega_m \equiv 1 - \Omega_\Lambda$ is also held constant, and thus $\Omega_b = \Omega_m - \Omega_c$ is decreased (increased) when Ω_c is increased (decreased). This decrease (increase) in the baryon fraction leads to a corresponding decrease (increase) in the tSZ power spectrum amplitude, as discussed in the previous item. The slight ℓ-dependence of the $\Omega_c h^2$ variations may be due to the associated change in h required to keep Ω_m constant, which leads to a change in the angular diameter distance to each cluster, and hence a change in the angular scale associated with a given physical scale. Increasing (decreasing) $\Omega_c h^2$ requires increasing (decreasing) h in order to leave Ω_m unchanged, which decreases (increases) the distance to each cluster, shifting a given physical scale in the spectrum to lower (higher) multipoles. However, it is hard to completely disentangle all of the effects described here, and in any case the overall influence of $\Omega_c h^2$ is quite small.

- Ω_Λ (2.7): An increase (decrease) in Ω_Λ has several effects which all tend to decrease (increase) the amplitude of the tSZ power spectrum. First, Ω_m is decreased (increased), which leads to fewer (more) halos, although this effect is compensated by the change in the comoving volume as described above. Second, for fixed Ω_b/Ω_c, this decrease (increase) in Ω_m leads to fewer (more) baryons in clusters, and thus less (more) tSZ power. Third, more (less) vacuum energy leads to more (less) suppression of late-time structure formation due to the decaying of gravitational potentials, and thus less (more) tSZ power. All of these effects combine coherently to produce the fairly large changes caused by Ω_Λ seen in Fig. 2.7. The slight ℓ-dependence may be due to the associated change in h required to keep $\Omega_c h^2$ and $\Omega_b h^2$ constant, similar (though in the opposite direction) to that discussed in the $\Omega_c h^2$ case above. Regardless, this effect is clearly subdominant to the amplitude shift caused by Ω_Λ, which is only slightly smaller on large angular scales than that caused by σ_8 (for a 1% change in either parameter).

- σ_8 (2.8): Increasing (decreasing) σ_8 leads to a significant overall increase (decrease) in the amplitude of the tSZ power spectrum, as has been known for many years (e.g. [19]). The effect is essentially ℓ-independent and appears in both the one- and two-halo terms.

- n_s (2.9): An increase (decrease) in n_s leads to more (less) power in the primordial spectrum at wavenumbers above (below) the pivot, which we set at the WMAP value $k_{piv} = 0.002$ Mpc^{-1} (no h). Since the halo mass function on cluster scales probes much smaller scales than the pivot (i.e., much higher wavenumbers $k \sim 0.1 - 1$ h/Mpc), an increase (decrease) in n_s should lead to more (fewer) halos at late times. However, since we require σ_8 to remain constant while increasing (decreasing) n_s, we must decrease (increase) A_s in order to compensate for the change in power on small scales. This is similar to the situation for M_ν described above. Thus, an increase (decrease) in n_s actually leads to a small decrease (increase) in the tSZ power spectrum on most scales, at least for the one-halo term. The cross-over in the two-halo term is likely related to

the pivot scale after it is weighted by the kernel in Eq. (2.21), but this is somewhat non-trivial to estimate. Regardless, the overall effect of n_s on the tSZ power spectrum is quite small.

- C_{P_0} (2.10): Since C_{P_0} sets the overall normalization of the ICM pressure profile (or, equivalently, the zero-point of the $Y - M$ relation), the tSZ power spectrum simply goes like $C_{P_0}^2$.

- C_β (2.11): Since C_β sets the logarithmic slope of the ICM pressure profile at large radii (see Eq. (2.22)), it significantly influences the total integrated thermal energy of each cluster, and thus the large-angular-scale behavior of the tSZ power spectrum. An increase (decrease) in C_β leads to a decrease (increase) in the pressure profile at large radii, and therefore a corresponding decrease (increase) in the tSZ power spectrum on angular scales corresponding to the cluster outskirts and beyond. On smaller angular scales, the effect should eventually vanish, since the pressure profile on small scales is determined by the other slope parameters in the pressure profile. This trend is indeed seen at high-ℓ in Fig. 2.11. Note that a 1% change in C_β leads to a much larger change in the tSZ power spectrum at nearly all angular scales than a 1% change in C_{P_0}, suggesting that simply determining the zero-point of the $Y-M$ relation may not provide sufficient knowledge of the ICM physics to break the long-standing ICM-cosmology degeneracy in tSZ power spectrum measurements. It appears that constraints on the shape of the pressure profile itself will be necessary.

2.5 Experimental Considerations

In this section we estimate the noise in the measurement of the tSZ power spectrum. The first ingredient is instrumental noise. We describe it for the Planck experiment and for an experiment with the same specifications as the proposed PIXIE satellite [24]. The second

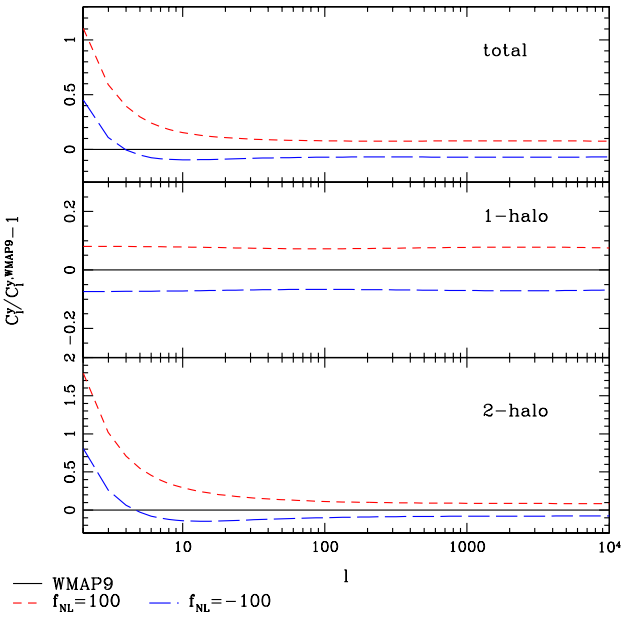

Figure 2.3: The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $f_{\rm NL} = \pm 100$.

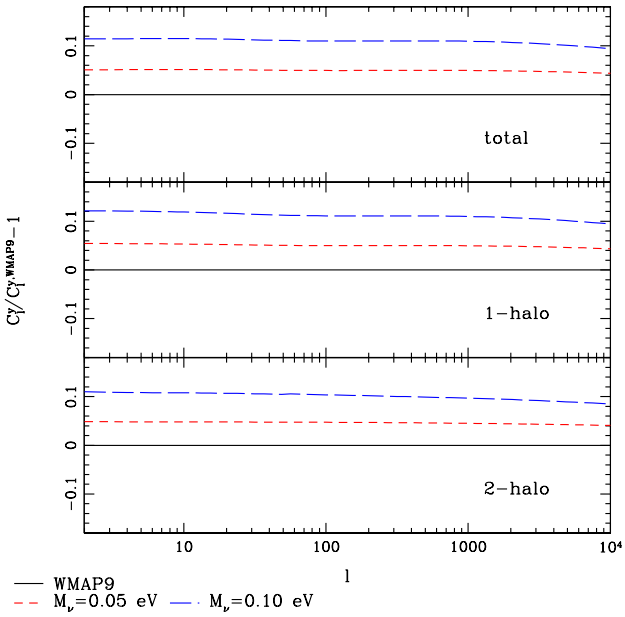

Figure 2.4: The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $M_\nu = 0.05$ eV and 0.10 eV.

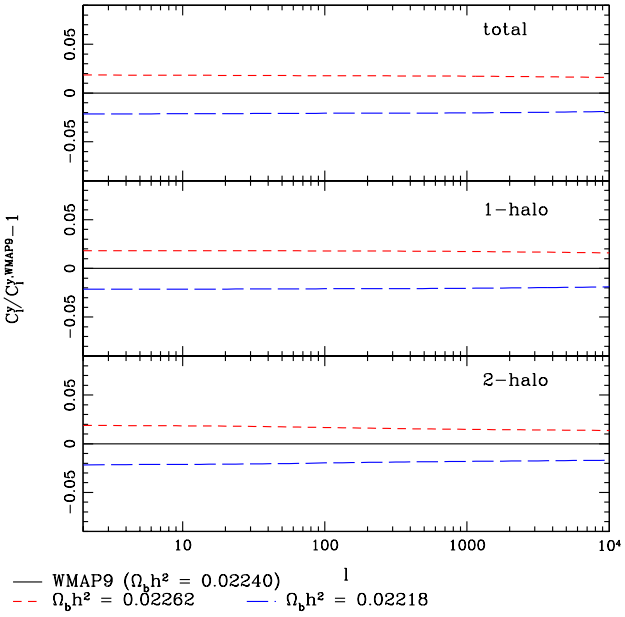

Figure 2.5: The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $\Omega_b h^2 = 0.02262$ and 0.02218.

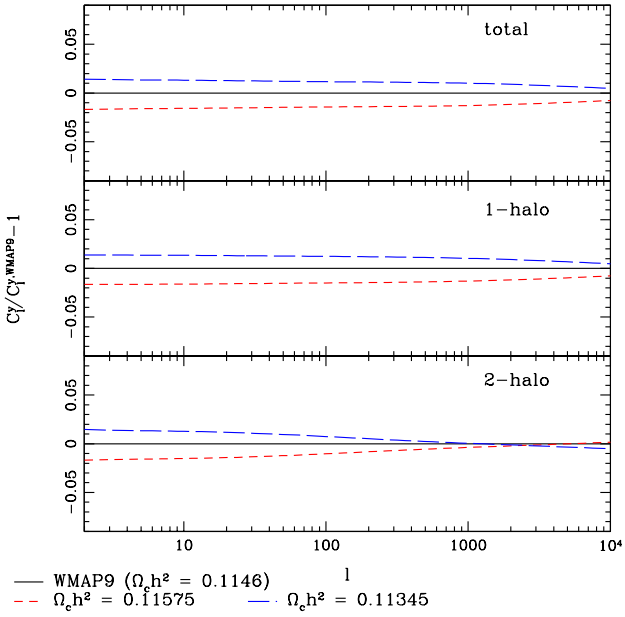

Figure 2.6: The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $\Omega_c h^2 = 0.11575$ and 0.11345.

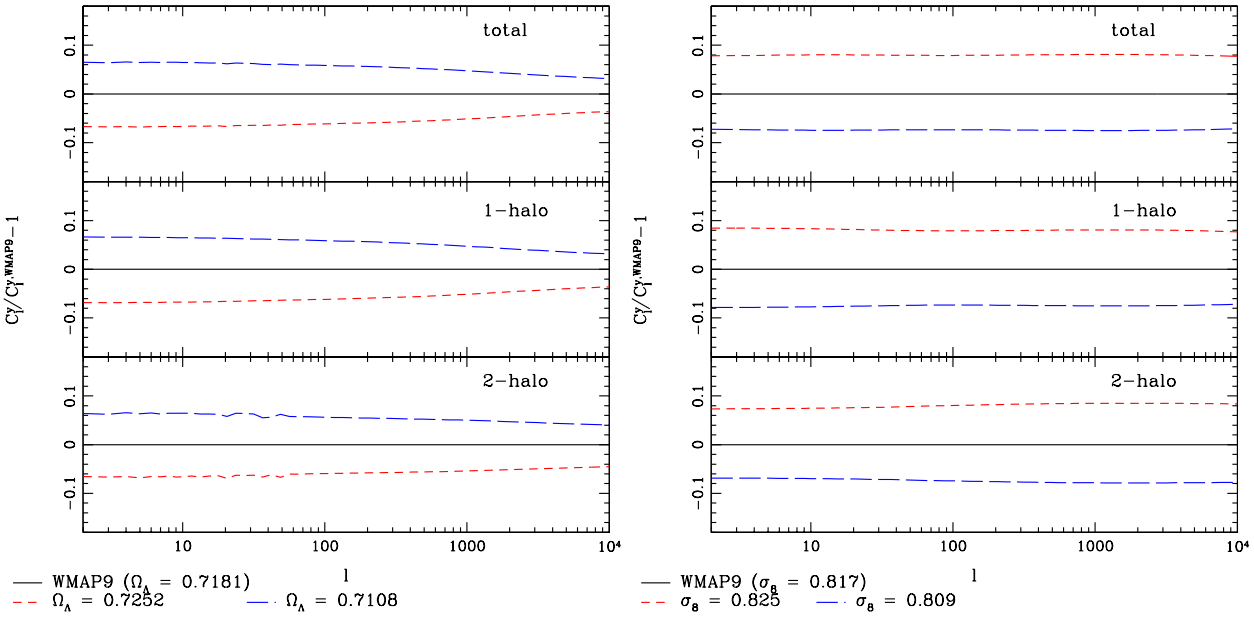

Figure 2.7: The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $\Omega_\Lambda = 0.7252$ and 0.7108.

Figure 2.8: The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $\sigma_8 = 0.825$ and 0.809.

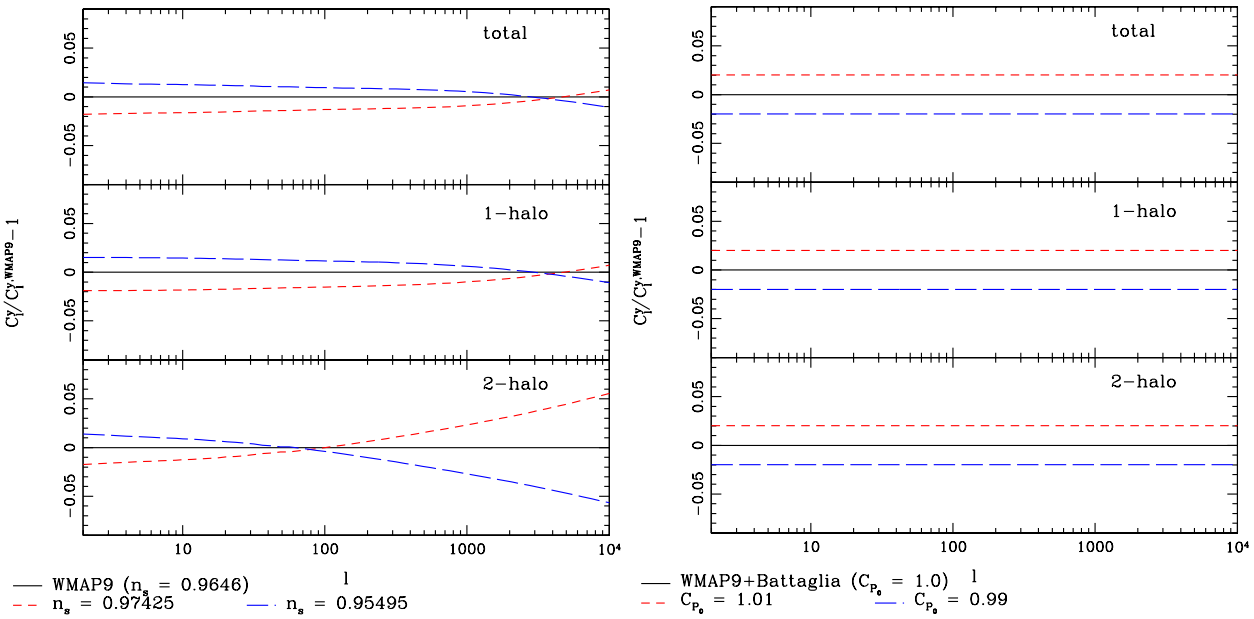

Figure 2.9: The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $n_s = 0.97425$ and 0.95495.

Figure 2.10: The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $C_{P_0} = 1.01$ and 0.99.

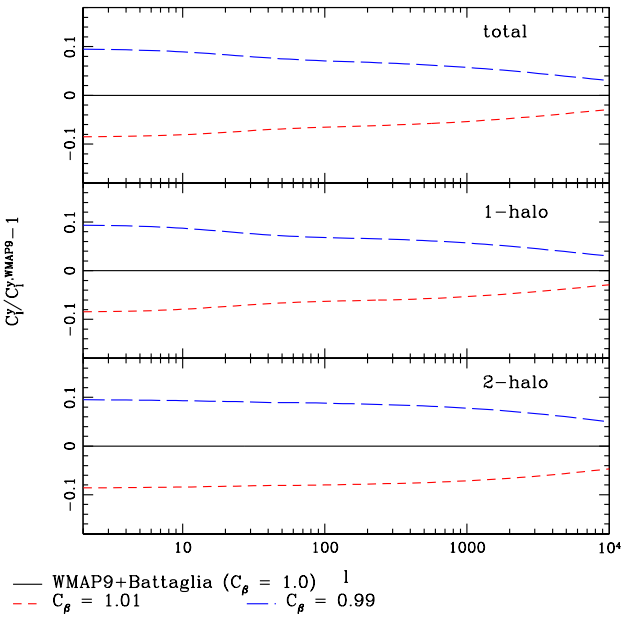

Figure 2.11: The fractional difference between the tSZ power spectrum computed using our fiducial model and power spectra computed for $C_\beta = 1.01$ and 0.99.

ingredient is foregrounds[7]. We try to give a rather complete account of all these signals and study how they can be handled using multifrequency subtraction. Our final results are in Fig. 2.12. Because of the several frequency channels, Planck and to a much larger extent PIXIE can remove all foregrounds and have a sensitivity to the tSZ power spectrum mostly determined by instrumental noise.

2.5.1 Multifrequency Subtraction

We discuss and implement multifrequency subtraction[8] along the lines of [47, 48]. The main idea is to find a particular combination of frequency channels that minimize the variance of

[7]To be precise we will consider both *foregrounds*, e.g. from our galaxy, and *backgrounds*, e.g. the CMB. On the other hand, in order to avoid repeating the cumbersome expression "foregrounds and backgrounds" we will collectively refer to all these contributions as foregrounds, sacrificing some semantic precision for the sake of an easier read.

[8]We are thankful to K. Smith for pointing us in this direction.

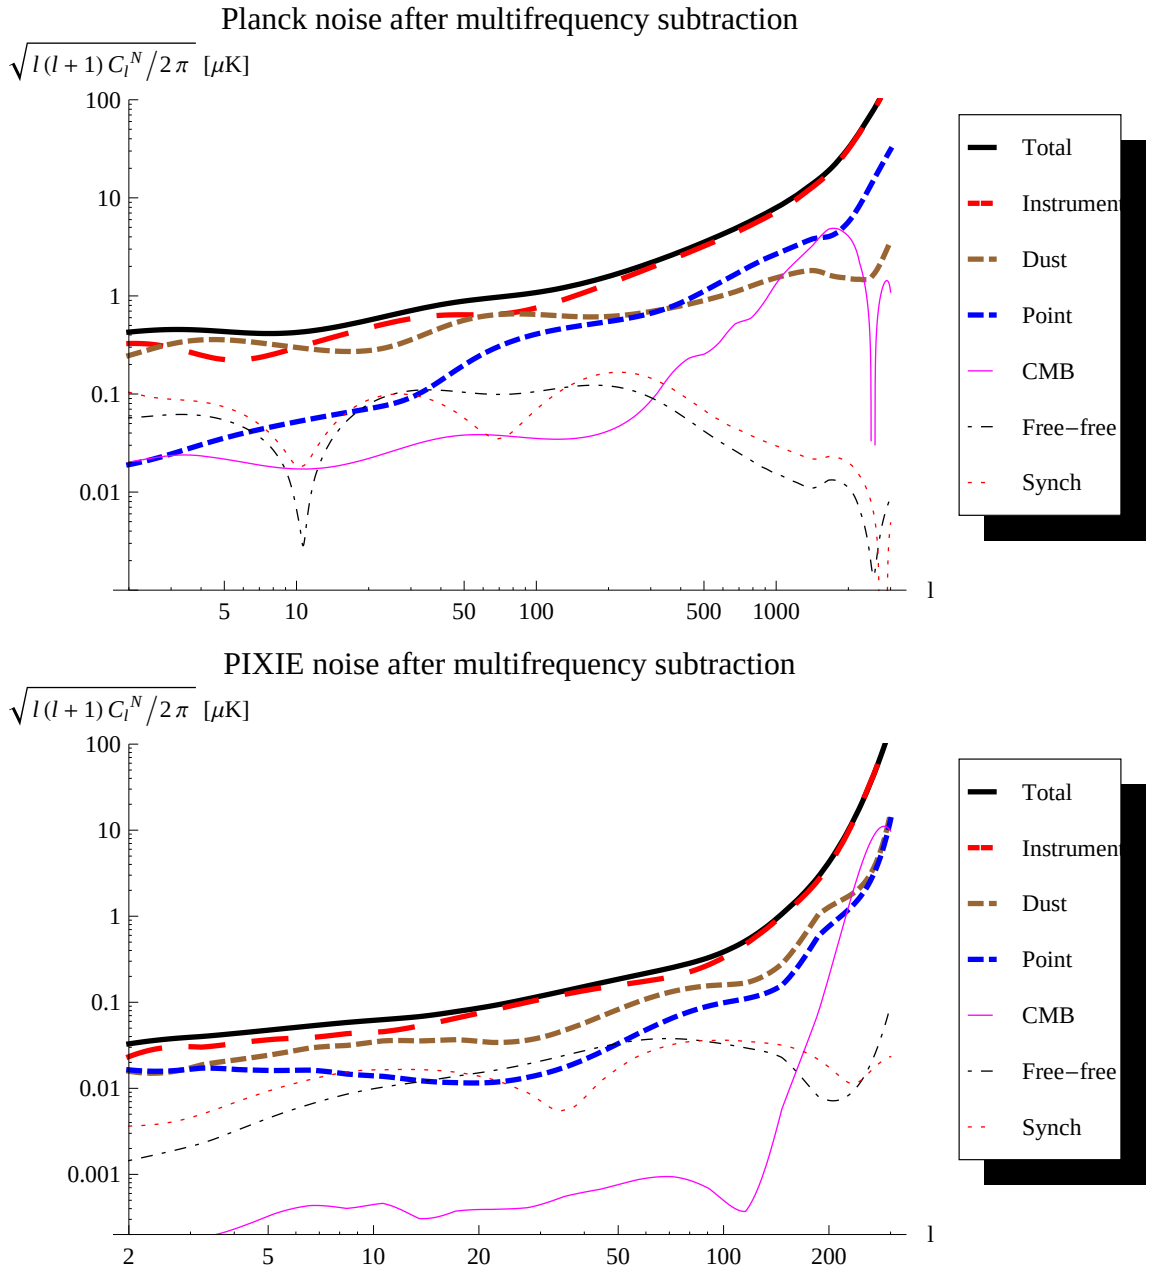

Figure 2.12: The two plots show the various contributions to the total noise per ℓ and m after multifrequency subtraction for Planck (top panel) and PIXIE (bottom panel). Because of the many frequency channels both experiments can subtract the various foregrounds and the total noise is not significantly different from the instrumental noise alone.

some desired signal, in our case the tSZ power spectrum. We hence start from

$$\hat{a}^{SZ}_{\ell m} = \sum_{\nu_i} \frac{w_i a_{\ell m}(\nu_i)}{g_{\nu_i}}\,, \tag{2.31}$$

where a^{SZ} refers to our estimator for the tSZ signal at 150 GHz (the conversion to a different frequency is straightforward), ν_i are the different frequency channels relevant for a given experiment, w_i are the weights for each channel, $a_{\ell m}(\nu_i)$ are spherical harmonic coefficients of the total measured temperature anisotropies at each frequency and finally g_{ν_i} is the tSZ spectral function defined in Section 2.4, allowing us to convert from Compton-y to ΔT. We can decompose the total signal according to $a_{\ell m} = a^{SZ}_{\ell m} + \sum_f a^f_{\ell m}$ with f enumerating all other contributions. We will assume that $\langle a^f_{\ell m} a^{f'}_{\ell m}\rangle \propto \delta_{ff'}$, i.e. different contributions are uncorrelated with each other. Dropping for the moment the ℓ and m indices, the variance of $\hat{a}^{SZ}$ is then found to be

$$\langle \hat{a}^{SZ}\hat{a}^{SZ}\rangle = C^{SZ}\left(\sum_{\nu_i} w_i\right)^2 + \sum_{\nu_i\nu_j} w_i w_j \sum_f \frac{C^f(\nu_i,\nu_j)}{g_{\nu_i} g_{\nu_j}}\,, \tag{2.32}$$

where $C^{SZ} = C^{SZ}_\ell$ is the tSZ power spectrum at 150 GHz as given in Eq. (2.16) and $C^f(\nu_i,\nu_j)$ (again the ℓ index is implicit) is the cross-correlation at different frequencies of the $a^f_{\ell m}$ of each foreground component (we will enumerate and describe these contributions shortly). To simplify the notation, in the following we will use

$$C(\nu_i,\nu_j) = C_{ij} \equiv \sum_f \frac{C^f(\nu_i,\nu_j)}{g_{\nu_i} g_{\nu_j}}\,. \tag{2.33}$$

We now want to minimize $\langle(\hat{a}^{SZ})^2\rangle$ with the constraint that the weights describe a unit response to a tSZ signal, i.e., $\sum_i w_i = 1$. This can be done using a Langrange multiplier λ

and solving the system

$$\partial_i \left[\langle (\hat{a}^{SZ})^2 \rangle + \lambda \left(\sum_i w_i - 1 \right) \right] = \partial_\lambda \left[\langle (\hat{a}^{SZ})^2 \rangle + \lambda \left(\sum_i w_i - 1 \right) \right] = 0 \,. \tag{2.34}$$

Because of the constraint $\sum_i w_i = 1$, the C^{SZ} term in $\langle (\hat{a}^{SZ})^2 \rangle$ is independent of w_i (alternatively one can keep this term and see that it drops out at the end of the computation). Then the solution of the first equation can be written as

$$w_i = -\lambda (C^{-1})_{ij} e_j = 0 \,, \tag{2.35}$$

where $e_j = 1$ is just a vector with all ones and $(C^{-1})_{ij}$ is the inverse of C_{ij} in Eq. (2.33). This solution can then be plugged back into the constraint $\sum_i w_i = 1$ to give

$$w_i = \frac{(C^{-1})_{ij} e_j}{e_k (C^{-1})_{kl} e_l} \,, \tag{2.36}$$

which is our final solution for the minimum-variance weights. From Eq. (2.32) we see that the total noise in each $\hat{a}^{SZ}_{\ell m}$ after multifrequency subtraction is

$$N_\ell = w_i C_{ij} w_j \,, \tag{2.37}$$

and the partial contributions to N_ℓ from each foreground can be obtained by substituting C with C^f (recall that there is an implicit ℓ index on C_{ij}). Notice that averaging over all m's for each ℓ and assuming that a given experiment covers only a fraction f_{sky} of the sky, the final noise in each ℓ is $N_\ell / (f_{sky}(2\ell + 1))$.

ν [GHz]	30	44	70	100	143	217	353	545	857
FWHM [arcmin]	33	24	14	10	7.1	5.0	5.0	5.0	5.0
$10^6 \Delta T/T_{CMB}$	2.0	2.7	4.7	2.5	2.2	4.8	14.7	147	6700

Table 2.1: For the nine frequency bands for Planck we report the central frequency (in GHz), the Full Width at Half Maximum (FWHM, in arcminutes, to be converted into radians in the noise computation) of each pixel, and the 1σ sensitivity to temperature per square pixel [23].

2.5.2 Foregrounds

We will consider the following sources of noise: instrumental noise (N), CMB (CMB), synchrotron ($Synch$), free-free (ff), radio and IR point sources ($Radio$ and IR) and thermal dust ($Dust$). We now discuss each of them in turn.

We assume that the noise is Gaussian with a covariance matrix diagonal in l-space and uncorrelated between different frequencies. Then [129]

$$C_\ell^N(\nu,\nu') = \delta_{\nu\nu'} \Delta T(\nu)^2 e^{\ell(\ell+1)\theta(\nu)^2} \left(8\ln 2\right) \theta(\nu)^2 \,, \tag{2.38}$$

where the beam size in radians at each frequency is $\theta(\nu) = \mathrm{FWHM}(\nu)(8\ln 2)^{-1/2} \times \pi/(180 \times 60)$. The frequency channels ν, the FWHM(ν) (Full Width at Half Maximum) and $\Delta T(\nu)$ depend on the experiment. In the following we consider the Planck satellite with specifications given in Table 2.1 and the proposed PIXIE satellite [24]. The latter is a fourth generation CMB satellite targeting primordial tensor modes through the polarization of the CMB. PIXIE will cover frequencies between 30 GHz and 6 THz with an angular resolution of 1°.6 Gaussian FWHM corresponding to $\ell_{max} \equiv \theta^{-1} \simeq 84$. The frequency coverage will be divided into 400 frequency channels each with a typical sensitivity of $\Delta I = 4 \times 10^{-24}\,\mathrm{W\,m^{-2} sr^{-1} Hz^{-1}}$ in each of 49152 sky pixels. In order to get ΔT we can use Planck's law with respect to CMB temperature

$$I(\nu, T_{CMB}) = \frac{2h}{c^2} \frac{\nu^3}{e^{\nu/(56.8GHz)} - 1} \quad \Rightarrow \quad \Delta T(\nu) = \left[\frac{\partial I(\nu,T)}{\partial T}\right]^{-1}_{T_{CMB}} \Delta I \,, \tag{2.39}$$

where we used the numerical value of fundamental constants and $T_{CMB} = 2.725\,\mathrm{K}$ to write $h\nu/(k_B T_{CMB}) = \nu/(56.8\mathrm{GHz})$. For example one finds $\Delta T(150\mathrm{GHz}) \simeq 1.00\,\mu\mathrm{K}$.

For all foregrounds except point sources we use the models and parameters discussed in [47]. We assume that different components are uncorrelated and for each component f we define

$$C^f(\nu_i, \nu_j) = \frac{\Theta^f(\nu_i)\Theta^f(\nu_j)}{\Theta^f(\nu_0)^2}\, R(\nu_i, \nu_j)\, C_\ell^f\,, \tag{2.40}$$

where $\Theta^f(\nu)$ encodes the frequency dependence, C_ℓ^f provides the ℓ-dependence and normalization at some fiducial frequency ν_0 (which will be different for different components) and finally $R(\nu_i, \nu_j)$ accounts for the frequency coherence. The latter ingredient was used in [47, 48] and discussed in [49]. The general picture is that the auto-correlation of some contribution f at two different frequencies might not be perfect. Instrumental noise is an extreme case of this in which two different frequency channels have completely uncorrelated noise, i.e. $R(\nu_i, \nu_j) = \delta_{ij}$. The CMB sits at the opposite extreme in that it follows a blackbody spectrum to very high accuracy, hence being perfectly coherent between any two frequencies: $R(\nu_i, \nu_j) = 1$ for any i, j. All other foregrounds lie in between these two extrema, having an $R(\nu_i, \nu_j)$ that starts at unity for $i = j$ and goes to zero as the frequencies are taken apart from each other. To model this Tegmark [49] proposed using

$$R(\nu_i, \nu_j) = \exp\left\{-\frac{1}{2}\left[\frac{\log\left(\nu_i/\nu_j\right)}{\xi^f}\right]^2\right\}\,, \tag{2.41}$$

where ξ^f depends on the foreground and can be estimated as $\xi^f \sim \left(\sqrt{2}\Delta\alpha\right)^{-1}$ with $\Delta\alpha$ being the variance across the sky of the spectral index of that particular component f. In the following we will write the frequency covariance as $R(\nu_i, \nu_j, \Delta\alpha)$, e.g. for the we CMB we will have $R(\nu_i, \nu_j, 0)$ while for instrumental noise $R(\nu_i, \nu_j, \infty)$.

We will parameterize the frequency dependences of the various components as

$$\Theta^{CMB}(\nu) = 1\,, \tag{2.42}$$

$$\Theta^{ff}(\nu) = \nu^{-2.15} c(\nu)\,, \tag{2.43}$$

$$\Theta^{Dust}(\nu) = \frac{c(\nu)\tilde{c}(\nu)\nu^{3+1.7}}{e^{\frac{\nu}{56.8\mathrm{GHz}}\frac{2.725K}{18K}} - 1}\,, \tag{2.44}$$

$$\Theta^{Synch}(\nu) = \nu^{-2.8} c(\nu)\,, \tag{2.45}$$

$$\Theta^{Radio}(\nu) = \nu^{-0.5}\left[\frac{\partial I(\nu,T)}{\partial T}\right]^{-1}_{T_{CMB}}\,, \tag{2.46}$$

$$\Theta^{IR}(\nu) = \nu^{2.1} I(\nu, 9.7K)\left[\frac{\partial I(\nu,T)}{\partial T}\right]^{-1}_{T_{CMB}}\,. \tag{2.47}$$

where we used Eq. (2.39) and

$$c(\nu) \equiv \left[\frac{2\sinh(x/2)}{x}\right]^2\,, \quad \tilde{c}(\nu) \propto \nu^{-2}\,. \tag{2.48}$$

Adding the information about the angular scale dependence we get

$$C_\ell^{CMB}(\nu_i,\nu_j) = C_\ell^{CMB} \tag{2.49}$$

$$C_\ell^{ff}(\nu_i,\nu_j) = \frac{\Theta^{ff}(\nu_i)\Theta^{ff}(\nu_j)}{\Theta^{ff}(31\mathrm{GHz})^2}(70\mu\mathrm{K})^2\ell^{-3}R(\nu_i,\nu_j,0.02) \tag{2.50}$$

$$C_\ell^{Dust}(\nu_i,\nu_j) = \frac{\Theta^{Dust}(\nu_i)\Theta^{Dust}(\nu_j)}{\Theta^{Dust}(90\mathrm{GHz})^2}(24\mu\mathrm{K})^2\ell^{-3}R(\nu_i,\nu_j,0.3) \tag{2.51}$$

$$C_\ell^{Synch}(\nu_i,\nu_j) = \frac{\Theta^{Synch}(\nu_i)\Theta^{Synch}(\nu_j)}{\Theta^{Synch}(19\mathrm{GHz})^2}(101\mu\mathrm{K})^2\ell^{-2.4}R(\nu_i,\nu_j,0.15) \tag{2.52}$$

$$C_\ell^{Radio}(\nu_i,\nu_j) = \frac{\Theta^{Radio}(\nu_i)\Theta^{Radio}(\nu_j)}{\Theta^{Radio}(31\mathrm{GHz})^2}(\sqrt{3}\mu\mathrm{K})^2\frac{2\pi}{\ell(\ell+1)}\left(\frac{\ell}{3000}\right)^2 \times R(\nu_i,\nu_j,0.5) \tag{2.53}$$

$$C_l^{IR}(\nu_i,\nu_j) = \frac{\Theta^{IR}(\nu_i)\Theta^{IR}(\nu_j)}{\Theta^{IR}(31\mathrm{GHz})^2}\frac{2\pi}{\ell(\ell+1)} \times \left[\left(\frac{\ell}{3000}\right)^2 7\mu\mathrm{K}^2 + \left(\frac{\ell}{3000}\right)^{2-1.2} 5.7\mu\mathrm{K}^2\right] R(\nu_i,\nu_j,0.3)\,. \tag{2.54}$$

For the CMB, C_ℓ^{CMB} is obtained using CAMB with the parameters of our fiducial cosmology. The parameters in the free-free, synchrotron and thermal dust components have been taken from the Middle Of the Road values in [47] (their Table 2 and text). The parameterization of the IR and Radio point sources follows [130].

2.5.3 Noise After Multifrequency Subtraction

Using the formulae in the last two sections we can estimate what the total variance in $\hat{a}^{SZ}$ will be after multifrequency subtraction. We denote the final result by N_ℓ for the total noise and by N_ℓ^f for each foreground component (see around Eq. (2.37)). Then we plot $[\ell(\ell+1)N_\ell/2\pi]^{1/2}$ and $\left[\ell(\ell+1)N_\ell^f/2\pi\right]^{1/2}$ in units of μK for Planck and PIXIE in Fig. 2.12. With this choice we can compare directly with the results of [48] and see that they agree for Planck once one accounts for the fact that we are constraining tSZ at 150 GHz while there the tSZ in the Rayleigh-Jeans tail is considered, which brings a factor of about 4 difference in N_ℓ.

The results for PIXIE are new. The proposed PIXIE design features 400 logarithmically-spaced frequency channels. This would require one to work with a very large multifrequency matrix which quickly becomes computationally expensive. Also, since C is very close to a singular matrix, the numerical inversion introduces some unavoidable error that becomes too large for matrices larger than about 35×35. For this reason, we decide to perform the computation binning the initial 400 channel into a smaller more manageable number. As we decrease the number of bins (i.e. bin more and more channels together) we expect two main effects to influence the final result. First, when the number of channels become comparable with the number of foregrounds that we want to subtract, the multifrequency subtraction will become very inefficient. Since we stay well away from this limit of very heavy binning, this is not an issue for us. Second, as we decrease the number of bins, the separation in frequency between adjacent bins grows larger. Because of the frequency decoherence (see discussion around Eq. (2.41)), when the bins are very far apart, they are contaminated by

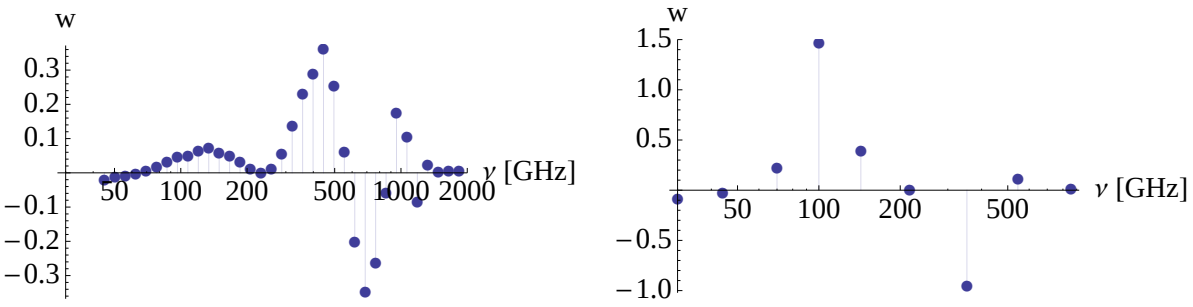

Figure 2.13: The plots show the weights appearing in Eq. (2.31) for PIXIE (left) and Planck (right) as a function of ν, for $\ell = 30$.

uncorrelated foregrounds and again the subtraction becomes inefficient. For a rough estimate of when this happens we take

$$\left[\frac{b}{400}\log\left(\frac{6000\text{GHz}}{30\text{GHz}}\right)\sqrt{2}\Delta\alpha\right]^2\frac{1}{2}\leq 1\,, \tag{2.55}$$

where b is the number of channels that we put in a bin, $(6000/30)^{1/400}$ is the logarithmic spacing and for $\Delta\alpha$ we take the largest one appearing in the foregrounds, i.e. $\Delta\alpha \sim .5$ for radio sources (dust and IR point sources have a comparable value). Then one finds that Eq. (2.55) starts being violated around $b \sim 8$, which is what we take in our analysis. The last point is that if we want to cover the frequency most relevant for tSZ with 35 bins each containing 8 channels, we cannot start from the lowest frequency covered by PIXIE, namely 30 GHz. We decide instead to start at 45 GHz, since the signal at lower frequencies is swamped anyhow by synchrotron and free-free radiation. Summarizing, we take 35 logarithmically spaced frequency channels between 45 GHZ and 1836 GHz and use them for the multifrequency subtraction. Given the arguments above we do not expect that using more channels will improve the final noise appreciably. The final noise for PIXIE after multifrequency subtraction is shown in the bottom panel of Fig. 2.12.

We show the weights w_i at $\ell = 30$ (as an example) for Planck and PIXIE in Fig. 2.13. As expected in both cases the weights are close to zero at 217 GHz, which is the null of the

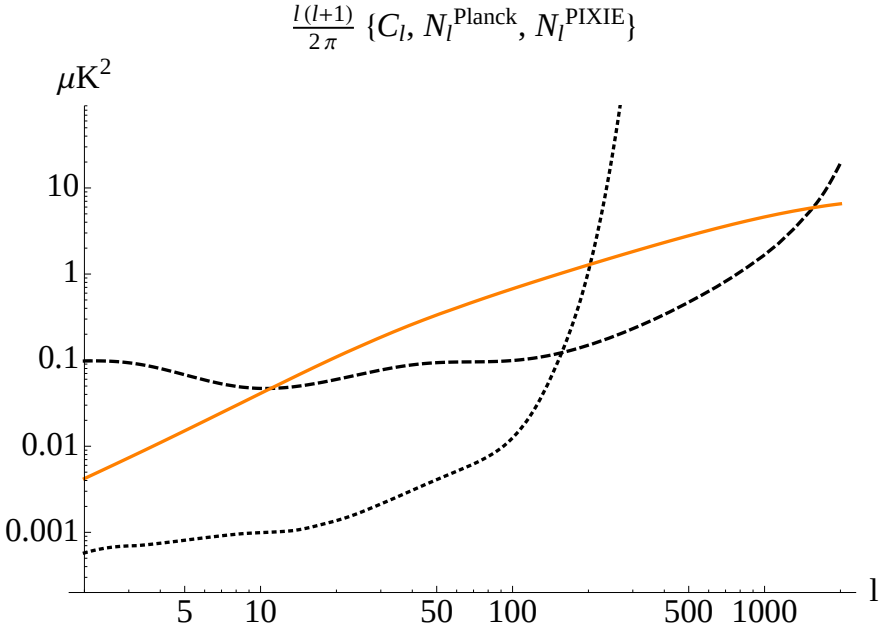

Figure 2.14: The plot compares the noise $N_\ell[f_{sky}(2\ell+1)]^{-1/2}$ for Planck (black dashed line) and PIXIE (black dotted line) with the expected tSZ power spectrum $C_\ell^{SZ} \equiv C_\ell^y$ (continuous orange line) at 150 GHz with $f_{sky} = 0.7$.

tSZ signal. Also in both cases, very low and very high frequencies have very small weights. Finally in Fig. 2.14 we compare the total noises for Planck and PIXIE after summing over m's and for a partial sky coverage $f_{sky} = 0.7$, i.e. $N_\ell[f_{sky}(2\ell+1)]^{-1/2}$, with the expected tSZ signal. PIXIE leads to an improvement of more than two orders of magnitude at low ℓ's, which is where the effect of the scale- dependent bias due to primordial non-Gaussianity arises. Since the PIXIE beam corresponds to $\ell_{max} \sim 84$, the PIXIE noise becomes very large beyond $\ell \sim$ few hundred where Planck is still expected to have signal-to-noise greater than one.

2.6 Covariance Matrix of the tSZ Power Spectrum

In order to forecast parameter constraints and the detection SNR of the tSZ power spectrum, we must compute its covariance matrix. The covariance matrix contains a Gaussian contribution from the total (signal+noise) tSZ power spectrum observed in a given experiment,

as well as a non-Gaussian cosmic variance contribution from the tSZ angular trispectrum. We compute the covariance matrix for three different experiments: Planck, PIXIE, and a future cosmic variance (CV)-limited experiment. The experimental noise after foreground subtraction is computed for Planck and PIXIE using the methods described in Section 2.5. For PIXIE, we assume a maximum multipole $\ell_{max} = 300$, while for Planck and the CV-limited experiment we assume a maximum multipole $\ell_{max} = 3000$. In the PIXIE and Planck cases, these values are well into the noise-dominated regime, so there is no reason to go to higher multipoles. For the CV-limited experiment, one can clearly compute up to as high a multipole as desired; however, it is unrealistic to imagine a satellite experiment being launched in the foreseeable future with noise levels better than PIXIE and angular resolution better than Planck, so we choose to adopt the semi-realistic value of $\ell_{max} = 3000$ for the CV-limited experiment. In all cases, we assume that the total available sky fraction used in the analysis is $f_{sky} = 0.7$, i.e., 30% of the sky is masked due to unavoidable contamination from foregrounds in our Galaxy.

In the remainder of this section, we outline the halo model-based calculations used to compute the tSZ power spectrum covariance matrix and then discuss in detail the different masking scenarios that we consider to reduce the level of cosmic variance error in the results.

Halo Model Formalism

We compute the tSZ power spectrum covariance matrix using the halo model approach, as was used for the power spectrum itself in Section 2.4.1. We provide additional background on these calculations in Chapter 1. The total tSZ power spectrum covariance matrix, $M^y_{\ell\ell'}$, is given by Eq. (1.29):

$$M^y_{\ell\ell'} = \frac{1}{4\pi f_{\rm sky}} \left(\frac{4\pi (C^y_\ell + N_\ell)^2}{\ell + 1/2} \delta_{\ell\ell'} + T^y_{\ell\ell'} \right) , \tag{2.56}$$

where C_ℓ^y is the tSZ power spectrum given by Eq. (2.16), N_ℓ is the noise power spectrum after foreground removal given by Eq. (2.37), and $T_{\ell\ell'}^y$ is the tSZ angular trispectrum. Note that we have neglected an additional term in the covariance matrix that arises from the so-called "halo sample variance" (HSV) effect (e.g., see Eq. (18) in [138] — although that result is for the weak lensing power spectrum, the tSZ result is directly analogous). The HSV term becomes negligible in the limit of a full-sky survey, which is all we consider in this chapter; thus, we do not expect this approximation to affect our results. Furthermore, we approximate the trispectrum contribution in Eq. (2.56) by the one-halo term only, which has been shown to dominate the trispectrum on nearly all angular scales [59]. We also restrict ourselves to the flat-sky limit, although the exact result is given in Chapter 1. The tSZ trispectrum is thus given by Eq. (1.28):

$$T_{\ell\ell'}^{y,1h} \approx \int dz \frac{d^2V}{dz d\Omega} \int dM \frac{dn}{dM} \left|\tilde{y}_\ell(M,z)\right|^2 \left|\tilde{y}_{\ell'}(M,z)\right|^2 . \tag{2.57}$$

Justifications for our approximations are given in Chapter 1. We compute Eq. (2.56) for our fiducial WMAP9 cosmology using each of the masking scenarios discussed in the following section. These results are then combined with the parameter variations discussed in Section 2.4.3 in order to compute Fisher matrix forecasts in Section 2.7.

Masking

For all of the experiments that we consider, the tSZ power spectrum covariance matrix in Eq. (2.56) is dominated by the trispectrum contribution over at least part of the multipole range relevant to that experiment. This issue is worse for PIXIE than for Planck, due to its lower noise levels, but even the Planck tSZ covariance is dominated by the trispectrum at some multipoles ($10 \lesssim \ell \lesssim 100$). Fortunately, the trispectrum contribution can be significantly reduced by simply masking the massive, nearby galaxy clusters that dominate the signal, especially at low-ℓ [20].

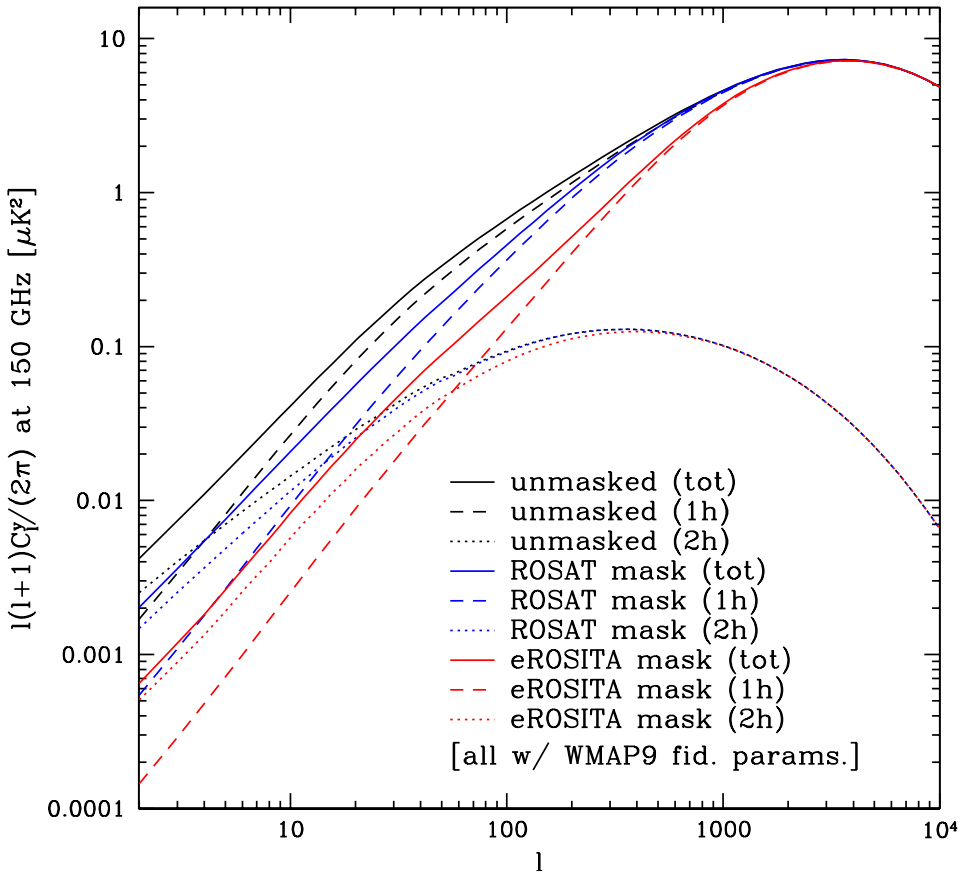

Figure 2.15: This plot shows the effect of different masking scenarios on the tSZ power spectrum calculated for our fiducial model. The amplitude of the one-halo term is suppressed relative to the two-halo term, leading to an increase in the multipole where their contributions are equal. In addition there is an overall decrease in the total signal at $\ell \lesssim 200$, as one would expect after masking the massive, nearby clusters that dominate the one-halo term in this regime.

However, there are some complications in this procedure. It is not obvious *a priori* that one should mask as many clusters as possible, since at some point the power spectrum signal itself will begin to decrease enough that the overall SNR decreases, despite the decrease in the noise. In addition, one must take care to use a very pure and complete cluster sample with a well-known selection function to do the masking; otherwise, it will be extremely difficult to properly account for the masking in the corresponding theoretical computations of the tSZ power spectrum. Finally, there is the unavoidable problem of scatter in the cluster mass-observable relation (e.g., L_X–M or Y–M), which will introduce additional uncertainty in the theoretical calculation of the masked tSZ power spectrum. Moreover, this uncertainty may be hard to precisely quantify.

In this chapter, we consider a set of masking scenarios that approximately correspond to catalogs from existing and future all-sky X-ray surveys. At present, X-ray cluster surveys likely possess the most well-understood selection functions, as compared to those derived from optical, SZ, or weak lensing data (e.g., [60, 64]). Furthermore, the scatter in the mass-observable relation for the X-ray quantity Y_X (a measure of the integrated gas pressure) is believed to be quite small, possibly $< 10\%$ [124, 101]. Although existing X-ray catalogs are only complete for fairly high masses and low redshifts (due to the steep decrease in X-ray surface brightness with redshift), these clusters are exactly the ones that need to be masked to suppress the tSZ trispectrum. Furthermore, the future catalogs from eROSITA, an upcoming X-ray satellite designed for an all-sky survey, will be complete to much lower masses and very high redshifts. Overall, it seems that X-ray-based masking is the most robust option at present.

In order to simplify our calculations and avoid the need to specify the details of any individual X-ray survey, we compute the effects of masking by removing all clusters in the tSZ calculations that lie above a mass threshold M_{mask} *and* below a redshift cutoff z_{mask}. This method also circumvents the issue of modeling scatter in the mass-observable relation.

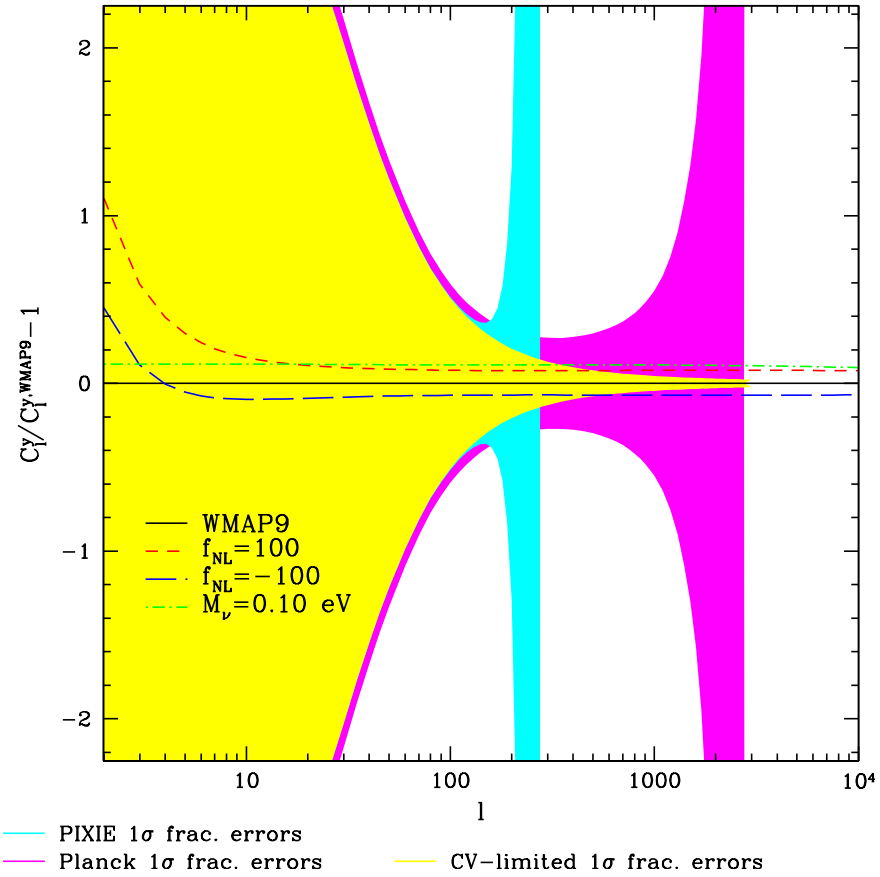

Figure 2.16: This plot shows the fractional difference between the tSZ power spectrum of our fiducial model and three different parameter variations, as labeled in the figure. In addition, we show the 1σ fractional errors on the tSZ power spectrum for each of the three experiments we consider, computed from the square root of the diagonal of the covariance matrix. In this unmasked calculation, the overwhelming influence of cosmic variance due to the tSZ trispectrum is clearly seen at low-ℓ.

Clearly a more sophisticated approach will be needed for the analysis of real data, but these choices allow us to explore several possibilities fairly quickly.

First, we consider a masking criterion based on the results of the ROSAT All-Sky Survey (RASS) [61]. The catalogs derived from RASS (e.g., BCS/eBCS [62, 63], REFLEX [64]) are nominally complete to a flux limit $f_X > 3\times10^{-12}$ erg/cm^2/s (0.5–2.0 keV band). At $z = 0.05$, this flux corresponds to a luminosity $L_X \approx 9 \times 10^{42}\ h^2$ erg/s. Using the scaling relations from [65] (for either optical or X-ray masses) to convert from L_X to mass, this luminosity corresponds to $M \approx 10^{14}\ M_\odot/h$, where M is the virial mass defined in Section 2.3.1 (we

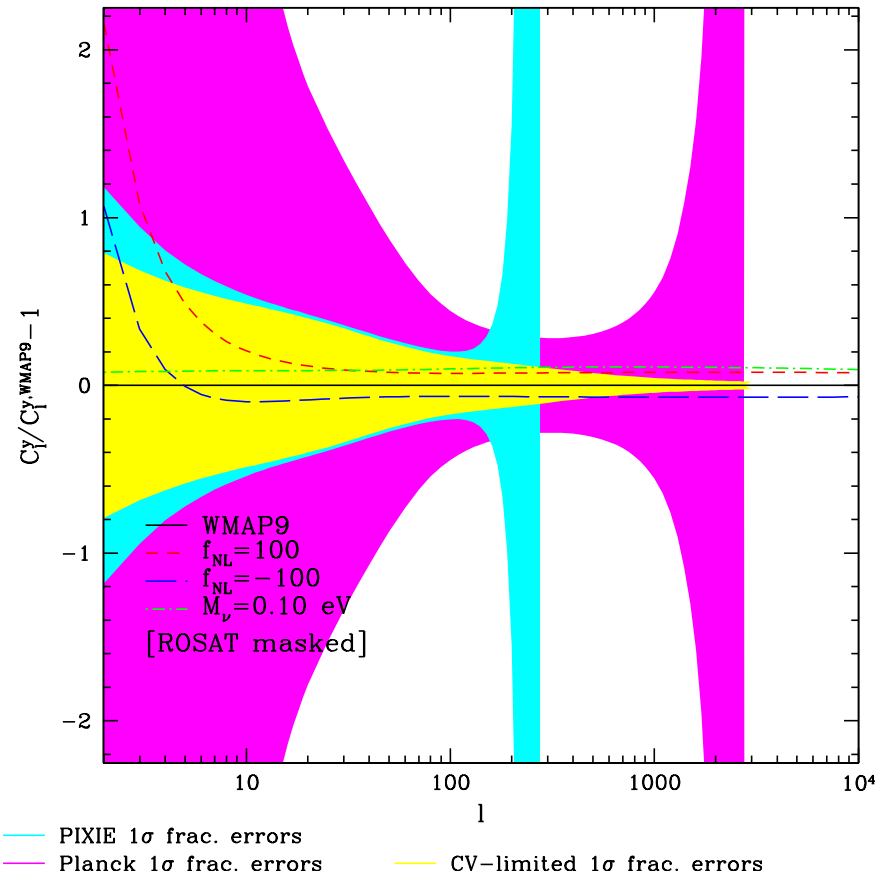

Figure 2.17: This plot is identical to Fig. 2.16 except that it has been computed for the ROSAT-masked scenario (see Section 2.6). Compared to Fig. 2.16, it is clear that the errors at low-ℓ for PIXIE and the CV-limited case have been dramatically reduced by the masking procedure. The errors from Planck are less affected because it is dominated by instrumental noise over most of its multipole range.

have also used the NFW profile [66] and the concentration-mass relation from [67] in this conversion). Thus, the catalogs derived from RASS should be $\sim 100\%$ complete for clusters with $M \gtrsim$ few $\times 10^{14}\ M_\odot/h$ at $z < 0.05$. To be conservative, we set $M_{mask} = 5 \times 10^{14}\ M_\odot/h$, which corresponds to a luminosity and flux of $L_X \approx 7 \times 10^{44}\ h^2$ erg/s and $f_X \approx 2 \times 10^{-11}$ erg/cm^2/s at $z = 0.05$, well above the flux limit given above. We will refer to this masking choice ($M_{mask} = 5 \times 10^{14}\ M_\odot/h$, $z_{mask} = 0.05$) as "ROSAT-masked" in the remainder of the chapter.

Second, we consider a masking criterion based on the upcoming all-sky survey conducted by eROSITA. The eROSITA cluster catalogs are expected to be nominally complete to a flux

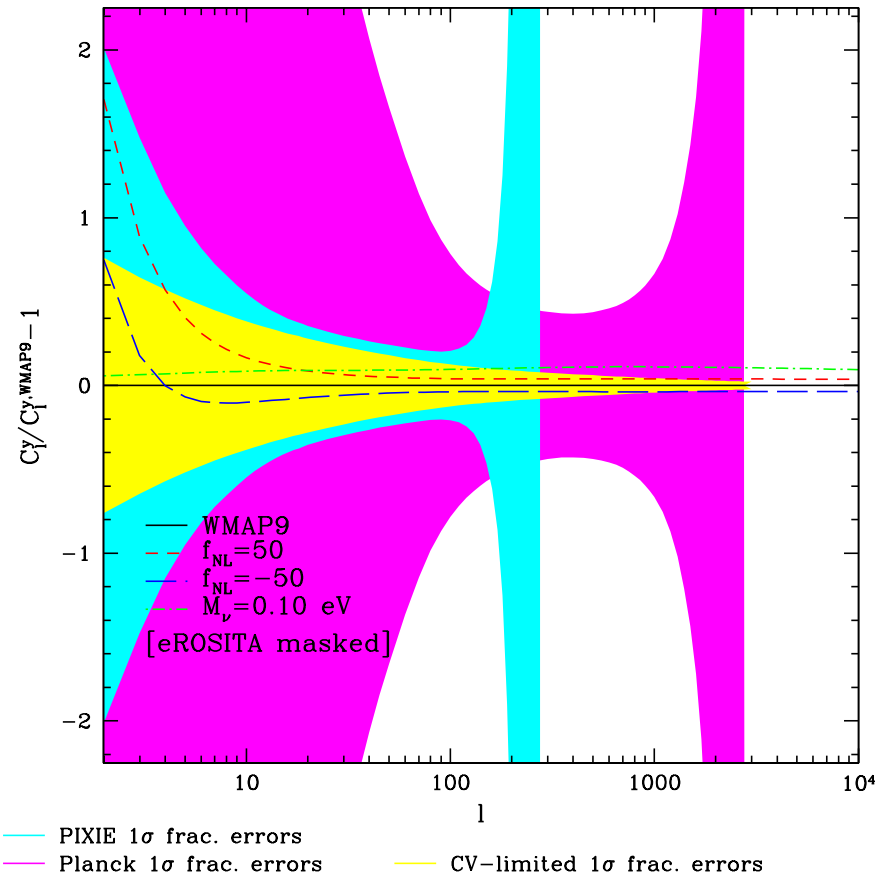

Figure 2.18: This plot is identical to Fig. 2.16 except that it has been computed for the eROSITA-masked scenario (see Section 2.6). Compared to Fig. 2.16, it is clear that the errors have been significantly reduced at low-ℓ, though again the effect is minimal for Planck. The comparison to the ROSAT-masked case in Fig. 2.17 is more complicated: it is clear that the errors are further reduced for the CV-limited case, but the PIXIE fractional errors actually increase at low-ℓ as a result of the reduction in the amplitude of the signal there due to the masking. However, the PIXIE fractional errors at higher multipoles are in fact slightly reduced compared to the ROSAT-masked case, although it is hard to see by eye in the plot.

limit $f_X > 4 \times 10^{-14}$ erg/cm^2/s (0.5–2.0 keV band) [68]. This limit is low enough that it will be possible to essentially mask clusters arbitrarily at low redshifts. At $z = 0.05$, this flux corresponds to $L_X \approx 10^{41}\ h^2$ erg/s, far below the emission from even a $\sim 10^{13}\ M_\odot/h$ group, for which $L_X \sim 10^{42}\ h^2$ erg/s (extrapolating the scaling relations from [65]). Thus, there is a wide range of possible masking choices based on the eROSITA catalogs. In principle, it would be best to fully explore the (M_{mask}, z_{mask}) parameter space and find the values that optimize the SNR for the tSZ power spectrum, or perhaps optimize the constraints on some particular parameter, such as $f_{\rm NL}$. This optimization is beyond the scope of this chapter. We anticipate that masking heavily above $z \sim 0.1$–0.2 will eventually begin to decrease the tSZ power spectrum signal too significantly, so we make a reasonably conservative cut at $z_{mask} = 0.15$. As mentioned, at these redshifts eROSITA will detect essentially all clusters (at $z = 0.15$, the flux limit corresponds to $L_X \approx 1.2 \times 10^{42}\ h^2$ erg/s, which roughly scales to $M \approx 2 \times 10^{13}\ M_\odot/h$ using the same scaling relations as above). Thus, we can choose the mass threshold at essentially any value. Our final values are ($M_{mask} = 2 \times 10^{14}\ M_\odot/h$, $z_{mask} = 0.15$); we will refer to this masking choice as "eROSITA-masked" in the remainder of the chapter.

Finally, we also consider a completely unmasked calculation, both due to its theoretical simplicity and as a means to assess how much the choice of masking can improve the SNR on the tSZ power spectrum, as well as how the forecasted parameter constraints change. Throughout the chapter, unless figures or tables are labelled with a masking choice, they have been calculated in the unmasked case.

In Fig. 2.15 we demonstrate the effects of masking on the tSZ power spectrum computed using our fiducial parameters. The main effects are a decrease in the amplitude of the one-halo term relative to the two-halo term at low-ℓ and an overall decrease in the amplitude of the power spectrum at low-ℓ. In the unmasked case, the one-halo term and two-halo term are roughly equal at $\ell \simeq 4$, while this cross-over multipole increases to $\simeq 14$ for the ROSAT-masked case and $\simeq 56$ for the eROSITA-masked case. Note that the enhancement

of the two-halo term relative to the one-halo term increases the sensitivity of the power spectrum to $f_{\rm NL}$ through the effect of the scale-dependent bias shown in Fig. 2.1. Although the amplitude of the total signal is decreased by masking, the suppression of the cosmic variance error due to the trispectrum is much larger, as seen in Figs. 2.16–2.21.

Figs. 2.16–2.18 demonstrate the significant reduction in error due to masking. These figures show the fractional difference with respect to our fiducial WMAP9 results for the tSZ power spectrum computed in three different cosmologies ($f_{\rm NL} = 100$ (or 50), $f_{\rm NL} = -100$ (or -50), and $M_\nu = 0.10$ eV). Note that these power spectra are computed using the same masking criterion as used in the covariance matrix. The figures also show the 1σ fractional errors on the tSZ power spectrum for each of the three experiments we consider. These fractional errors have been computed from the diagonal of the covariance matrix, although we emphasize that the entire covariance matrix is used in all calculations presented in the chapter. The overall implication of Figs. 2.16–2.18 is that masking is in general highly beneficial for tSZ power spectrum measurements. Although the improvements for Planck are marginal due to the fact that it is instrumental noise-dominated (see the subsequent figures), the improvements for PIXIE and the CV-limited case are dramatic. We note, however, that the reduction in the signal for the eROSITA-masked case is large enough at low-ℓ that the fractional errors for PIXIE are actually larger than in the ROSAT-masked case, because the signal at low-ℓ is becoming smaller than the PIXIE instrumental noise. Nonetheless, at higher multipoles the PIXIE fractional errors are in fact slightly smaller for the eROSITA-masked case than ROSAT-masked. This trend suggests that the optimal choice of masking for a given experiment depends on the noise levels of the experiment and the multipole range that one would like to measure with the highest SNR. For the CV-limited case, it is perhaps unsurprising that the fractional errors continue to decrease with heavier masking, as the trispectrum term is further and further suppressed relative to the Gaussian term in Eq. (2.56). It is unclear how far one can push this masking before the SNR starts to

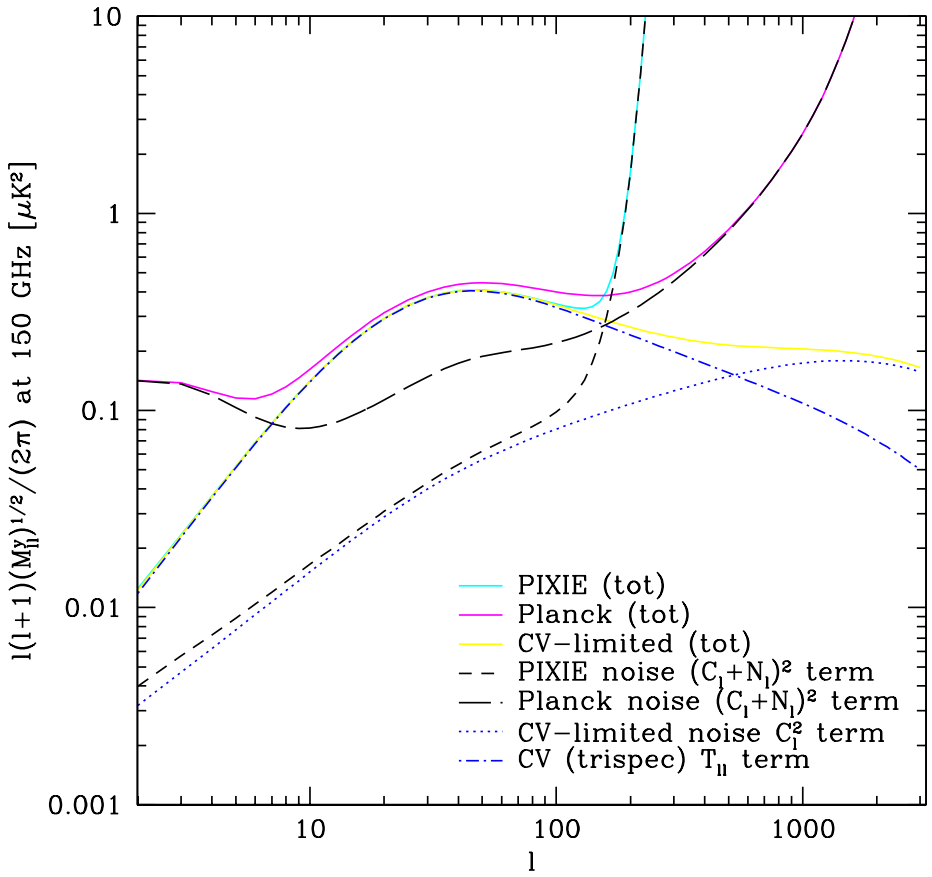

Figure 2.19: This plot shows the square root of the diagonal elements of the covariance matrix of the tSZ power spectrum, as well as the contributions from the Gaussian and non-Gaussian terms in Eq. (2.56). Results are shown for the total in each experiment, as well as the different Gaussian terms in each experiment and the trispectrum contribution that is identical for all three. See the text for discussion.

decrease due to the reduction in the signal, but Fig. 2.21 (discussed in the next paragraph) suggests that our eROSITA masking scenario is not too far from this limit.

Figs. 2.19–2.21 show the contributions to (and total of) the diagonal of the covariance matrix from the Gaussian and non-Gaussian terms in Eq. (2.56) for each of the experimental and masking scenarios. From Fig. 2.19, one can see that Planck is dominated by instrumental noise over most of its multipole range, the exception being a window from $10 \lesssim \ell \lesssim 100$ where the trispectrum contribution dominates. Fig. 2.19 also demonstrates that PIXIE is near the CV-limited case over essentially its entire multipole range. This explains the large decrease in the fractional errors for PIXIE after applying the ROSAT-masking in Fig. 2.17.

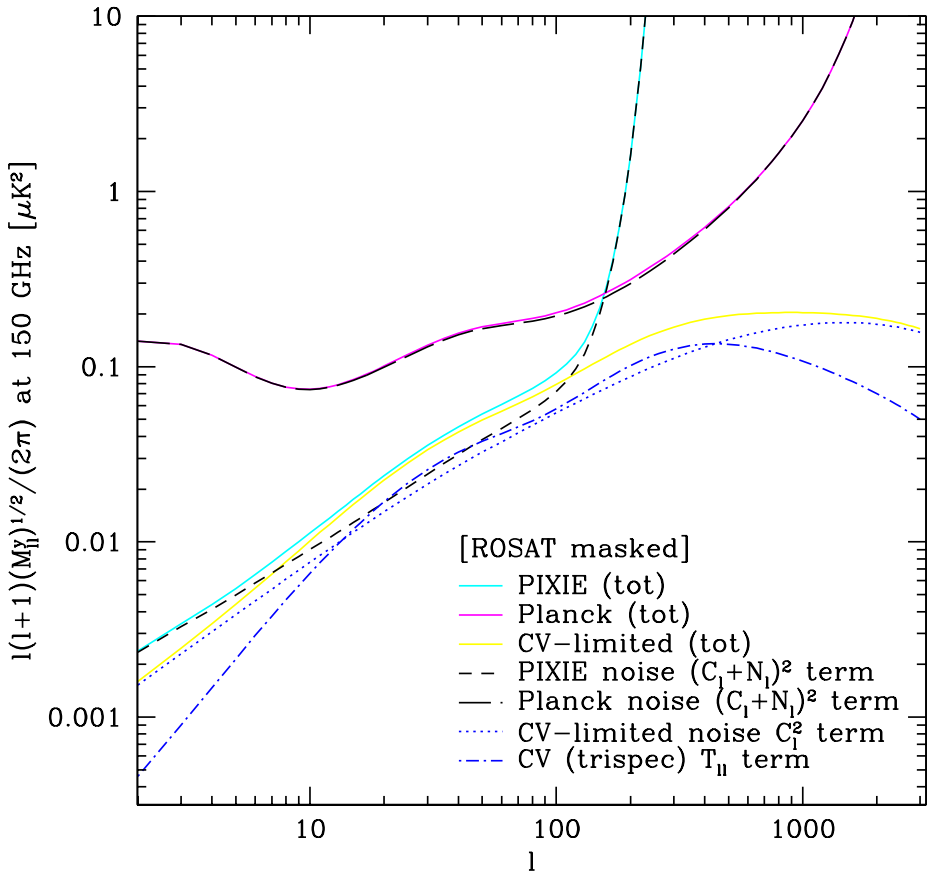

Figure 2.20: This plot is identical to Fig. 2.19 except that it has been computed for the ROSAT-masked scenario (see Section 2.6). Results are shown for the total in each experiment, as well as the different Gaussian terms in each experiment and the trispectrum contribution that is identical for all three. See the text for discussion.

Fig. 2.20 shows the covariance matrix contributions in the ROSAT-masked case, where it is clear that Planck is now completely dominated by instrumental noise at all multipoles. Thus, further masking is not beneficial for Planck. PIXIE is still near the CV-limited case for the ROSAT-masked scenario, except at low-ℓ where its instrumental noise starts to become significant. This suggests that further masking will decrease the SNR at low multipoles for PIXIE, as indeed can be seen in the eROSITA-masked scenario of Fig. 2.18. Finally, Fig. 2.21 shows the covariance matrix contributions in the eROSITA-masked case, where one can see that the trispectrum term has now been suppressed below the Gaussian term at all multipoles, even for the CV-limited experiment. This trend likely suggests that further

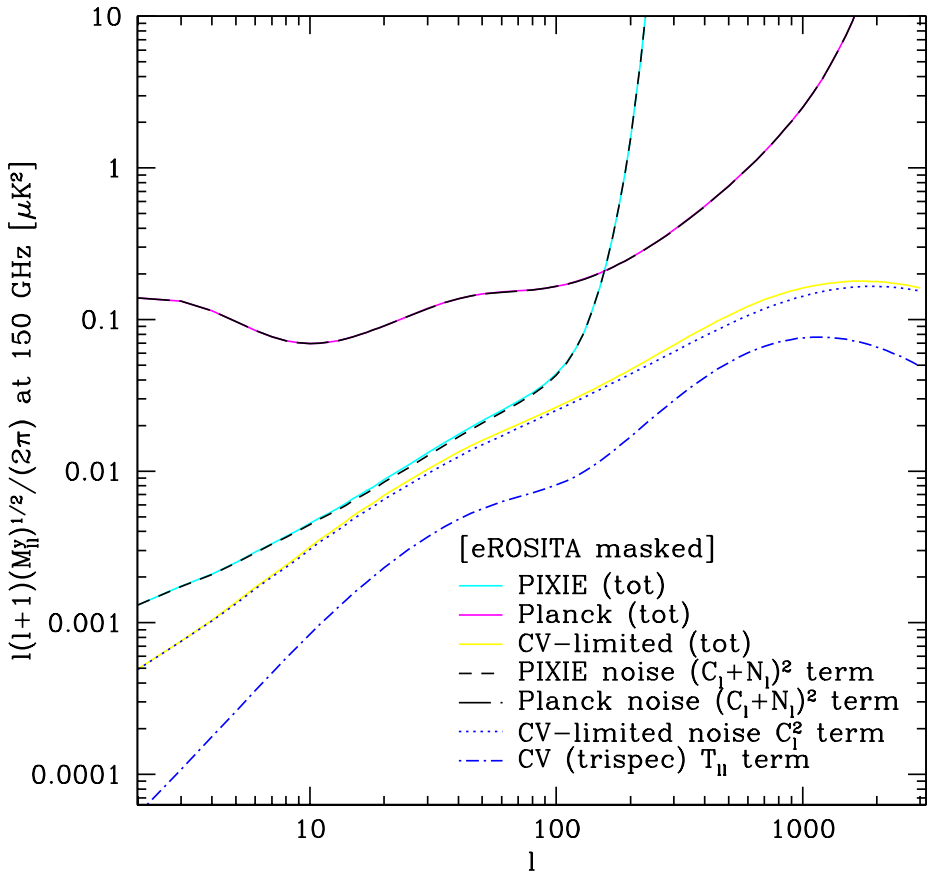

Figure 2.21: This plot is identical to Fig. 2.19 except that it has been computed for the eROSITA-masked scenario (see Section 2.6). Results are shown for the total in each experiment, as well as the different Gaussian terms in each experiment and the trispectrum contribution that is identical for all three. See the text for discussion.

masking will begin to lead to a reduction in the SNR even for the CV-limited case, although we have not computed this precisely, as we have no masking cases beyond eROSITA. Both PIXIE and Planck are clearly dominated by instrumental noise for the eROSITA-masked case, suggesting that further masking is unlikely to lead to significant improvements in SNR for these experiments.

2.7 Forecasted Constraints

Having described our model for the signal and its covariance matrix, we now use the Fisher matrix formalism [131, 129, 132, 133] to forecast the expected constraints on the parameters

listed in Eq. (2.29) for a total of nine different experimental specifications (Planck, PIXIE and CV-limited) and masking choices (unmasked, ROSAT-masked, eROSITA-masked). We compute derivatives of the tSZ power spectrum signal with respect to each parameter around the fiducial values given in Eq. (2.30). The Fisher matrix is given by

$$F_{ij} = \frac{\partial C_\ell^y}{\partial p_i} \left(M^{-1}\right)_{\ell\ell'} \frac{\partial C_{\ell'}^y}{\partial p_j}\,, \tag{2.58}$$

where p_i is the $i^{\rm th}$ parameter in Eq. (2.29), C_ℓ^y is the tSZ power spectrum given by Eq. (2.16), and $(M^{-1})_{\ell\ell'}$ is the inverse of the covariance matrix given by Eq. (2.56). Note that we only consider $f_{\rm NL} \neq 0$ and $M_\nu > 0$ cosmologies separately; in the interest of simplicity, we only seek to constrain minimal one-parameter extensions of the ΛCDM standard model. Thus, our Fisher matrices are eight-by-eight, containing the five relevant ΛCDM parameters, the two ICM physics parameters (C_{P_0} and C_β) and either $f_{\rm NL}$ or M_ν . The unmarginalized 1σ error on parameter p_i is given by $1/\sqrt{F_{ii}}$, and it describes the best possible error one can obtain when all other parameters are known exactly. The marginalized 1σ error is given by $\sqrt{(F^{-1})_{ii}}$, and it describes the case in which all other parameters are constrained from the same set of data. In the following sections, we quote the unmarginalized 1σ errors on each of the parameters from the tSZ power spectrum alone, as well as marginalized 1σ errors computed by the following methods. The results are summarized in Fig. 2.22. In addition, we provide the complete Fisher matrices for these calculations online[9].

In order to compute marginalized constraints, we find that it is necessary to include external data (in addition to basic priors, which are described below), because the degeneracies between the various parameters are too strong for the tSZ signal alone to provide meaningful constraints. For this purpose, we include a Fisher matrix computed for the imminent results from the Planck measurements of the primordial CMB temperature fluctuation power spectrum[10]. This Fisher matrix includes our five primary cosmological parameters

[9] `http://www.astro.princeton.edu/~jch/tSZFisher/`

[10] We are thankful to M. Takada for providing a computation of the Planck CMB Fisher matrix.

$\{\Omega_b h^2, \Omega_c h^2, \Omega_\Lambda, \sigma_8, n_s\}$, as well as the optical depth to reionization, τ (note that it does not include any non-ΛCDM parameters, i.e., $f_{\rm NL}$ and M_ν are not included). We marginalize over τ, since the tSZ signal is insensitive to this parameter. In addition, we include a prior on H_0 from the results of [134], who find a $\approx 3.3\%$ constraint on this parameter from cosmological distance ladder measurements. Although their central value is slightly discrepant from our fiducial value, we simply include the statistical power of this measurement as a representative measure of current constraints on h. To include this prior, we transform to a parameter set in which h lies on the diagonal, add a Gaussian prior of width $\Delta h = 0.033 \times 0.697 = 0.023$, where $h = 0.697$ is the value for our fiducial cosmology, and then transform back to our original parameters. We add the resulting "Planck CMB + H_0" Fisher matrix to the appropriate five-dimensional sub-matrix of our tSZ Fisher matrices in order to break degeneracies between the various parameters, and investigate to what extent the tSZ signal can improve on the upcoming Planck CMB measurements of the primary cosmological parameters.

In addition to the Planck CMB prior matrix for the primary cosmological parameters, we also place simple priors on the ICM physics parameters C_{P_0} and C_β. We adopt 1σ Gaussian priors of $\Delta C_{P_0} = 0.2 = \Delta C_\beta$ (recall that these parameters are normalized versions of the GNFW pressure profile amplitude and outer logarithmic slope). These values correspond roughly to the variances determined for these parameters due to the scatter between cluster pressure profiles in the simulations from which the fiducial model was obtained[11]. They also encompass values obtained from various X-ray and SZ observations [119, 3, 120], although we note that our use of $P_{200,c}$ in Eq. (2.22) makes direct comparisons of the P_0 parameter somewhat nontrivial between our model and others in the literature. Finally, these priors guarantee that the unphysical values $C_{P_0} = 0$ and $C_\beta = 0$ are highly disfavored ($> 5\sigma$). We find that the Fisher forecasts computed below are often strong enough to provide constraints on the ICM parameters well below the width of these priors (at least for C_β, which is currently

[11] N. Battaglia, priv. comm.

the less well-constrained of the two), and thus we conclude that the resulting errors are robust.

We place no priors on $f_{\rm NL}$ or M_ν as our goal is to assess the detectability of these parameters using information in the tSZ power spectrum, with no external constraints (apart from those necessary to break degeneracies amongst the primary cosmological parameters, as implemented in our Planck CMB+H_0 prior matrix).

In the following subsections we describe our main parameter forecast results, summarized in Fig. 2.22, for the two separate cases in which the standard cosmological and astrophysical parameters (see Eq. (2.29)) have been supplemented by either $f_{\rm NL}$ or M_ν .

2.7.1 $f_{\rm NL}$

We discuss local primordial non-Gaussianity and its effects on the tSZ power spectrum in Sections 2.3 and 2.4.3. The current strongest bounds come from WMAP9 and correspond to $-3 < f_{\rm NL} < 77$ at 95% CL [37], but Planck will likely improve on these limits by a factor of ≈ 3. From the top table in Fig. 2.22, one can see that at the unmarginalized level the tSZ power spectrum is quite sensitive to $f_{\rm NL}$, with an unmarginalized CV limit of $\Delta f_{\rm NL} \approx 7$. On the other hand, $f_{\rm NL}$ is very degenerate with all the other parameters and the marginalized bounds are far weaker than the current bounds (see the central right table in Fig. 2.22). Adding the tSZ power spectrum constraint to the current bounds in quadrature improves the overall constraint beyond the current bounds by at most a few percent, even in the most optimistic scenario (CV-limited experiment with eROSITA masking). We have checked that no single parameter is driving the marginalized error on $f_{\rm NL}$ away from the unmarginalized value, which would be comparable with constraints from the Planck primordial temperature bispectrum measurements. On the contrary, in order to get closer to the unmarginalized bound, all other parameters, both cosmological and astrophysical, would need to be constrained much better. We point out that if one ever hopes to constrain $f_{\rm NL}$ using the tSZ power spectrum, it is crucial to obtain high SNR on the lowest possible multipoles,

Unmargin.	f_{NL}	M_ν	C_β	C_{P0}	n_s	$h^2\,\Omega_b$	$h^2\,\Omega_c$	Ω_Λ	σ_8
Planck	40.7	0.0341	0.00547	0.0151	0.04	0.000359	0.00303	0.00441	0.00309
Planck ROSAT	40.2	0.0337	0.00541	0.0149	0.0394	0.000356	0.003	0.00438	0.00306
Planck eROSITA	37.2	0.0322	0.00433	0.0143	0.0426	0.000341	0.00293	0.00423	0.00291
PIXIE	212.	0.196	0.0268	0.0866	0.162	0.00202	0.0158	0.0216	0.018
PIXIE ROSAT	108.	0.135	0.0152	0.0559	0.122	0.00132	0.0107	0.0146	0.0123
PIXIE eROSITA	47.4	0.0569	0.00558	0.0224	0.111	0.000541	0.00498	0.00653	0.00511
CV	7.65	0.00668	0.00138	0.00284	0.017	0.0000704	0.00074	0.00105	0.00059
CV ROSAT	7.6	0.00665	0.00138	0.00283	0.0169	0.0000701	0.000738	0.00104	0.000588
CV eROSITA	7.13	0.00633	0.00128	0.00267	0.0147	0.0000664	0.00071	0.000993	0.000559

Margin.	M_ν	C_β	C_{P0}
Planck	0.668	0.0755	0.2
Planck ROSAT	0.474	0.0461	0.191
Planck eROSITA	0.376	0.0225	0.183
PIXIE	1.44	0.184	0.2
PIXIE ROSAT	0.802	0.0864	0.198
PIXIE eROSITA	0.651	0.0661	0.197
CV	0.525	0.0209	0.189
CV ROSAT	0.416	0.0165	0.151
CV eROSITA	0.2	0.00832	0.0699

Margin.	f_{NL}	C_β	C_{P0}
Planck	561.	0.0619	0.182
Planck ROSAT	464.	0.0451	0.172
Planck eROSITA	400.	0.0226	0.157
PIXIE	452.	0.0832	0.19
PIXIE ROSAT	176.	0.0584	0.186
PIXIE eROSITA	78.8	0.0508	0.19
CV	309.	0.0139	0.114
CV ROSAT	128.	0.0132	0.0588
CV eROSITA	42.9	0.00822	0.0383

Planck TT	f_{NL}	M_ν	C_β	C_{P0}	n_s	$h^2\,\Omega_b$	$h^2\,\Omega_c$	Ω_Λ	σ_8
Unmarg.	∞	∞	0.2	0.2	0.00195	0.000127	0.000479	0.00386	0.00402
Marg.	∞	∞	0.2	0.2	0.00434	0.000173	0.00143	0.0094	0.00723

Figure 2.22: The tables show the estimated unmarginalized and marginalized 1σ error bars for the indicated parameters and for a total of nine different experimental specifications (Planck, PIXIE, and CV-limited) and cluster masking choices (no masking, ROSAT masking, and eROSITA masking). The unmarginalized errors (top table) are derived using only the tSZ power spectrum. The marginalized errors (two central tables) are derived by adding as an external prior the forecasted Planck constraints from the CMB TT power spectrum plus a prior on H_0 from [134]. Since adding the tSZ power spectrum leads to effectively no improvement in the errors on the ΛCDM parameters with respect to the Planck TT priors, we do not show those numbers in the two central tables (marginalized tSZ plus Planck CMB+H_0 priors). They are equal to the numbers in the bottom table, which show the marginalized Planck CMB+H_0 priors, to $\approx 1\%$ precision. Note that all constraints on M_ν are in units of eV, while the other parameters are dimensionless.

$f_{NL}^{fid}=37$	f_{NL} unmarg.	f_{NL} marg	C_β marg.	C_{P0} marg.
Planck	39.3	524.	0.0614	0.18
Planck ROSAT	38.8	421.	0.045	0.166
Planck eROSITA	35.8	361.	0.0227	0.149
PIXIE	99.4	115.	0.068	0.19
PIXIE ROSAT	49.4	55.9	0.0548	0.186
PIXIE eROSITA	26.2	29.7	0.0486	0.189
CV	7.35	91.2	0.0138	0.0495
CV ROSAT	7.21	38.3	0.0132	0.04
CV eROSITA	6.14	13.1	0.00822	0.0354

Figure 2.23: The same as in the right central table of Fig. 2.22, but for a fiducial $f_{\rm NL} = 37$, the central value of WMAP9 [37]. A CV-limited experiment could lead to a 3σ detection.

where the signature of the scale-dependent bias breaks the degeneracy between $f_{\rm NL}$ and other parameters that influence the overall amplitude of the power spectrum. This fact explains why PIXIE would achieve much tighter marginalized constraints on $f_{\rm NL}$ using the tSZ power spectrum than Planck would (as seen in the central right table in Fig. 2.22), despite Planck's larger multipole range. Nevertheless, even the PIXIE constraints are unlikely to be competitive with those from other probes of primordial non-Gaussianity.

Thus far we have only considered constraints around our fiducial cosmology with $f_{\rm NL} = 0$. If instead we assume a fiducial $f_{\rm NL} = 37$, corresponding to the central value of WMAP9 [37], we find as before that Planck and PIXIE would have less than a 2σ detection. However, for a CV-limited experiment using eROSITA masking, we find a marginalized error $\Delta f_{\rm NL} \simeq 13$. Thus, if $f_{\rm NL}$ turns out to be of this magnitude, one could obtain a 3σ detection using methods completely independent of the primordial CMB temperature bispectrum. The full results of these calculations are given in Fig. 2.23.

2.7.2 M_ν

We discuss massive neutrinos and their effects on the tSZ power spectrum in Sections 2.3 and 2.4.3. The current strongest bounds are in the range $M_\nu \lesssim 0.3$ eV [12, 101, 42, 43]. However, a 3σ detection near this mass scale was recently claimed in [44]. From the top table in Fig. 2.22, it is clear that the unmarginalized constraints on M_ν from the tSZ power

spectrum alone are quite strong. The Planck tSZ power spectrum unmarginalized error $\Delta M_\nu \simeq 0.03$ eV is slightly smaller than the lower bounds from neutrino oscillations, and the unmarginalized error for a CV-limited experiment would lead to a very robust detection of M_ν. The bounds from PIXIE are much weaker due to its lower angular resolution; the change induced by M_ν in the tSZ power spectrum is effectively an overall amplitude shift, as seen in Fig. 2.4, and thus one can gain much more leverage on this parameter by going to higher multipoles. For PIXIE, one must mask heavily in order to reduce the CV errors to a sufficient level at the low multipoles where it observes in order to measure the effect of small neutrino masses. Masking makes much less of a difference in the Planck and CV-limited constraints on M_ν because most of their constraining power comes from higher multipoles where the masking procedure does not significantly reduce the error.

After adding the external Planck CMB+H_0 priors on the ΛCDM parameters (which are summarized in the bottom table of Fig. 2.22), the marginalized error for Planck with eROSITA masking is $\Delta M_\nu \simeq 0.37$ eV, comparable with current bounds. This could be useful to strengthen current bounds and confirm or reject a detection at this level. For a CV-limited experiment, the error is $\Delta M_\nu \simeq 0.20$ eV, a factor of two better than Planck. Note that we have derived these bounds after fully marginalizing over the primary ΛCDM parameters and the ICM gas physics parameters, so they are unlikely to be overly optimistic. In addition, we have only used the information on M_ν contained the tSZ power spectrum; we have not used any information from the Planck CMB temperature power spectrum regarding M_ν . Including primordial CMB constraints could significantly tighten the forecasted Planck errors to a level well below the current constraints. Given the imminent release of the Planck sky maps, we leave this as an avenue to be pursued with data.

Thus far we have only considered constraints around our fiducial cosmology with $M_\nu = 0$. If instead we assume a fiducial $M_\nu = 0.1$ eV, roughly corresponding to the minimum value in the inverted neutrino hierarchy [39], we find that the marginalized Planck bound after eROSITA masking becomes slightly tighter than the current upper limits in the literature

$M_\nu^{fid}=0.1$	M_ν unmarg.	M_ν marg	C_β marg.	C_{P0} marg.
Planck	0.0231	0.498	0.0836	0.199
Planck ROSAT	0.0229	0.319	0.0469	0.186
Planck eROSITA	0.0218	0.246	0.0225	0.177
PIXIE	0.131	0.981	0.188	0.2
PIXIE ROSAT	0.0906	0.524	0.0851	0.197
PIXIE eROSITA	0.0391	0.542	0.0765	0.197
CV	0.0047	0.399	0.0342	0.176
CV ROSAT	0.00468	0.271	0.0225	0.122
CV eROSITA	0.00447	0.107	0.00821	0.0506

Figure 2.24: The same as in the left central table of Fig. 2.22, but for a fiducial $M_\nu = 0.1$ eV, similar to the minimum allowed value in the inverted neutrino hierarchy [39]. Note that all constraints on M_ν are in units of eV.

$\simeq 0.3$ eV. Note that the forecasted errors depend on the fiducial neutrino mass assumed; if the actual neutrino mass is near the current upper limits, the eROSITA-masked Planck 1σ error would lie below the actual mass, suggesting a possible marginal detection. However, we have focused on forecasts for M_ν near the minimum allowed values in order to be conservative.

More interestingly, we find that the primary degeneracy of M_ν is with the ICM physics parameters C_{P_0} and C_β. These degeneracies are comparable to but much stronger than that of M_ν with any of the ΛCDM parameters (after imposing the Planck CMB+H_0 prior), which suggests that strong external constraints on the ICM pressure profile would be of great use in constraining M_ν using tSZ power spectrum measurements. We investigate this issue in detail in Fig. 2.25, which shows the forecasted 1σ error on M_ν as a function of the width of the Gaussian prior placed on C_{P_0} and C_β, which is 0.2 in our standard case. We provide plots for both the fiducial $M_\nu = 0$ case and the case where the fiducial $M_\nu = 0.1$ eV. The latter plot indicates that strengthening the prior on C_{P_0} and C_β by a factor of two would lead to a $\simeq 40\%$ decrease in the forecasted error on M_ν for the eROSITA-masked Planck experiment. Although the improvement eventually saturates as the degeneracies with other parameters become important for very tight priors on C_{P_0} and C_β, these results demonstrate the importance of strong external constraints on the ICM physics in tightening the bounds on cosmological parameters from tSZ measurements. Fortunately, it will be possible to derive

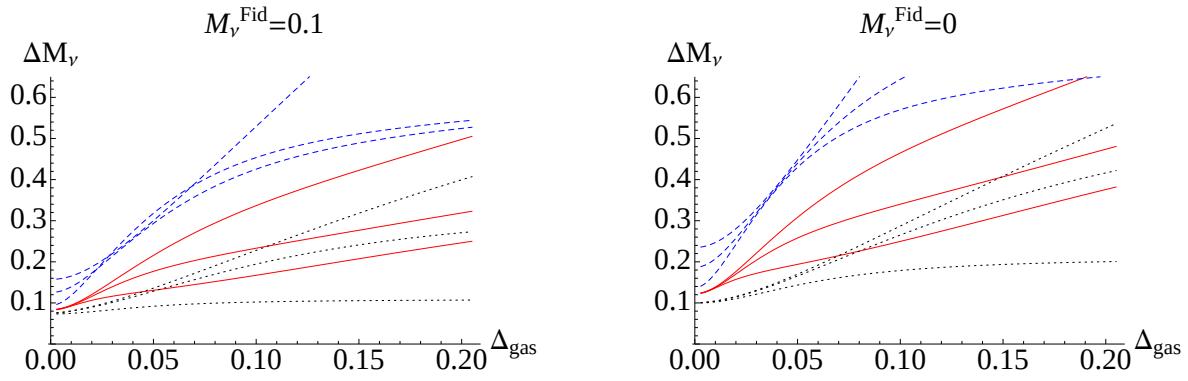

Figure 2.25: These plots show the effect of strengthening the prior on the ICM gas physics parameters C_{P_0} and C_β on our forecasted 1σ error on the sum of the neutrino masses (vertical axis, in units of eV). The left panel assumes a fiducial neutrino mass of $M_\nu = 0.1$ eV, while the right panel assumes a fiducial value of $M_\nu = 0$ eV. The solid red curves show the results for Planck, assuming (from top to bottom) no masking, ROSAT masking, or eROSITA masking. The blue dashed curves show the results for PIXIE, with the masking scenarios in the same order. The dotted black curves show the results for a CV-limited experiment, again with the masking scenarios in the same order. The horizontal axis in both plots is the value of the 1σ Gaussian prior $\Delta C_{P_0} = \Delta C_\beta \equiv \Delta_{gas}$ placed on the gas physics parameters (our fiducial value is 0.2). Note that we vary the priors simultaneously, keeping them fixed to the same value. The plots demonstrate that even modest improvements in the external priors on C_{P_0} and C_β could lead to significant decreases in the expected error on M_ν from tSZ power spectrum measurements.

such external constraints from detailed studies of the ICM pressure profile with Planck (e.g., [120]) and eROSITA themselves, which portends a bright future for tSZ-based cosmological constraints.

We conclude that the tSZ power spectrum is a promising probe for the neutrino masses and could lead to interesting results already using Planck data. We find $\Delta M_\nu \simeq 0.3$ eV around a fiducial $M_\nu = 0$ using the tSZ power spectrum alone to probe M_ν (including other data in the M_ν constraint would tighten this bound). Considering larger values of the fiducial M_ν and/or imposing stronger priors on the ICM gas physics parameters C_{P_0} and C_β can shrink this bound considerably, perhaps to the level of a 2σ measurement using Planck data alone.

2.7.3 Other Parameters

Although our focus in this chapter has been on constraining currently unknown extensions to the ΛCDM standard model, we also compute the forecasted constraints on all of the standard parameters in our model. Unfortunately, we find that adding the tSZ power spectrum data to the priors from the Planck CMB+H_0 Fisher matrix yields essentially no improvement in the forecasted errors, which are already very small using Planck+H_0 alone. Even for a CV-limited experiment with eROSITA masking, the marginalized constraints on the five ΛCDM parameters in our model only improve by $\approx 1-2\%$ beyond the marginalized priors from Planck+H_0 (which are given in the bottom table in Fig. 2.22). The primary cosmological utility of tSZ measurements is in providing a low-redshift probe of the amplitude of fluctuations, which permits a constraint on M_ν when combined with a high-redshift probe of this amplitude from the CMB.

In contrast to the results for the cosmological parameters, the forecasted constraints on the ICM physics parameters C_{P_0} and C_β are more encouraging, as seen in the central tables in Fig. 2.22, as well as Figs. 2.23 and 2.24. This is especially true of the outer slope parametrized by C_β, for which we forecast a $\approx 6-8\%$ constraint using the unmasked Planck data, which decreases to nearly the percent level after masking with eROSITA. Most current observational measurements of this parameter do not have reported error bars, but the analysis of SPT stacked SZ profiles in [3] found a $\sim 40\%$ uncertainty in this parameter after marginalization, which is probably a representative value. We thus expect this error to decrease dramatically very shortly. However, it will be difficult for Planck or PIXIE to constrain the overall normalization of the ICM pressure profile using the tSZ power spectrum alone due to its strong degeneracy with the cosmological parameters (e.g., σ_8). This is reflected in the fact that the marginalized constraints on C_{P_0} in Figs. 2.22–2.24 are often near the bound imposed by our prior, the exceptions being either of the masked Planck measurements or any of the CV-limited measurements. The best approach to obtaining constraints on C_{P_0} is likely using cross-correlations between the tSZ signal and lensing maps

of the dark matter distribution. However, it is also worth noting that we have not considered the minimal ΛCDM case in these calculations, in which both $f_{\rm NL}$ and M_ν are fixed to zero. In such a scenario, the bounds on the ICM physics parameters would be stronger than those quoted here. However, given that we know $M_\nu > 0$ in our universe, it is perhaps most reasonable to look at the marginalized constraints in that case as an example of future constraints on C_{P_0} and C_β.

2.7.4 Forecasted SNR

The cumulative SNR on the tSZ power spectrum using multipoles $\ell < \ell_{max}$ is given by

$$\mathrm{SNR}_{cumul}(\ell < \ell_{max}) = \sqrt{\sum_{\ell,\ell'=2}^{\ell_{max}} C_\ell^y (M_{\ell\ell'}^y)^{-1} C_{\ell'}^y}\,, \tag{2.59}$$

where $(M_{\ell\ell'}^y)^{-1}$ refers to the inverse of the ℓ_{max}-by-ℓ_{max} submatrix of the full covariance matrix. The cumulative SNR provides a simple way to assess the constraining power of a given experimental and masking choice on the tSZ power spectrum, without regard to constraints on particular parameters. We compute Eq. (2.59) for our fiducial cosmology using each of the experiment (Planck, PIXIE, CV-limited) and masking options (unmasked, ROSAT-masked, eROSITA-masked) considered in the previous sections. Note that the covariance matrix in Eq. (2.59) includes all contributions from cosmic variance (Gaussian and non-Gaussian) and experimental noise after foreground removal, as discussed in Section 2.6. Note that our approach includes the trispectrum (or sample variance) contribution to the covariance matrix in calculating the SNR, which is perhaps more conservative than an approach in which only the Gaussian errors are considered in assessing the SNR.

The results of these calculations are shown in Fig. 2.26. We find that Planck can detect the tSZ power spectrum with a cumulative SNR ≈ 35 using multipoles $\ell < 3000$ (see Section 2.9 for a comparison with the recently-released initial tSZ results from the Planck collaboration [137]). This result is essentially independent of the masking scenario, although

masking more heavily can lead to greater cumulative SNR using lower values of ℓ_{max}, as compared to the unmasked case. However, masking leads to significant improvement in the PIXIE results: the unmasked PIXIE cumulative SNR using $\ell < 300$ is ≈ 5.8, while the ROSAT- and eROSITA-masked results are ≈ 8.9 and 22, respectively. These results follow from the fact that PIXIE is nearly CV-limited for $\ell < 100$, as seen in the unmasked curves shown in Fig. 2.26 (and discussed earlier in Section 2.6). For the masked cases, the CV errors are reduced sufficiently that the PIXIE noise starts to become important at $\ell \lesssim 10$.

There is one subtlety of the masking procedure that can be understood by considering Fig. 2.26 in combination with Figs. 2.19–2.21. Looking at the ROSAT-masked case for Planck in Fig. 2.20, it appears that the Planck errors are dominated by the Gaussian instrumental noise term at all ℓ, and hence that masking further for Planck should be harmful rather than helpful; this appears to be confirmed by the fact that the Planck fractional errors in the eROSITA-masked case in Fig. 2.18 are indeed larger than in the ROSAT-masked case in Fig. 2.17. However, it is clear in Fig. 2.26 that the eROSITA-masked Planck case has a larger total SNR than the ROSAT-masked case. The resolution of this apparent discrepancy lies in the fact that the masking continues to suppress the off-diagonal terms in the covariance matrix, which arise solely from the trispectrum term in Eq. (2.56), even as the on-diagonal fractional errors begin to increase. The plots in Figs. 2.16–2.21 only show the diagonal entries in the covariance matrix, and thus one may not realize the effect of the masking on the off-diagonal terms in the covariance matrix without examining the cumulative SNR (this result is also implied by the improved parameter constraints for the eROSITA-masked Planck case given in the previous sections). These considerations imply that there is likely an optimal masking choice for a given experimental noise level and survey specifications, but obtaining the precise answer to this question lies beyond the scope of this chapter.

Overall, Fig. 2.26 indicates that near-term data promises highly significant detections of the tSZ power spectrum on larger angular scales than have been probed thus far by ACT and SPT.

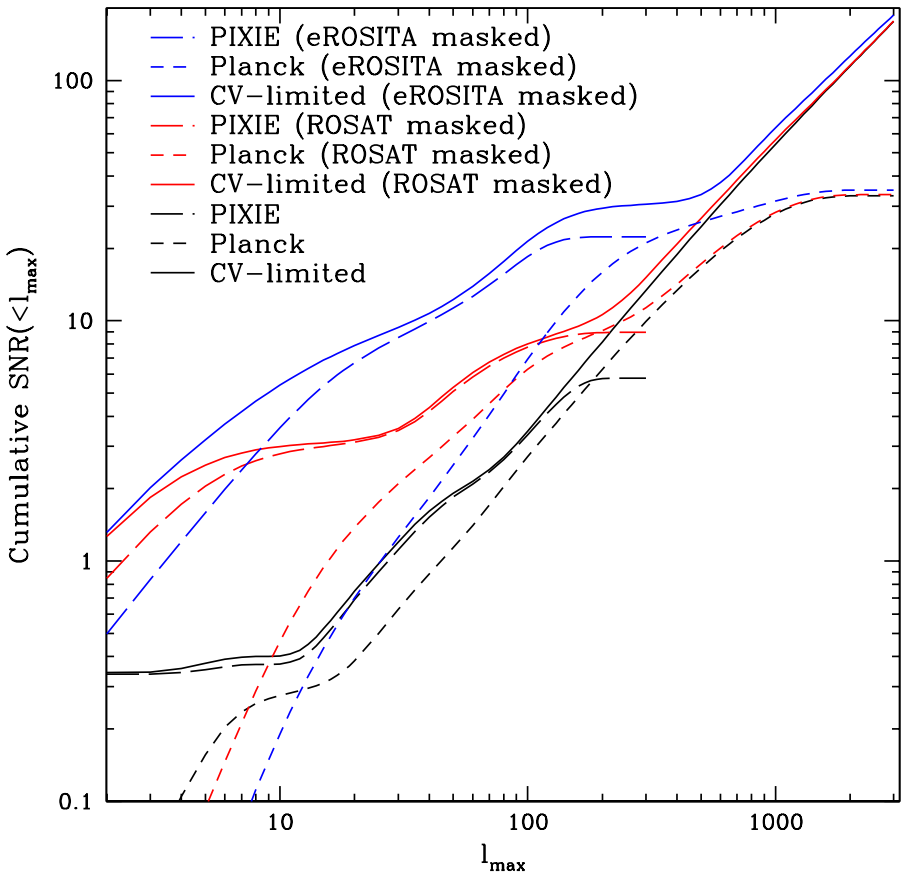

Figure 2.26: This plot shows the cumulative SNR achievable on the tSZ power spectrum for each of nine different experimental and making scenarios. The solid curves display results for a CV-limited experiment, the short dashed curves show results for Planck, and the long dashed curves show results for PIXIE. The different colors correspond to different masking options, as noted in the figure. Note that PIXIE is close to the CV limit over its signal-dominated multipole range. The total SNR using the imminent Planck data is ≈ 35, essentially independent of the masking option used (note that masking can increase the cumulative SNR up to lower multipoles, however, as compared to the unmasked case).

2.8 Discussion and Outlook

In this chapter we have performed a comprehensive analysis of the possible constraints on cosmological and astrophysical parameters achievable with measurements of the tSZ power spectrum from upcoming full-sky CMB observations, with a particular focus on extensions to the ΛCDM standard model parametrized by $f_{\rm NL}$ and M_ν . We have included all of the important physical effects due to these additional parameters, including the change to the halo mass function and the scale-dependent halo bias induced by primordial non-Gaussianity. Our halo model calculations of the tSZ power spectrum include both the one- and two-halo terms, and we use the exact expressions where necessary to obtain accurate results on large angular scales. We model the ICM pressure profile using parameters that have been found to agree well with existing constraints, and furthermore we model the uncertainty in the ICM physics by freeing two of these parameters. We also include a realistic treatment of the instrumental noise for the Planck and PIXIE experiments, accounting for the effects of foregrounds by using a multifrequency subtraction technique. Our calculations of the covariance matrix of the tSZ power spectrum include both the Gaussian noise terms and the non-Gaussian cosmic variance term due to the tSZ trispectrum. We investigate two masking scenarios motivated by the ROSAT and eROSITA all-sky surveys, which significantly reduce the large errors that would otherwise be induced by the trispectrum term, especially at low-ℓ. Finally, we use these calculations to forecast constraints on $f_{\rm NL}$, M_ν , the primary ΛCDM parameters, and two parameters describing the ICM pressure profile.

Our primary findings are as follows:

- The tSZ power spectrum can be detected with a total SNR > 30 using the imminent Planck data up to $\ell = 3000$, regardless of masking (see Section 2.9 for a comparison with the initial tSZ power spectrum results released by the Planck collaboration while this manuscript was under review [137]) ;

- The tSZ power spectrum can be detected with a total SNR between ≈ 6 and 22 using the future PIXIE data up to $\ell = 300$, with the result being sensitive to the level of masking applied to remove massive, nearby clusters;

- Adding the tSZ power spectrum information to the forecasted constraints from the Planck CMB temperature power spectrum and existing H_0 data is unlikely to significantly improve constraints on the primary cosmological parameters, but may give interesting constraints on the extensions we consider:

  - If the true value of $f_{\rm NL}$ is near the WMAP9 ML value of ≈ 37, a future CV-limited experiment combined with eROSITA-masking could provide a 3σ detection, completely independent of the primordial CMB temperature bispectrum; alternatively, PIXIE could give $1-2\sigma$ evidence for such a value of $f_{\rm NL}$ with this level of masking;

  - If the true value of M_ν is near 0.1 eV, the Planck tSZ power spectrum with eROSITA masking can provide upper limits competitive with the current upper bounds on M_ν ; with stronger external constraints on the ICM physics, Planck with eROSITA masking could provide $1-2\sigma$ evidence for massive neutrinos from the tSZ power spectrum, depending on the true neutrino mass;

- Regardless of the cosmological constraints, Planck will allow for a very tight constraint on the logarithmic slope of the ICM pressure profile in the outskirts of galaxy clusters, and may also provide some information on the overall normalization of the pressure profile (which sets the zero point of the $Y-M$ relation).

Our results are subject to a few caveats. We have made the usual Fisher matrix approximation that the likelihood function is nearly Gaussian around our fiducial parameter values, but this should be safe for small variations, which are all that we consider (in particular, σ_8 is tightly constrained by the external Planck CMB prior, and it would be most likely to have a non-Gaussian likelihood). We have also neglected any tSZ signal from the intergalactic

medium, filaments, or other diffuse structures, but the comparison between simulations and halo model calculations in [116] indicates that this approximation should be quite good. We have also assumed that the mass function parameters are perfectly well known, while in reality some uncertainties remain, especially in the exponential tail. Given that our most optimistic results involve masking nearly all of the clusters that live in the exponential tail, we believe that our forecasts should be fairly robust to the mass function uncertainties, in contrast to cluster count calculations which are highly sensitive to small changes in the tail of the mass function. Finally, we have only included the flat-sky version of the one-halo term in our computations of the tSZ power spectrum covariance matrix. For the masked calculations, the flat-sky result should suffice, since massive, nearby clusters are removed; however, it is possible that the two-, three-, or four-halo terms could eventually become relevant in the masked calculations. These would be largest at low-ℓ, however, where our primary interest is in constraining $f_{\rm NL}$, which does not appear very optimistic in any case. Thus, we neglect these corrections for our purposes.

There are many future extensions of this work involving higher-order tSZ statistics and cross-correlations with other tracers of large-scale structure. Recent work on the tSZ bispectrum and skewness [15, 21, 22, 17] indicates that significantly stronger constraints on both cosmology and the ICM physics can be obtained by using higher-order statistics. These may also be a better place to look for $f_{\rm NL}$ constraints, as the additional powers of the halo bias could lead to a larger signal at low-ℓ than in the power spectrum (controlling systematics will be of paramount importance, as will masking to reduce the very large cosmic variance due to the tSZ six-point function). Determining an optimal strategy to extract the neutrino mass through combinations of tSZ statistics and cross-spectra with other tracers is also work in progress. The key factor remains breaking the degeneracy with the ICM physics, or, more optimistically, simultaneously constraining both the ICM and cosmological parameters using tSZ measurements.

2.9 Comparison with Planck Results

While this manuscript was under review, the Planck team released its initial set of cosmological results, including the construction of a Compton-y map and estimation of the tSZ power spectrum from this map [137]. In this section, we provide a brief comparison of the publicly released Planck results with the forecasts in our work. Based on the analysis presented in Fig. 2.16, the Planck tSZ power spectrum should be signal-dominated over roughly the multipole range $100 \lesssim \ell \lesssim 1500$. Comparing with Fig. 15 in [137], this prediction is in very good agreement. The bandpowers and associated error bars presented in Table 3 of [137] imply a detection of the tSZ power spectrum with SNR ≈ 12.3, assuming a diagonal covariance matrix (no off-diagonal terms are presented in the Planck results). Our basic Planck forecast predicts a SNR ≈ 35. There are several reasons behind the difference in SNR between our forecast and the initial Planck result:

- The usable fraction of sky in the Planck analysis ($f_{sky} \approx 0.5$) is found to be somewhat lower than that used in our analysis ($f_{sky} = 0.7$) — this is primarily due to heavier masking of Galactic dust contamination than we anticipated;

- The number of frequency channels used in the Planck analysis (six HFI channels only) is lower than that used in our analysis (all nine of the HFI and LFI channels);

- The Planck analysis explicitly accounts for uncertainties in the contributions of the relevant foreground components (clustered CIB, IR point sources, and radio point sources) to the derived tSZ power spectrum, and finds that these uncertainties dominate the overall errors on the power spectrum; in our analysis, we have used reasonable models for the foregrounds to compute their contributions to the tSZ power spectrum, but have not explicitly propagated through uncertainties in these models to the final error bars. Our choice on this issue is partly driven by the fact that it is hard to quantify these uncertainties — in the Planck analysis, simulations are used to provide an estimate of the amplitude of each residual spectrum in the derived Compton-y power

spectrum, but a 50% uncertainty remains. This uncertainty dominates the derived errors on the tSZ bandpowers, which is likely the main reason the Planck analysis SNR is significantly lower than our forecast.

In addition to these differences, we also note that the Planck analysis does not consider the possibility of masking nearby, massive clusters to reduce the sample variance in the tSZ power spectrum — however, it appears that the angular trispectrum contribution to the covariance matrix may not have been included in the Planck analysis at all, in which case masking would not be relevant. Regardless, this is another difference between our forecasts and the Planck results. Ultimately, the Planck analysis constrains $\sigma_8 \left(\Omega_m/0.28\right)^{3.2/8.1} = 0.784 \pm 0.016$. These constraints are obtained in a ΛCDM framework with all other cosmological parameters fixed, using the pressure profile of [119] with a hydrostatic mass bias of 20%, and without including the Planck constraints from the primordial CMB temperature power spectrum. In addition, the amplitudes of the foreground contributions to the tSZ power spectrum are allowed to vary, and are included as nuisance parameters. Given that this framework is rather different from ours, it is difficult to compare directly our forecasted parameter constraints with those obtained in the Planck analysis. Regardless, it is clear that the tSZ power spectrum is a useful cosmological probe, especially of the low-redshift amplitude of fluctuations, provided that uncertainties related to cluster gas physics and foreground contamination are treated with care.

2.10 Acknowledgments

We are thankful to Masahiro Takada for providing a computation of the Planck CMB Fisher matrix. We are also grateful to Kendrick Smith for his suggestion of using the multifrequency subtraction techniques to remove foregrounds, as well as many other insightful conversations. We thank Nick Battaglia, Bruce Draine, Eiichiro Komatsu, Marilena LoVerde, Blake Sherwin, David Spergel, and Matias Zaldarriaga for a number of helpful exchanges.

JCH is supported by NASA Theory Grant NNX12AG72G. EP is supported in part by the Department of Energy grant DE-FG02-91ER-40671.

Bibliography

[1] Sunyaev, R. A., & Zeldovich, Y. B. 1970, Ap&SS, 7, 3

[2] Reese, E. D., Mroczkowski, T., Menanteau, F., et al. 2011, arXiv:1108.3343

[3] Plagge, T. J., Marrone, D. P., Abdulla, Z., et al. 2012, arXiv:1203.2175

[4] Lancaster, K., Birkinshaw, M., Gawroński, M. P., et al. 2011, Mon. Not. R. Astron. Soc, 418, 1441

[5] AMI Consortium, Shimwell, T. W., Rodriguez-Gonzalvez, C., et al. 2011, arXiv:1101.5590

[6] Marriage, T. A., Acquaviva, V., Ade, P. A. R., et al. 2011, ApJ, 737, 61

[7] Williamson, R., Benson, B. A., High, F. W., et al. 2011, ApJ, 738, 139

[8] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2011, A&A, 536, A8

[9] Hasselfield, M., Hilton, M., Marriage, T. A., et al. 2013, arXiv:1301.0816

[10] Reichardt, C. L., Stalder, B., Bleem, L. E., et al. 2013, ApJ, 763, 127

[11] Dunkley, J., Hlozek, R., Sievers, J., et al. 2011, ApJ, 739, 52

[12] Sievers, J. L., Hlozek, R. A., Nolta, M. R., et al. 2013, arXiv:1301.0824

[13] Reichardt, C. L., Shaw, L., Zahn, O., et al. 2012, ApJ, 755, 70

[14] Story, K. T., Reichardt, C. L., Hou, Z., et al. 2012, arXiv:1210.7231

[15] Wilson, M. J., Sherwin, B. D., Hill, J. C., et al. 2012, Phys. Rev. D, 86, 122005

[16] Rubiño-Martín, J. A., & Sunyaev, R. A. 2003, Mon. Not. R. Astron. Soc, 344, 1155

[17] Crawford, T. M., Schaffer, K. K., Bhattacharya, S., et al. 2013, arXiv:1303.3535

[18] Cole, S., & Kaiser, N. 1988, Mon. Not. R. Astron. Soc, 233, 637

[19] Komatsu, E., & Seljak, U. 2002, Mon. Not. R. Astron. Soc, 336, 1256

[20] Komatsu, E., & Kitayama, T. 1999, ApJ, 526, L1

[21] Hill, J. C., & Sherwin, B. D. 2013, Phys. Rev. D, 87, 023527

[22] Bhattacharya, S., Nagai, D., Shaw, L., Crawford, T., & Holder, G. P. 2012, ApJ, 760, 5

[23] Planck Blue Book. http://www.rssd.esa.int/SA/PLANCK/ docs/Bluebook-ESA-SCI%282005%291_V2.pdf

[24] Kogut, A., Fixsen, D. J., Chuss, D. T., et al. 2011, J. Cosmology Astropart. Phys., 7, 25

[25] Acquaviva, V., Bartolo, N., Matarrese, S., & Riotto, A. 2003, Nuclear Physics B, 667, 119

[26] Maldacena, J. 2003, Journal of High Energy Physics, 5, 13

[27] Creminelli, P., & Zaldarriaga, M. 2004, J. Cosmology Astropart. Phys., 10, 6

[28] Salopek, D. S., & Bond, J. R. 1990, Phys. Rev. D, 42, 3936

[29] Gangui, A., Lucchin, F., Matarrese, S., & Mollerach, S. 1994, ApJ, 430, 447

[30] Komatsu, E., & Spergel, D. N. 2001, Phys. Rev. D, 63, 063002

[31] Linde, A., & Mukhanov, V. 1997, Phys. Rev. D, 56, 535

[32] Lyth, D. H., & Wands, D. 2002, Physics Letters B, 524, 5

[33] Lyth, D. H., Ungarelli, C., & Wands, D. 2003, Phys. Rev. D, 67, 023503

[34] Buchbinder, E. I., Khoury, J., & Ovrut, B. A. 2008, Physical Review Letters, 100, 171302

[35] Creminelli, P., & Senatore, L. 2007, J. Cosmology Astropart. Phys., 11, 10

[36] Lehners, J.-L., & Steinhardt, P. J. 2008, Phys. Rev. D, 77, 063533

[37] Bennett, C. L., Larson, D., Weiland, J. L., et al. 2012, arXiv:1212.5225

[38] Giannantonio, T., Ross, A. J., Percival, W. J., et al. 2013, arXiv:1303.1349

[39] McKeown, R. D., & Vogel, P. 2004, Physics Reports, 394, 315

[40] Abazajian, K. N., Calabrese, E., Cooray, A., et al. 2011, Astroparticle Physics, 35, 177

[41] Lesgourgues, J., & Pastor, S. 2006, Physics Reports, 429, 307

[42] Mantz, A., Allen, S. W., & Rapetti, D. 2010, Mon. Not. R. Astron. Soc, 406, 1805

[43] de Putter, R., Mena, O., Giusarma, E., et al. 2012, ApJ, 761, 12

[44] Hou, Z., Reichardt, C. L., Story, K. T., et al. 2012, arXiv:1212.6267

[45] Lueker, M., Reichardt, C. L., Schaffer, K. K., et al. 2010, ApJ, 719, 1045

[46] Trac, H., Bode, P., & Ostriker, J. P. 2011, ApJ, 727, 94

[47] Tegmark, M., Eisenstein, D. J., Hu, W., & de Oliveira-Costa, A. 2000, ApJ, 530, 133

[48] Cooray, A., Hu, W., & Tegmark, M. 2000, ApJ, 540, 1

[49] Tegmark, M. 1998, ApJ, 502, 1

[50] Remazeilles, M., Aghanim, N., & Douspis, M. 2013, Mon. Not. R. Astron. Soc, 430, 370

[51] Remazeilles, M., Delabrouille, J., & Cardoso, J.-F. 2011, Mon. Not. R. Astron. Soc, 410, 2481

[52] Leach, S. M., Cardoso, J.-F., Baccigalupi, C., et al. 2008, A&A, 491, 597

[53] Sadeh, S., Rephaeli, Y., & Silk, J. 2007, Mon. Not. R. Astron. Soc, 380, 637

[54] Roncarelli, M., Moscardini, L., Branchini, E., et al. 2010, Mon. Not. R. Astron. Soc

[55] Shimon, M., Sadeh, S., & Rephaeli, Y. 2011, Mon. Not. R. Astron. Soc, 412, 1895

[56] Shimon, M., Rephaeli, Y., Itzhaki, N., Dvorkin, I., & Keating, B. G. 2012, Mon. Not. R. Astron. Soc, 427, 828

[57] Taburet, N., Douspis, M., & Aghanim, N. 2010, Mon. Not. R. Astron. Soc, 404, 1197 as, 402, 923

[58] Mak, D. S. Y., & Pierpaoli, E. 2013, arXiv:1303.2081

[59] Cooray, A. 2001, Phys. Rev. D, 64, 063514

[60] Melin, J.-B., Bartlett, J. G., & Delabrouille, J. 2005, A&A, 429, 417

[61] Voges, W., Aschenbach, B., Boller, T., et al. 1999, A&A, 349, 389

[62] Ebeling, H., Edge, A. C., Bohringer, H., et al. 1998, Mon. Not. R. Astron. Soc, 301, 881

[63] Ebeling, H., Edge, A. C., Allen, S. W., et al. 2000, Mon. Not. R. Astron. Soc, 318, 333

[64] Böhringer, H., Schuecker, P., Guzzo, L., et al. 2001, A&A, 369, 826

[65] Popesso, P., Biviano, A., Böhringer, H., Romaniello, M., & Voges, W. 2005, A&A, 433, 431

[66] Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, ApJ, 490, 493

[67] Duffy, A. R., Schaye, J., Kay, S. T., & Dalla Vecchia, C. 2008, Mon. Not. R. Astron. Soc, 390, L64

[68] Merloni, A., Predehl, P., Becker, W., et al. 2012, arXiv:1209.3114

[69] Hinshaw, G., Larson, D., Komatsu, E., et al. 2012, arXiv:1212.5226

[70] Bryan, G. L., & Norman, M. L. 1998, ApJ, 495, 80

[71] Press, W. H., & Schechter, P. 1974, ApJ, 187, 425

[72] Sheth, R. K., & Tormen, G. 1999, Mon. Not. R. Astron. Soc, 308, 119

[73] Sheth, R. K., Mo, H. J., & Tormen, G. 2001, Mon. Not. R. Astron. Soc, 323, 1

[74] Jenkins, A., Frenk, C. S., White, S. D. M., et al. 2001, Mon. Not. R. Astron. Soc, 321, 372

[75] Tinker, J., Kravtsov, A. V., Klypin, A., et al. 2008, ApJ, 688, 709

[76] Warren, M. S., Abazajian, K., Holz, D. E., & Teodoro, L. 2006, ApJ, 646, 881

[77] Bhattacharya, S., Heitmann, K., White, M., et al. 2011, ApJ, 732, 122

[78] Mantz, A., Allen, S. W., Rapetti, D., & Ebeling, H. 2010, Mon. Not. R. Astron. Soc, 406, 1759

[79] Lucchin, F., & Matarrese, S. 1988, ApJ, 330, 535

[80] Colafrancesco, S., Lucchin, F., & Matarrese, S. 1989, ApJ, 345, 3

[81] Chiu, W. A., Ostriker, J. P., & Strauss, M. A. 1998, ApJ, 494, 479

[82] Robinson, J., Gawiser, E., & Silk, J. 2000, ApJ, 532, 1

[83] Koyama, K., Soda, J., & Taruya, A. 1999, Mon. Not. R. Astron. Soc, 310, 1111

[84] Verde, L., Jimenez, R., Kamionkowski, M., & Matarrese, S. 2001, Mon. Not. R. Astron. Soc, 325, 412

[85] Matarrese, S., Verde, L., & Jimenez, R. 2000, ApJ, 541, 10

[86] Lo Verde, M., Miller, A., Shandera, S., & Verde, L. 2008, J. Cosmology Astropart. Phys., 4, 14

[87] D'Amico, G., Musso, M., Noreña, J., & Paranjape, A. 2011, J. Cosmology Astropart. Phys., 2, 1

[88] LoVerde, M., & Smith, K. M. 2011, J. Cosmology Astropart. Phys., 8, 3

[89] Grossi, M., Verde, L., Carbone, C., et al. 2009, Mon. Not. R. Astron. Soc, 398, 321

[90] Pillepich, A., Porciani, C., & Hahn, O. 2010, Mon. Not. R. Astron. Soc, 402, 191

[91] Wagner, C., Verde, L., & Boubekeur, L. 2010, J. Cosmology Astropart. Phys., 10, 22

[92] Hoyle, B., Jimenez, R., & Verde, L. 2011, Phys. Rev. D, 83, 103502

[93] Cayón, L., Gordon, C., & Silk, J. 2011, Mon. Not. R. Astron. Soc, 415, 849

[94] Mortonson, M. J., Hu, W., & Huterer, D. 2011, Phys. Rev. D, 83, 023015

[95] Enqvist, K., Hotchkiss, S., & Taanila, O. 2011, J. Cosmology Astropart. Phys., 4, 17

[96] Harrison, I., & Hotchkiss, S. 2012, arXiv:1210.4369

[97] Hoyle, B., Jimenez, R., Verde, L., & Hotchkiss, S. 2012, J. Cosmology Astropart. Phys., 2, 9

[98] Brandbyge, J., Hannestad, S., Haugb olle, T., & Wong, Y. Y. Y. 2010, J. Cosmology Astropart. Phys., 9, 14

[99] Marulli, F., Carbone, C., Viel, M., Moscardini, L., & Cimatti, A. 2011, Mon. Not. R. Astron. Soc, 418, 346

[100] Ichiki, K., & Takada, M. 2012, Phys. Rev. D, 85, 063521

[101] Vikhlinin, A., Kravtsov, A. V., Burenin, R. A., et al. 2009, ApJ, 692, 1060

[102] Bardeen, J. M., Bond, J. R., Kaiser, N., & Szalay, A. S. 1986, ApJ, 304, 15

[103] Mo, H. J., & White, S. D. M. 1996, Mon. Not. R. Astron. Soc, 282, 347

[104] Smith, R. E., Scoccimarro, R., & Sheth, R. K. 2007, Phys. Rev. D, 75, 063512

[105] Tinker, J. L., Robertson, B. E., Kravtsov, A. V., et al. 2010, ApJ, 724, 878

[106] Dalal, N., Doré, O., Huterer, D., & Shirokov, A. 2008, Phys. Rev. D, 77, 123514

[107] Slosar, A., Hirata, C., Seljak, U., Ho, S., & Padmanabhan, N. 2008, J. Cosmology Astropart. Phys., 8, 31

[108] Matarrese, S., & Verde, L. 2008, ApJ, 677, L77

[109] Giannantonio, T., & Porciani, C. 2010, Phys. Rev. D, 81, 063530

[110] Desjacques, V., Seljak, U., & Iliev, I. T. 2009, Mon. Not. R. Astron. Soc, 396, 85

[111] Grossi, M., Dolag, K., Branchini, E., Matarrese, S., & Moscardini, L. 2007, Mon. Not. R. Astron. Soc, 382, 1261

[112] Smith, K. M., Ferraro, S., & LoVerde, M. 2012, J. Cosmology Astropart. Phys., 3, 32

[113] Nozawa, S., Itoh, N., Suda, Y., & Ohhata, Y. 2006, Nuovo Cimento B Serie, 121, 487

[114] Cooray, A., & Sheth, R. 2002, Physics Reports, 372, 1

[115] Limber, D. N. 1954, ApJ, 119, 655

[116] Battaglia, N., Bond, J. R., Pfrommer, C., & Sievers, J. L. 2012, ApJ, 758, 75

[117] Battaglia, N., Bond, J. R., Pfrommer, C., Sievers, J. L., & Sijacki, D. 2010, ApJ, 725, 91

[118] Bode, P., Ostriker, J. P., Cen, R., & Trac, H. 2012, arXiv:1204.1762

[119] Arnaud, M., Pratt, G. W., Piffaretti, R., et al. 2010, A & A, 517, A92

[120] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, A&A, 550, A131

[121] Battaglia, N., Bond, J. R., Pfrommer, C., & Sievers, J. L. 2012, ApJ, 758, 74

[122] Komatsu, E., & Seljak, U. 2001, Mon. Not. R. Astron. Soc, 327, 1353

[123] Shaw, L. D., Nagai, D., Bhattacharya, S., & Lau, E. T. 2010, ApJ, 725, 1452

[124] Vikhlinin, A., Burenin, R. A., Ebeling, H., et al. 2009, ApJ, 692, 1033

[125] Nagai, D., Kravtsov, A. V., & Vikhlinin, A. 2007, ApJ, 668, 1

[126] Kaiser, N. 1986, Mon. Not. R. Astron. Soc, 222, 323

[127] Voit, G. M. 2005, Reviews of Modern Physics, 77, 207

[128] Sun, M., Sehgal, N., Voit, G. M., et al. 2011, ApJ, 727, L49

[129] Knox, L.. 1995, Phys. Rev. D, 52, 4307

[130] Dunkley, J., Calabrese, E., Sievers, J., et al. 2013, arXiv:1301.0776

[131] Fisher, R. A.. 1935, J. Roy. Stat. Soc., 98, 39

[132] Tegmark, M., Taylor, A. N., & Heavens, A. F. 1997, ApJ, 480, 22

[133] Jungman, G., Kamionkowski, M., Kosowsky, A., & Spergel, D. N. 1996, Phys. Rev. D, 54, 1332

[134] Riess, A. G., Macri, L., Casertano, S., et al. 2011, ApJ, 730, 119

[135] Loverde, M., & Afshordi, N. 2008, Phys. Rev. D, 78, 123506

[136] Pillepich, A., Porciani, C., & Reiprich, T. H. 2012, Mon. Not. R. Astron. Soc, 422, 44

[137] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5081

[138] Sato, M., Hamana, T., Takahashi, R., et al. 2009, ApJ, 701, 945

Chapter 3

Detection of Thermal SZ – CMB Lensing Cross-Correlation in Planck Nominal Mission Data

3.1 Abstract

The nominal mission maps from the Planck satellite contain a wealth of information about secondary anisotropies in the cosmic microwave background (CMB), including those induced by the thermal Sunyaev-Zel'dovich (tSZ) effect and gravitational lensing. As both the tSZ and CMB lensing signals trace the large-scale matter density field, the anisotropies sourced by these processes are expected to be correlated. We report the first detection of this cross-correlation signal, which we measure at 6.2σ significance using the Planck data. We take advantage of Planck's multifrequency coverage to construct a tSZ map using internal linear combination techniques, which we subsequently cross-correlate with the publicly-released Planck CMB lensing potential map. The cross-correlation is subject to contamination from the cosmic infrared background (CIB), which is known to correlate strongly with CMB lensing. We correct for this contamination via cross-correlating our tSZ map with the Planck

857 GHz map and confirm the robustness of our measurement using several null tests. We interpret the signal using halo model calculations, which indicate that the tSZ – CMB lensing cross-correlation is a unique probe of the physics of intracluster gas in high-redshift, low-mass groups and clusters. Our results are consistent with extrapolations of existing gas physics models to this previously unexplored regime and show clear evidence for contributions from both the one- and two-halo terms, but no statistically significant evidence for contributions from diffuse, unbound gas outside of collapsed halos. We also show that the amplitude of the signal depends rather sensitively on the amplitude of fluctuations (σ_8) and the matter density (Ω_m), scaling as $\sigma_8^{6.1}\Omega_m^{1.5}$ at $\ell = 1000$. We constrain the degenerate combination $\sigma_8(\Omega_m/0.282)^{0.26} = 0.824 \pm 0.029$, a result that is in less tension with primordial CMB constraints than some recent tSZ analyses. We also combine our measurement with the Planck measurement of the tSZ auto-power spectrum to demonstrate a technique that can in principle constrain both cosmology and the physics of intracluster gas simultaneously. Our detection is a direct confirmation that hot, ionized gas traces the dark matter distribution over a wide range of scales in the universe (~ 0.1–$50\,\mathrm{Mpc}/h$).

3.2 Introduction

The primordial anisotropies in the cosmic microwave background radiation (CMB) have been a powerful source of cosmological information over the past two decades, establishing the ΛCDM standard model and constraining its parameters to nearly percent-level precision [1, 2, 3, 4, 5]. However, as CMB photons propagate from the surface of last scattering, they are affected by a number of physical processes that produce secondary anisotropies. These processes include gravitational lensing of CMB photons by intervening large-scale structures along the line-of-sight (LOS) [6] and the Sunyaev-Zel'dovich (SZ) effect due to inverse Compton scattering of CMB photons off free electrons along the LOS [7, 8]. The SZ effect contains two distinct contributions: one due to the thermal motion of hot electrons,

primarily located in the intracluster medium (ICM) of galaxy clusters (the thermal SZ effect), and one due to the bulk motion of electrons along the LOS (the kinetic SZ effect). The thermal SZ (tSZ) effect is roughly an order of magnitude larger than the kinetic SZ (kSZ) effect for typical massive galaxy clusters [9] and is now recognized as a robust method with which to find and characterize clusters in blind surveys of the microwave sky [10, 11, 12]. The first detection of the kSZ effect was only achieved recently due to its much smaller amplitude [13]. We will consider only the tSZ effect in this work.

Measurements of CMB lensing have improved dramatically in recent years, from first detections using cross-correlation techniques [14, 15] to precision measurements using CMB data alone [16, 17, 46]. The lensing effect leads to $\approx$ 2–3 arcminute coherent distortions of $\sim$degree-sized regions in the CMB temperature field. This distortion produces statistical anisotropy in the small-scale CMB fluctuations, which allows the lensing potential (or convergence) to be reconstructed (e.g., [83]). The CMB lensing signal is a probe of the integrated mass distribution out to the surface of last scattering at $z \approx 1100$ — the most distant source plane possible. The redshift kernel for the CMB lensing signal peaks around $z \sim 2$, although it receives significant contributions over a wide redshift range ($0.1 \lesssim z \lesssim 10$). The CMB lensing power spectrum is a robust cosmological probe as it is primarily sourced by Fourier modes that are still in the linear regime, allowing accurate theoretical predictions to be computed, at least for multipoles $\ell \lesssim 1000$ [69]. In short, the CMB lensing field is an excellent tracer of the large-scale matter density field over a broad redshift range.

The tSZ signal is predominantly sourced by hot, ionized electrons located in the deep potential wells of massive galaxy clusters. These halos trace the large-scale matter density field — in fact, they are highly biased tracers ($b \sim$ 3–4). Thus, the tSZ and CMB lensing signals must be correlated. We report the first detection of this cross-correlation in this chapter. The level of correlation is sensitive to the particular details of how ICM gas traces the dark matter overdensity field, on both large scales (the "two-halo" term) and small scales (the "one-halo term"). In addition, due to the rare nature of the objects responsible for the

tSZ signal, the amplitude of the tSZ – CMB lensing cross-correlation is quite sensitive to the amplitude of fluctuations (σ_8) and the matter density (Ω_m). We show below that the normalized cross-correlation coefficient between the tSZ and CMB lensing fields is $\approx$ 30–40% at $\ell \lesssim 1000$ for our fiducial model, with a stronger correlation on smaller angular scales. This value is somewhat less than the normalized correlation between the cosmic infrared background (CIB) and the CMB lensing field, which reaches values as high as $\approx$ 80% [48], because the redshift kernels of the tSZ and CMB lensing fields are not as well-matched as those of the CIB and CMB lensing signals. Physically, this arises because the clusters responsible for the tSZ signal have not formed until recent epochs ($z \lesssim 1.5$), while the dusty star-forming galaxies responsible for most of the CIB emission formed much earlier ($z \gtrsim 2$–3), providing a stronger overlap with the CMB lensing redshift kernel.

The tSZ – CMB lensing cross-correlation is a bispectrum of the CMB temperature field, as the CMB lensing potential (or convergence) is reconstructed from quadratic combinations of the temperature fluctuations in multipole space [83], while the tSZ field is reconstructed from linear combinations of the temperature maps (see Section 3.4.2). To our knowledge, this bispectrum was first estimated in [18], where its signal-to-noise ratio (SNR) was forecasted for the (then-forthcoming) WMAP experiment, and compared to the signal from the integrated Sachs-Wolfe (ISW) – CMB lensing bispectrum. The authors estimated a SNR $\approx$ 3 for the total tSZ+ISW – CMB lensing bispectrum in WMAP, with most of the signal arising from the tSZ – CMB lensing term. These estimates were later updated using more detailed halo model calculations in [19], who considered contamination to the tSZ – CMB lensing bispectrum from reionization-induced bispectra. Shortly thereafter, [20] investigated the use of multifrequency subtraction techniques to reconstruct the tSZ signal using forecasted WMAP and Planck data, and estimated the SNR with which various tSZ statistics could be detected. Using a simplified model for the ICM gas physics (which they stated was likely only valid at best to the order-of-magnitude level), they estimated a SNR $\sim \mathcal{O}(10)$ for the tSZ – CMB lensing cross-spectrum in Planck data, depending on the maximum multipole used

in the analysis. Finally, [21] computed predictions for the cross-correlation between the tSZ signal and weak lensing maps constructed from forecasted Sloan Digital Sky Survey (SDSS) data, although no SNR estimates were calculated. Predictions for the cross-correlation of SDSS galaxies (not lensing) with the tSZ signal from WMAP were computed in [22], and a first detection of this signal was made in [23] at SNR ≈ 3.1. In the intervening decade since these early papers, our knowledge of the cosmological parameters, ICM gas physics, and contaminating signals from other sources (e.g., the CIB) has improved immensely, making the initial tSZ – CMB lensing cross-correlation SNR estimates somewhat obsolete. However, they are roughly accurate at the order-of-magnitude level, and it is indeed the case that the SNR of the tSZ – CMB lensing cross-spectrum is higher than that of the ISW – CMB lensing cross-spectrum: we detect the former at 6.2σ significance in this chapter, while only 2.5–2.6σ evidence has been found for the latter in Planck collaboration analyses [24, 25].

We measure the tSZ – CMB lensing cross-correlation by computing the cross-power spectrum of the publicly released Planck CMB lensing potential map [46] and a tSZ (or "Compton-y") map that we construct from a subset of the Planck channel maps. As shown for more general CMB lensing-induced bispectra in Appendix B of [14], the optimal estimator for the tSZ – CMB lensing temperature bispectrum factors into the individual steps of performing lens reconstruction and correlating with a y-map that utilizes the spectral signature of the tSZ effect. Thus, although we do not claim optimality in this analysis, our approach should be close to the optimal estimator (technically this statement holds only in the limit of weak non-Gaussianity; future high-SNR measurements of this cross-correlation may require an improved estimator). A possible goal for future analyses may be the simultaneous measurement of the tSZ – CMB lensing and CIB – CMB lensing bispectra, rather than measuring each individually while treating the other as a systematic (as done for the CIB in [48], and in this chapter for the tSZ). Also, it is worth noting that the tSZ – CMB lensing bispectrum (and CIB – CMB lensing as well) can contaminate measurements of primordial non-Gaussianity; these bispectra have thus far been neglected in such constraints

(e.g., [24]). However, the smooth shapes of these late-time bispectra are rather distinct from the acoustic oscillations in the primordial bispectrum arising from the transfer function, and thus the contamination to standard "shapes" considered in non-Gaussianity analyses may not be significant.

In this chapter, we demonstrate that the tSZ – CMB lensing cross-correlation signal is a unique probe of the ICM physics in high-redshift ($z \sim 1$), relatively low-mass ($M \sim 10^{13}$–$10^{14}\ M_{\odot}/h$) groups and clusters. In fact, this signal receives contributions from objects at even higher redshifts and lower masses than the tSZ auto-power spectrum, which is well-known for its dependence on the ICM physics in high-z, low-mass groups and clusters that have been unobserved thus far with direct X-ray or optical observations (e.g., [29, 79]). In a broader sense, this cross-correlation presents a new method with which to constrain the pressure–mass relation, which remains the limiting factor in cosmological constraints based on tSZ measurements [12, 44, 26, 27]. Although any individual tSZ statistic, such as the tSZ auto-power spectrum or the tSZ – CMB lensing cross-power spectrum, is subject to a complete degeneracy between the normalization of the pressure–mass relation and the amplitude of cosmic density fluctuations (i.e., σ_8), it is possible to break this degeneracy by combining multiple such statistics with different dependences on the ICM physics and the background cosmology. A simple version of this idea is applied to measurements of the tSZ auto-power spectrum and the tSZ skewness/bispectrum in Chapter 3 [60] to show that the physics of intracluster gas could be constrained in a manner nearly independent of the background cosmology. We demonstrate a similar technique in this chapter, which will soon become much more powerful with higher SNR detections of tSZ statistics.

The remainder of this chapter is organized as follows. In Section 3.3, we describe the theory underlying our halo model calculations of the tSZ – CMB lensing cross-power spectrum, including our modeling of the ICM physics, and compute the mass and redshift contributions to the signal. Section 3.4 presents our Compton-y map constructed from the Planck channel maps using internal linear combination (ILC) techniques. We also briefly discuss

the tSZ auto-power spectrum. In Section 3.5, we describe our measurement of the tSZ – CMB lensing cross-power spectrum, including a number of null tests and a small correction for leakage of CIB emission into the Compton-y map. In Section 3.6, we use our results to constrain the physics of the ICM and the cosmological parameters σ_8 and Ω_m. We also demonstrate that by combining multiple tSZ statistics — in this case, the tSZ – CMB lensing cross-power spectrum and the tSZ auto-power spectrum — it is possible to break the ICM–cosmology degeneracy and simultaneously constrain both. We discuss our results and conclude in Section 3.7.

We assume a flat ΛCDM cosmology throughout, with parameters set to their WMAP9+eCMB+BAO+H_0 maximum-likelihood values [1] unless otherwise specified. In particular, $\Omega_m = 0.282$ and $\sigma_8 = 0.817$ are our fiducial values for the matter density and the rms amplitude of linear density fluctuations on 8 Mpc/h scales at $z = 0$, respectively. All masses are quoted in units of $M_\odot/h$, where $h \equiv H_0/(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}})$ and H_0 is the Hubble parameter today. All distances and wavenumbers are in comoving units of Mpc/h.

3.3 Theory

3.3.1 Thermal SZ Effect

The tSZ effect is a frequency-dependent change in the observed CMB temperature due to the inverse Compton scattering of CMB photons off of hot electrons along the LOS, e.g., ionized gas in the ICM of a galaxy cluster. The temperature change ΔT at angular position $\vec{\theta}$ with respect to the center of a cluster of mass M at redshift z is given by [8]:

$$\begin{aligned} \frac{\Delta T(\vec{\theta}, M, z)}{T_{\mathrm{CMB}}} &= g_\nu y(\vec{\theta}, M, z) \\ &= g_\nu \frac{\sigma_T}{m_e c^2} \int_{\mathrm{LOS}} P_e\left(\sqrt{l^2 + d_A^2 |\vec{\theta}|^2}, M, z\right) dl\,, \end{aligned} \tag{3.1}$$

where $g_\nu = x\coth(x/2) - 4$ is the tSZ spectral function with $x \equiv h\nu/k_B T_{\rm CMB}$, $y(\vec{\theta}, M, z)$ is the Compton-y parameter, σ_T is the Thomson scattering cross-section, m_e is the electron mass, and $P_e(\vec{r}, M, z)$ is the ICM electron pressure at (three-dimensional) location $\vec{r}$ with respect to the cluster center. We have neglected relativistic corrections to the tSZ spectral function in Eq. (3.1) (e.g., [28]), as these effects are only non-negligible for the most massive clusters in the universe ($\gtrsim 10^{15}\ M_\odot/h$). The tSZ – CMB lensing cross-spectrum is dominated by contributions from clusters well below this mass scale (see Figs. 3.3 and 3.4), and thus we do not expect relativistic corrections to be important in our analysis. Our measurements and theoretical calculations will be given in terms of the Compton-y parameter, which is the frequency-independent quantity characterizing the tSZ signal (in the absence of relativistic corrections).

Our theoretical calculations assume a spherically symmetric pressure profile, i.e., $P_e(\vec{r}, M, z) = P_e(r, M, z)$ in Eq. (3.1). Note that the integral in Eq. (3.1) is computed along the LOS such that $r^2 = l^2 + d_A(z)^2\theta^2$, where $d_A(z)$ is the angular diameter distance to redshift z and $\theta \equiv |\vec{\theta}|$ is the angular distance between $\vec{\theta}$ and the cluster center in the plane of the sky. A spherically symmetric pressure profile implies that the Compton-y profile is azimuthally symmetric in the plane of the sky, i.e., $\Delta T(\vec{\theta}, M, z) = \Delta T(\theta, M, z)$. Finally, note that the electron pressure $P_e(\vec{r}, M, z)$ is related to the thermal gas pressure via $P_{th} = P_e(5X_H + 3)/2(X_H + 1) = 1.932 P_e$, where $X_H = 0.76$ is the primordial hydrogen mass fraction.

We define the mass M in Eq. (3.1) to be the virial mass, which is the mass enclosed within a radius $r_{\rm vir}$ [72]:

$$r_{\rm vir} = \left(\frac{3M}{4\pi\Delta_{cr}(z)\rho_{cr}(z)}\right)^{1/3}, \tag{3.2}$$

where $\Delta_{cr}(z) = 18\pi^2 + 82(\Omega(z)-1) - 39(\Omega(z)-1)^2$ and $\Omega(z) = \Omega_m(1+z)^3/(\Omega_m(1+z)^3+\Omega_\Lambda)$. For some calculations, we require the spherical overdensity (SO) mass contained within some

radius, defined as follows: $M_{\delta c}$ ($M_{\delta d}$) is the mass enclosed within a sphere of radius $r_{\delta c}$ ($r_{\delta d}$) such that the enclosed density is δ times the critical (mean matter) density at redshift z. For clarity, c subscripts refer to masses defined with respect to the critical density at redshift z, $\rho_{cr}(z) = 3H^2(z)/8\pi G$ with $H(z)$ the Hubble parameter at redshift z, whereas d subscripts refer to masses defined with respect to the mean matter density at redshift z, $\bar{\rho}_m(z) \equiv \bar{\rho}_m$ (this quantity is constant in comoving units). We convert between the virial mass M and other SO masses (e.g., M_{200c} or M_{200d}) using the Navarro-Frenk-White (NFW) density profile [70] and the concentration-mass relation from [71].

Our fiducial model for the ICM pressure profile is the parametrized fit to the "AGN feedback" simulations given in [29], which fully describes the pressure profile as a function of cluster mass (M_{200c}) and redshift. The simulations on which the "AGN feedback" model is based include virial shock heating, radiative cooling, and sub-grid prescriptions for star formation and feedback from supernovae and active galactic nuclei (AGN) [30]. Also, the smoothed particle hydrodynamics method used for the simulations captures the effects of bulk motions and turbulence, which provide non-thermal pressure support in the cluster outskirts. While this non-thermal pressure support suppresses the tSZ signal in the outer regions of groups and clusters (by allowing the thermal pressure to decrease while still maintaining balance against the dark matter-sourced gravitational potential), the AGN feedback heats and expels gas from the inner regions of the cluster, leading to flatter pressure profiles in the center and higher temperatures and pressures in the outskirts compared to calculations with cooling only (no AGN feedback) [30]. However, the feedback can also lower the cluster gas fraction by blowing gas out of the cluster (see, e.g., [31, 32] for other cosmological simulations incorporating AGN feedback in the ICM). All of these effects are accounted for in the phenomenological "GNFW" fit provided in [29]. This profile has been found to agree with a number of recent X-ray and tSZ studies [62, 65, 66, 58].

We also consider the parametrized fit to the "adiabatic" simulations given in [29]. These simulations include only heating from virial shocks; no cooling or feedback prescriptions of

any kind are included. Thus, this fit predicts much more tSZ signal from a cluster of a given mass and redshift than our fiducial "AGN feedback" model does. It is already in tension with many observations (e.g., [60, 67]), but we include it here as an extreme example of the ICM physics possibilities.

For comparison purposes, we also show results below computed using the "universal pressure profile" (UPP) of Arnaud et al. [62], which consists of a similar "GNFW"-type fitting function that specifies the pressure profile as a function of mass (M_{500c}) and redshift. Within r_{500c}, this profile is derived from observations of local ($z < 0.2$), massive ($10^{14}\,M_{\odot} < M_{500c} < 10^{15}\,M_{\odot}$) clusters in the X-ray band with XMM-Newton. Beyond r_{500c}, this profile is a fit to various numerical simulations, which include radiative cooling, star formation, and supernova feedback, but do not include AGN feedback [33, 34, 35]. The overall normalization of this profile is subject to uncertainty due to the so-called "hydrostatic mass bias", $(1-b)$, because the cluster masses used in the analysis are derived from X-ray observations under the assumption of hydrostatic equilibrium (HSE), which is not expected to be exactly valid in actual clusters. This bias is typically expressed via $M_{500c}^{\rm HSE} = (1-b)M_{500c}$. Typical values are expected to be $(1-b) \approx 0.8$–0.9 [62, 68], but recent Planck results require much more extreme values ($(1-b) \approx 0.5$–0.6) in order to reconcile observed tSZ cluster counts with predictions based on cosmological parameters from the Planck observations of the primordial CMB [12]. We note that $(1-b)$ should not be thought of as a single number, valid for all clusters — rather, it is expected to be a function of cluster mass and redshift, and likely exhibits scatter about any mean relation as well. Statements in the literature about this bias thus only reflect comparisons to the massive, low-redshift population of clusters from which the UPP of [62] was derived. As a rough guide, our fiducial "AGN feedback" model corresponds approximately to $(1-b) \approx 15\%$ for this population of clusters, though this varies with radius (see Fig. 2 of [30]). Beyond r_{500c}, where the UPP of [62] is not based on X-ray data, the "AGN feedback" fit predicts much larger pressures than the UPP, likely due to the newer simulations' inclusion of AGN feedback and the earlier simulations' neglect of

this effect. Finally, we stress that although we will show the results of calculations using the UPP of [62], it is derived from observations of clusters at much lower redshifts and higher masses than those that dominate the tSZ – CMB lensing cross-power spectrum, and hence these calculations are an extrapolation of this model. Nonetheless, our results can constrain the mean value of $(1-b)$ averaged over the cluster population, and can further be combined with the results of [12] or [44] to obtain even tighter constraints.

Finally, we also show results below computed directly from a cosmological simulation described in Sehgal et al. [63][1]. The simulation is a large dark matter-only N-body simulation that is post-processed to include gas and galaxies according to various phenomenological prescriptions. We use the full-sky CMB lensing and tSZ maps derived from these simulations to compute the tSZ – CMB lensing cross-power spectrum. The tSZ auto-power spectrum of this simulation has been studied previously (e.g., [63, 64]) and lies higher than the results from ACT and SPT at $\ell = 3000$, likely due to missing non-thermal physics in the baryonic post-processing applied to the N-body simulation. We find a similar result for the tSZ – CMB lensing cross-power spectrum in Section 3.6. Note that the cosmological parameters used in the simulation ($\sigma_8 = 0.80$ and $\Omega_m = 0.264$) are consistent with the WMAP5 cosmology, but in order to facilitate comparison with our halo model calculations, we rescale the tSZ – CMB lensing cross-power spectrum derived from the simulation to our fiducial WMAP9 values ($\sigma_8 = 0.817$ and $\Omega_m = 0.282$) using the dependences on these parameters computed in Section 3.3.3. We emphasize that this rescaling is not computed an overall rescaling of the entire cross-spectrum, but rather as an ℓ-dependent calculation. As shown later, we find a general agreement in shape between the tSZ – CMB lensing cross-power spectrum derived from our halo model calculations and from this simulation (see Fig. 3.15); the amplitudes differ as a result of the differing gas physics treatments.

[1] `http://lambda.gsfc.nasa.gov/toolbox/tb_sim_ov.cfm`

3.3.2 CMB Lensing

Gravitational lensing of the CMB causes a re-mapping of the unlensed temperature field (e.g., [69]):

$$\begin{aligned} T(\hat{n}) &= T^{\rm un}(\hat{n} + \nabla\phi(\hat{n})) \\ &= T^{\rm un}(\hat{n}) + \nabla^i\phi(\hat{n})\nabla_i T^{\rm un}(\hat{n}) + \mathcal{O}(\phi^2)\,, \end{aligned} \tag{3.3}$$

where $T^{\rm un}$ is the unlensed primordial temperature and $\phi(\hat{n})$ is the CMB lensing potential:

$$\phi(\hat{n}) = -2\int_0^{\chi_*} d\chi \frac{\chi_* - \chi}{\chi_*\chi}\Psi(\chi\hat{n}, \chi)\,, \tag{3.4}$$

where we have specialized to the case of a flat universe. In this equation, $\chi(z)$ is the comoving distance to redshift z, χ_* is the comoving distance to the surface of last scattering at $z_* \approx 1100$, and $\Psi(\chi\hat{n}, \chi)$ is the gravitational potential. Note that the lensing convergence is given by $\kappa(\hat{n}) = -\nabla^2\phi(\hat{n})/2$ (where ∇^2 is now the two-dimensional Laplacian on the sky), or $\kappa_\ell = \ell(\ell+1)\phi_\ell/2$ in multipole space.

Analogous to the Compton-y profile for a halo of mass M at redshift z defined in Eq. (3.1), we can define a CMB lensing convergence profile, $\kappa(\vec{\theta}, M, z)$:

$$\kappa(\vec{\theta}, M, z) = \Sigma_{\rm crit}^{-1}(z)\int_{\rm LOS} \rho\left(\sqrt{l^2 + d_A^2|\vec{\theta}|^2}, M, z\right) dl\,, \tag{3.5}$$

where $\rho(\vec{r}, M, z)$ is the halo density profile and $\Sigma_{\rm crit}(z)$ is the critical surface density for lensing at redshift z assuming a source plane at $z_* \approx 1100$:

$$\Sigma_{\rm crit}^{-1}(z) = \frac{4\pi G\chi(z)\left(\chi_* - \chi(z)\right)}{c^2\chi_*(1+z)}\,. \tag{3.6}$$

We adopt the NFW density profile [70] (which is spherically symmetric, i.e., $\rho(\vec{r}, M, z) = \rho(r, M, z)$) and the concentration-mass relation from [71] when calculating Eq. (3.5). Also,

note that Eq. (3.5) describes the lensing convergence profile, but we will work in terms of the lensing potential below, as this is the quantity directly measured in the publicly released Planck lensing map. The convergence and potential are trivially related in multipole space, as mentioned above: $\phi_\ell = 2\kappa_\ell/(\ell(\ell+1))$.

3.3.3 Power Spectra

Given models for the Compton-y and lensing potential signals from each halo of mass M and redshift z, we compute the tSZ – CMB lensing cross-spectrum in the halo model framework (see [36] for a review). The following expressions are directly analogous to those derived in Chapters 1 and 2; the only change is that now one factor of Compton-y will be replaced by the lensing potential ϕ. We work in the flat-sky and Limber approximations in this chapter, since we only consider multipoles $\ell \gtrsim 100$ in the cross-spectrum analysis (complete full-sky derivations can be found in Chapter 1).

The tSZ – CMB lensing cross-power spectrum, $C_\ell^{y\phi}$, is given by the sum of the one-halo and two-halo terms:

$$C_\ell^{y\phi} = C_\ell^{y\phi,1h} + C_\ell^{y\phi,2h} \,. \tag{3.7}$$

The one-halo term arises from correlations between the Compton-y and lensing potential ϕ profiles of the same object. In the flat-sky limit, the one-halo term is given by (e.g., [37, 39]):

$$C_\ell^{y\phi,1h} = \int dz \frac{d^2V}{dz d\Omega} \int dM \frac{dn(M,z)}{dM} \tilde{y}_\ell(M,z) \tilde{\phi}_\ell(M,z) \,, \tag{3.8}$$

where

$$\tilde{y}_\ell(M,z) = \frac{\sigma_T}{m_e c^2} \frac{4\pi r_{s,y}}{\ell_{s,y}^2} \int dx_y \, x_y^2 \frac{\sin((\ell+1/2)x_y/\ell_{s,y})}{(\ell+1/2)x_y/\ell_{s,y}} P_e(x_y r_{s,y}, M, z) \tag{3.9}$$

and

$$\begin{aligned}\tilde{\phi}_\ell(M,z) &= \frac{2}{\ell(\ell+1)}\tilde{\kappa}_\ell(M,z) \\ &= \frac{2}{\ell(\ell+1)}\frac{4\pi r_{s,\phi}}{\ell_{s,\phi}^2}\int dx_\phi\, x_\phi^2 \frac{\sin((\ell+1/2)x_\phi/\ell_{s,\phi})}{(\ell+1/2)x_\phi/\ell_{s,\phi}}\frac{\rho(x_\phi r_{s,\phi},M,z)}{\Sigma_{\rm crit}(z)}\,. \quad (3.10)\end{aligned}$$

In Eq. (3.8), $d^2V/dzd\Omega$ is the comoving volume element per steradian and $dn(M,z)/dM$ is the halo mass function, i.e., the comoving number density of halos per unit mass as a function of redshift. All calculations in this chapter use the mass function from [73], but with the updated parameters provided in Eqs. (8)–(12) of [74], which explicitly enforce the physical constraint that the mean bias of all matter in the universe at a fixed redshift must equal unity. (This constraint was not enforced in the original fits in [73], and it is relevant when calculating quantities that receive contributions from very low-mass halos, such as the CMB lensing power spectrum.) We work with the mass function fit given for the SO mass M_{200d}. Further details on the mass function calculations can be found in Chapter 2.

In Eq. (3.9), $r_{s,y}$ is a characteristic scale radius (not the NFW scale radius) of the ICM pressure profile, $\ell_{s,y} = a(z)\chi(z)/r_{s,y} = d_A(z)/r_{s,y}$ is the associated multipole moment, and $x_y \equiv r/r_{s,y}$ is a dimensionless radial integration variable. For the fiducial "AGN feedback" pressure profile from [29] used in our calculations, the natural scale radius is r_{200c}. For the UPP of [62], the natural scale radius is r_{500c}.

Similarly, in Eq. (3.10), $r_{s,\phi}$ is a characteristic scale radius of the halo density profile, which in this case is indeed the NFW scale radius. Likewise, $\ell_{s,\phi} = a(z)\chi(z)/r_{s,\phi} = d_A(z)/r_{s,\phi}$ is the associated multipole moment and $x_\phi \equiv r/r_{s,\phi}$ is a dimensionless radial integration variable.

The two-halo term arises from correlations between the Compton-y and lensing potential ϕ profiles of two separate objects. In the Limber approximation [38], the two-halo term is

given by (e.g., [37, 75]):

$$C_\ell^{y\phi,2h} = \int dz \frac{d^2V}{dz d\Omega} \left[\int dM_1 \frac{dn(M_1,z)}{dM_1} b(M_1,z) \tilde{y}_\ell(M_1,z) \right] \times \left[\int dM_2 \frac{dn(M_2,z)}{dM_2} b(M_2,z) \tilde{\phi}_\ell(M_2,z) \right] P_{\rm lin}\left(\frac{\ell+1/2}{\chi(z)}, z \right) , \quad (3.11)$$

where $b(M,z)$ is the linear halo bias and $P_{\rm lin}(k,z)$ is the linear theory matter power spectrum. We use the fitting function for the halo bias from Table 2 of [74], which matches our mass function to ensure a mean bias of unity at a given redshift, as mentioned above. We compute the matter power spectrum for a given set of cosmological parameters using `CAMB`[2].

In this work we model the two-halo term using the full expression in Eq. (3.11). However, in the large-scale limit ($\ell \to 0$) this equation can be simplified to match the standard expression used in analyses of cross-correlation between CMB lensing and various mass tracers (e.g., Eq. (48) of [46] or Eq. (6) of [76]). The simplification rests on the fact that on very large scales ($k \to 0$) the 3D Fourier transform of the halo density profile is simply the total mass M of the halo, i.e., $\tilde{\rho}(k,M,z) \to M$ as $k \to 0$. Using this fact and the previously noted result that the mean bias of all matter at a fixed redshift is unity, i.e., $\int dM \frac{dn}{dM} b(M,z) \frac{M}{\bar{\rho}_m} = 1$, one obtains

$$\begin{aligned} C_{\ell\to 0}^{y\phi,2h} &\approx \frac{2}{\ell(\ell+1)} \int d\chi W^\kappa(\chi) P_{\rm lin}\left(\frac{\ell+1/2}{\chi(z)}, z \right) \int dM \frac{dn(M,z)}{dM} b(M,z) \tilde{y}_\ell(M,z) \\ &\approx \frac{2}{\ell(\ell+1)} \int d\chi W^\kappa(\chi) P_{\rm lin}\left(\frac{\ell+1/2}{\chi(z)}, z \right) \times \\ &\quad \int dM \frac{dn(M,z)}{dM} b(M,z) \tilde{y}_0(M,z) , \end{aligned} \quad (3.12)$$

where $W^\kappa(\chi) = \frac{3\Omega_m H_0^2 (1+z)}{2c^2} \chi \left(\frac{\chi_* - \chi}{\chi_*} \right)$ is the lensing convergence projection kernel (assuming a flat universe) and we have changed integration variables from z to χ to match the standard CMB lensing literature (e.g., Eq. (48) of [46]). In obtaining the second line of Eq. (3.12), we have noted that the $\ell \to 0$ limit corresponds to angular scales much larger than the

[2] http://www.camb.info

scale of individual clusters, and thus one can further approximate $\tilde{y}_\ell(M,z) \approx \tilde{y}_0(M,z)$, which is proportional to the mean Compton-y signal from that cluster. In this limit, one can thus constrain an effective "Compton-y" bias corresponding to the product of the halo bias and Compton-y signal. In contrast, when cross-correlating a redshift catalog of objects (e.g., quasars or galaxy clusters) with the CMB lensing field, one directly constrains the typical halo bias (and hence mass) of those objects. Our measurements are all at multipoles $\ell > 100$, and thus we work with the full expression for the two-halo term in Eq. (3.11), in which it is unfortunately not straightforward to disentangle the influence of the bias and the ICM physics. However, we can still probe the ICM by varying the gas physics model while holding all other ingredients in the calculation fixed; since the gas physics is likely the most uncertain ingredient (especially at these redshifts and masses — see Figs. 3.3 and 3.4), this is a reasonable approach.

The fiducial integration limits in our calculations are $0.005 < z < 10$ for all redshift integrals and $10^5 M_\odot/h < M < 5 \times 10^{15} M_\odot/h$ for all mass integrals. We note that this involves an extrapolation of the halo mass and bias functions from [73, 74] to mass and redshift ranges in which they were not explicitly measured in the simulations. However, the bulk of the tSZ – CMB lensing cross-spectrum comes from mass and redshift ranges in which the fitting functions are indeed measured (see Figs. 3.3 and 3.4), so this extrapolation should not have a huge effect; it is primarily needed in order to compute the CMB lensing auto-spectrum for comparison calculations, as this signal does indeed receive contributions from very high redshifts and low masses. In addition, we must define a boundary at which to cut off the integrals over the pressure and density profiles in Eqs. (3.9) and (3.10). We determine this boundary by requiring our halo model calculation of the CMB lensing auto-spectrum to agree with the standard method in which one simply integrates over the linear theory matter power spectrum:

$$C_\ell^{\phi\phi,2h} \approx \frac{4}{\ell^2(\ell+1)^2} \int d\chi \left(\frac{W^\kappa(\chi)}{\chi}\right)^2 P_{\rm lin}\left(\frac{\ell+1/2}{\chi(z)}, z\right) , \tag{3.13}$$

which is quite accurate for $\ell \lesssim 1000$ (see Fig. 2 of [69]). It is straightforward to derive Eq. (3.13) from the CMB lensing analogue of Eq. (3.11) by considering the $\ell \to 0$ limit and using the procedure described in the previous paragraph. However, since we are using a variety of fitting functions from the literature, this formal derivation does not hold precisely when applied to the numerical calculations. The primary obstacle is the logarithmic divergence of the enclosed mass in the NFW profile as $r \to \infty$ [77]. By testing various values of the outer cut-off in Eq. (3.10), we find that $r_{\rm out} = 1.5 r_{\rm vir}$ leads to better than 5% agreement between our halo model calculation of the CMB lensing potential power spectrum and the linear theory calculation in Eq. (3.13) for all multipoles $\ell \lesssim 900$, above which nonlinear corrections to Eq. (3.13) become important. Thus, we adopt $r_{\rm out} = 1.5 r_{\rm vir}$ when computing Eqs. (3.9) and (3.10). Future work requiring more accurate predictions will necessitate detailed simulations and better understanding of the halo model approximations in this context.

We use the model described thus far to compute the tSZ – CMB lensing cross-power spectrum for our fiducial cosmology. The results are shown in Fig. 3.1. As expected, the two-halo term dominates at low ℓ, while the one-halo term dominates at high ℓ; the two terms are roughly equal at $\ell \approx 450$. As one might expect, this behavior lies between the extremes of the CMB lensing power spectrum, which is dominated by the two-halo term for $\ell \lesssim 1000$ [69], and the tSZ power spectrum, which is dominated by the one-halo term for $\ell \gtrsim 10$ [37]. Fig. 3.1 also shows the cross-power spectrum computed for $\pm 5\%$ variations in σ_8 and Ω_m around our fiducial cosmology. Using these variations, we compute approximate power-law scalings of the cross-power spectrum with respect to these parameters. The scalings vary as a function of ℓ — for example, the scaling with σ_8 is strongest near $\ell \sim 1000$–2000, where the one-halo term dominates but the power is not yet being sourced by the variations within the pressure profile itself, which occurs at very high ℓ. The scaling with Ω_m is strongest near

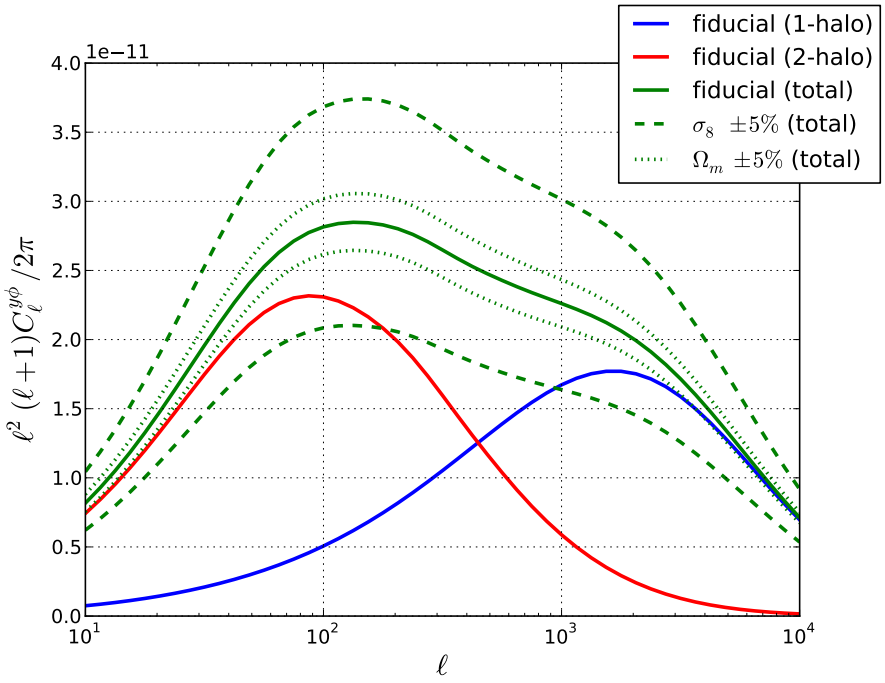

Figure 3.1: The tSZ – CMB lensing cross-power spectrum computed for our fiducial model (WMAP9 cosmological parameters and the "AGN feedback" ICM pressure profile fit from [29]). The total signal is the green solid curve, while the one-halo and two-halo contributions are the blue and red solid curves, respectively. The plot also shows the total signal computed for variations around our fiducial model: ±5% variations in σ_8 are the dashed green curves, while ±5% variations in Ω_m are the dotted green curves. In both cases, an increase (decrease) in the parameter's value yield an increase (decrease) in the amplitude of the cross-spectrum.

$\ell \sim 800$–1500. Representative scalings are:

$$\begin{aligned} C^{y\phi}_{\ell=100} &\propto \left(\frac{\sigma_8}{0.817}\right)^{5.7}\left(\frac{\Omega_m}{0.282}\right)^{1.4} \\ C^{y\phi}_{\ell=1000} &\propto \left(\frac{\sigma_8}{0.817}\right)^{6.1}\left(\frac{\Omega_m}{0.282}\right)^{1.5} . \end{aligned} \tag{3.14}$$

Over the ℓ range where we measure $C_\ell^{y\phi}$ (see Section 3.5) — $100 < \ell < 1600$ — the average values of the scalings are 6.0 and 1.5 for σ_8 and Ω_m, respectively. These scalings provide the theoretical degeneracy between these parameters, which will be useful in Section 3.6. Note that the dependence of the tSZ – CMB lensing cross-power spectrum on these parameters is not as steep as the dependence of the tSZ auto-power spectrum [39, 37] or other higher-order

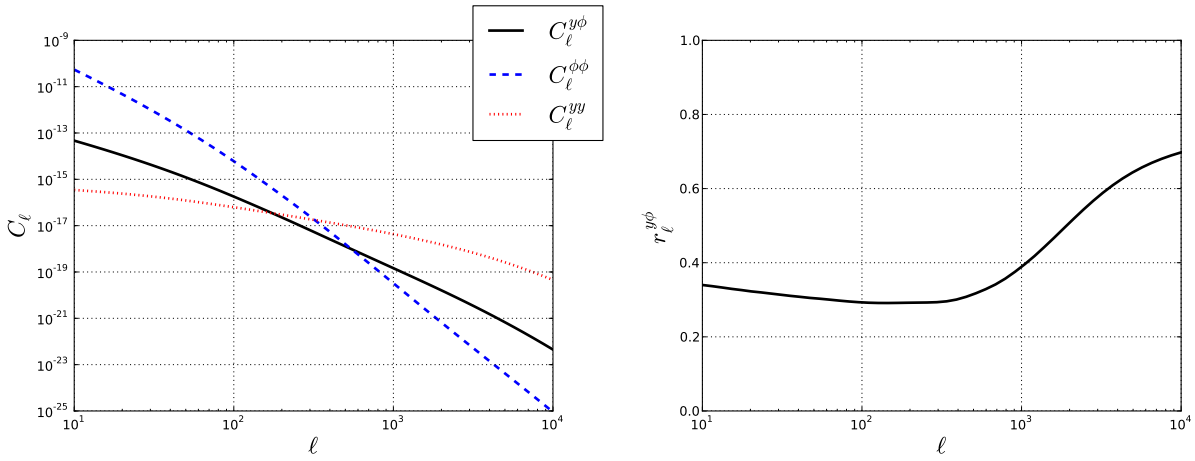

Figure 3.2: The left panel shows the power spectra of the tSZ and CMB lensing potential fields for our fiducial model, as labeled in the figure. The right panel shows the normalized cross-correlation coefficient of the tSZ and CMB lensing potential signals as a function of multipole (see Eq. (3.15)). The typical normalized correlation is 30–40% over the ℓ range where we measure $C_\ell^{y\phi}$ ($100 < \ell < 1600$). The strength of the correlation increases at smaller angular scales, where both signals are dominated by the one-halo term.

tSZ statistics, such as the skewness [59, 60] or bispectrum [78]. The dependence is stronger for the tSZ auto-statistics because their signals arise from more massive, rare clusters that lie further in the exponential tail of the mass function than those that source the tSZ – CMB lensing cross-spectrum (see below). The cross-spectrum is, however, uncharted territory for measurements of the ICM pressure profile, as its signal comes from groups and clusters at much higher redshifts and lower masses than those that source the tSZ auto-statistics.

Before assessing the origin of the cross-spectrum signal in detail, it is useful to examine the predicted strength of the correlation between the tSZ and CMB lensing potential fields. A standard method to assess the level of correlation is through the normalized cross-correlation coefficient, $r_\ell^{y\phi}$:

$$r_\ell^{y\phi} = \frac{C_\ell^{y\phi}}{\sqrt{C_\ell^{yy} C_\ell^{\phi\phi}}}\,, \tag{3.15}$$

where C_ℓ^{yy} and $C_\ell^{\phi\phi}$ are the tSZ and CMB lensing auto-power spectra, respectively. Fig. 3.2 shows the result of this calculation, as well as a comparison of the auto-spectra to the cross-

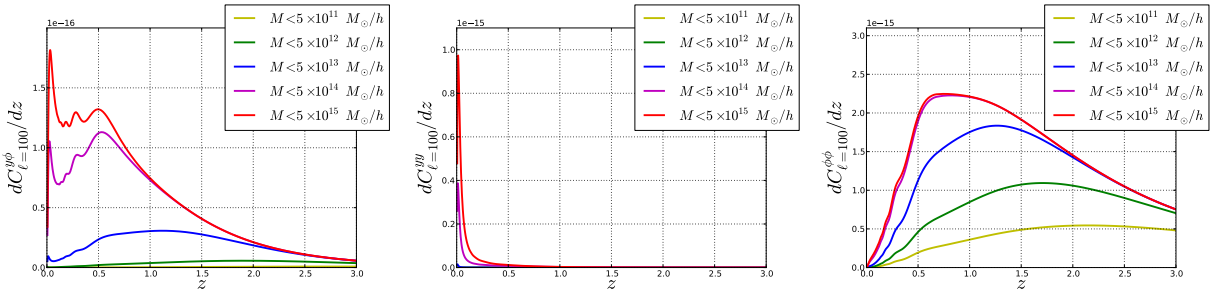

Figure 3.3: Mass and redshift contributions to the tSZ – CMB lensing cross-spectrum (left panel), tSZ auto-spectrum (middle panel), and CMB lensing potential auto-spectrum (right panel), computed at $\ell = 100$, where the two-halo term dominates the cross-spectrum. Each curve includes contributions from progressively higher mass scales, as labelled in the figures, while the vertical axis encodes the differential contribution from each redshift. At this multipole, the tSZ auto-spectrum is heavily dominated by massive, low-redshift halos, while the CMB lensing auto-spectrum receives contributions from a much wider range of halo masses and redshifts. The cross-spectrum, as expected, lies between these extremes. Note that the sharp features in the cross-spectrum curves are an artifact of combining the various fitting functions used in our calculation. See the text for further discussion.

spectrum. The behavior of the power spectra is consistent with expectations: $C_\ell^{\phi\phi}$ is much larger at low ℓ than high ℓ; C_ℓ^{yy} is close to Poisson for $\ell < 1000$, above which intracluster structure contributes power; and $C_\ell^{y\phi}$ lies between these two extremes. The normalized cross-correlation $r_\ell^{y\phi} \approx 30$–$40\%$ over the ℓ range where we measure the tSZ – CMB lensing cross-spectrum ($100 < \ell < 1600$); the average value of the normalized correlation within this range is 33.2%. At smaller angular scales, the correlation reaches values as high as 60–70%. This increase occurs because both the tSZ and CMB lensing signals are dominated by the one-halo term from individual objects at these scales, and the gas traces the dark matter quite effectively within halos. Overall, the tSZ signal is a reasonably strong tracer of the large-scale matter distribution, although not nearly as strong as the CIB [48], at least on large scales. The tSZ – CMB lensing cross-correlation is, however, interesting as a probe of a previously unstudied population of high-redshift, low-mass groups and clusters.

We plot the redshift and mass kernels for the tSZ – CMB lensing cross-spectrum and each auto-spectrum at $\ell = 100$ and 1000 in Figs. 3.3 and 3.4, respectively. These multipoles probe different regimes of the cross-spectrum: at $\ell = 100$, the cross-spectrum and CMB

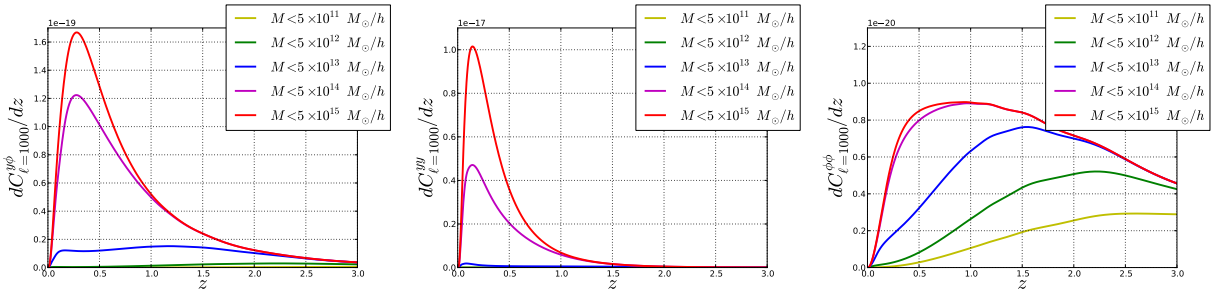

Figure 3.4: Identical to Fig. 3.3, but computed at $\ell = 1000$, where the one-halo term dominates the cross-spectrum. The tSZ auto-spectrum contributions here extend to lower masses and higher redshifts than at $\ell = 100$, leading to a stronger expected cross-correlation with the CMB lensing signal. See the text for further discussion.

lensing auto-spectrum are dominated by the two-halo term, while the tSZ auto-spectrum is dominated by the one-halo term; at $\ell = 1000$, the cross-spectrum and tSZ auto-spectrum are dominated by the one-halo term, while the CMB lensing auto-spectrum receives significant contributions from both terms. It is clear that at $\ell = 100$ the mass and redshift kernels of the tSZ and CMB lensing auto-spectra are not as well-matched as they are at $\ell = 1000$, which explains why $r_\ell^{y\phi}$ is larger at $\ell = 1000$ than at $\ell = 100$.

The most striking feature of Fig. 3.3 is the difference in the mass and redshift kernels for the tSZ auto-spectrum and the tSZ – CMB lensing cross-spectrum. At these low multipoles, the auto-spectrum is dominated by massive nearby clusters, while the cross-spectrum receives significant contributions out to very high redshift ($z \lesssim 2$) and low masses ($5 \times 10^{12}\, M_\odot/h$). This arises from the fact that the CMB lensing redshift and mass kernels upweight these scales in the cross-spectrum. Similar behavior is also seen at $\ell = 1000$ in Fig. 3.4, though the difference between the tSZ auto-spectrum and tSZ – CMB lensing cross-spectrum is somewhat less dramatic in this case.

These results are further illustrated in Fig. 3.5, which shows the cumulative contribution to each power spectrum as a function of mass (integrated over all redshifts) at $\ell = 100$ and 1000. The primary takeaway is that the tSZ – CMB lensing cross-spectrum is indeed sourced by much less massive halos than those that source the tSZ auto-spectrum — the difference is especially dramatic at $\ell = 100$, but also seen at $\ell = 1000$. At $\ell = 100$ (1000),

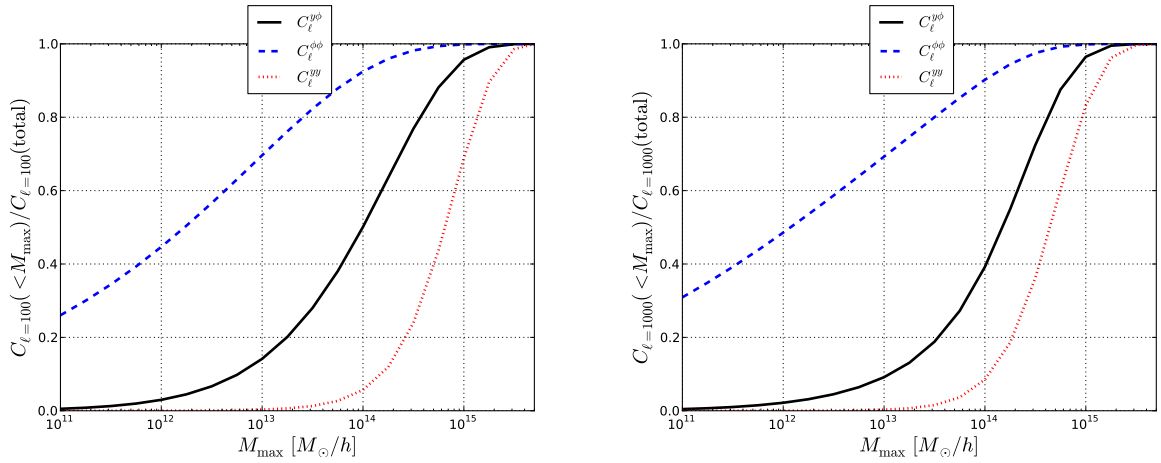

Figure 3.5: Cumulative contributions to the tSZ – CMB lensing cross-spectrum (black), CMB lensing auto-spectrum (blue), and tSZ auto-spectrum (red) as a function of halo mass, at $\ell = 100$ (left panel) and $\ell = 1000$ (right panel). These results are integrated over all redshifts. Note that the cross-spectrum receives significant contributions from much lower mass scales than the tSZ auto-spectrum. See the text for further discussion.

50% (40%) of the cross-spectrum signal comes from masses below $10^{14}\,M_\odot/h$, while 6% (9%) of the tSZ auto-spectrum signal comes from these masses[3]. Our calculations for the tSZ auto-spectrum contributions agree with earlier results [29, 79]. Overall, the tSZ – CMB lensing cross-spectrum provides a new probe of the ICM at mass and redshift scales that are presently unexplored.

As a final point regarding our theoretical calculations, we note that the cosmological parameter analysis presented in Section 3.6 accounts not only for statistical errors in our measurements, but also sample variance arising from the angular trispectrum of the tSZ – CMB lensing signals. Considering only the one-halo term in the flat-sky limit (which dominates the trispectrum at the scales we consider [40]), the trispectrum contribution to the tSZ – CMB lensing cross-power spectrum covariance matrix is [39, 37]

$$T_{\ell\ell'}^{y\phi} = \int dz \frac{d^2V}{dz d\Omega} \int dM \frac{dn(M,z)}{dM} \left|\tilde{y}_\ell(M,z)\right|^2 \left|\tilde{\phi}_{\ell'}(M,z)\right|^2 . \tag{3.16}$$

[3]Note that masking massive, nearby clusters could change the tSZ auto-spectrum contributions significantly [37, 80], but we will not consider this possibility here.

We also compute the trispectrum contribution to the tSZ auto-power spectrum when including it in our constraints, using an analogous expression to Eq. (3.16) (note that the tSZ trispectrum is not mentioned in [44]). Due to computational constraints, we only compute the diagonal elements of $T^{y\phi}_{\ell\ell'}$ and $T^{yy}_{\ell\ell'}$, as the trispectra must be re-calculated for each variation in the cosmological parameters or ICM physics model. This choice also allows us to work in the limit in which the multipole bins in our measurement are uncorrelated. Based on the results of Chapter 2, the neglect of the off-diagonal elements should only have an effect for very large angular scales ($\ell \lesssim 200$), and is likely only important for the tSZ auto-spectrum (not the cross-spectrum). Future high-precision measurements of the tSZ auto-spectrum (and possibly the tSZ – CMB lensing cross-spectrum) for which the sample variance is comparable to the statistical errors will require more careful treatment and more efficient computational implementation of the trispectrum.

3.4 Thermal SZ Reconstruction

3.4.1 Data and Cuts

Our thermal SZ reconstruction is based on the nominal mission maps from the first 15.5 months of operation of the Planck satellite [41]. In particular, our ILC pipeline uses the data from the 100, 143, 217, 353, and 545 GHz channels of the Planck High-Frequency Instrument (HFI) [42]. Note that these frequencies span the zero point of the tSZ spectral function at 217 GHz (the spectral function is given after Eq. (3.1)). We use the HFI 857 GHz map as an external tracer of dust emission, both from the Galaxy and from the CIB. The FWHM of the beams in the HFI channels ranges from 9.66 arcmin at 100 GHz to 4.63 arcmin at 857 GHz. We do not use the data from the Low-Frequency Instrument (LFI) channels due to their significantly larger beams [43]. We work with the Planck maps in `HEALPix`[4] format at the provided resolution of $N_{\rm side} = 2048$, for which the pixels have a typical width of 1.7

[4] `http://www.healpix.jpl.nasa.gov`

arcmin. At all HFI frequencies, we use the Zodi-corrected maps provided by the Planck collaboration. Also, when needed for the 545 and 857 GHz maps, we convert from MJy/sr to K_{CMB} using the values provided in the Planck explanatory supplement.[5]

Our analysis pipeline is designed to follow the approach of [44] to a large degree. As in [44], we approximate the Planck beams as circular Gaussians, with FWHM values given in Table 1 of [44]. The first step of our analysis pipeline is to convolve all of the HFI maps to a common resolution of 10 arcmin, again following [44]. This value is dictated by the angular resolution of the 100 GHz map, and is necessary in order to apply the ILC algorithm described in the following section.

A key difference between our analysis and [44] is that our ILC algorithm relies on the use of cross-correlations between independent "single-survey" maps, which prevents any auto-correlations due to noise in the input channel maps leaking into the output Compton-y map — effectively, the ILC minimizes "non-tSZ" signal in the reconstructed y-map, rather than "non-tSZ + instrumental noise". Thus, instead of using the full Planck HFI nominal mission maps, we work with the "survey 1" (S1) and "survey 2" (S2) maps; we also include the associated masks for the S1 and S2 maps in our analysis below, as neither survey completely covers the full sky. The only exception to this choice is our use of the full nominal mission 857 GHz map, which, as mentioned earlier, is used only as an external dust tracer rather than as a component in the ILC pipeline.

We use the 857 GHz map to construct a mask that excludes regions of the sky with the brightest dust emission. This mask is primarily meant to remove Galactic dust emission, but also removes a small fraction of the CIB emission. We construct the mask by simply thresholding the 857 GHz map until the desired sky fraction remains. Our fiducial results use a mask in which 70% of the sky is removed, i.e., $f_{\rm sky} = 0.30$. We also consider $f_{\rm sky} = 0.20$ and 0.40 cases as tests for our primary results.

[5]We use the "545-avg" and "857-avg" values from
`http://www.sciops.esa.int/wikiSI/planckpla/index.php?title=UC_CC_Tables`.

We construct a point source mask using the same approach as in [44]. We take the union of the individual point source masks provided at each of the LFI and HFI frequencies [45]. In order to mask sources as thoroughly as possible, we use the 5σ catalogs rather than the 10σ catalogs. The authors of [44] verify that this procedure removes all resolved radio sources, as well as an unknown number of IR sources. Unresolved radio and IR point sources will still be present in the derived Compton-y map, but this masking prevents their emission from significantly biasing the ILC algorithm.

In our tSZ reconstruction pipeline, we take the union of the 857 GHz Galactic dust mask, the point source mask, and all of the masks associated with the S1 and S2 channel maps. We will refer to this mask as the "y-map mask". For our baseline results, which use the $f_{\rm sky} = 0.30$ Galactic dust mask, the total sky fraction left in the y-map mask is $f_{\rm sky} = 0.25180$.

When estimating the tSZ – CMB lensing cross-spectrum, we also must account for the mask associated with the Planck CMB lensing potential map. The construction of this mask is described in full detail in [46]. In addition to masking Galactic dust and point sources, this mask also removes regions contaminated by CO emission, as well as extended nearby objects, such as the Magellanic Clouds. Note that although some tSZ clusters are masked in the 143 GHz lensing reconstruction, they are not masked in the 217 GHz reconstruction, and thus in the publicly released map — a minimum-variance combination of 143 and 217 GHz — these tSZ clusters are not masked. (The noise levels at the location of these clusters will be somewhat higher than elsewhere in the lensing potential map, since the signal there is only reconstructed from one of the two channels, but we neglect this small effect in our analysis.) The y-map mask is sufficiently thorough that the inclusion of the lensing mask only covers a small amount of additional sky — the total sky fraction left in the combined mask for our fiducial results is $f_{\rm sky} = 0.25177$.

We apodize all masks used in the analysis by smoothing them with a Gaussian beam of FWHM 10 arcmin. The apodization prevents excessive ringing or other artifacts when we

compute power spectra. The effective sky fraction is reduced by $< 0.01\%$ by the apodization (compared to using an unapodized mask), but we nonetheless take this into account by computing $f_{\rm sky}$ via

$$f_{\rm sky} = \frac{1}{4\pi} \int d^2\hat{n} M(\hat{n}) \,, \tag{3.17}$$

where $M(\hat{n})$ is the apodized mask. Note that when estimating power spectra (see below), we correct for the effects of the mask using the average value of the square of the apodizing mask, rather than its mean [46]:

$$f_{\rm sky,2} = \frac{1}{4\pi} \int d^2\hat{n} M^2(\hat{n}) \,. \tag{3.18}$$

3.4.2 ILC

In order to construct a tSZ map over a large fraction of the sky using the Planck HFI maps, we implement a slightly modified version of the standard ILC technique. In the ILC approach, one assumes that the observed temperature T at each pixel p, of which there are $N_{\rm pix}$, in all $N_{\rm obs}$ maps (indexed by i in the following) can be written as a linear combination of the desired signal (y) and noise (n):

$$T_i(p) = a_i y(p) + n_i(p) \,, \tag{3.19}$$

where a_i is the product of $T_{\rm CMB} = 2.726$ K and the tSZ spectral function at the $i^{\rm th}$ frequency, which is computed by integrating over the relevant bandpass for each channel. We use the values of a_i given in Table 1 of [44]; we obtain values consistent with these when integrating the tSZ spectral function over the publicly released Planck HFI bandpasses [47]. The effect of the bandpass integration compared to simply using the central frequency of each channel is fairly small except for the 217 GHz channel, because it spans the tSZ null frequency (as emphasized in [48]). For the 217 GHz channel, a naïve calculation gives a result with the

wrong sign and an amplitude that is incorrect by roughly an order of magnitude; thus, accounting for the bandpass in quite important at this frequency.

The standard ILC approach computes an estimate of the desired signal in each pixel $\hat{y}(p)$ by constructing the minimum-variance linear combination of the observed maps that simultaneously satisfies the constraint of unit response to the signal of interest. Defining the ILC weights w_i via $\hat{y}(p) = w_i T_i(p)$, we thus seek to minimize the variance $\sigma_y^2 = N_{\rm pix}^{-1} \sum_p \left(\hat{y}(p) - \langle \hat{y} \rangle\right)^2$ while enforcing the constraint $w_i a_i = 1$ (summation over repeated indices is implied throughout). A simple derivation using Lagrange multipliers yields the desired result [51]:

$$w_j = \frac{a_i (\hat{R}^{-1})_{ij}}{a_k (\hat{R}^{-1})_{kl} a_l} \,, \tag{3.20}$$

where $\hat{R}_{ij} = N_{\rm pix}^{-1} \sum_p \left(T_i(p) - \langle T_i \rangle\right) \left(T_j(p) - \langle T_j \rangle\right)$ is the empirical covariance matrix of the (masked) observed maps. Note that the weights w_i have units of $\mathrm{K}_{\rm CMB}^{-1}$ in this formulation. Also note that it is important to mask the Galaxy and point sources *before* applying the ILC algorithm, as otherwise the strong emission from these sources can heavily bias the derived weights. We apply the full y-map mask described in the previous section before running our ILC algorithm.

This standard ILC approach can be extended to explicitly prevent any signal from the primordial CMB leaking into the derived y-map. This step is facilitated by the fact that the CMB spectrum is known to be a blackbody to extremely good precision. Eq. (3.19) is now modified to explicitly include the CMB as a second signal, in addition to the tSZ signal:

$$T_i(p) = a_i y(p) + b_i s(p) + n_i(p) \,, \tag{3.21}$$

where $s(p)$ is the CMB signal in pixel p and b_i is the CMB spectrum evaluated at each map's frequency, which is simply unity for maps in $\mathrm{K}_{\rm CMB}$ units (as ours are). Imposing the condition that the ILC y-map has zero response to the CMB signal, i.e., $w_i b_i = 0$,

and solving the system with Lagrange multipliers leads to the following modified version of Eq. (3.20) [52]:

$$w_j = \frac{\left(b_k(\hat{R}^{-1})_{kl}b_l\right) a_i(\hat{R}^{-1})_{ij} - \left(a_k(\hat{R}^{-1})_{kl}b_l\right) b_i(\hat{R}^{-1})_{ij}}{\left(a_k(\hat{R}^{-1})_{kl}a_l\right)\left(b_m(\hat{R}^{-1})_{mn}b_n\right) - \left(a_k(\hat{R}^{-1})_{kl}b_l\right)^2} \,. \tag{3.22}$$

Following [52], we will refer to this approach as the "constrained" ILC (CILC). This method ensures that no CMB signal leaks into the derived y-map. We implement Eq. (3.22) in our tSZ reconstruction pipeline, with additional modifications described in the following.

Our method modifies the standard ILC or CILC approach in two ways. First, we note that the variance σ_y^2 can be written as a sum over the power spectrum C_ℓ^y as

$$\sigma_y^2 = \sum_{\ell=0}^{\infty} \frac{2\ell+1}{4\pi} C_\ell^y \,. \tag{3.23}$$

Thus, the minimization of the variance in the ILC map is the minimization of this sum, taken over all ℓ. Our goal is to produce an ILC y-map to be used for cross-correlation with the Planck CMB lensing potential map (and possibly other extragalactic maps). It is therefore most important to minimize the extragalactic contamination, especially from the CIB and IR sources, rather than contamination from Galactic dust. Hence, we modify the ILC to minimize only a restricted sum over the power spectrum:

$$\tilde{\sigma}_y^2 = \sum_{\ell=\ell_a}^{\ell_b} \frac{2\ell+1}{4\pi} C_\ell^y \,. \tag{3.24}$$

We choose $\ell_a = 300$ and $\ell_b = 1000$, as Galactic dust emission is subdominant to the CIB over this ℓ range (after heavily masking the Galaxy), but the tSZ signal is still significant. In general, even after thorough masking, Galactic dust dominates over the tSZ at low-ℓ, while at high-ℓ the CIB takes over [44]. We verify that our results are stable to modest variations in these values.

In addition to minimizing a restricted sum over the power spectrum, we implement a second modification to the ILC algorithm which prevents instrumental noise from contributing to the ILC weights. Instead of minimizing $\tilde{\sigma}_y^2$ computed from the auto-spectra of the channel maps, we compute this quantity using cross-spectra of the S1 and S2 channel maps. The minimized quantity is now a cross-statistic:

$$\tilde{\sigma}_{y_{12}}^2 = \sum_{\ell=\ell_a}^{\ell_b} \frac{2\ell+1}{4\pi} C_\ell^{y_1 y_2} \,, \tag{3.25}$$

where the "1" and "2" subscripts refer to S1 and S2. Note that this approach now implies that we construct both an S1 and S2 ILC y-map, but using the same weights w_i, i.e., $y_1(p) = w_i T_i^1(p)$ and $y_2(p) = w_i T_i^2(p)$, where T_i^1 and T_i^2 are the $i^{\rm th}$ S1 and S2 channel maps, respectively.

Our ILC approach thus consists of finding the linear combination of maps that minimizes $\tilde{\sigma}_{y_{12}}^2$ in Eq. (3.25) while simultaneously requiring unit response to a tSZ spectrum and zero response to a CMB spectrum. Straightforward algebra using Lagrange multipliers yields the final expression for the weights — it is identical to the CILC weights in Eq. (3.22), except that the covariance matrix $\hat{R}_{ij}$ is now replaced by a slightly modified quantity:

$$\hat{\tilde{R}}_{ij} = \sum_{\ell=\ell_a}^{\ell_b} \frac{2\ell+1}{2\pi} C_{\ell,12}^{ij} \,, \tag{3.26}$$

where $C_{\ell,12}^{ij}$ is the cross-power spectrum of the $i^{\rm th}$ S1 channel map and the $j^{\rm th}$ S2 channel map. Eqs. (3.22) and (3.26) define our modified CILC approach to constructing a y-map.

As mentioned above, the weights determined by Eqs. (3.22) and (3.26) actually yield two y-maps: one constructed by applying the weights to the S1 channel maps and one from the S2 channel maps. In Section 3.4.3 we compute the tSZ auto-power spectrum by taking the cross-power spectrum of the S1 and S2 y-maps. In order to compute the cross-power spectrum of the tSZ and CMB lensing signals, we co-add the S1 and S2 y-maps to obtain

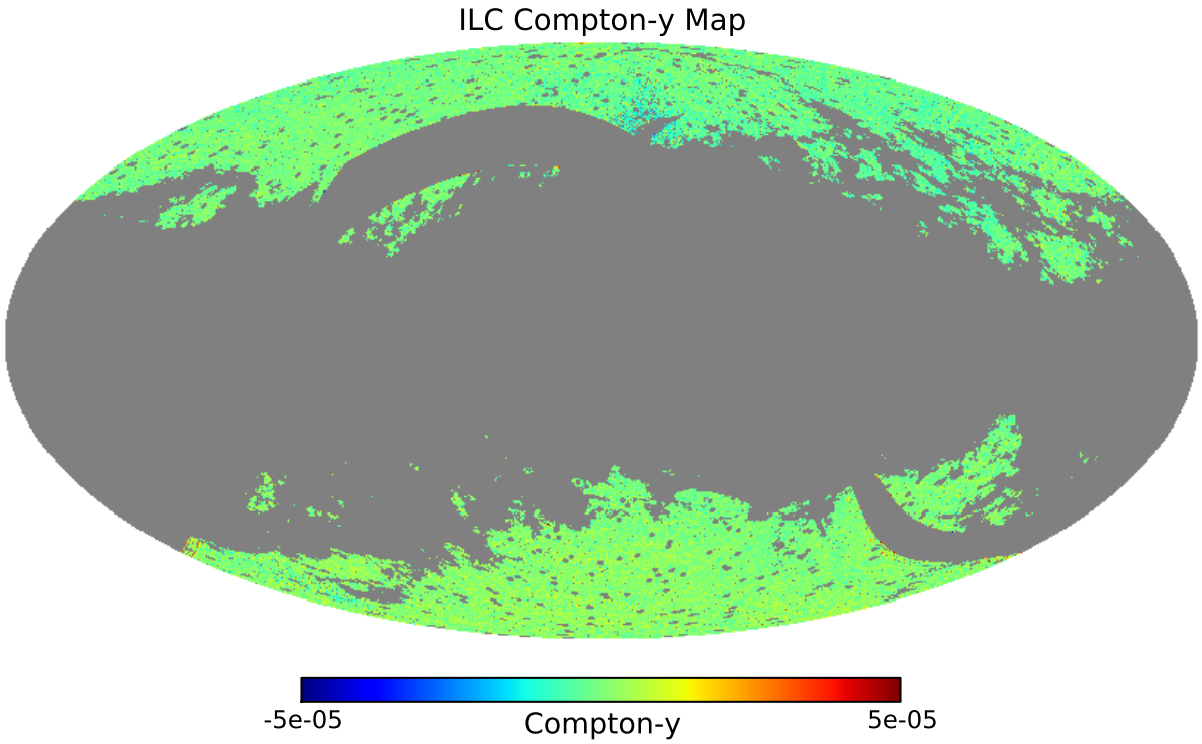

Figure 3.6: The Compton-y map reconstructed by our ILC pipeline for the fiducial $f_{\rm sky} = 0.30$ case, plotted in Galactic coordinates. Note that the mean of the map has been removed before plotting for visual clarity.

a final y-map which we then cross-correlate with the CMB lensing potential map. The co-addition of the S1 and S2 y-maps is performed with an inverse variance weighting to obtain the minimum-variance combination of the two maps.

Fig. 3.6 shows the final co-added y-map obtained from our modified CILC pipeline for the fiducial $f_{\rm sky} = 0.30$ (note that the actual final $f_{\rm sky} = 0.25180$ after accounting for the S1 and S2 masks and the point source mask, as mentioned in Section 3.4.1). The mean has been removed from this map before plotting, in order to allow better visual clarity. Fig. 3.7 shows the histogram of Compton-y values in this map. Apart from the dominant Gaussian noise component, there is a clear non-Gaussian tail extending to positive Compton-y values, which provides some evidence that there is indeed tSZ signal in the map. A smaller non-Gaussian tail extending to negative values is likely caused by residual Galactic dust in the map. Interpreting the moments of this histogram [59, 60, 61] is an interesting prospect for future

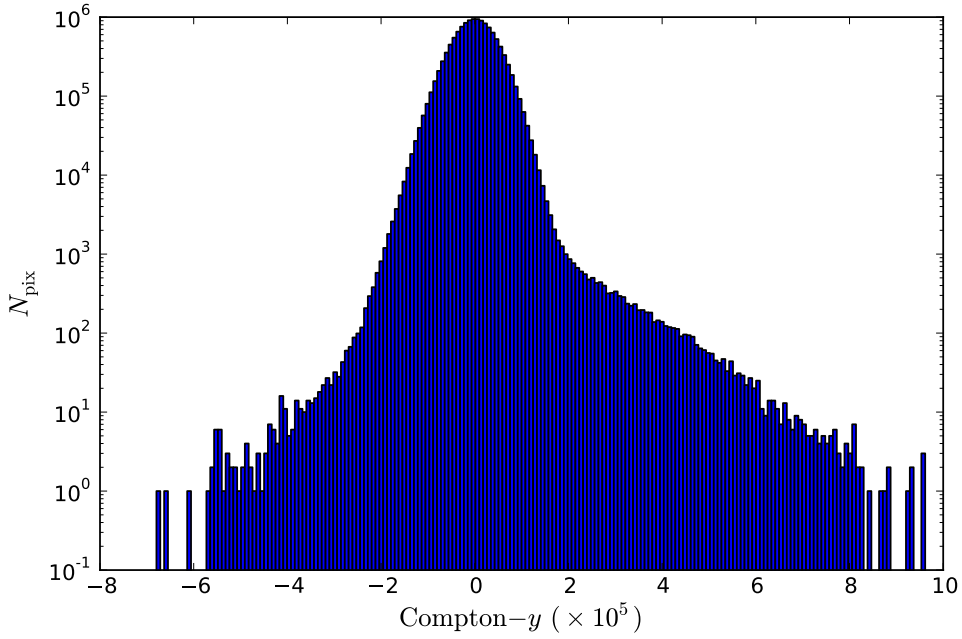

Figure 3.7: Histogram of Compton-y values in our fiducial ILC y-map. The non-Gaussian tail extending to positive y-values provides some evidence that the map does indeed contain tSZ signal. The negative tail is likely due to residual Galactic dust in the map. See [59, 60, 61] for related work on interpreting the moments of Compton-y histograms.

work, but will require careful consideration of the non-tSZ components, which is beyond the scope of our analysis here. As further evidence of the success of our reconstruction, we show a sub-map of a region centered on the Coma cluster in Fig. 3.8. Coma is the most significant tSZ cluster in the Planck sky [58]. The central pixel's y-value in Coma in our (mean-subtracted) ILC map is $\approx 7 \times 10^{-5}$, which is fairly consistent with the determination in [58] (see their Figs. 4 or 5). Although we cannot easily estimate the error on this quantity without simulations, this result provides additional evidence for the overall success of our ILC reconstruction.

The primary disadvantage of the ILC technique is its assumption that the signal of interest (y in Eq. (3.19)) is completely uncorrelated with the noise and foregrounds at each frequency (n_i in Eq. (3.19)). In the case of the tSZ signal, this assumption is likely violated by the correlation between the tSZ and IR sources (e.g., dusty star-forming galaxies) [53],

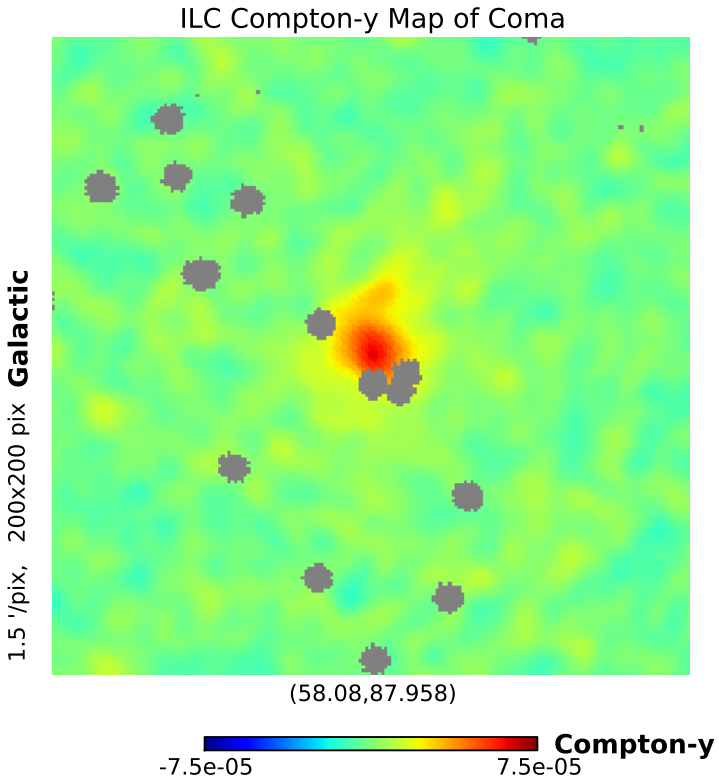

Figure 3.8: Sub-map of a 5°-by-5° region centered on the Coma cluster in the ILC y-map (after removing the map mean). The masked areas are the locations of point sources in this field.

for which some evidence has recently been found [54, 55]. However, given the indirect nature of this evidence and difficulty in assessing the amplitude of the correlation, we neglect it for the present time (as in [44]). Future tSZ reconstructions at higher SNR — for example, using the complete Planck data set — may need to consider its implications.

Finally, we note that our ILC approach differs somewhat from those employed in the Planck y-map analysis [44]. In particular, we do not implement a method that combines reconstructions at varying angular scales, such as the Needlet ILC (NILC) [56]. Since the Planck HFI maps are smoothed to a common 10 arcmin resolution before the tSZ reconstruction is performed in both [44] and this work, it seems difficult to take advantage of the additional power of a multi-scale approach. In future work we plan to combine the novel

elements of our ILC pipeline described above with a multi-scale reconstruction pipeline, allowing the simultaneous use of data from Planck (HFI and LFI), WMAP, and ground-based experiments such as ACT and SPT.

Although the Compton-y map reconstructed from our pipeline is not free of residual contamination, as evident in the histogram and power spectrum (see the following section), we believe that it should nonetheless be useful for cross-correlation studies if treated carefully. Thus, we make available to the community a `HEALPix` version of our fiducial $f_{\rm sky} = 0.30$ map[6].

3.4.3 Thermal SZ Auto-Power Spectrum

We compute the tSZ auto-power spectrum by cross-correlating the S1 and S2 Compton-y-maps obtained with the pipeline described in Section 3.4.2. Given a reconstructed Compton-y map, $\hat{y}(\hat{n})$, we perform a spherical harmonic transform via $\hat{y}_{\ell m} = \int d^2\hat{n}\, Y^*_{\ell m}(\hat{n})\, \hat{y}(\hat{n})$ and compute the tSZ power spectrum using a simple pseudo-C_ℓ estimator:

$$\hat{C}^{yy}_\ell = \frac{f^{-1}_{\rm sky,2}}{2\ell+1} \sum_m \hat{y}^1_{\ell m} \hat{y}^{2*}_{\ell m}\,, \tag{3.27}$$

where $f_{\rm sky,2}$ is given by Eq. (3.18), computed using only the y-map mask (clearly the lensing potential is not involved in this analysis), and y^1 and y^2 refer to the S1 and S2 y-maps, respectively. Our power spectrum estimate also accounts for the deconvolution of the 10 arcmin smoothing applied to all channel maps before the ILC reconstruction pipeline. We "whiten" the estimated power spectrum by multiplying by $\ell(\ell+1)/(2\pi)$ and then bin it using bins identical to those chosen in [44] in order to allow a direct comparison of the results. We follow [46, 48, 44] in neglecting any correlations arising in the error bars due to mode coupling induced by the mask; we simply correct for the power lost through masking with the factor $f^{-1}_{\rm sky,2}$.

[6] `http://www.astro.princeton.edu/~jch/ymapv1/`

Fig. 3.9 shows the tSZ power spectrum estimated from the cross-spectrum of our fiducial $f_{\rm sky} = 0.30$ S1 and S2 ILC Compton-y maps, as well as the Compton-y power spectrum estimated in [44] before and after subtracting residual contributions to the power spectrum from clustered CIB, IR point sources, and radio point sources. Note that we have not attempted to subtract these residual contributions from our tSZ power spectrum shown in Fig. 3.9. Our data points (red circles) should be compared directly to the uncorrected Planck points shown as blue circles; moreover, they should only be compared over the region in which our ILC is designed to remove contamination, i.e., $300 < \ell < 1000$ (as delineated in the figure). Within this region, our results are in reasonable agreement with the uncorrected Planck results. The green squares show the Planck points after power from residual foregrounds is subtracted. In general, the uncorrected tSZ power spectra are contaminated by power from residual Galactic dust at low multipoles and from residual CIB, IR sources, and radio sources at high multipoles [44]. Simple tests varying the sky fraction used in our ILC Compton-y reconstruction give power spectra broadly consistent with that shown in Fig. 3.9; the $f_{\rm sky} = 0.20$ results appear to have slightly more CIB leakage and less Galactic dust, while the opposite is true for the $f_{\rm sky} = 0.40$ results.

It is also important to note that the error bars shown on our data points and the uncorrected Planck data points in Fig. 3.9 are statistical errors only:

$$\left(\Delta\hat{C}_\ell^{yy}\right)^2 = \frac{1}{f_{\rm sky,2}} \frac{2}{(2\ell+1)\Delta\ell} \left(\hat{C}_\ell^{yy}\right)^2 , \tag{3.28}$$

where $\Delta\ell$ is the width of a given multipole bin centered at ℓ. Note that $\hat{C}_\ell^{yy}$ is the measured tSZ auto-power spectrum, i.e., it includes the noise bias. We show both the statistical-only and statistical+foreground errors on the corrected Planck points. Neither set of error bars includes contributions from sample variance, although we will include this effect in our cosmological analysis in Section 3.6. The uncertainties arising from the subtraction of residual foreground contributions dominate the total errors on the corrected Planck data

points [44]. We emphasize again that the visual discrepancy between the red and green points in Fig. 3.9 is not a sign that our y-map is significantly more contaminated than that in [44], but only a reflection of the fact that we have not subtracted residual foreground contributions to the auto-spectrum of the y-map. Within $300 < \ell < 1000$, the raw power spectrum of our ILC y reconstruction is quite similar to that of [44], as seen in the similarity between the red and blue points over this multipole range in Fig. 3.9. Our aim is to measure the cross-correlation of the tSZ and CMB lensing signals, rather than the tSZ auto-spectrum; the latter requires a higher threshold for removing foregrounds and contamination. We do not strive to directly repeat the foreground analysis of [44] here, and instead will simply use their tSZ power spectrum in combination with our tSZ – CMB lensing cross-power spectrum in the cosmological interpretation of our results in Section 3.6.

3.5 Thermal SZ – CMB Lensing Cross-Power Spectrum

3.5.1 CMB Lensing Potential Map

To lowest order in the CMB lensing potential ϕ, gravitational lensing introduces a correlation between the lensed temperature field and the gradient of the unlensed temperature field (see Eq. (3.3)). This correlation leads to statistical anisotropy in the observed (lensed) CMB temperature field on the arcminute angular scales where lensing effects become detectable (the RMS deflection of a CMB photon due to lensing is $\sim$ 2–3 arcmin). Using this property of the lensed CMB, it is possible to reconstruct a map of the lensing potential itself, $\phi(\hat{n})$ [83, 82]. The Planck team performed a detailed study of these effects in the nominal mission data, including reconstructions of the lensing potential at 100, 143, and 217 GHz and a 25σ detection of the power spectrum of the lensing potential [46].

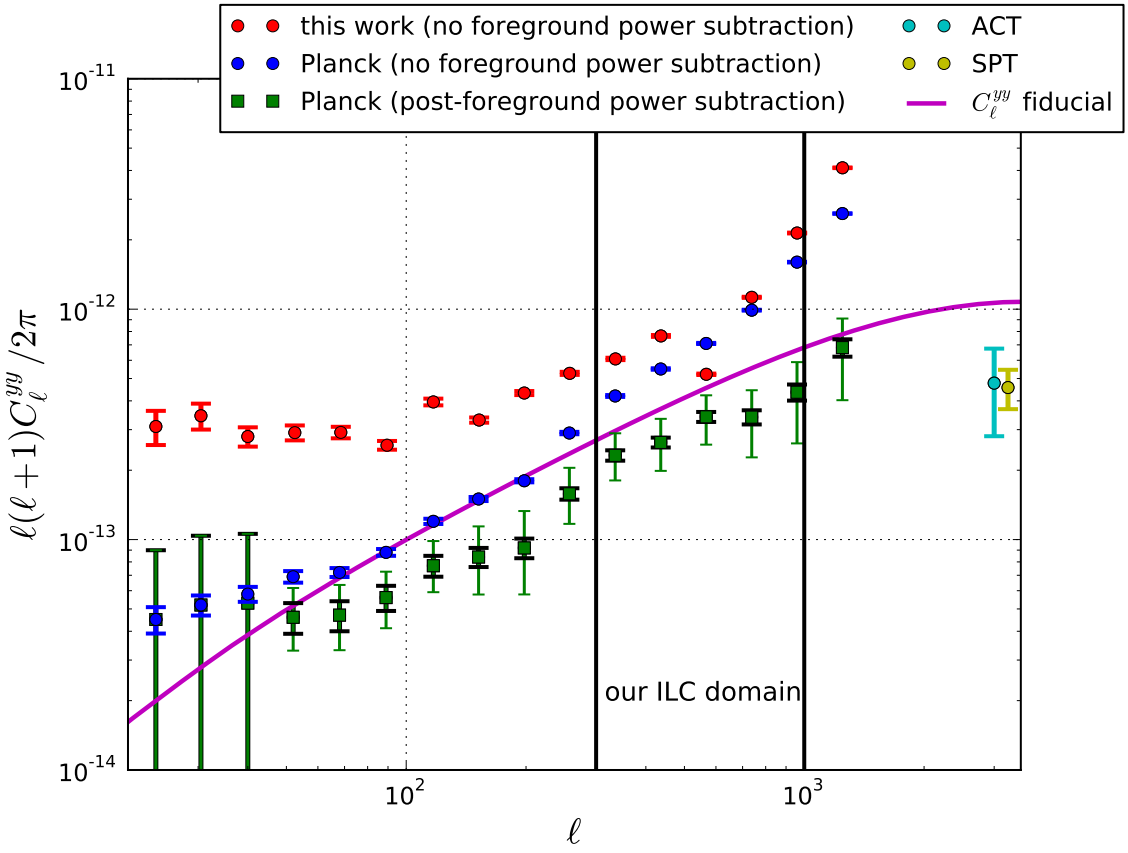

Figure 3.9: The tSZ power spectrum estimated from our fiducial S1 and S2 Compton-y maps (red circles), with statistical errors only. The Planck results from [44] are shown as blue circles and green squares — the green squares have been corrected to subtract residual power from leakage of CIB and IR and radio point sources into the y-map constructed in [44], while the blue circles are the uncorrected results (with statistical errors only). We have not attempted to subtract the residual power due to these foregrounds; the red circles are the raw cross-spectrum of our S1 and S2 ILC y-maps, which are in reasonable agreement with the uncorrected Planck points (blue circles) over the region in which our ILC algorithm minimizes residual contamination ($300 < \ell < 1000$, as delineated on the plot; see Section 3.4.2). The thick black error bars are the statistical errors alone on the foreground-corrected Planck points, while the thin green error bars are the combined statistical and foreground-subtraction uncertainties (from Table 3 of [44]). The solid magenta curve shows the tSZ power spectrum predicted by our fiducial model. The cyan and yellow points are the most recent ACT [4] and SPT [57] constraints on the tSZ power spectrum amplitude at $\ell = 3000$, respectively (the SPT point is slightly offset for visual clarity). Note that the ACT and SPT constraints are derived in a completely different data analysis approach, in which the tSZ signal is constrained through its indirect contribution to the total CMB power measured at $\ell = 3000$ by these experiments; in other words, a y-map is not constructed. We use the corrected Planck points with their full errors when deriving cosmological and astrophysical constraints in Section 3.6.

In our measurement of the tSZ – CMB lensing cross-power spectrum, we utilize the publicly released Planck CMB lensing potential map, which is a minimum-variance combination of the 143 and 217 GHz reconstructions. As described in Section 3.4.1, we combine the mask associated with the CMB lensing potential map with the mask associated with our ILC y-map when computing the tSZ – CMB lensing cross-power spectrum. We refer the reader to [46] for a complete description of the Planck CMB lensing reconstruction. We have verified that our pipeline produces a lensing potential power spectrum estimate from the publicly released map which is broadly consistent with that presented in [46].

Note that the inclusion of the 217 GHz channel in the publicly released map is crucial, as it allows reconstruction to be performed in regions of the sky centered on tSZ clusters that must be masked in the 143 GHz reconstruction. The 143 GHz masking of the tSZ-detected clusters does imply that some small effects could be present at these locations in the combined 143+217 GHz map, such as changes in the effective noise level and estimator normalization. However, given that only the most massive clusters in the universe are masked in this procedure ($\sim$ 500–600 objects[7]), we do not expect it to significantly affect our measurement of the tSZ – CMB lensing cross-spectrum, which is dominated by clusters well below the Planck detection threshold (see Section 3.3.3). Fig. 3.5 indicates that the clusters masked in the 143 GHZ lensing reconstruction contribute $\approx 10\%$ of the tSZ – CMB lensing cross-power spectrum signal. The change in the effective normalization of the lensing estimator in the combined 143+217 GHz map at these locations is a factor of ≈ 2 due to the masking, and thus it is possible that our treatment has missed an error of roughly this factor on $\approx 10\%$ of the signal. This level of error is well below the statistical error on our detection (see the next two sections), and thus we do not consider these effects further. Future high-signal-to-noise detections of this cross-spectrum will need to carefully consider the effects of tSZ masking in the lensing reconstruction — performing the lensing reconstruction on either

[7]These are SNR > 5 clusters detected with the MMF1 or MMF3 pipelines [50].

a 217 GHz map or an ILC CMB map constructed to have zero tSZ signal would likely solve this problem, but simulations should be used to completely characterize the effects.

3.5.2 Measurement

We use the co-added y-map described in Section 3.4.2 and the CMB lensing potential map described in Section 3.5.1 to measure the tSZ – CMB lensing cross-power spectrum for our fiducial $f_{\rm sky} = 0.3$ mask. We compute the cross-power spectrum using a simple pseudo-C_ℓ estimator analogous to that in Eq. (3.27):

$$\hat{C}_\ell^{y\phi} = \frac{f_{\rm sky,2}^{-1}}{2\ell+1} \sum_m \hat{y}_{\ell m} \hat{\phi}^*_{\ell m} \,, \tag{3.29}$$

where $f_{\rm sky,2} = 0.24281$ is given by Eq. (3.18), computed using the union of the y-map mask and the lensing map mask. We multiply the cross-power spectrum by $\ell^2(\ell+1)/(2\pi)$ and bin it using twelve linearly spaced bins over the range $100 < \ell < 1600$. These bins are identical to those chosen in [48], with the exception that our upper multipole limit is slightly lower than theirs (1600 instead of 2000) — we conservatively choose a lower multipole cutoff due to our simplified treatment of the Planck beams. As in Section 3.4.3, we follow [46, 48, 44] in neglecting any correlations arising in the error bars due to mode coupling induced by the mask, and simply correct for the power lost through masking with the factor $f_{\rm sky,2}^{-1}$.

We compute statistical error bars on the cross-power spectrum using the standard analytic prescription (e.g., [48]):

$$\left(\Delta \hat{C}_\ell^{y\phi}\right)^2 = \frac{1}{f_{\rm sky,2}} \frac{1}{(2\ell+1)\Delta\ell} \left(\hat{C}_\ell^{yy} \hat{C}_\ell^{\phi\phi} + \left(C_\ell^{y\phi}\right)^2 \right) \,, \tag{3.30}$$

where $\hat{C}_\ell^{yy}$ and $\hat{C}_\ell^{\phi\phi}$ are the measured auto-spectra of the co-added y-map- and the ϕ-map, respectively (i.e., these power spectra include the noise bias), and $C_\ell^{y\phi}$ is the fiducial theoretical cross-spectrum. The contribution from the second term is typically 3–5 orders of magnitude smaller than the contribution from the first — thus any uncertainty in the the-

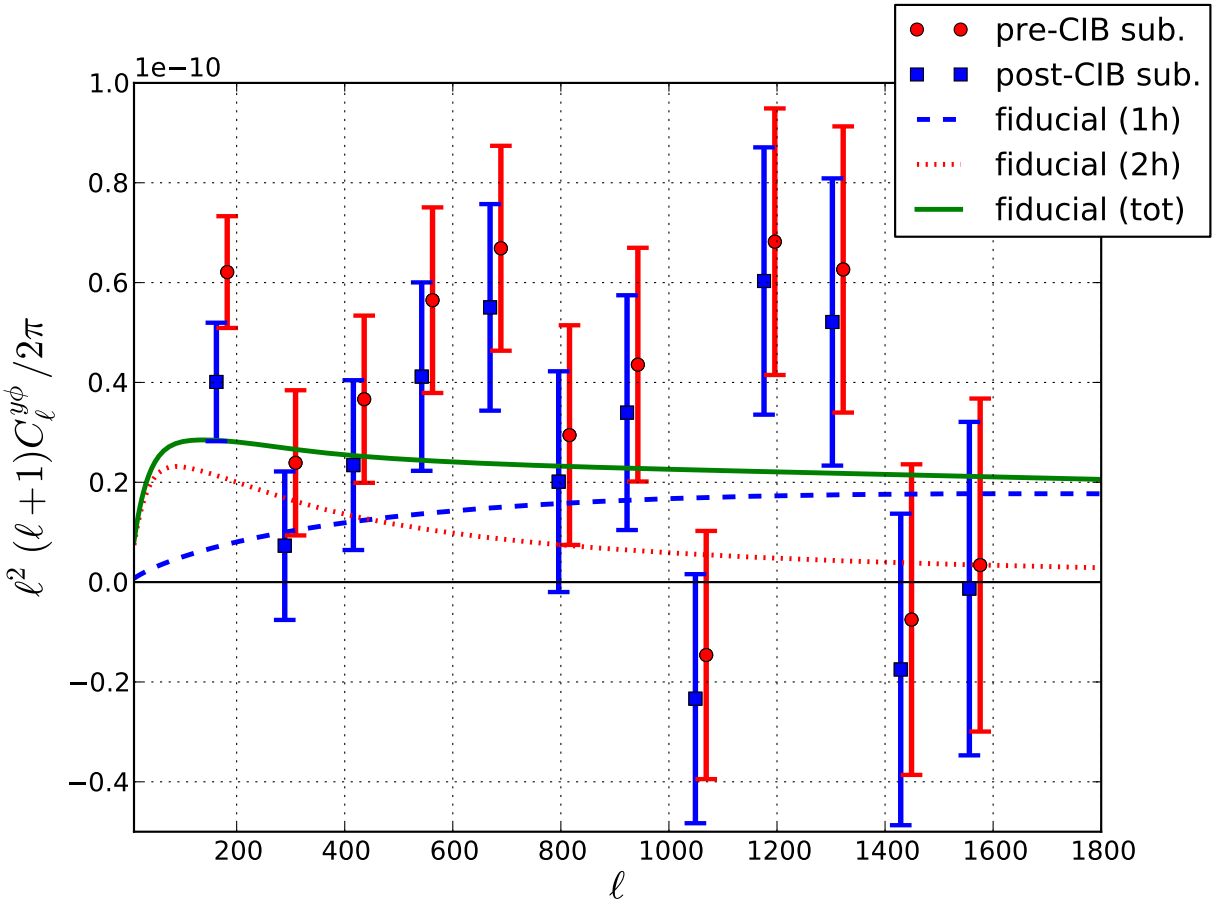

Figure 3.10: The tSZ – CMB lensing cross-power spectrum estimated from our fiducial co-added Compton-y map and the publicly released Planck lensing potential map is shown in the red circles, with statistical errors computed via Eq. (3.30). The blue squares show the cross-power spectrum after a correction for CIB leakage into the y-map has been applied (see Section 3.5.3 for a description of the subtraction procedure). The errors on the blue squares include statistical uncertainty and additional systematic uncertainty arising from the CIB subtraction. The final significance of the CIB-corrected measurement of $C_\ell^{y\phi}$ is 6.2σ. The solid green curve shows the theoretical prediction of our fiducial model for the tSZ – CMB lensing cross-spectrum (see Section 3.3) — we emphasize that this curve is not a fit to the data, but rather an independent prediction. It agrees well with our measurement: $\chi^2 = 14.6$ for 12 degrees of freedom. The dashed blue and dotted red curves show the one- and two-halo contributions to the fiducial theoretical cross-spectrum, respectively. Our measurement shows clear evidence of contributions from both terms.

oretical modeling of $C_\ell^{y\phi}$ is completely negligible for the purposes of Eq. (3.30). This fact arises because the individual y- and ϕ-maps are quite noisy, and thus the approximation that they are nearly uncorrelated (which is implicit in Eq. (3.30)) is valid. Finally, we note that five additional error calculation methods were tested in [48] (see their Appendix A) using both simulations and data, and all were found to give results consistent with Eq. (3.30). Given the similarity between our analysis and theirs, we expect that Eq. (3.30) should be quite accurate for our measurement as well. In future work we will assess this issue more carefully using forthcoming public lensing simulations from the Planck team.

Fig. 3.10 shows our initial measurement of the tSZ – CMB lensing power spectrum (red circles). A strong signal is clearly detected. We quantify the detection significance according to:

$$\mathrm{SNR} = \sqrt{\sum_{i=1}^{12} \left(\frac{\hat{C}_i^{y\phi}}{\Delta C_i^{y\phi}} \right)^2}, \tag{3.31}$$

where i indexes each of the twelve multipole bins in our measurement. The SNR of our initial detection is 8.7σ. However, as described in detail in the following section, the measurement is subject to contamination from CIB leakage into the y-map. The CIB-subtracted results are shown as blue squares in Fig. 3.10.

3.5.3 CIB Contamination Correction

The primary systematic affecting our measurement of the tSZ – CMB lensing cross-power spectrum is leakage of CIB emission into the ILC y-map. Not only is the CIB emission difficult to completely remove from the y-map [44], it also correlates very strongly with the CMB lensing signal due to the similar mass and redshift kernels sourcing the two fields [84, 48]. Thus, it is clear that we need to investigate the possible contamination of our tSZ – CMB lensing cross-spectrum measurement due to residual CIB emission correlating with the CMB lensing potential.

We adopt a primarily empirical approach based on the Planck 857 GHz map to address this contamination, although we also test an alternative method in Section 3.5.4 and find that it gives consistent results. We use the 857 GHz map as a tracer of dust emission from both the Galaxy and the CIB. We cross-correlate the y-map constructed in Section 3.4.2 with the 857 GHz map and use the resulting cross-power spectrum to assess the level of CIB leakage into the y-map. Specifically, we assume that the 857 GHz signal $T_{857}(p)$ in pixel p is given by

$$T_{857}(p) = T_{\rm CIB}(p) + T_{\rm Gal}(p)\,, \tag{3.32}$$

where $T_{\rm CIB}(p)$ and $T_{\rm Gal}(p)$ are the CIB and Galactic dust emission in pixel p. (The CMB, tSZ, and other sky components are completely overwhelmed by dust emission at 857 GHz [48, 49].) Neglecting any other contaminants in the ILC y-map, we assume that the observed signal $\hat{y}(p)$ in pixel p is given by

$$\hat{y}(p) = y(p) + \alpha_{\rm CIB} T_{\rm CIB}(p) + \alpha_{\rm Gal} T_{\rm Gal}(p)\,, \tag{3.33}$$

where $y(p)$ is the "true" uncontaminated Compton-y signal and $\alpha_{\rm CIB}$ and $\alpha_{\rm Gal}$ are free parameters that describe the mean leakage of CIB and Galactic dust emission into the reconstructed y-map, respectively. Note that the definitions in Eqs. (3.32) and (3.33) imply that we have normalized the frequency dependence of the CIB and Galactic dust emission by their values at 857 GHz. An important assumption implicit in the following is that the CIB emission at the lower HFI frequencies is perfectly correlated with that at 857 GHz; while this correlation is not perfect in reality, it is in fact very strong ($\approx$ 80–90% — see Fig. 13 in [48]), which suffices for our purposes here.[8]

Using Eq. (3.32) and assuming that the CIB and Galactic dust emission is uncorrelated (which physically must be true given their disparate origins), the auto-power spectrum of

[8]We thank U. Seljak for emphasizing this point.

the 857 GHz map is

$$C_\ell^{857} = C_\ell^{\rm CIB} + C_\ell^{\rm Gal}\,, \tag{3.34}$$

where $C_\ell^{\rm CIB}$ and $C_\ell^{\rm Gal}$ are the power spectra of the CIB and Galactic dust emission, respectively. Similarly, the cross-power spectrum between our ILC y-map and the Planck 857 GHz map is

$$C_\ell^{y\times 857} = \alpha_{\rm CIB} C_\ell^{\rm CIB} + \alpha_{\rm Gal} C_\ell^{\rm Gal}\,, \tag{3.35}$$

where we have neglected any possible correlation between the CIB emission and the tSZ signal, as in Section 3.4.2 (following [44]). Such a correlation likely does exist [53], but in this analysis it would be very difficult to separate from the correlation between CIB leakage in the y-map and the CIB emission in the 857 GHz map, given the similar shape of the signals in multipole space. The $\alpha_{\rm CIB}$ parameter essentially captures a combination of the physical tSZ – CIB correlation and the spurious correlation due to leakage of CIB into the y-map. For simplicity, we assume that the spurious correlation dominates. Given a determination of $\alpha_{\rm CIB}$, it is straightforward to correct our measurement of the tSZ – CMB lensing cross-spectrum:

$$\hat{C}_\ell^{y\phi,{\rm corr}} = \hat{C}_\ell^{y\phi} - \alpha_{\rm CIB}\hat{C}_\ell^{\phi\times 857}\,, \tag{3.36}$$

where $\hat{C}_\ell^{\phi\times 857}$ is the cross-power spectrum of the CMB lensing potential and the CIB emission at 857 GHz, which was measured at 16σ in [48] (including all statistical and systematic errors). If we considered both the physical tSZ – CIB correlation and the spurious correlation due to leakage of CIB into the y-map in this analysis, then the correction computed in Eq. (3.36) would be smaller than what we find below. The results of Section 3.6 suggest that it is unlikely that we have overestimated the CIB contamination correction, and thus our

current neglect of the physical tSZ – CIB correlation seems reasonable. Nevertheless, this is clearly an area in need of improvement in future tSZ analyses.

The first step in the CIB correction pipeline is measuring the auto-power spectrum of the 857 GHz map and separating the CIB and Galactic components using their different shapes in multipole space. We adopt a model for the CIB emission at 857 GHz from [81], which is based on a simultaneous fit to power spectrum measurements from Planck, Herschel, ACT, and SPT, as well as number count measurements from Spitzer and Herschel. We refer the reader to [81] for a complete description of the model — the values used in this work are shown in the "Planck 857 GHz" panel of Fig. 1 in [81][9]. We subtract the CIB model power spectrum from the measured 857 GHz power spectrum to obtain an estimate of the power spectrum of the Galactic dust emission. The results of this procedure are shown in the left panel of Fig. 3.11 for three different sky cuts, including our fiducial $f_{\rm sky} = 0.3$ case (the other cases will be considered when we perform null tests in Section 3.5.4). It is clear in the figure that the total power at 857 GHz increases as the sky fraction increases, simply because more Galactic dust emission is present. (Note that these power spectra have been corrected with the $f_{\rm sky,2}$ factor described in Eq. (3.18).) The only case in which the CIB power partially dominates the total measured power is for $f_{\rm sky} = 0.2$, primarily around $\ell = 1000$; however, the CIB contribution is clearly non-negligible in all cases. The key result from this analysis is an estimate for both $C_\ell^{\rm CIB}$ and $C_\ell^{\rm Gal}$ in Eq. (3.34).

The second step in the CIB correction pipeline is applying these results to the cross-spectrum of the ILC y-map and the 857 GHz map in order to measure $\alpha_{\rm CIB}$ and $\alpha_{\rm Gal}$, as described in Eq. (3.35). The results of this procedure are shown in the right panel of Fig. 3.11, for the fiducial $f_{\rm sky} = 0.3$ case only. The red circles in the plot are the measured cross-power spectrum between our ILC y-map and the Planck 857 GHz map, while the curves show the best-fit result when applying Eq. (3.35) to these measurements. For our fiducial $f_{\rm sky} = 0.3$

[9]We are grateful to G. Addison for providing the CIB model fitting results from [81].

analysis, we measure

$$\begin{aligned} \alpha_{\rm CIB} &= (6.7 \pm 1.1) \times 10^{-6}\ {\rm K}_{\rm CMB}^{-1} \\ \alpha_{\rm Gal} &= (-5.4 \pm 0.4) \times 10^{-6}\ {\rm K}_{\rm CMB}^{-1}\,, \end{aligned} \tag{3.37}$$

with marginalized 1σ uncertainties given for both parameters. The model provides a reasonable fit to the data, with $\chi^2 = 16.9$ for $12 - 2 = 10$ degrees of freedom. Note that the results of the fit also imply that our ILC weights yield a positive response when applied to a CIB-like spectrum, and a negative response when applied to a Galactic-dust-like spectrum, a result which was also found in [44, 50]. Finally, we note that while this approach is sufficient for correcting CIB leakage into the tSZ – CMB lensing cross-power spectrum, it is unlikely to suffice for correcting CIB leakage into quadratic or higher-order tSZ statistics computed from the y-map, such as the tSZ auto-power spectrum. The crucial point is that the tSZ – CMB lensing cross-spectrum is linear in y, and hence only an estimate of the mean CIB leakage into the y-map is needed to compute the correction, which our cross-correlation method provides. Fluctuations around this mean leakage will clearly contribute residual power to the tSZ auto-spectrum at a level which is not constrained in this approach. Physically, this arises because the spectrum of the CIB emission varies from source to source; thus, although we can remove the mean leakage into the y-map with our approach, fluctuations around this mean will still exist and contribute to the auto-power spectrum and higher-order statistics. The Planck team used simulations to assess the residual amount of non-tSZ power in their y auto-spectrum [44]; it seems likely that a combination of simulations and more refined empirical methods will be useful in future studies. For now we do not attempt to provide a cleaned tSZ auto-power spectrum from our y-maps.

We use the measured value of $\alpha_{\rm CIB}$ to correct the tSZ – CMB lensing cross-power spectrum following Eq. (3.36). We propagate the uncertainty resulting from the fit, as well as the uncertainties on $\hat{C}_\ell^{\phi\times 857}$ given in [48], into the final error bars on $\hat{C}_\ell^{y\phi,{\rm corr}}$. Our final, CIB-

corrected measurement of the tSZ – CMB lensing cross-power spectrum is shown in Fig. 3.10 as the blue squares. The final detection significance is 6.2σ after the CIB correction. The solid green curve in Fig. 3.10 shows the theoretical prediction of our fiducial model described in Section 3.3. This curve is not a fit to the measurements, but already agrees well with the data: $\chi^2 = 14.6$ for 12 degrees of freedom. The figure also shows the one- and two-halo terms for the fiducial model; it is clear that our measurement shows evidence of contributions from both terms.

The detection significance is somewhat lower for the tSZ – CMB lensing cross-power spectrum than for either of the auto-spectra, which are detected at 12.3σ [44] and 25σ [46], respectively. If the y and ϕ fields were perfectly correlated, i.e., $r_\ell^{y\phi} = 1$, and if the contributions to their individual SNR were over identical multipole ranges, then we could expect a $\sqrt{12 \times 25} \approx 17\sigma$ significance for the detection of $C_\ell^{y\phi}$. However, since $\langle r_\ell^{y\phi} \rangle \approx 0.33$ over the relevant multipole range for our measurement, we obtain a detection significance of $17 \times 0.33 \approx 6\sigma$. The fact that this estimate is in rough agreement with our actual measurement indicates that the multipole ranges over which the tSZ and CMB lensing auto-spectra have significant SNR are fairly similar.

Table 3.1 gives our final CIB-corrected bandpowers and the associated uncertainties.

3.5.4 Null Tests and Systematic Errors

We perform a series of null tests designed to search for any residual systematic errors that might affect our measurement of the tSZ – CMB lensing cross-power spectrum. Our calculation of the statistical errors on both the fiducial measurement and the null tests follows Eq. (3.30) throughout. We refer the reader to [46, 48] for a complete discussion of possible systematics that can affect the CMB lensing reconstruction. Essentially all of these effects are far smaller than our error bars, the smallest of which is $\approx 30\%$ (this is the fractional error on our lowest ℓ-bin). A particular cause for concern, however, is spurious lensing signal induced by other sources of statistical anisotropy in the temperature maps, such as the

$\ell_{\rm mean}$	$\ell^2(\ell+1)C_\ell^{y\phi}/(2\pi) \times 10^{11}$	$\Delta(\ell^2(\ell+1)C_\ell^{y\phi}/(2\pi)) \times 10^{11}$
163	4.01	1.19
290	0.73	1.49
417	2.34	1.70
543	4.12	1.89
670	5.50	2.07
797	2.01	2.21
923	3.39	2.35
1050	−2.34	2.49
1177	6.03	2.68
1303	5.21	2.88
1430	−1.75	3.12
1557	−0.13	3.34

Table 3.1: Measured tSZ – CMB lensing cross-spectrum bandpowers. All values are dimensionless. These bandpowers have been corrected for contamination from CIB emission following the procedure described in Section 3.5.3; they correspond to the blue squares shown in Fig. 3.10. Uncertainties due to the CIB subtraction have been propagated into the final errors provided here. Note that the bins have been chosen to match those in [48].

Galactic and point source masks. We mitigate this "mean-field" bias by discarding all data below $\ell = 100$ in our analysis, following [48].

The first null test we perform is to cross-correlate a null Compton-y map with the CMB lensing potential map. The null y-map is constructed by subtracting, rather than co-adding, the S1 and S2 y-maps that are produced by our ILC pipeline, as described in Section 3.4.2. The subtraction is performed using the same inverse-variance combination that is applied in the co-addition, guaranteeing that the noise level in the subtracted map is identical to that in the co-added map. The null map should be free of any astrophysical emission, including tSZ, CIB, or Galactic dust signals. The cross-power spectrum of the null y-map with the CMB lensing potential map is shown in Fig. 3.12 (black points with solid black error bars). The test is consistent with a null signal: $\chi^2 = 13.2$ for 12 degrees of freedom. For reference, Fig. 3.12 also shows the 1σ error bars on our fiducial signal in cyan. The consistency between the errors, as well as the reasonable value of χ^2, provides additional support for the robustness of our statistical error calculation.

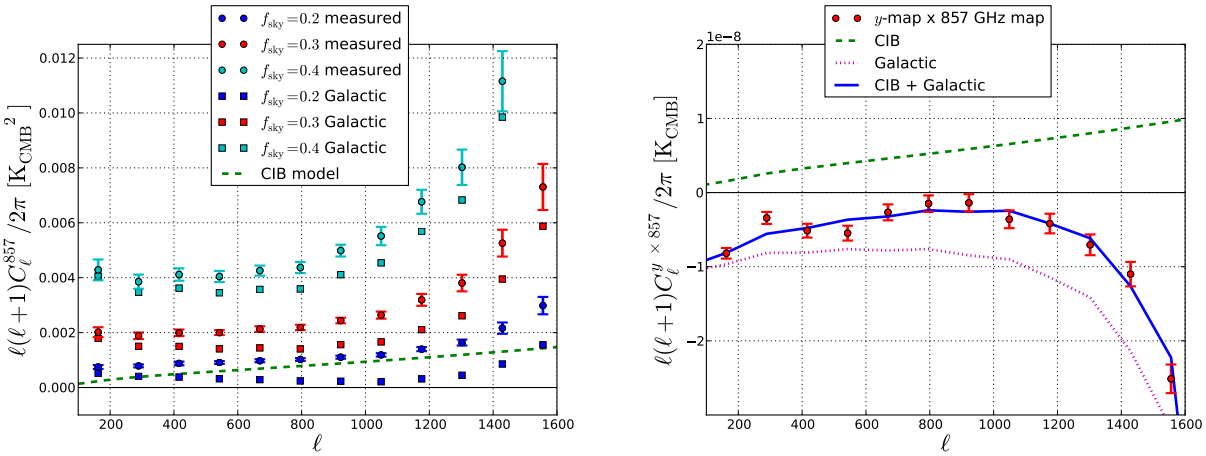

Figure 3.11: The left panel shows the auto-power spectrum of the Planck 857 GHz map (circles with error bars), computed for three sky cuts — note that the fiducial case is $f_{\rm sky} = 0.3$ (red circles). The right panel shows the cross-power spectrum between our fiducial co-added y-map and the Planck 857 GHz map (for $f_{\rm sky} = 0.3$ only). These measurements provide a method with which to assess the level of residual contamination in the y-map due to dust emission from the CIB and from the Galaxy. We model the 857 GHz map as a combination of CIB emission (computed from the best-fit model in [81], shown as a dashed green curve in the left panel) and Galactic dust emission (computed by subtracting the CIB model from the measured 857 GHz power spectrum, shown as squares in the left panel). We then model the y-map – 857 GHz cross-power spectrum as a linear combination of these emission sources using two free amplitude parameters (see the text for further information). We fit these parameters using the measured cross-power spectrum; the results of the fit are shown as dotted magenta (Galactic) and dashed green (CIB curves), with the sum of the two in solid blue. The model provides a reasonable fit to the cross-power spectrum ($\chi^2 = 16.9$ for $12 - 2 = 10$ degrees of freedom) and allows us to effectively constrain the level of CIB leakage into the y-map.

The second null test we perform is designed to search for residual contamination from the Galaxy. If Galactic emission (e.g., synchrotron, free-free, or dust) leaks into both the y-map and the CMB lensing potential map, it could produce a spurious correlation signal in our analysis. We investigate this possibility by running our entire analysis pipeline described thus far — including the combination with the point source mask, the ILC y reconstruction, the cross-power spectrum estimation, and the CIB correction — on two additional sky masks, one more conservative than our fiducial choice, and one more aggressive. These masks are constructed by thresholding the 857 GHz map in the same manner as our fiducial $f_{\rm sky} = 0.3$ mask (see Section 3.4.1); the conservative mask has $f_{\rm sky} = 0.2$, while the aggressive mask

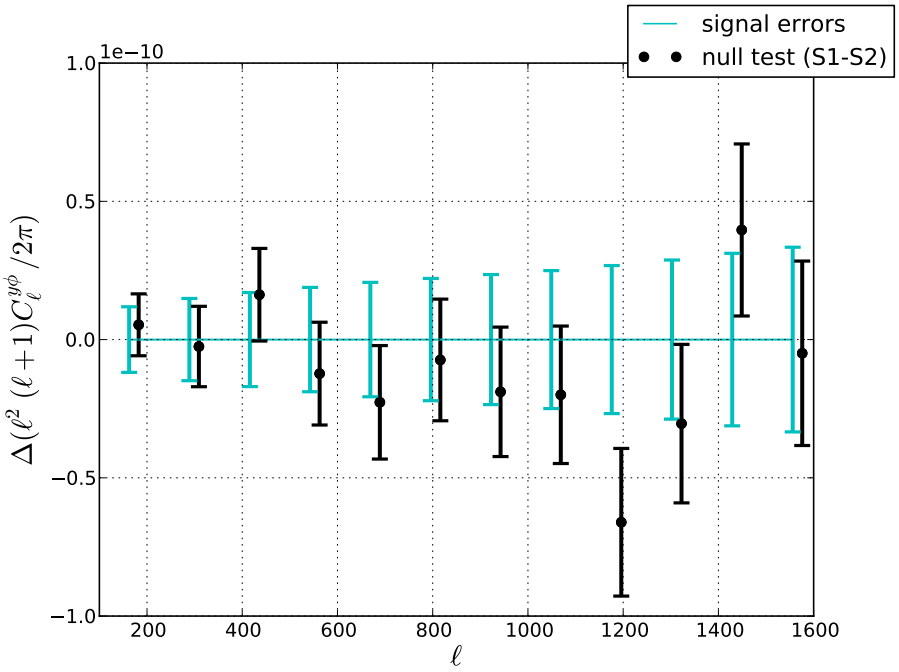

Figure 3.12: Cross-power spectrum between the null Compton y-map and the CMB lensing potential map (black points). The null y-map is constructed by subtracting (rather than co-adding) the S1 and S2 season y-maps output by our ILC pipeline in Section 3.4.2. The error bars, shown in black, are computed by applying Eq. (3.30) to the null y-map and CMB lensing potential map. The result is consistent with a null signal: $\chi^2 = 13.2$ for 12 degrees of freedom. The cyan error bars show the 1σ errors on our fiducial tSZ – CMB lensing cross-spectrum (i.e., the error bars on the blue squares in Fig. 3.10), which allow a quick assessment of the possible importance of any systematic effect in our measurement.

has $f_{\rm sky} = 0.4$. As seen in the left panel of Fig. 3.11, the Galactic emission varies strongly as a function of masking, and thus if our results are significantly contaminated by the Galaxy, a strong variation in the signal should be seen. The results of these tests are shown in Fig. 3.13 — the left panel shows the difference between $\hat{C}_\ell^{y\phi}$ computed for $f_{\rm sky} = 0.2$ and for the fiducial $f_{\rm sky} = 0.3$, while the right panel shows the difference between the $f_{\rm sky} = 0.4$ result and $f_{\rm sky} = 0.3$ result. In each case, the error bars shown in black are the errors on the $f_{\rm sky} = 0.2$ or 0.4 analyses, while the cyan errors are those for the fiducial $f_{\rm sky} = 0.3$ case. No significant variation is seen for either $f_{\rm sky} = 0.2$ or $f_{\rm sky} = 0.4$. For the former, we find $\chi^2 = 9.7$ (12 degrees of freedom) when comparing the difference spectrum to a null signal; for the latter, we find $\chi^2 = 2.8$ (12 degrees of freedom). The χ^2 value for the $f_{\rm sky} = 0.4$ test

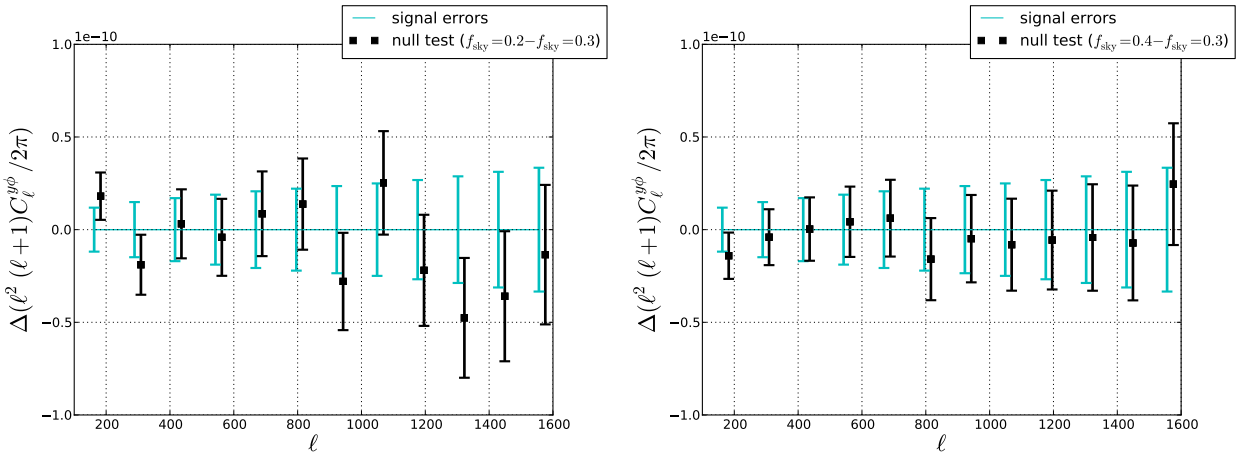

Figure 3.13: Difference between the tSZ – CMB lensing cross-power spectrum measured using our fiducial $f_{\rm sky} = 0.3$ Galactic mask and using an $f_{\rm sky} = 0.2$ mask (black squares in left panel) or $f_{\rm sky} = 0.4$ mask (black squares in right panel). The error bars on the black points are computed by applying Eq. (3.30) to the $f_{\rm sky} = 0.2$ or 0.4 case, i.e., using the same procedure as for the fiducial $f_{\rm sky} = 0.3$ analysis. In both cases, the difference between the fiducial cross-power spectrum and the $f_{\rm sky}$ variation cross-power spectrum is consistent with null: $\chi^2 = 9.7$ (12 degrees of freedom) for the $f_{\rm sky} = 0.2$ case, and $\chi^2 = 2.8$ (12 degrees of freedom) for the $f_{\rm sky} = 0.4$ case. The errors on the fiducial signal are shown in cyan, as in Fig. 3.12. Note that the entire pipeline described in Sections 3.4 and 3.5, including the CIB correction, is re-run to obtain the $f_{\rm sky} = 0.2$ and 0.4 results. The ILC weights for each sky cut are given in Table 3.2.

is somewhat low because the error bars on $\hat{C}_\ell^{y\phi}$ do not decrease significantly as $f_{\rm sky}$ increases, likely because the foreground-dominated noise in the y-map does not decrease, which more than compensates for the beneficial effect of having more modes in the map. Regardless, the null test is passed in both the $f_{\rm sky} = 0.2$ and 0.4 cases. For reference, we also provide the ILC channel map weights derived in each case in Table 3.2. The weights change slightly as $f_{\rm sky}$ is increased in order to compensate for the increased Galactic dust emission, but are generally fairly stable around the fiducial $f_{\rm sky} = 0.3$ case.

The third null test we perform uses an alternative method for calculating the CIB correction to the tSZ – CMB lensing measurement. Our standard approach is based on cross-correlating the ILC y-map with the 857 GHz map, as described in Section 3.5.3. As an alternative, we consider a method based on the best-fit model to the CIB – CMB lensing cross-power spectrum measurements obtained in [48]. We simply weight their $C_\ell^{\rm CIB\times\phi}$

$f_{\rm sky}$	100 GHz	143 GHz	217 GHz	353 GHz	545 GHz
0.2	0.41782	−0.93088	0.50931	0.00701	−0.00325
0.3 (fiducial)	0.45921	−1.00075	0.54397	0.00032	−0.00275
0.4	0.51380	−1.11722	0.62638	−0.02113	−0.00182

Table 3.2: Channel map weights computed by our ILC pipeline for three different sky cuts (see Eqs. (3.22) and (3.26)). All values are in units of $\mathrm{K}_{\rm CMB}^{-1}$. The fiducial case used throughout this chapter is $f_{\rm sky} = 0.3$. Note that these $f_{\rm sky}$ values are computed by simply thresholding the 857 GHz map (see Section 3.4.1); the final sky fractions used in each case are smaller than these values due to the inclusion of the point source mask and S1 and S2 season map masks. The weights change slightly as $f_{\rm sky}$ is increased as the ILC adjusts to remove the increased Galactic dust emission.

result at each HFI frequency between 100 and 545 GHz by the ILC relevant weight computed with our fiducial $f_{\rm sky} = 0.3$ pipeline (as given in Table 3.2), and then sum the results to obtain the total contamination to our $C_\ell^{y\phi}$ measurement. We work with the model fitting results (in particular, the "j reconstruction" results shown in Fig. 15 of [48][10]) rather than the direct measurements of $C_\ell^{\mathrm{CIB}\times\phi}$ at each frequency because the noise in the 100 and 143 GHz measurements renders them too imprecise for our purposes. We compute and subtract the CIB contamination from our initial tSZ – CMB lensing cross-spectrum measurement (the red circles in Fig. 3.10) to obtain a CIB-corrected measurement independent of our standard CIB-correction pipeline. Fig. 3.14 shows the difference between the CIB-corrected cross-spectrum obtained using this alternative procedure and the CIB-corrected cross-spectrum obtained using our standard procedure (described in Section 3.5.3). The errors on the alternative-correction points do not include uncertainty from the subtraction procedure, as we do not have the full errors on the CIB model fits from [48], but this is likely a very small correction. The difference between the cross-spectra estimated from the two subtraction procedures is consistent with null: $\chi^2 = 9.4$ for 12 degrees of freedom. A slightly significant difference is seen between the cross-spectra in the lowest two ℓ-bins; however, Fig. 15 in [48] indicates that the $C_\ell^{\mathrm{CIB}\times\phi}$ model fits severely underpredict the measured results at 100, 143, and 217 GHz for these low ℓ values. Modest variations in the low-ℓ

[10]We are grateful to O. Doré for providing the CIB – CMB lensing model fits from [48].

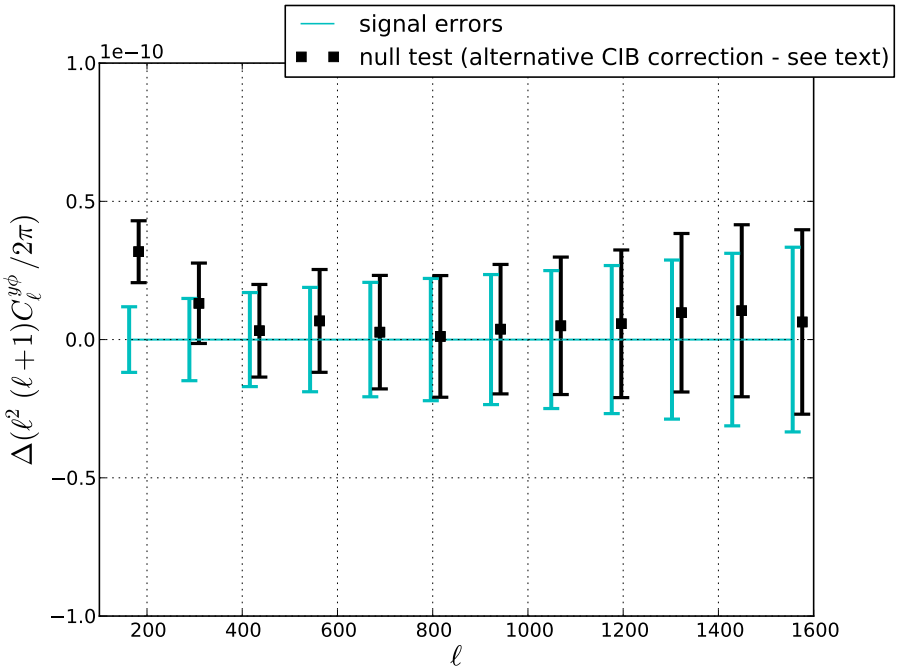

Figure 3.14: Difference between the tSZ – CMB lensing cross-power spectrum corrected for CIB contamination using our fiducial approach and corrected using an alternative approach based on the model fitting results presented in [48] (black squares). In the alternative approach, we simply weight the $C_\ell^{\mathrm{CIB}\times\phi}$ results at each HFI frequency by the relevant ILC weight computed by our Compton-y reconstruction pipeline (given in Table 3.2), and then sum the results to get the total leakage into the $C_\ell^{y\phi}$ measurement. The difference between our fiducial CIB-corrected cross-spectrum and the cross-spectrum corrected using this method is consistent with null: $\chi^2 = 9.4$ for 12 degrees of freedom. The slight disagreement at low-ℓ is likely due to discrepancies between the model fit from [48] and the measured CIB – CMB lensing cross-spectrum at 100, 143, and 217 GHz — see the text for discussion. The errors on our fiducial tSZ – CMB lensing cross-spectrum measurement are shown in cyan, as in Figs. 3.12 and 3.13.

$C_\ell^{\mathrm{CIB}\times\phi}$ values at these frequencies can easily explain the discrepancy seen in Fig. 3.14, and thus we do not consider it a cause for concern. The general consistency between this CIB subtraction procedure and our standard approach based on the y – 857 GHz cross-correlation gives us confidence that the signal in Fig. 3.10 is indeed the tSZ – CMB lensing cross-power spectrum.

These three null tests address the primary possible systematic errors in our measurement. However, in principle, any component of the CMB temperature field that produces a non-zero $\langle T(\ell)T(\ell-\ell')T(-\ell')\rangle$ could contaminate our measurement. Fortunately, all such bispectra

have been examined in detail in [48] for each of the HFI channels. One possible worry is contamination from extragalactic point sources, either due to radio emission at the lower HFI frequencies or infrared dust emission at the higher HFI frequencies. The contamination of this shot-noise bispectrum to the measurement of $C_\ell^{\mathrm{CIB}\times\phi}$ is shown in Fig. 11 of [48], where it can be seen that the induced bias is completely negligible at $\ell \lesssim 1500$ for every HFI frequency. Thus, we neglect it for the purposes of our analysis, which essentially lies completely below this multipole. Similarly, we neglect any contamination from the clustered CIB bispectrum (recently measured in [57, 49]), as this was also found to be negligible in [48], even after a dedicated search for its effects. For such a bispectrum to affect our results, the CIB emission would have to leak into both the y-map and CMB lensing potential map; the ILC used to construct the y-map should suppress such leakage below the raw level estimated in [48] for the $C_\ell^{\mathrm{CIB}\times\phi}$ measurements.

The integrated Sachs-Wolfe (ISW) effect in the CMB presents another possible source of contamination, as it traces the same large-scale matter density fluctuations as the CMB lensing potential [85]. Crucially, however, the ISW-induced fluctuations in the CMB have the same frequency dependence as the primordial CMB fluctuations, i.e., a blackbody spectrum. Since our ILC y-map construction explicitly removes all signal with a CMB blackbody spectrum (see Section 3.4.2), no ISW – CMB lensing correlation should contaminate our measurement of $C_\ell^{y\phi}$.

A final worry is the direct leakage of tSZ signal into the CMB lensing potential reconstruction, which would of course correlate with the tSZ signal in the ILC y-map. This leakage was recently considered in [86, 87], and we utilize their results here. In particular, it is shown in [87] that the bias induced in the CMB lensing power spectrum by tSZ leakage is $\lesssim 10\%$ for a completely unmasked reconstruction over the multipole range we consider (the bias is more problematic at high multipoles, e.g., $\ell > 2000$). However, this is the bias in the CMB lensing auto-power spectrum; the bias in the lensing reconstruction itself is much smaller ($\lesssim \sqrt{10}\%$). Furthermore, given that tSZ-detected clusters have been masked in the

143 GHz Planck CMB lensing reconstruction, and given that we work with the publicly released minimum-variance combination of the 143 and 217 GHz lensing reconstructions, the tSZ-induced bias in the CMB lensing potential map should be completely negligible for the purposes of our analysis. Follow-up analyses with detailed simulations (such as those used in [87]) will easily test this assumption.

Having passed a number of null tests and demonstrated that essentially all non-CIB sources of systematic error are negligible, we are confident in the robustness of our detection of the tSZ – CMB lensing cross-power spectrum. The blue squares in Fig. 3.10 show our final CIB-corrected results, which are also presented in Table 3.1 with the associated uncertainties. The final detection significance is 6.2σ.

3.6 Interpretation

As demonstrated in Section 3.3, the tSZ – CMB lensing cross-correlation signal is a probe of both the background cosmology (primarily via σ_8 and Ω_m) and the ICM physics in the groups and clusters responsible for the signal. In this section, we first fix the cosmological parameters and interpret our measurement in terms of the ICM physics. We then fix the ICM physics to our fiducial model and constrain the cosmological parameters. Finally, we demonstrate a method to simultaneously constrain the ICM and cosmology by combining our measurement with the Planck measurement of the tSZ auto-power spectrum.

3.6.1 ICM Constraints

Fixing the background cosmological parameters to their WMAP9 values (in particular, $\sigma_8 = 0.817$ and $\Omega_m = 0.282$, i.e., our fiducial model), we derive constraints on ICM physics models from our measurement of the tSZ – CMB lensing cross-power spectrum. The ICM models are described in Section 3.3, and include our fiducial "AGN feedback" model, the "adiabatic" model from [29], the UPP of Arnaud et al. [62], for which we consider various

values of the HSE mass bias $(1-b)$ (see Section 3.3), and the simulation results of [63]. The results are shown in Fig. 3.15, in which we have re-binned the points from Fig. 3.10 for visual clarity (note that all fitting results and χ^2 values are computed using the 12 bins in Fig. 3.10). We find that the fiducial model is the best fit to the data of the five models shown in Fig. 3.15, with $\chi^2 = 14.6$ for 12 degrees of freedom. The adiabatic model is in some tension with our measurement, with $\chi^2 = 17.9$. Fig. 3.15 also shows three versions of the UPP with differing amounts of HSE mass bias: $(1-b) = 0.9$ is a reasonable fit to the data ($\chi^2 = 15.7$), while $(1-b) = 0.8$ is in some tension ($\chi^2 = 17.8$). Perhaps most interestingly, $(1-b) = 0.55$ is in serious tension with our measurement, with $\chi^2 = 25.1$, i.e., $\Delta\chi^2 = 10.5$ with respect to our fiducial model. This value of $(1-b)$ was found to be necessary in [12] in order to reconcile the Planck SZ cluster count-derived cosmological parameters with those derived in the Planck primordial CMB analysis. In the context of a fixed WMAP9 background cosmology, our results highly disfavor such an extreme value of the HSE mass bias for the UPP, although we are probing a lower-mass and higher-redshift subset of the cluster population than that used in the counts analysis. Finally, the cross-spectrum derived from the Sehgal et al. simulation [63] lies somewhat high but is a reasonable fit to the data, with $\chi^2 = 15.4$. Note that this simulation result has been rescaled to our fiducial WMAP9 cosmological parameters, as described in Section 3.3.

If we instead fix the background cosmology to the Planck + WMAP polarization parameters from [3] ($\sigma_8 = 0.829$ and $\Omega_m = 0.315$), we obtain similar results. In this case, the UPP with $(1-b) = 0.9$ is slightly preferred ($\chi^2 = 14.4$) over the fiducial "AGN feedback" model ($\chi^2 = 15.0$). The UPP with $(1-b) = 0.8$ is also a reasonable fit, with $\chi^2 = 15.3$. The "adiabatic" model is now highly disfavored, with $\chi^2 = 28.7$, while the Sehgal et al. simulation is also in tension, with $\chi^2 = 20.7$. Most strikingly, the UPP with $(1-b) = 0.55$ is still in serious tension with the data, with $\chi^2 = 22.4$. Thus, regardless of the choice of background cosmology (WMAP9 or Planck), we do not find evidence for a value of the HSE mass bias in excess of the standard values predicted in numerical simulations, i.e., $(1-b) \approx 0.8$–0.9.

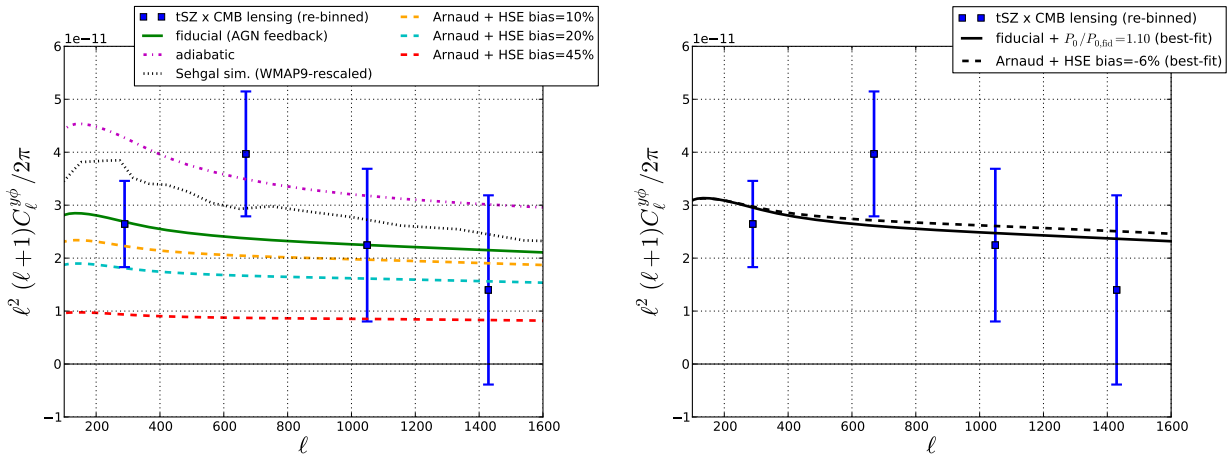

Figure 3.15: ICM model comparison results. In both panels, the blue squares show a re-binned version of the tSZ – CMB lensing cross-power spectrum results presented in Fig. 3.10. The re-binning was performed using inverse-variance weighting, and is only done for visual purposes; all model fitting results are computed using the twelve bins shown in Fig. 3.10. The left panel shows a comparison between the predictions of several ICM physics models, all computed using a WMAP9 background cosmology. The fiducial "AGN feedback" model [29] is the best fit, with $\chi^2 = 14.6$ for 12 degrees of freedom (see the text for comparisons to the other models shown here). The right panel shows the result of a fit for the amplitude of the pressure–mass relation P_0 with respect to its standard value in our fiducial model (solid black curve), as well as the result of a fit for the HSE mass bias in the UPP of Arnaud et al. [62] (dashed black curve). Our measurements are consistent with the fiducial value of P_0 ($P_0/P_{0,\mathrm{fid}} = 1.10 \pm 0.22$) in our standard "AGN feedback" model, and similarly with standard values of the HSE mass bias for the UPP, $(1-b) = 1.06^{+0.11}_{-0.14}$ (all results at 68% C.L.).

In addition to these general ICM model comparison results, we can use our tSZ – CMB lensing measurements to fit an overall amplitude parameter for a given ICM model, holding all of the other parameters in the model (and the background cosmology) fixed. For the fiducial "AGN feedback" model from [29], we fit the overall amplitude of the normalization of the pressure–mass relation (keeping its mass and redshift dependence fixed), $P_0/P_{0,\mathrm{fid}}$, where $P_{0,\mathrm{fid}} = 18.1$ is the fiducial amplitude (see [29] for full details). In the context of this model and a WMAP9 cosmology, we find

$$P_0/P_{0,\mathrm{fid}} = 1.10 \pm 0.22 \;\; (68\%\mathrm{C.L.}) \tag{3.38}$$

using our tSZ – CMB lensing cross-power spectrum results, with $\chi^2 = 14.4$ for 12 degrees of freedom. Thus, the fiducial model ($P_0/P_{0,\rm fid} = 1$) is consistent with our results. Similarly, we constrain the HSE mass bias $(1-b)$ in the UPP [62]. Holding all other UPP parameters fixed and using a WMAP9 cosmology, we find

$$(1-b) = 1.06^{+0.11}_{-0.14} \ (68\%\text{C.L.}) \tag{3.39}$$

with $\chi^2 = 14.4$ for 12 degrees of freedom. This constraint rules out extreme values of the HSE mass bias, if interpreted in the context of a fixed WMAP9 background cosmology. In particular, we find $0.60 < (1-b) < 1.38$ at the 99.7% confidence level.

For comparison, we also perform the same analysis using the Planck collaboration measurement of the tSZ auto-power spectrum [44] (see their Table 3 for the relevant data). For the fiducial "AGN feedback" model and a WMAP9 cosmology, we find

$$P_0/P_{0,\rm fid} = 0.80 \pm 0.05 \ (68\%\text{C.L.}) \tag{3.40}$$

using the Planck tSZ auto-power spectrum results, with $\chi^2 = 2.5$ for 16 degrees of freedom (the low value of χ^2 could be an indication that the additional error due to uncertainties in residual foreground power subtraction in [44] could be too large) . The fiducial model ($P_0/P_{0,\rm fid} = 1$) is in some tension with this constraint. However, the fiducial ICM model can be brought within the 2σ allowed region (i.e., $P_0/P_{0,\rm fid} = 0.90 \pm 0.05$) if one fixes $\sigma_8 = 0.793$ rather than $\sigma_8 = 0.817$, which is within the 2σ allowed region of the WMAP9+eCMB+BAO+H_0 result [1]. This is a clear manifestation of the usual ICM–cosmology degeneracy in tSZ constraints (see Section 3.6.3 for a method to break this degeneracy). For the HSE mass bias $(1-b)$ in the UPP, we find (again, for a fixed WMAP9 cosmology)

$$(1-b) = 0.78^{+0.03}_{-0.04} \ (68\%\text{C.L.}) \tag{3.41}$$

with $\chi^2 = 2.2$ for 16 degrees of freedom (again, the low χ^2 could be an indication that the foreground subtraction errors are over-estimated). This result, as in Eq. (3.39), rules out extreme values of $(1 - b)$, although it is more dependent on the assumption of the fixed background WMAP9 cosmology, as the tSZ auto-power spectrum depends more sensitively on σ_8 and Ω_m than the tSZ – CMB lensing cross-power spectrum. Overall, the constraints in Eqs. (3.38) and (3.39) from our measurement of the tSZ – CMB lensing cross-power spectrum are consistent with those obtained in Eqs. (3.40) and (3.41) from the Planck collaboration measurement of the tSZ auto-power spectrum. The lower values favored by the latter are a reflection of the low amplitude of the measured tSZ auto-power spectrum compared to fiducial theoretical expectations. The low amplitude could be due to ICM physics, cosmology, or observational systematics — this is discussed in more detail in the next section.

The slightly high amplitude of the cross-spectrum results in Eqs. (3.38) and (3.39) could be explained in many different ways. For example, a small amount of residual CIB contamination could be present in the tSZ – CMB lensing cross-correlation measurement, though clearly neither constraint provides statistically significant evidence for this claim. Alternatively, the results could suggest that perhaps our model has not fully accounted for all of the hot, ionized gas present in the universe — in particular, highly diffuse, unbound gas is not included in the halo model computations described in Section 3.3. However, it is worth noting that in the simulations from which our fiducial ICM model is extracted, the baryon fraction within $r_{\rm vir}$ is found to be $\approx 80\%$ (at the group scale, i.e, $M_{200c} \approx 10^{14}\, M_\odot/h$) to 95% (at the massive cluster scale, i.e., $M_{200c} \approx 10^{15}\, M_\odot/h$) of the cosmic mean, with the remainder of the gas located within 2–3$r_{\rm vir}$ (see Fig. 10 in [90]). Thus, the "missing baryons" in this model constitute only a small fraction of the cosmic baryon budget. The fact that the measured cross-power spectrum at low-ℓ lies slightly above the halo model predictions computed using our fiducial model might suggest a signal from these "missing baryons", but our results provide no statistically significant support for this claim. We do clearly measure both the one- and two-halo terms in the tSZ – CMB lensing cross-power spectrum

(see Fig. 3.10), but by no means does this imply that the measurement shows signatures of diffuse, unbound gas. Comparisons to full cosmological hydrodynamics simulations (such as those used to extract the pressure profile used in our fiducial model [29]) will be needed to assess the relative importance of halo-bound gas and diffuse, unbound gas in the tSZ – CMB lensing cross-power spectrum, which likely has a strong angular scale dependence. On the largest angular scales at which we measure the tSZ – CMB lensing cross-power spectrum ($\ell \sim 100$, or, $\theta \sim 1^\circ$), the results are sensitive to the correlation between ionized gas and dark matter on comoving distance scales as large as $\sim 50\,\mathrm{Mpc}/h$ (see the kernel in the left panel of Fig. 3.3, which receives significant contributions at redshifts as large as $z \sim 2$). In the halo model framework in which we interpret the measurement, these large-scale correlations physically arise from supercluster-sized dark matter structures that host many group- and cluster-sized collapsed halos. The large, linear-regime dark matter structures imprint the CMB lensing distortions, while the hot electrons sourcing the tSZ signal are bound in the collapsed halos. Whether the measured correlation also implies the existence of diffuse, unbound gas located outside of these collapsed halos is a question that will require further interpretation with simulations.

3.6.2 Cosmological Constraints

Fixing the ICM physics prescription to our fiducial "AGN feedback" model [29], we derive constraints on the cosmological parameters σ_8 and Ω_m using our measurement of the tSZ – CMB lensing cross-power spectrum. All other cosmological parameters are held fixed to their WMAP9+eCMB+BAO+H_0 maximum-likelihoood values [1]. Note that we include the tSZ – CMB lensing angular trispectrum (Eq. (3.16)) in the covariance matrix when deriving these constraints, as discussed in Section 3.3.3. The error bars given in Eq. (3.30) thus include an additional term (e.g., [39, 37]):

$$\left(\Delta\hat{C}_\ell^{y\phi}\right)^2 = \frac{1}{f_{\mathrm{sky},2}}\frac{1}{(2\ell+1)\Delta\ell}\left(\hat{C}_\ell^{yy}\hat{C}_\ell^{\phi\phi} + \left(C_\ell^{y\phi}\right)^2\right) + \frac{1}{4\pi f_{\mathrm{sky},2}}T_{\ell\ell}^{y\phi}\,. \tag{3.42}$$

The trispectrum term is highly subdominant to the noise arising from the y and ϕ auto-spectra: for our fiducial ICM model and WMAP9 cosmology, the second term in Eq. (3.42) is $\lesssim 0.2\%$ of the first term. However, we include it here for completeness. Since the covariance matrix is now a function of cosmological parameters (i.e., $T_{\ell\ell}^{y\phi}$ depends on σ_8 and Ω_m), simple χ^2 minimization is no longer the correct procedure to use when deriving constraints (e.g., [88]). The likelihood function is:

$$\mathcal{L}(\sigma_8, \Omega_m) = \frac{1}{\sqrt{(2\pi)^{N_b}\det(\mathrm{Cov})}} \exp\left(-\frac{1}{2}(C_\ell^{y\phi} - \hat{C}_\ell^{y\phi})(\mathrm{Cov}^{-1})_{\ell\ell'}(C_{\ell'}^{y\phi} - \hat{C}_{\ell'}^{y\phi})\right), \quad (3.43)$$

where $N_b = 12$ is the number of bins in the measurement and $\mathrm{Cov}_{\ell\ell'} = \left(\Delta\hat{C}_\ell^{y\phi}\right)^2 \delta_{\ell\ell'}$ is the covariance matrix, with $\left(\Delta\hat{C}_\ell^{y\phi}\right)^2$ given in Eq. (3.42). Both $C_\ell^{y\phi}$ and $\mathrm{Cov}_{\ell\ell'}$ are computed as a function of σ_8 and Ω_m.

Using the likelihood in Eq. (3.43), we compute constraints on σ_8 and Ω_m. We most tightly constrain the combination $\sigma_8(\Omega_m/0.282)^{0.26}$:

$$\sigma_8\left(\frac{\Omega_m}{0.282}\right)^{0.26} = 0.824 \pm 0.029 \ \ (68\%\mathrm{C.L.}). \quad (3.44)$$

Note that here and elsewhere throughout the chapter we report the best-fit value as the mean of the marginalized likelihood, while the lower and upper error bounds correspond to the 16% and 84% points in the marginalized cumulative distribution, respectively. Marginalizing the likelihood in Eq. (3.43) over Ω_m, we obtain $\sigma_8 = 0.826^{+0.037}_{-0.036}$ at 68% C.L. If we instead marginalize over σ_8, we find $\Omega_m = 0.281 \pm 0.033$ at 68% C.L.

We also re-analyze the Planck collaboration measurement of the tSZ auto-power spectrum presented in [44]. We use the bandpowers and (statistical + foreground) errors given in Table 3 of [44]. However, we also include the tSZ angular trispectrum in the covariance matrix, calculated via an expression analogous to Eq. (3.16). The trispectrum contribution was neglected (or at least not discussed) in [44], but it is non-negligible on large angular scales, as shown in Chapter 2 (see Fig. 2.19). For our fiducial ICM model and WMAP9 cosmology,

we find that the tSZ trispectrum term is as much as ~ 50 times larger than the statistical + foreground error term in the covariance matrix (in terms of the error bar itself, this translates to a factor of $\sqrt{50} \approx 7$); this maximum occurs for the fourth lowest multipole bin shown in Fig. 3.9, centered at $\ell \approx 52$, where the measurement errors are fairly small but the trispectrum contribution is enormous. This effect is highly ℓ-dependent: for the highest multipole bin shown in Fig. 3.9, the trispectrum term is only $\approx 1\%$ as large as the statistical + foreground error term in the covariance matrix, and is essentially negligible. Note that the tSZ trispectrum contribution to the covariance matrix can be heavily suppressed by masking nearby, massive clusters when measuring the tSZ power spectrum [80, 37]; however, since this procedure was not used in [44], we do not apply it either, in order to obtain results that can be directly compared with those from [44]. It would likely be beneficial to consider such a masking procedure in future large-angular scale measurements of the tSZ power spectrum.

In addition to our inclusion of the tSZ trispectrum term in the covariance matrix, our re-analysis of the Planck collaboration measurement of the tSZ auto-power spectrum also differs in our choice of fiducial ICM model — as mentioned throughout this chapter, we use the "AGN feedback" model from [29], whereas [44] used the UPP of [62] with a fixed HSE mass bias $(1-b) = 0.80$.

Implementing a likelihood for the tSZ power spectrum analogous to that in Eq. (3.43), we compute constraints on σ_8 and Ω_m. Again, we find that the combination $\sigma_8(\Omega_m/0.282)^{0.26}$ is most tightly constrained:

$$\sigma_8 \left(\frac{\Omega_m}{0.282}\right)^{0.26} = 0.771^{+0.012}_{-0.011} \ (68\%\text{C.L.}). \tag{3.45}$$

This result is consistent with that presented in [44], although the Planck team found that a slightly different parameter combination was best constrained, obtaining $\sigma_8 \left(\Omega_m/0.28\right)^{0.395} = 0.784 \pm 0.016$. We obtain a slightly lower value due to our choice of ICM model with which to interpret the data — our model (from [29]) predicts slightly more tSZ power than the

UPP+20% HSE bias model used in [44], and thus requires a lower value of σ_8 to fit the data. It is also worth noting that the errors bars in Eq. (3.45) are slightly smaller than those from [44], despite our addition of the tSZ trispectrum to the covariance matrix. This result is due to the fact that the likelihood in [44] included two additional foreground parameters (for point sources and clustered CIB), which were marginalized over in the determination of cosmological constraints. There is also possibly a small difference due to the fact that the trispectrum provides another cosmology-dependent quantity in the likelihood function in addition to the power spectrum itself (see [88] for a discussion of this effect in the context of the weak lensing power spectrum).

Comparison of the constraints presented in Eqs. (3.44) and (3.45) clearly indicates that the tSZ auto-power spectrum measurement from [44] favors a lower value of this parameter combination than our measurement of the tSZ – CMB lensing cross-power spectrum. The most plausible explanation for this result involves systematics, in particular the subtraction of residual power due to leakage of CIB and IR and radio point sources into the y-map in [44]. As can be seen by eye in Fig. 3.9, the unsubtracted (blue) points agree quite well with our fiducial theoretical model over a fairly wide multipole range ($50 \lesssim \ell \lesssim 300$), although contamination is clearly present at higher multipoles. The residual power from non-tSZ contaminants is assessed with simulations in [44], and an additional 50% error is included in the final results due to uncertainties in the subtraction of this power. If the subtraction procedure has removed too much power, clearly $\sigma_8 \, (\Omega_m/0.282)^{0.26}$ will be biased low. A detailed characterization of all the foregrounds is needed to carefully assess this issue, but a rough estimate (assuming that the tSZ auto-power spectrum amplitude scales as σ_8^8) indicates that a value of $\sigma_8 \approx 0.80$ would be obtained if the foreground power subtraction has been over-estimated by $\approx$ 20–25%. Further study of the CIB, Galactic dust, and point source power spectra will be useful to clarify the accuracy of the subtraction procedure. Along the same lines, it is also possible that our result from the tSZ – CMB lensing cross-power spectrum is biased slightly high due to residual CIB emission in the y-map; however,

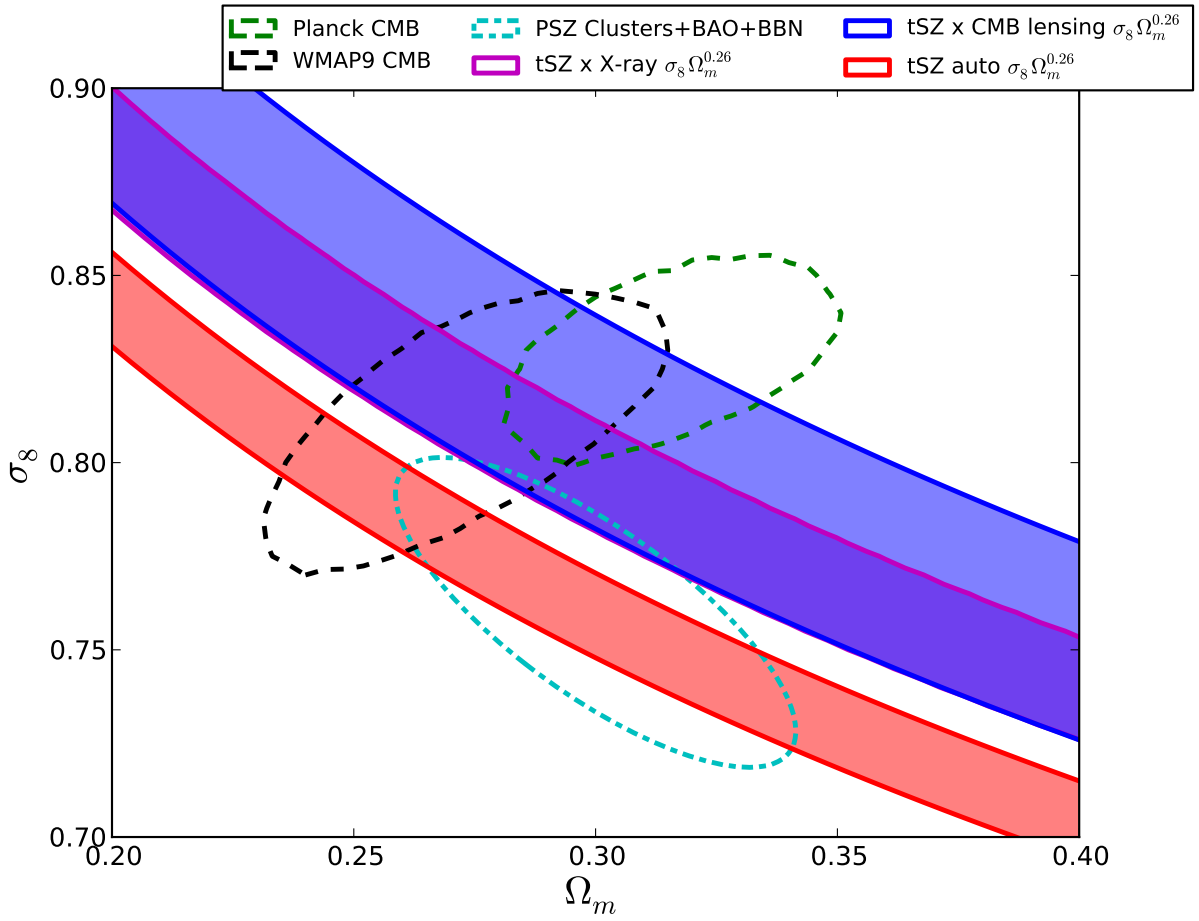

Figure 3.16: Constraints on σ_8 and Ω_m from the tSZ – CMB lensing cross-power spectrum measurement presented in this work (blue shaded), the tSZ auto-power spectrum measurement from [44] re-analyzed in this work (red shaded), the Planck+WMAP polarization CMB analysis [3] (green dashed), WMAP9 CMB analysis [1] (black dashed), number counts of Planck tSZ clusters [12] (cyan dashed), and a recent cross-correlation of tSZ signal from Planck and WMAP with an X-ray "δ"-map based on ROSAT data [67] (magenta shaded). All contours shown are 68% confidence intervals only (for visual clarity). See the text for further discussion.

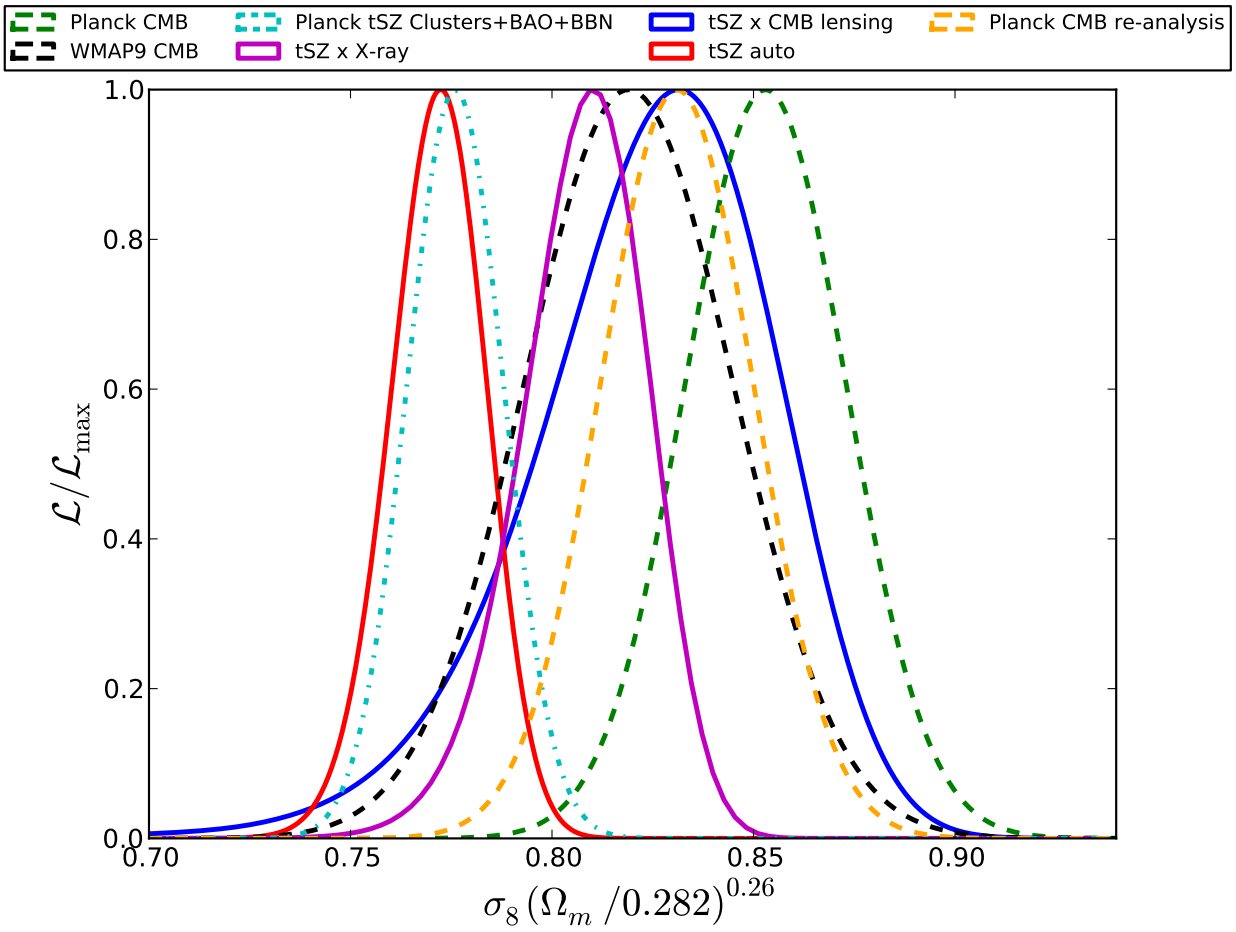

Figure 3.17: Normalized one-dimensional likelihoods for the degenerate parameter combination $\sigma_8(\Omega_m/0.282)^{0.26}$ that is best constrained by the tSZ auto- and cross-power spectra measurements. The probes shown are identical to those in Fig. 3.16, except for an additional result from a recent re-analysis of the Planck CMB + WMAP polarization data [93]. See the text for further discussion.

the thorough systematics analysis in Section 3.5.4 gives us confidence that we have treated this issue with sufficient precision.

The other possible explanation for the discrepancy between Eqs. (3.44) and (3.45) lies in the physics of the ICM. However, this explanation requires an unexpected mass dependence for the relative importance of non-thermal pressure support or feedback effects in the ICM (i.e., the HSE mass bias). Recall that the tSZ – CMB lensing cross-power spectrum is sourced by lower mass, higher redshift objects than the tSZ auto-power spectrum (see Section 3.3). The low value of Eq. (3.45) compared to Eq. (3.44) implies that a larger value

of the HSE mass bias is required to reconcile σ_8 and Ω_m inferred from the tSZ auto-power spectrum measurement from [44] with the WMAP9 or Planck values of these parameters than is required for the measurement of the tSZ – CMB lensing cross-power spectrum presented in this chapter. Eq. (3.39) directly demonstrates that our results do not require an extreme value of the HSE mass bias in order to remain consistent with a WMAP9 cosmology. However, given the mass- and redshift-dependences of the tSZ auto- and tSZ – CMB lensing cross-spectra, Eqs. (3.39) and (3.41) imply that the HSE mass bias (or the influence of non-thermal ICM physics) is larger for more massive, lower redshift groups and clusters. This mass dependence is the opposite of the general expectation from simulations or simple theoretical arguments (e.g., [68]), which posit that the most massive clusters in the universe are dominated by their gravitational potential energy, with energetic feedback and non-thermal effects playing a smaller role than in less massive systems (but see [89] for an effect that does have this type of mass dependence). Thus, it seems more plausible that the discrepancy between Eqs. (3.44) and (3.45) arises from systematic effects such as those discussed in the previous paragraph.

We present these constraints on σ_8 and Ω_m and compare them to results from various cosmological probes in Figs. 3.16 and 3.17. We show results from the following analyses:

- Planck+WMAP polarization+high-ℓ (ACT+SPT) CMB analysis [3]: $\sigma_8 = 0.829 \pm 0.012$ and $\Omega_m = 0.315^{+0.016}_{-0.018}$, or, $\sigma_8(\Omega_m/0.27)^{0.3} = 0.87 \pm 0.02$;

- WMAP9 CMB-only analysis (note that this is CMB-only, and hence the best-fit parameters are slightly different than our fiducial model, which is WMAP9+eCMB+BAO+H_0) [1]: $\sigma_8 = 0.821 \pm 0.023$ and $\Omega_m = 0.279 \pm 0.025$;

- Planck tSZ cluster number counts + BAO + BBN [12]: $\sigma_8(\Omega_m/0.27)^{0.3} = 0.782 \pm 0.010$, or, $\sigma_8 = 0.77 \pm 0.02$ and $\Omega_m = 0.29 \pm 0.02$;

- Cross-correlation of tSZ signal from Planck and WMAP with an X-ray "δ"-map based on ROSAT data [67]: $\sigma_8(\Omega_m/0.30)^{0.26} = 0.80 \pm 0.02$, i.e., $\sigma_8(\Omega_m/0.282)^{0.26} = 0.81 \pm 0.02$;

- Re-analysis of Planck+WMAP polarization CMB [93] (Fig. 3.17 only): $\sigma_8(\Omega_m/0.30)^{0.26} = 0.818 \pm 0.019$, i.e., $\sigma_8(\Omega_m/0.282)^{0.26} = 0.831 \pm 0.019$.

The tSZ auto-power spectrum results shown are those derived in our re-analysis of the measurements from [44], as detailed above. Also, the contours shown for the tSZ auto- and cross-power spectra in Fig. 3.16 assume that only a single parameter ($\sigma_8 \, (\Omega_m/0.282)^{0.26}$) is fit to the data. Note that all contours shown in Fig. 3.16 are 68% confidence intervals only. Fig. 3.16 indicates that there is general concordance between the various probes, with the exception of the tSZ auto-power spectrum and tSZ cluster count constraints. The discordance of these two probes is more visually apparent in Fig. 3.17, which shows the one-dimensional constraints on $\sigma_8 \, (\Omega_m/0.282)^{0.26}$ from each probe. The low amplitude of the tSZ auto-power spectrum is discussed in the previous two paragraphs. A thorough investigation of the low σ_8 values inferred from the tSZ cluster counts is beyond the scope of this chapter, but it is possible that the ICM physics discussed in the previous paragraph is again responsible, or that systematic effects due to e.g. the complicated selection function arising from the inhomogeneous noise in the Planck maps are responsible. Theoretical uncertainties in the exponential tail of the mass function could also be an issue, as the clusters used in the counts analysis are much more massive (and hence rarer) than those which dominate the various tSZ statistics shown in Figs. 3.16 and 3.17.

Finally, it is worth mentioning that at some level the value of σ_8 inferred from low-redshift tSZ measurements must disagree with that inferred from primordial CMB measurements, due to the non-zero masses of neutrinos. Neutrino oscillation experiments imply a minimum mass for the most massive of the three neutrino species of ≈ 0.05 eV; thus the energy density in neutrinos is at least $\approx 1\%$ of the energy density in baryons [92]. While this contribution seems small, it is enough to reduce the number of $z > 1, 10^{15} \, M_\odot/h$ clusters by $\approx 25\%$ and yield a $\approx 1.5\%$ decrease in the value of σ_8 inferred from low-redshift large-scale structure measurements compared to that inferred from $z = 1100$ CMB measurements [91]. For example, if the true CMB-inferred value is $\sigma_8 = 0.817$ (WMAP9), then the true tSZ-

inferred value should be $\sigma_8 \approx 0.805$. While this effect does not completely explain the current discrepancy between the results shown in Figs. 3.16 and 3.17, it is worth keeping in mind as the statistical precision of these measurements continues to improve.

3.6.3 Simultaneous Constraints on Cosmology and the ICM

In the previous two sections, we have constrained ICM pressure profile models while fixing the background cosmology, and constrained cosmological parameters while holding the ICM model fixed. We now demonstrate the feasibility of simultaneously constraining both the ICM and cosmological parameters by combining our measurement of the tSZ – CMB lensing cross-power spectrum with the tSZ auto-power spectrum measurement from [44]. Heuristically, since the cross-power spectrum amplitude scales roughly as $\sigma_8^6 \Omega_m^{1.5} P_0$, while the auto-power spectrum amplitude scales roughly as $\sigma_8^8 \Omega_m^{2.1} P_0^2$, it is clear that by combining the different probes we can begin to break the cosmology–ICM degeneracy that currently hinders cosmological constraints from tSZ measurements. The parameter P_0 here represents the overall normalization of the pressure–mass relation, as discussed prior to Eq. (3.38). It is important to note that this approach to simultaneously constraining the ICM and cosmological parameters is only valid for an ICM model without scatter (we do not consider scatter in our prescription for the pressure profile as a function of halo mass and redshift; we simply implement the fitting function from the "AGN feedback" simulations of [29]). If scatter were included, the situation would be slightly more complicated, as the tSZ auto-spectrum would probe $\langle P_0^2 \rangle = P_0^2 + \sigma_{P_0}^2$, where $\sigma_{P_0}^2$ is the scatter in the normalization of the pressure–mass relation, whereas the cross-spectrum would probe $\langle P_0 \rangle = P_0$. For our purposes, $\sigma_{P_0}^2 = 0$, and hence both power spectra probe P_0. Including scatter would not invalidate this framework, but would simply require fitting for an additional parameter in the analysis[11]. This is likely a useful direction to pursue in future tSZ statistical analyses.

[11] We thank M. Becker and E. Rozo for emphasizing these points to us.

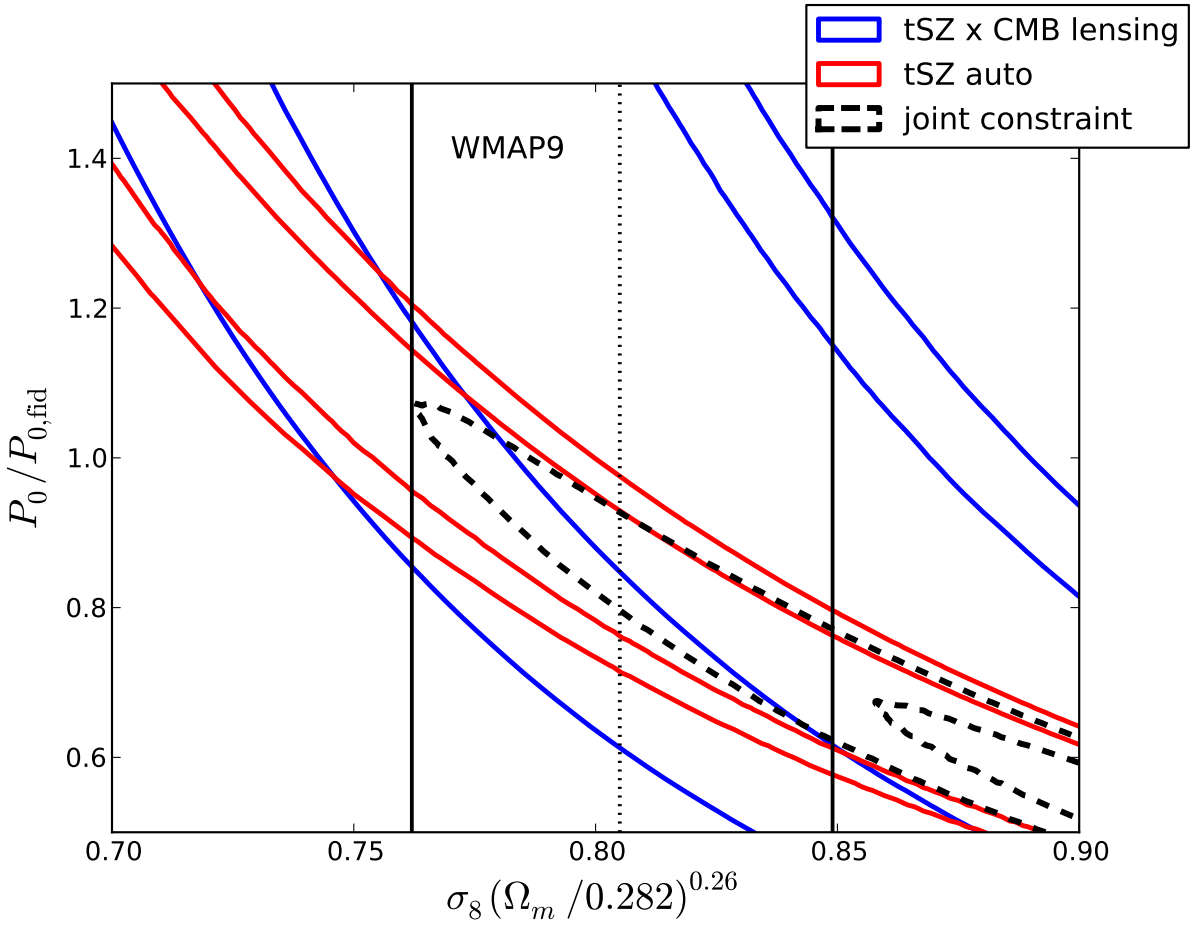

Figure 3.18: Constraints on $\sigma_8(\Omega_m/0.282)^{0.26}$ and P_0 (the amplitude of the pressure–mass relation normalized to its fiducial value in our ICM model [29]) from the tSZ – CMB lensing cross-power spectrum measurement presented in this work (blue), the tSZ auto-power spectrum measurement from [44] re-analyzed in this work (red), and the combination of the two probes (black dashed). The different dependences of the two probes on these parameters allow the degeneracy between the ICM and cosmology to be broken (albeit weakly in this initial application) using tSZ measurements alone. The contours show the 68% and 95% confidence intervals for the tSZ probes; in addition, the vertical lines show the 68% confidence interval for the WMAP9 constraint on this cosmological parameter combination, which is obviously independent of P_0. Standard values for the pressure–mass normalization ($P_0 \approx 0.8$–1.0) are compatible with all of the probes.

We implement a two-parameter likelihood for the joint constraint, consisting of P_0 and $\sigma_8(\Omega_m/0.282)^{0.26}$. Note that all other parameters in the ICM model are held fixed; only the overall normalization is varied. We also neglect any covariance between the tSZ – CMB lensing cross-power spectrum and the tSZ auto-power spectrum. Since the two signals are sourced by somewhat different populations of clusters (see Figs. 3.3–3.5) and the noise in the auto-power measurement is dominated by foreground uncertainties, this seems reasonable in the current analysis, but for future studies with higher SNR this covariance should be accurately quantified. The results of the joint analysis are shown in Fig. 3.18. Using either the tSZ – CMB lensing cross-power spectrum or tSZ auto-power spectrum alone, it is clear that the parameters are completely degenerate. However, by combining the two probes, the allowed region dramatically shrinks — the degeneracy is broken. At the current signal-to-noise levels for these measurements, the degeneracy-breaking power is fairly weak, but this situation will improve with forthcoming data from Planck, ACTPol, and SPTPol. As shown in Fig. 3.18, the current combination of the two measurements leads to a preference for very low values of P_0 and high values of $\sigma_8(\Omega_m/0.282)^{0.26}$. This result is driven by the low amplitude of the measured tSZ auto-power spectrum from [44], which was discussed in detail in the previous section. Given current uncertainties in the origin of this low amplitude (e.g., systematics due to residual foreground power subtraction), we refrain from quoting precise parameter constraints from this joint analysis. Regardless, the 95% confidence interval in Fig. 3.18 includes regions of parameter space consistent with both standard ICM models ($P_0 \approx 0.8$–1.0) and standard cosmological parameters from other probes ($\sigma_8(\Omega_m/0.282)^{0.26} \approx 0.8$). For clear visual comparison, Fig. 3.18 shows the 68% confidence interval for $\sigma_8(\Omega_m/0.282)^{0.26}$ from the WMAP9 CMB-only analysis, which overlaps significantly with the allowed region from the combination of the tSZ probes. Fig. 3.18 thus represents a clear proof-of-principle for this technique. By simultaneously analyzing multiple tSZ statistics, the ICM–cosmology degeneracy can be broken and robust constraints on both can be obtained [60, 61].

3.7 Discussion and Outlook

In this chapter we have presented the first detection of the correlation between the tSZ and CMB lensing fields. Based on the theoretical calculations presented in Section 3.3, this cross-correlation signal is sensitive to the ICM physics in some of the highest redshift and least massive groups and clusters ever probed. The measurement shows clear signatures of both the one- and two-halo terms, probing the correlation between ionized gas and dark matter over scales ranging from $\approx 0.1\,\mathrm{Mpc}/h$ at $z \approx 0.05$ to $\approx 50\,\mathrm{Mpc}/h$ at $z \approx 2$. Interpreting the measurement with halo model calculations, we do not find evidence for the presence of diffuse, unbound gas that lies outside of collapsed halo regions (the "missing baryons" — see discussion in Section 3.6.1). Further interpretation of these measurements using cosmological hydrodynamics simulations will be extremely useful.

An additional area in clear need of future improvement is the characterization of foreground contamination in the Compton-y map reconstruction and the estimation of the tSZ auto-power spectrum, with the tSZ – CIB correlation a particular point of concern. We make the y-map constructed in this work publicly available, but stress that it contains signal from a number of non-tSZ sources, including Galactic dust, CIB, and unresolved point sources. However, it may prove a useful resource for further cross-correlation studies, provided that contamination from these other sources can be properly understood and mitigated (as done in this analysis for the CIB contamination).

Our measurement of the tSZ – CMB lensing cross-power spectrum is consistent with extrapolations of existing ICM physics models in the literature. In particular, we find that our fiducial model (the "AGN feedback" model from [29]) is consistent with the data, assuming a WMAP9 background cosmology. We do not find evidence for extreme values of the HSE mass bias in the UPP model of [62], obtaining $0.60 < (1-b) < 1.38$ at the 99.7% confidence level. Working in the context of our fiducial gas physics model, we constrain the cosmological parameters σ_8 and Ω_m, obtaining $\sigma_8(\Omega_m/0.282)^{0.26} = 0.824 \pm 0.029$ at 68% C.L. In addition, we re-analyze the tSZ auto-power spectrum measurement from [44], and combine

the results with our cross-power measurement to break the degeneracy between ICM physics and cosmological parameters using tSZ statistics alone. Within the 95% confidence interval, standard ICM models and WMAP9 or Planck CMB-derived cosmological parameters are consistent with the results (Fig. 3.18). Higher SNR measurements will soon greatly improve on these constraints, and may allow robust conclusions regarding σ_8 and massive neutrinos to be drawn from the tSZ signal, overcoming systematic uncertainties arising from the ICM physics. However, currently the tSZ data alone are driven to rather extreme regions of parameter space by the low amplitude of the measured tSZ auto-power spectrum. As discussed in Section 3.6.2, this low amplitude could be the result of over-subtraction of residual foreground power in [44], or could be due to ICM physics currently missing in the theoretical calculations. In this context, it is worth noting that the ACT and SPT constraints on the tSZ power at $\ell \approx 3000$ are also low compared to our fiducial model, although these constraints are obtained in a different approach based on fitting multifrequency measurements of the high-ℓ CMB power spectrum. However, at these very small angular scales, the tSZ power spectrum is more sensitive to the details of the gas distribution in lower-mass halos (e.g., gas blown out of groups by AGN feedback), and thus the low amplitude of the ACT and SPT values may indeed be due to ICM physics considerations. Upcoming data from ACTPol [94], SPTPol [95], and Planck will shed further light on these issues.

Overall, our measurement of the tSZ – CMB lensing cross-correlation is a direct confirmation that hot, ionized gas traces dark matter throughout the universe over a wide range of physical scales.

3.8 Acknowledgments

We thank Graeme Addison for providing the theoretical model of the CIB power spectrum at 857 GHz, as well as for a number of useful exchanges. We thank Duncan Hanson for guidance regarding the Planck CMB lensing potential map, Olivier Doré for providing the

CIB–CMB lensing cross-power spectrum modeling results from [48], and Jacques Delabrouille for providing an early release of the Planck Sky Model code. We thank Nick Battaglia and Amir Hajian for providing results from [67] for comparisons in Section 3.6, as well as many informative discussions. We are also grateful to Daisuke Nagai, Enrico Pajer, Uros Seljak, Blake Sherwin, and Kendrick Smith for many helpful conversations. While this manuscript was in preparation, we became aware of related work in [96], and subsequently exchanged correspondence with the authors of that study; we leave a detailed comparison of the results for future work. JCH and DNS are supported by NASA Theory Grant NNX12AG72G and NSF AST-1311756.

Bibliography

[1] Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, ApJS, 208, 19

[2] Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, ApJS, 208, 20

[3] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5076

[4] Sievers, J. L., Hlozek, R. A., Nolta, M. R., et al. 2013, J. Cosmology Astropart. Phys., 10, 60

[5] Story, K. T., Reichardt, C. L., Hou, Z., et al. 2012, arXiv:1210.7231

[6] Blanchard, A., & Schneider, J. 1987, A&A, 184, 1

[7] Zel'dovich, Y. B., & Sunyaev, R. A. 1969, Ap&SS, 4, 301

[8] Sunyaev, R. A., & Zeldovich, Y. B. 1970, Ap&SS, 7, 3

[9] Carlstrom, J. E., Holder, G. P., & Reese, E. D. 2002, ARA&A, 40, 643

[10] Marriage, T. A., Acquaviva, V., Ade, P. A. R., et al. 2011, ApJ, 737, 61

[11] Vanderlinde, K., Crawford, T. M., de Haan, T., et al. 2010, ApJ, 722, 1180

[12] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5080

[13] Hand, N., Addison, G. E., Aubourg, E., et al. 2012, Physical Review Letters, 109, 041101

[14] Smith, K. M., Zahn, O., & Doré, O. 2007, Phys. Rev. D, 76, 043510

[15] Hirata, C. M., Ho, S., Padmanabhan, N., Seljak, U., & Bahcall, N. A. 2008, Phys. Rev. D, 78, 043520

[16] Das, S., Sherwin, B. D., Aguirre, P., et al. 2011, Physical Review Letters, 107, 021301

[17] van Engelen, A., Keisler, R., Zahn, O., et al. 2012, ApJ, 756, 142

[18] Goldberg, D. M., & Spergel, D. N. 1999, Phys. Rev. D, 59, 103002

[19] Cooray, A., & Hu, W. 2000, ApJ, 534, 533

[20] Cooray, A., Hu, W., & Tegmark, M. 2000, ApJ, 540, 1

[21] Cooray, A. 2000, Phys. Rev. D, 62, 103506

[22] Peiris, H. V., & Spergel, D. N. 2000, ApJ, 540, 605

[23] Afshordi, N., Loh, Y.-S., & Strauss, M. A. 2004, Phys. Rev. D, 69, 083524

[24] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5084

[25] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5079

[26] Hasselfield, M., Hilton, M., Marriage, T. A., et al. 2013, J. Cosmology Astropart. Phys., 7, 8

[27] Reichardt, C. L., Stalder, B., Bleem, L. E., et al. 2013, ApJ, 763, 127

[28] Nozawa, S., Itoh, N., Suda, Y., & Ohhata, Y. 2006, Nuovo Cimento B Serie, 121, 487

[29] Battaglia, N., Bond, J. R., Pfrommer, C., & Sievers, J. L. 2012, ApJ, 758, 75

[30] Battaglia, N., Bond, J. R., Pfrommer, C., Sievers, J. L., & Sijacki, D. 2010, ApJ, 725, 91

[31] Sijacki, D., Springel, V., Di Matteo, T., & Hernquist, L. 2007, Mon. Not. R. Astron. Soc, 380, 877

[32] McCarthy, I. G., Schaye, J., Ponman, T. J., et al. 2010, Mon. Not. R. Astron. Soc, 406, 822

[33] Borgani, S., Murante, G., Springel, V., et al. 2004, Mon. Not. R. Astron. Soc, 348, 1078

[34] Piffaretti, R., & Valdarnini, R. 2008, A&A, 491, 71

[35] Nagai, D., Kravtsov, A. V., & Vikhlinin, A. 2007, ApJ, 668, 1

[36] Cooray, A., & Sheth, R. 2002, Physics Reports, 372, 1

[37] Hill, J. C., & Pajer, E. 2013, Phys. Rev. D, 88, 063526

[38] Limber, D. N. 1954, ApJ, 119, 655

[39] Komatsu, E., & Seljak, U. 2002, Mon. Not. R. Astron. Soc, 336, 1256

[40] Cooray, A. 2001, Phys. Rev. D, 64, 063514

[41] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5062

[42] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5067

[43] Planck Collaboration, Aghanim, N., Armitage-Caplan, C., et al. 2013, arXiv:1303.5063

[44] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5081

[45] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5088

[46] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5077

[47] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5070

[48] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5078

[49] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1309.0382

[50] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5089

[51] Eriksen, H. K., Banday, A. J., Górski, K. M., & Lilje, P. B. 2004, ApJ, 612, 633

[52] Remazeilles, M., Delabrouille, J., & Cardoso, J.-F. 2011, Mon. Not. R. Astron. Soc, 410, 2481

[53] Addison, G. E., Dunkley, J., & Spergel, D. N. 2012, Mon. Not. R. Astron. Soc, 427, 1741

[54] Mesinger, A., McQuinn, M., & Spergel, D. N. 2012, Mon. Not. R. Astron. Soc, 422, 1403

[55] Hincks, A. D., Hajian, A., & Addison, G. E. 2013, J. Cosmology Astropart. Phys., 5, 4

[56] Remazeilles, M., Aghanim, N., & Douspis, M. 2013, Mon. Not. R. Astron. Soc, 430, 370

[57] Crawford, T. M., Schaffer, K. K., Bhattacharya, S., et al. 2013, arXiv:1303.3535

[58] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, A&A, 554, A140

[59] Wilson, M. J., Sherwin, B. D., Hill, J. C., et al. 2012, Phys. Rev. D, 86, 122005

[60] Hill, J. C., & Sherwin, B. D. 2013, Phys. Rev. D, 87, 023527

[61] Hill, J. C., Sherwin, B. D., Smith, K. M., et al. in prep.

[62] Arnaud, M., Pratt, G. W., Piffaretti, R., et al. 2010, A & A, 517, A92

[63] Sehgal, N., Bode, P., Das, S., et al. 2010, ApJ, 709, 920

[64] Dunkley, J., Hlozek, R., Sievers, J., et al. 2011, ApJ, 739, 52

[65] Sun, M., Sehgal, N., Voit, G. M., et al. 2011, ApJ, 727, L49

[66] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, A&A, 550, A131

[67] Hajian, A., Battaglia, N., Spergel, D. N., et al. 2013, arXiv:1309.3282

[68] Shaw, L. D., Nagai, D., Bhattacharya, S., & Lau, E. T. 2010, ApJ, 725, 1452

[69] Lewis, A., & Challinor, A. 2006, Physics Reports, 429, 1

[70] Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, ApJ, 490, 493

[71] Duffy, A. R., Schaye, J., Kay, S. T., & Dalla Vecchia, C. 2008, Mon. Not. R. Astron. Soc, 390, L64

[72] Bryan, G. L., & Norman, M. L. 1998, ApJ, 495, 80

[73] Tinker, J., Kravtsov, A. V., Klypin, A., et al. 2008, ApJ, 688, 709

[74] Tinker, J. L., Robertson, B. E., Kravtsov, A. V., et al. 2010, ApJ, 724, 878

[75] Komatsu, E., & Kitayama, T. 1999, ApJ, 526, L1

[76] Sherwin, B. D., Das, S., Hajian, A., et al. 2012, Phys. Rev. D, 86, 083006

[77] Binney, J., & Tremaine, S. 2008, Galactic Dynamics: Second Edition, by James Binney and Scott Tremaine. ISBN 978-0-691-13026-2 (HB). Published by Princeton University Press, Princeton, NJ USA, 2008.,

[78] Bhattacharya, S., Nagai, D., Shaw, L., Crawford, T., & Holder, G. P. 2012, ApJ, 760, 5

[79] Trac, H., Bode, P., & Ostriker, J. P. 2011, ApJ, 727, 94

[80] Shaw, L. D., Zahn, O., Holder, G. P., & Doré, O. 2009, ApJ, 702, 368

[81] Addison, G. E., Dunkley, J., & Bond, J. R. 2013, Mon. Not. R. Astron. Soc, 436, 1896

[82] Hirata, C. M., & Seljak, U. 2003, Phys. Rev. D, 67, 043001

[83] Okamoto, T., & Hu, W. 2003, Phys. Rev. D, 67, 083002

[84] Song, Y.-S., Cooray, A., Knox, L., & Zaldarriaga, M. 2003, ApJ, 590, 664

[85] Sachs, R. K., & Wolfe, A. M. 1967, ApJ, 147, 73

[86] Osborne, S. J., Hanson, D., & Doré, O. 2013, arXiv:1310.7547

[87] van Engelen, A., Bhattacharya, S., Sehgal, N., et al. 2013, arXiv:1310.7023

[88] Eifler, T., Schneider, P., & Hartlap, J. 2009, A&A, 502, 721

[89] Rudd, D. H., & Nagai, D. 2009, ApJ, 701, L16

[90] Battaglia, N., Bond, J. R., Pfrommer, C., & Sievers, J. L. 2013, ApJ, 777, 123

[91] Ichiki, K., & Takada, M. 2012, Phys. Rev. D, 85, 063521

[92] McKeown, R. D., & Vogel, P. 2004, Physics Reports, 394, 315

[93] Spergel, D., Flauger, R., & Hlozek, R. 2013, arXiv:1312.3313

[94] Niemack, M. D., Ade, P. A. R., Aguirre, J., et al. 2010, Proc. SPIE, 7741

[95] Austermann, J. E., Aird, K. A., Beall, J. A., et al. 2012, Proc. SPIE, 8452

[96] Van Waerbeke, L., Hinshaw, G., & Murray, N. 2013, arXiv:1310.5721

Chapter 4

Cosmological Constraints from Moments of the Thermal Sunyaev-Zel'dovich Effect

4.1 Abstract

In this chapter, we explain how moments of the thermal Sunyaev-Zel'dovich (tSZ) effect can constrain both cosmological parameters and the astrophysics of the intracluster medium (ICM). As the tSZ signal is strongly non-Gaussian, higher moments of tSZ maps contain useful information. We first calculate the dependence of the tSZ moments on cosmological parameters, finding that higher moments scale more steeply with σ_8 and are sourced by more massive galaxy clusters. Taking advantage of the different dependence of the variance and skewness on cosmological and astrophysical parameters, we construct a statistic, $|\langle T^3\rangle|/\langle T^2\rangle^{1.4}$, which cancels much of the dependence on cosmology (i.e., σ_8) yet remains sensitive to the astrophysics of intracluster gas (in particular, to the gas fraction in low-mass clusters). Constraining the ICM astrophysics using this statistic could break the well-known degeneracy between cosmology and gas physics in tSZ measurements, allowing for tight con-

straints on cosmological parameters. Although detailed simulations will be needed to fully characterize the accuracy of this technique, we provide a first application to data from the Atacama Cosmology Telescope and the South Pole Telescope. We estimate that a *Planck*-like full-sky tSZ map could achieve a $\lesssim 1\%$ constraint on σ_8 and a 1σ error on the sum of the neutrino masses that is comparable to the existing lower bound from oscillation measurements.

4.2 Introduction

The thermal Sunyaev-Zel'dovich (tSZ) effect is a spectral distortion of the cosmic microwave background (CMB) caused by inverse Compton scattering of CMB photons off hot electrons in the intracluster medium (ICM) of galaxy clusters [1]. The tSZ effect, which is largest on arcminute angular scales, has traditionally been studied either through observations of individual clusters — both pointed measurements [2, 3, 4, 5] and recent detections in blind surveys [6, 7, 8] — or indirectly through its contribution to the CMB power spectrum [9, 10].

The goal of many direct studies is to characterize the masses and redshifts of a well-defined sample of clusters, and thereby reconstruct the high-mass end of the halo mass function, an important cosmological quantity that is sensitive to a number of parameters. These parameters include σ_8, the rms fluctuation amplitude on scales of $8\,h^{-1}$ Mpc. However, one can also constrain these parameters through their influence on the tSZ power spectrum, which is exceptionally sensitive to σ_8 in particular [11]. In recent years, many studies have obtained competitive constraints on σ_8 from the power spectrum [9, 10]. This approach is advantageous not only because the tSZ signal is extremely sensitive to σ_8, but also because it does not require the measurement of individual cluster masses, which is a difficult procedure. Unfortunately, it does require a precise understanding of the pressure profile of the ICM gas for clusters over a wide range of masses and redshifts. Consequently, systematic errors due to theoretical uncertainty in the astrophysics of the ICM have remained comparable to or

greater than statistical errors in tSZ power spectrum measurements [9, 10]. This situation has hindered the progress of tSZ measurements in providing cosmological constraints.

In this chapter, we propose a method to reduce the theoretical systematic uncertainty in tSZ-derived cosmological constraints by combining different moments of the tSZ effect. Although previous studies have investigated tSZ statistics beyond the power spectrum [12, 13, 14, 15], none have attempted to find a tSZ observable that isolates the dependence on either astrophysical or cosmological parameters. In §II, we compute the variance (second moment) and skewness (third moment) of the tSZ signal for a variety of ICM astrophysics scenarios while varying σ_8, which allows us to probe the mass and redshift dependence of each statistic, as well as to derive their dependence on cosmology. We find that the skewness is not only more sensitive to σ_8 than the variance, but is also dominated by contributions from higher-mass clusters, for which the ICM astrophysics is better-constrained by existing observations. In §III, we use these results to find a particular combination of the variance and skewness — the "rescaled skewness" — that only depends weakly on the underlying cosmology while remaining sensitive to the ICM astrophysics prescription. We test this statistic using a variety of models for the ICM, which should span the space of reasonable theoretical possibilities. In §IV, we apply this method to data from the Atacama Cosmology Telescope (ACT) [16, 17] and the South Pole Telescope (SPT) [18, 19] to demonstrate its feasibility and to obtain a first weak constraint on the ICM astrophysics. In §V, we estimate the extent to which this approach can increase the precision of future constraints on σ_8 and the sum of the neutrino masses through tSZ measurements. We also highlight applications to other parameters that affect the tSZ signal through the mass function.

We assume a concordance ΛCDM cosmology throughout, with parameters taking their maximum-likelihood WMAP5 values (WMAP+BAO+SN) [20] unless otherwise specified. All masses are quoted in units of $M_\odot/h$, where $h \equiv H_0/(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}})$ and H_0 is the Hubble parameter today.

4.3 Calculating tSZ Moments

4.3.1 Background

The tSZ effect leads to a frequency-dependent change in the observed CMB temperature in the direction of a galaxy group or cluster. Neglecting relativistic corrections (which are relevant only for the most massive systems) [21], the temperature change T at angular position $\boldsymbol{\theta}$ with respect to the center of a cluster of mass M at redshift z is given by [1]

$$\begin{aligned} \frac{T(\boldsymbol{\theta}, M, z)}{T_{\rm CMB}} &= g_\nu y(\boldsymbol{\theta}, M, z) \\ &= g_\nu \frac{\sigma_T}{m_e c^2} \int_{\rm LOS} P_e\left(\sqrt{l^2 + d_A^2 |\boldsymbol{\theta}|^2}, M, z\right) dl\,, \end{aligned} \tag{4.1}$$

where the tSZ spectral function $g_\nu = x\coth(x/2) - 4$ with $x \equiv h\nu/k_B T_{\rm CMB}$, y is the Compton-y parameter, σ_T is the Thomson scattering cross-section, m_e is the electron mass, and $P_e(\boldsymbol{r})$ is the ICM electron pressure at location $\boldsymbol{r}$ with respect to the cluster center. In this work we only consider spherically symmetric pressure profiles with $P_e(\boldsymbol{r}) = P_e(r)$. Also, we calculate all observables at $\nu = 150$ GHz, where the tSZ effect is observed as a decrement in the CMB temperature in the direction of a cluster (i.e., $g_\nu < 0$). Note that the integral in Eq. (4.1) is taken along the line of sight such that $r^2 = l^2 + d_A(z)^2\theta^2$, where $d_A(z)$ is the angular diameter distance to redshift z and $\theta \equiv |\boldsymbol{\theta}|$ is the angular distance from the cluster center in the plane of the sky. Given a spherically symmetric pressure profile, Eq. (4.1) implies that the temperature decrement profile is azimuthally symmetric in the plane of the sky, that is, $T(\boldsymbol{\theta}, M, z) = T(\theta, M, z)$. Finally, note that the electron pressure is related to the thermal gas pressure via $P_{th} = P_e(5X_H + 3)/2(X_H + 1) = 1.932 P_e$, where $X_H = 0.76$ is the primordial hydrogen mass fraction.

In order to calculate moments of the tSZ effect, we assume that the distribution of clusters on the sky can be adequately described by a Poisson distribution (and that contributions due to clustering are negligible, which is valid on sub-degree scales) [22, 23]. The $N^{\rm th}$ moment

at zero lag is then given by

$$\langle T^N \rangle = \int dz \frac{dV}{dz} \int dM \frac{dn(M,z)}{dM} \int d^2\theta \, T(\theta, M, z)^N \,, \tag{4.2}$$

where dV/dz is the comoving volume per steradian at redshift z, $dn(M,z)/dM$ is the comoving number density of halos of mass M at redshift z (the halo mass function), and $T(\theta, M, z)$ is given by Eq. (4.1). Our fiducial integration limits are $0.005 < z < 4$ and $5 \times 10^{11} M_\odot/h < M < 5 \times 10^{15} M_\odot/h$. Eq. (4.2) only includes the 1-halo contribution to the $N^{\rm th}$ moment at zero lag — as noted, we have neglected contributions due to clustering. Using the bias model of [24], we find that including the 2-halo term typically increases $\langle T^2 \rangle$ by only 1–2%. Moreover, the 2- and 3-halo contributions to $\langle T^3 \rangle$ should be even less significant, since the higher moments are dominated by regions increasingly near the center of each cluster (where $|T|$ is larger). Thus, we include only the 1-halo term in all calculations, as given in Eq. (4.2).

Note that we define M to be the virial mass, that is, the mass enclosed within a radius r_{vir} [25]:

$$r_{vir} = \left(\frac{3M}{4\pi \Delta_{cr}(z) \rho_{cr}(z)} \right)^{1/3} , \tag{4.3}$$

where $\Delta_{cr}(z) = 18\pi^2 + 82(\Omega(z) - 1) - 39(\Omega(z) - 1)^2$, $\Omega(z) = \Omega_m (1+z)^3 / (\Omega_m (1+z)^3 + \Omega_\Lambda)$, and $\rho_{cr}(z) = 3H^2(z)/8\pi G$ is the critical density at redshift z. However, some of the pressure profiles that we use below are specified as a function of the spherical overdensity mass rather than the virial mass. Thus, when necessary, we use the NFW density profile [26] and the concentration-mass relation from [27] to convert the virial mass to a spherical overdensity mass $M_{\delta,c/d}$, where $M_{\delta,c}$ ($M_{\delta,d}$) is the mass enclosed within a sphere of radius $r_{\delta,c}$ ($r_{\delta,d}$) such that the enclosed density is δ times the critical (mean matter) density at redshift z. To be completely explicit, this procedure involves solving the following non-linear equation for $r_{\delta,c}$

(or $r_{\delta,d}$):

$$\int_0^{r_{\delta,c}} 4\pi r'^2 \rho_{\rm NFW}(r', M_{vir}, c_{vir}) dr' = \frac{4}{3}\pi r_{\delta,c}^3 \rho_{cr}(z)\delta \tag{4.4}$$

where $c_{vir} \equiv r_{vir}/r_s$ is the concentration parameter (r_s is the NFW scale radius) and one replaces the critical density $\rho_{cr}(z)$ with the mean matter density $\rho_m(z)$ in this equation in order to obtain $r_{\delta,d}$ instead of $r_{\delta,c}$. After solving Eq. (4.4) to find $r_{\delta,c}$, $M_{\delta,c}$ is simply calculated via $M_{\delta,c} = \frac{4}{3}\pi r_{\delta,c}^3 \rho_{cr}(z)\delta$.

From Eqs. (4.1) and (4.2), it is clear that two ingredients are required for the tSZ moment calculation, given a specified cosmology: (1) the halo mass function $dn(M,z)/dM$ and (2) the electron pressure profile $P_e(r, M, z)$ for halos of mass M at redshift z. The calculation thus naturally divides into a cosmology-dependent component (the mass function) and an ICM-dependent component (the pressure profile). Because the small-scale baryonic physics responsible for the details of the ICM pressure profile is decoupled from the large-scale physics described by the cosmological parameters (σ_8, Ω_m, ...), it is conventional to determine the pressure profile from numerical cosmological hydrodynamics simulations (or observations of galaxy clusters), which are run for a fixed cosmology. One can then take this pressure profile and compute its predictions for a different cosmology by using the halo mass function appropriate for that cosmology. We follow this approach below.

For all the calculations in this chapter, we use the $M_{200,d}$ halo mass function of [28] with the redshift-dependent parameters given in their Eqs. (5)–(8); we will hereafter refer to this as the Tinker mass function. However, the tSZ moments are somewhat sensitive to the particular choice of mass function used in the calculation. As a test, we perform identical calculations using the $M_{400,d}$ halo mass function of [28]. Using our fiducial cosmology and the Battaglia pressure profile (see below), we find that the $M_{400,d}$ mass function gives a result for $\langle T^2 \rangle$ that is 11% higher than that found using the $M_{200,d}$ mass function, while the result for $\langle T^3 \rangle$ is 22% higher. In general, higher tSZ moments are more sensitive to changes in

the mass function, as they are dominated by progressively higher-mass halos that live in the exponential tail of the mass function. While these changes are non-negligible for upcoming high precision measurements, they are smaller than those caused by changes in the choice of pressure profile; hence, for the remainder of this chapter, we use the $M_{200,d}$ Tinker mass function for all calculations.

Our calculations include four different pressure profiles from [29], [30], and [11] (and additional results using the profile from [34]), which we briefly describe in the following. We consider two profiles derived from the simulations of [31], which are reported in [29]. The first, which we hereafter refer to as the Battaglia profile, is derived from hydrodynamical simulations that include radiative cooling, star formation, supernova feedback, and feedback from active galactic nuclei (AGN). These feedback processes tend to lower the gas fraction in low-mass clusters, as gas is blown out by the injection of energy into the ICM. The smoothed particle hydrodynamics used in these simulations also captures the effects of non-thermal pressure support due to bulk motions and turbulence. The second profile that we use from [29], which we hereafter refer to as the Battaglia Adiabatic profile[1], is derived from hydrodynamical simulations with all forms of feedback turned off. This setup leads to a non-radiative adiabatic cluster model with only formation shock heating present. Thus, in these simulations, the gas fraction in low-mass clusters is close to the cosmological value, Ω_b/Ω_m (as is the case in high-mass clusters). This leads to much higher predicted amplitudes for the tSZ moments (see below). Note that the Battaglia profile is specified as a function of $M_{200,c}$, while the Battaglia Adiabatic profile is given as a function of $M_{500,c}$.

In addition, we use the "universal" pressure profile derived in [30], which we hereafter refer to as the Arnaud profile. This profile is obtained from a combination of *XMM-Newton* observations of massive, $z < 0.2$ clusters and hydrodynamical simulations that include radiative cooling and some feedback processes (though not AGN feedback). The simulations are used to extend the profile beyond R_{500}, due to the lack of X-ray photons at large radii. Im-

[1] N. Battaglia, priv. comm. (now included in [29])

portantly, the normalization of this profile is obtained using hydrostatic equilibrium (HSE)-based estimates of the mass $M_{500,c}$, which are known to be biased low by $\approx 10-15\%$ [32]. Thus, following [33], we use a HSE-bias correction of 13% in our calculations with this profile, that is, we set $M_{500,c}^{HSE} = 0.87 M_{500,c}$. This correction slightly lowers the amplitude of the tSZ moments for this profile.

We also use the profile derived in [11], which we hereafter refer to as the Komatsu-Seljak (or K-S) profile. This profile is derived analytically under some simplifying assumptions, including HSE, gas tracing dark matter in the outer regions of clusters, and a constant polytropic equation of state for the ICM gas. In particular, it includes no feedback prescriptions or sources of non-thermal pressure support. The gas fraction in low-mass clusters (indeed, in all clusters) is thus equal to the cosmological value, leading to high predicted amplitudes for the tSZ moments. This profile is now known to over-predict the tSZ power spectrum amplitude [9, 10], but we include it here as an extreme example of the possible ICM physics scenarios. It is specified as a function of the virial mass M as defined in Eq. (4.3).

Finally, we present results computed using the fiducial profile from [34], which have been graciously provided by the authors of that study. This profile is based on that of [33], and thus we hereafter refer to this as the Shaw profile. It is derived from an analytic ICM model that accounts for star formation, feedback from supernovae and AGN, and non-thermal pressure support. Its tSZ predictions are generally similar to those of the Battaglia profile. In particular, it also predicts a suppression of the gas fraction in low-mass clusters due to feedback processes.

Overall, then, we have two profiles that include a variety of detailed feedback prescriptions (Battaglia and Shaw), one profile based primarily on an empirical fit to X-ray data (Arnaud), and two profiles for which the ICM is essentially in HSE (Battaglia Adiabatic and Komatsu-Seljak). The Arnaud, Battaglia, and Shaw models all agree reasonably well with X-ray observations of massive, low-redshift clusters [35]. The Komatsu-Seljak profile is somewhat discrepant for these high-mass systems, but disagrees more significantly with the predictions

Profile	A_2 [μK^2]	A_3 [μK^3]	α_2	α_3
Arnaud	20.5	-1790	7.9	11.5
Battaglia	22.6	-1660	7.7	11.2
Batt. Adiabatic	47.1	-3120	6.6	9.7
Komatsu-Seljak	53.0	-3040	7.5	10.6
Shaw	23.8	-1610	7.9	10.7

Table 4.1: Amplitudes and power-law scalings with σ_8 for the tSZ variance and skewness, as defined in Eq. (4.5). The first column lists the pressure profile used in the calculation (note that all calculations use the Tinker mass function). The amplitudes are specified at $\sigma_8 = 0.817$, the WMAP5 maximum-likelihood value. All results are computed at $\nu = 150$ GHz.

of the feedback-based profiles for low-mass groups and clusters. The predicted gas fraction in these low-mass objects is a major source (along with non-thermal pressure support) of the difference in the tSZ predictions between the feedback- or X-ray-based models (Arnaud, Battaglia, and Shaw) and the adiabatic models (Battaglia Adiabatic and Komatsu-Seljak).

We also analyze a tSZ simulation from [36] that covers an octant of the sky, providing a nontrivial test of our calculations. We hereafter refer to this as the Sehgal simulation. The simulated data was produced by populating a large dark matter N-body simulation with gas according to a polytropic cluster model that also includes some feedback prescriptions. However, the model requires all clusters — including low-mass objects — to contain enough gas to reach the cosmological value of the gas fraction. Note that the simulation results include tSZ signal from components that are not accounted for in the halo model-based calculations, including substructure within halos, deviations from the globally averaged pressure profile, and diffuse emission from the intergalactic medium. These effects are expected to contribute to the tSZ power spectrum at the $\approx 10 - 20\%$ level at high-ℓ [29]. The simulation was run using $\sigma_8 = 0.8$, with the other ΛCDM parameters taking values consistent with WMAP5. Since an average pressure profile has not been derived from this simulation, we cannot recompute its predictions for different cosmologies, and thus it is presented as a single data point in the figures throughout this chapter.

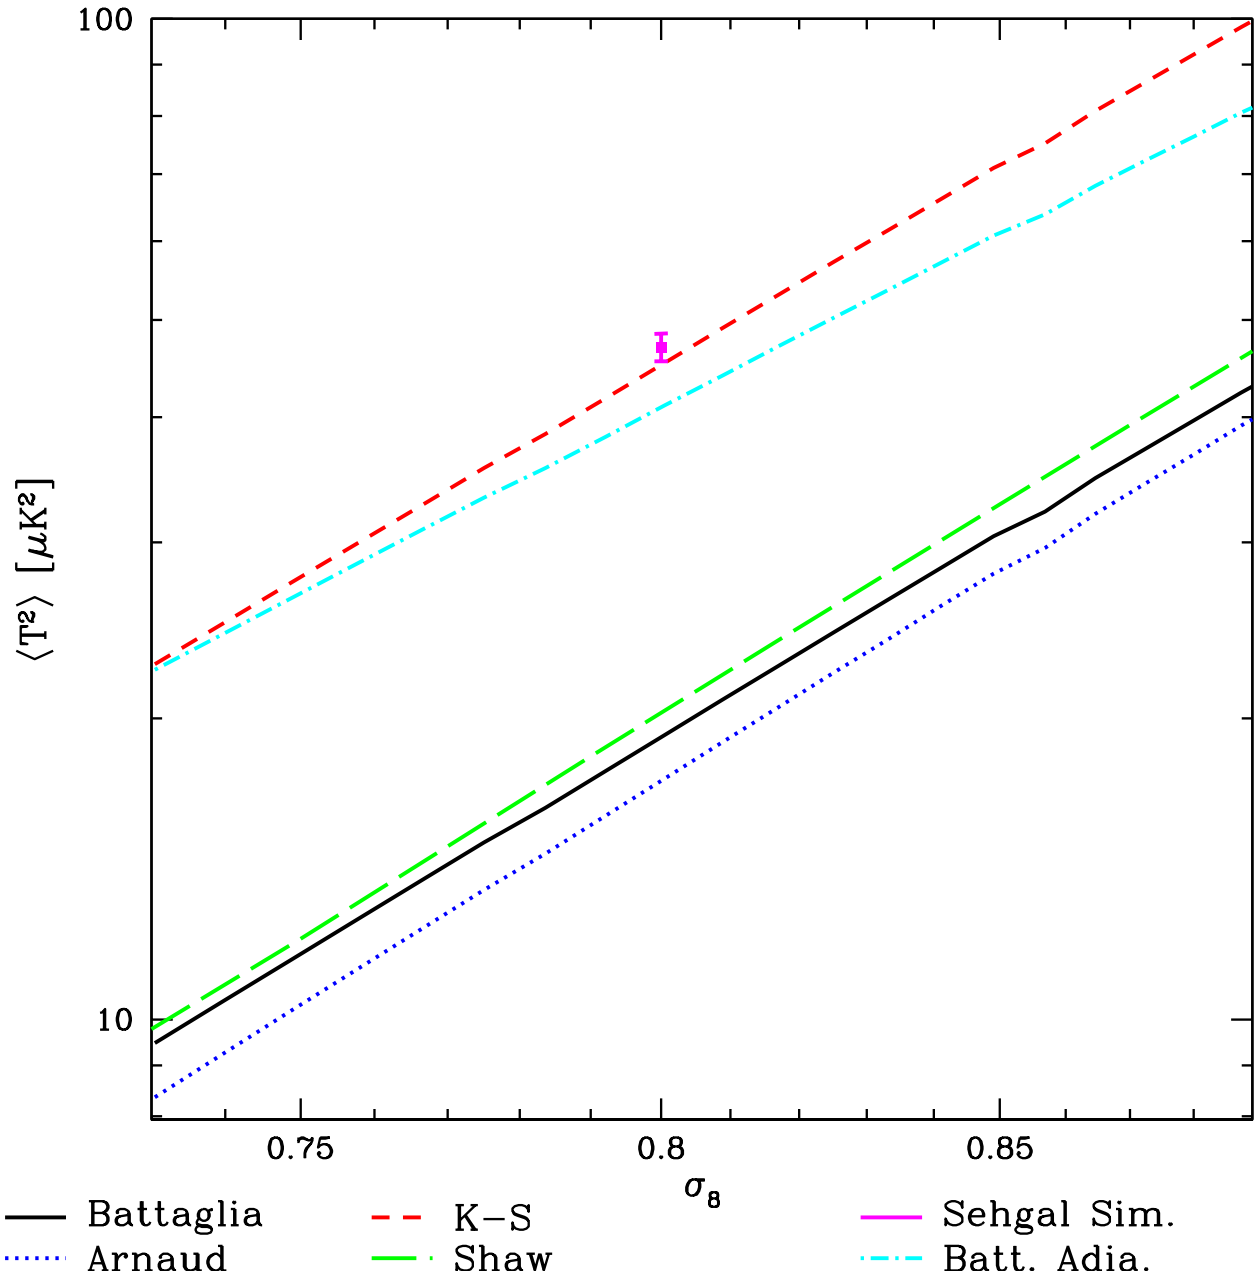

Figure 4.1: The tSZ variance versus σ_8 obtained from five different pressure profiles (using the Tinker mass function) and one direct simulation measurement. The scalings with σ_8 are similar for all the models: $\langle T^2 \rangle \propto \sigma_8^{6.6-7.9}$. It is evident that $\langle T^2 \rangle$ is very sensitive to σ_8, but the scatter due to uncertainties in the ICM astrophysics (which is greater than twice the signal for large σ_8) makes precise constraints from this quantity difficult.

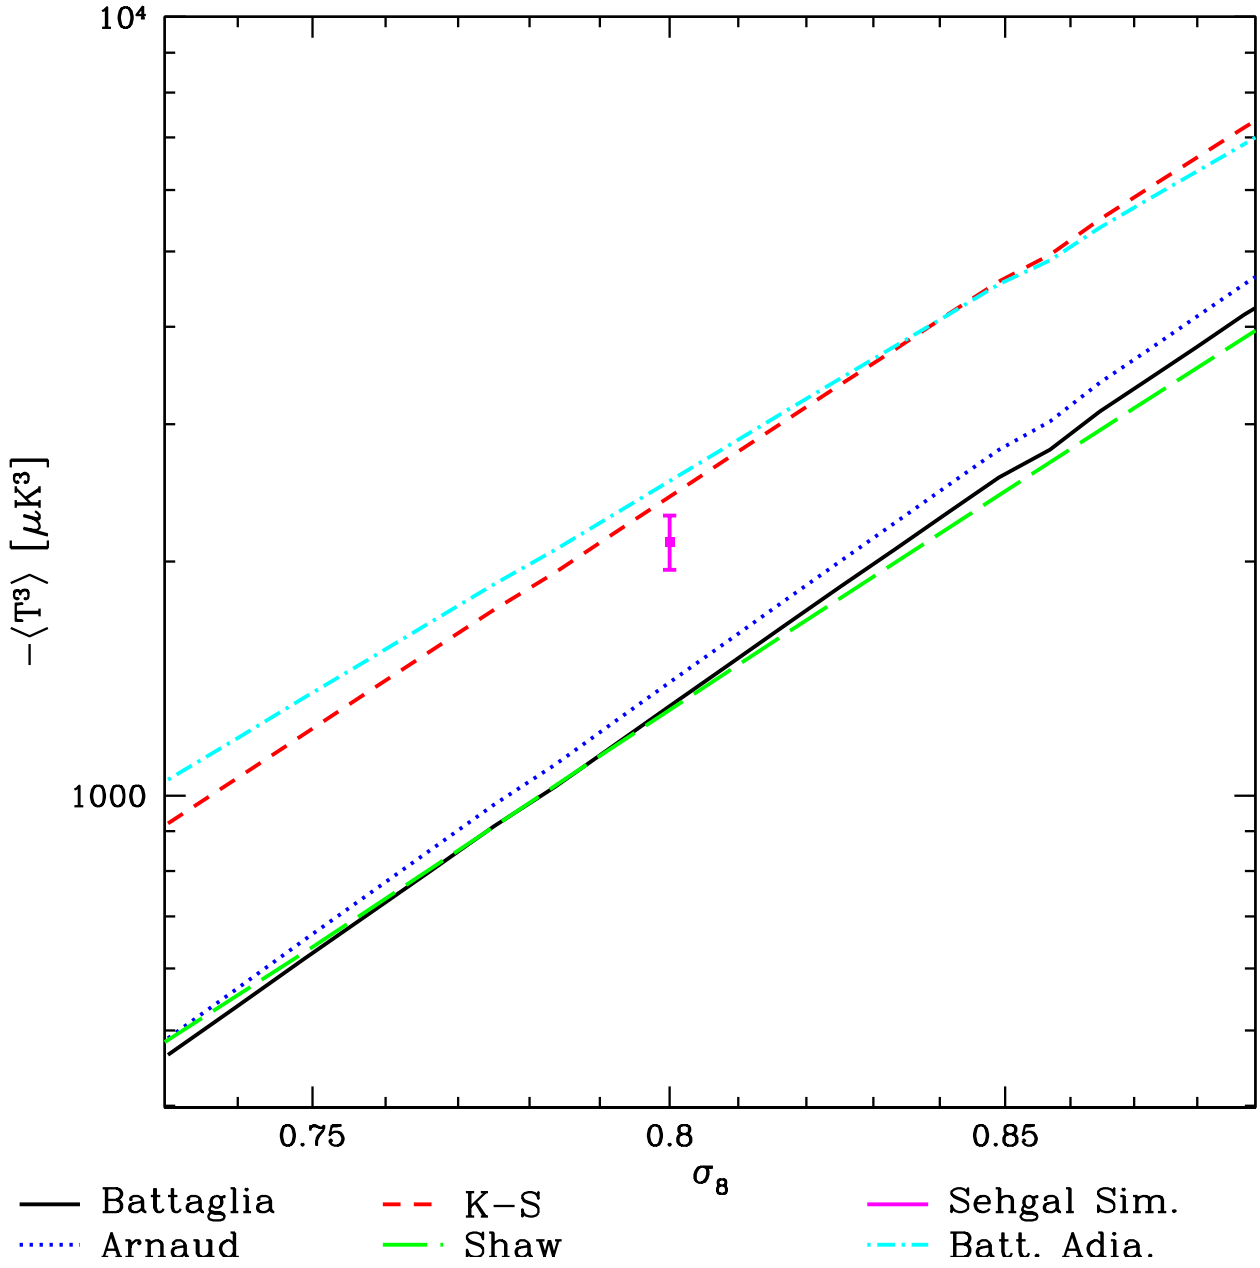

Figure 4.2: Similar to Fig. 1, but now showing the tSZ skewness. The scalings with σ_8 are again similar for all the models: $\langle T^3 \rangle \propto \sigma_8^{9.7-11.5}$. As for the variance, the sensitivity to σ_8 is pronounced, but degraded due to uncertainties in the ICM astrophysics (as represented by the different choices of pressure profile).

4.3.2 Results

We compute Eq. (4.2) for the variance ($N = 2$) and the unnormalized skewness ($N = 3$) for each of the profiles while varying σ_8 and keeping the other cosmological parameters fixed. The results are shown in Figs. 1 and 2. (See Appendix B for results involving the unnormalized kurtosis, i.e., $N = 4$.) We find that the scalings of the variance and skewness with σ_8 are well-described by power-laws for each of these profiles,

$$\langle T^{2,3} \rangle = A_{2,3} \left(\frac{\sigma_8}{0.817} \right)^{\alpha_{2,3}} , \tag{4.5}$$

where we have normalized to the WMAP5 value of σ_8. The amplitudes $A_{2,3}$ and scalings $\alpha_{2,3}$ for each of the pressure profiles are given in Table I. The scalings are similar for all the pressure profiles. Note that the slightly steeper scalings for the Battaglia and Shaw profiles are due to their inclusion of feedback processes that suppress the tSZ signal from low-mass clusters (the Arnaud profile also suppresses the signal from these objects, though it is primarily based on a fit to higher-mass clusters). Thus, $\langle T^2 \rangle$ and $\langle T^3 \rangle$ for these profiles are dominated by higher-mass (rarer) objects than for the Komatsu-Seljak or Battaglia Adiabatic profiles, and they are correspondingly more sensitive to σ_8. Despite the steep scalings, it is clear in Figs. 1 and 2 that the systematic uncertainty due to the unknown ICM astrophysics significantly degrades any potential constraints that could be derived from these observables. We address possible ways around this problem in the next section.

We also investigate the characteristic mass and redshift ranges responsible for the variance and skewness signals. Figs. 3 and 4 show the fraction of the variance and skewness signals contributed by clusters with virial mass $M < M_{\rm max}$. We find that the variance receives $\approx 50 - 60\%$ of its amplitude from clusters with $M < 2$–3×10^{14} $M_\odot/h$, while the skewness receives only $\approx 20 - 40\%$ of its amplitude from these less massive objects. These results vary for the different profiles, as the different feedback prescriptions suppress the amplitude contributed by low-mass clusters by different amounts. More massive clusters are domi-

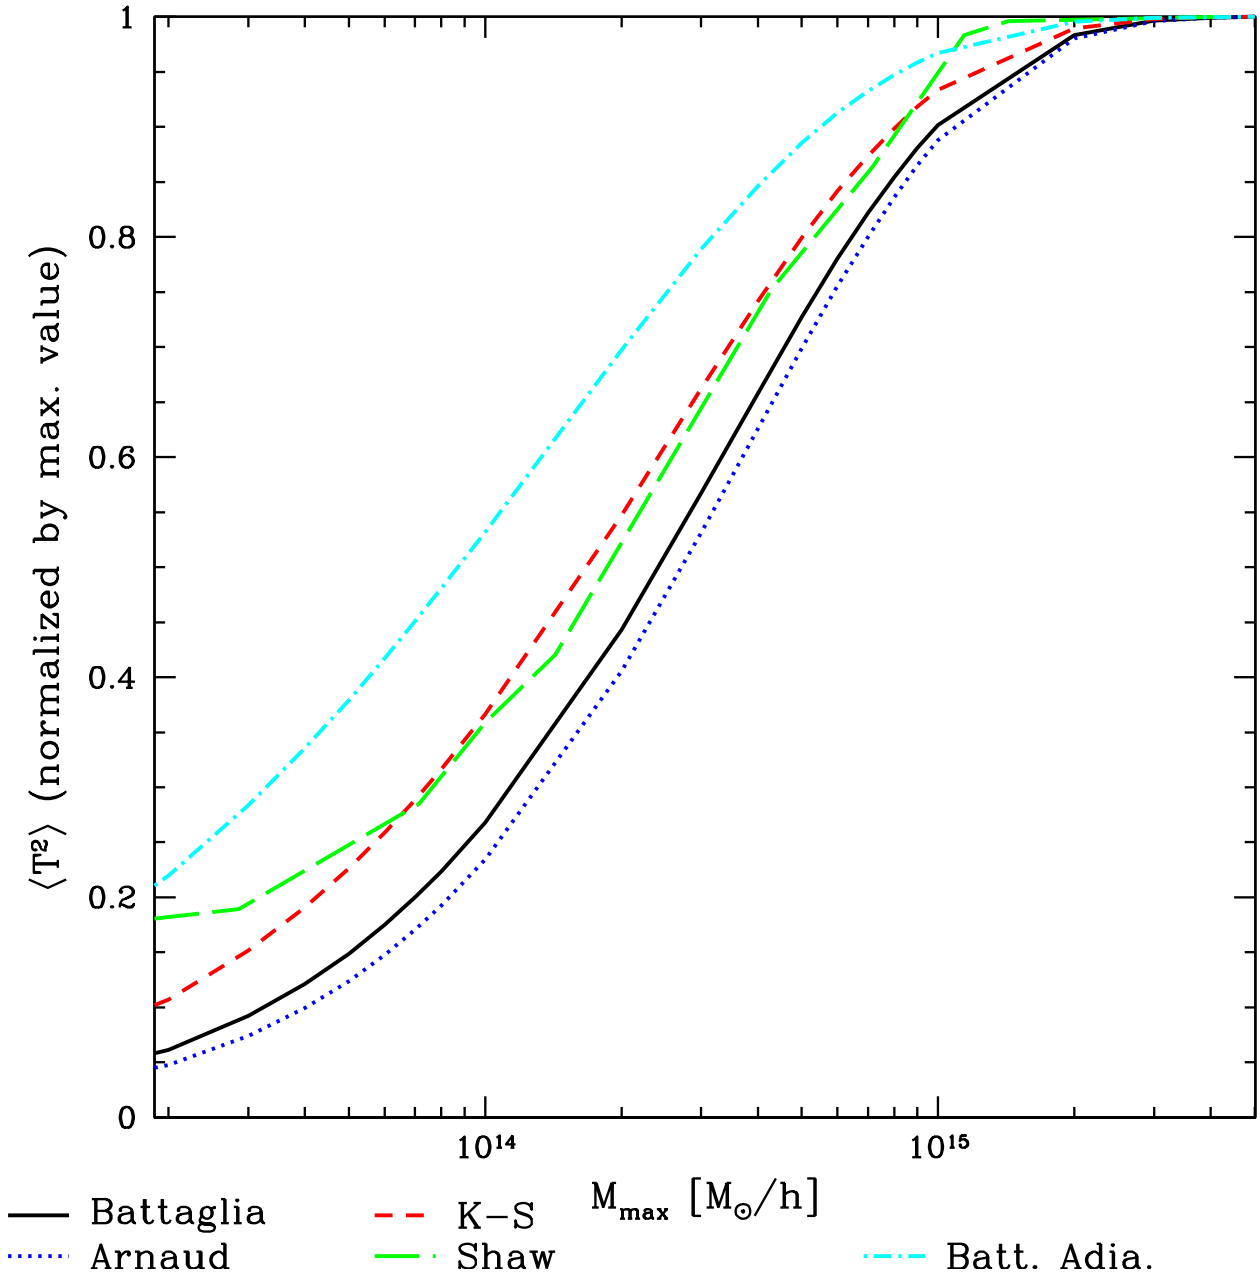

Figure 4.3: Fraction of the tSZ variance contributed by clusters with virial mass $M < M_{\rm max}$.

nated by gravitational heating and are thus less sensitive to energy input via feedback from AGN, turbulence, and other mechanisms [37, 33]. Since the Komatsu-Seljak and Battaglia Adiabatic profiles include no feedback, the amplitudes of their variance and skewness are somewhat more weighted toward low-mass objects. In all cases, though, these results indicate that the ICM astrophysics underlying the tSZ skewness is much better constrained by observation than that responsible for much of the variance (or power spectrum), as the skewness amplitude is dominated by significantly more massive objects which have been better studied observationally. This explains the smaller scatter seen between the different profiles in Fig. 2 as compared to Fig. 1: the skewness signal is dominated by objects for which the ICM astrophysics is less uncertain.

In addition, we find that the variance receives $\approx$ 25% (Arnaud/Battaglia) – 45% (Battaglia Adiabatic) of its amplitude from groups and clusters at $z > 1$, while the skewness receives only $\approx$ 15% (Arnaud/Battaglia) – 35% (Battaglia Adiabatic) of its amplitude from these high-redshift objects. (The Komatsu-Seljak results lie in the middle of these ranges.) These results imply that a greater fraction of the tSZ skewness than the variance originates from clusters that have been studied using X-rays, lensing, and other techniques. On the other hand, many of the objects comprising the variance (or power spectrum) have not been directly observed.

As it depends sensitively on σ_8 and is less affected by uncertainties in astrophysical modeling than the power spectrum, the tSZ skewness is a powerful cosmological probe. Nevertheless, despite the signal originating from characteristically higher-mass, lower-redshift clusters, the amplitude of the skewness is still quite uncertain due to the poorly understood astrophysics of the ICM, as can be seen in Fig. 2. We investigate methods to solve this problem in the remainder of this chapter.

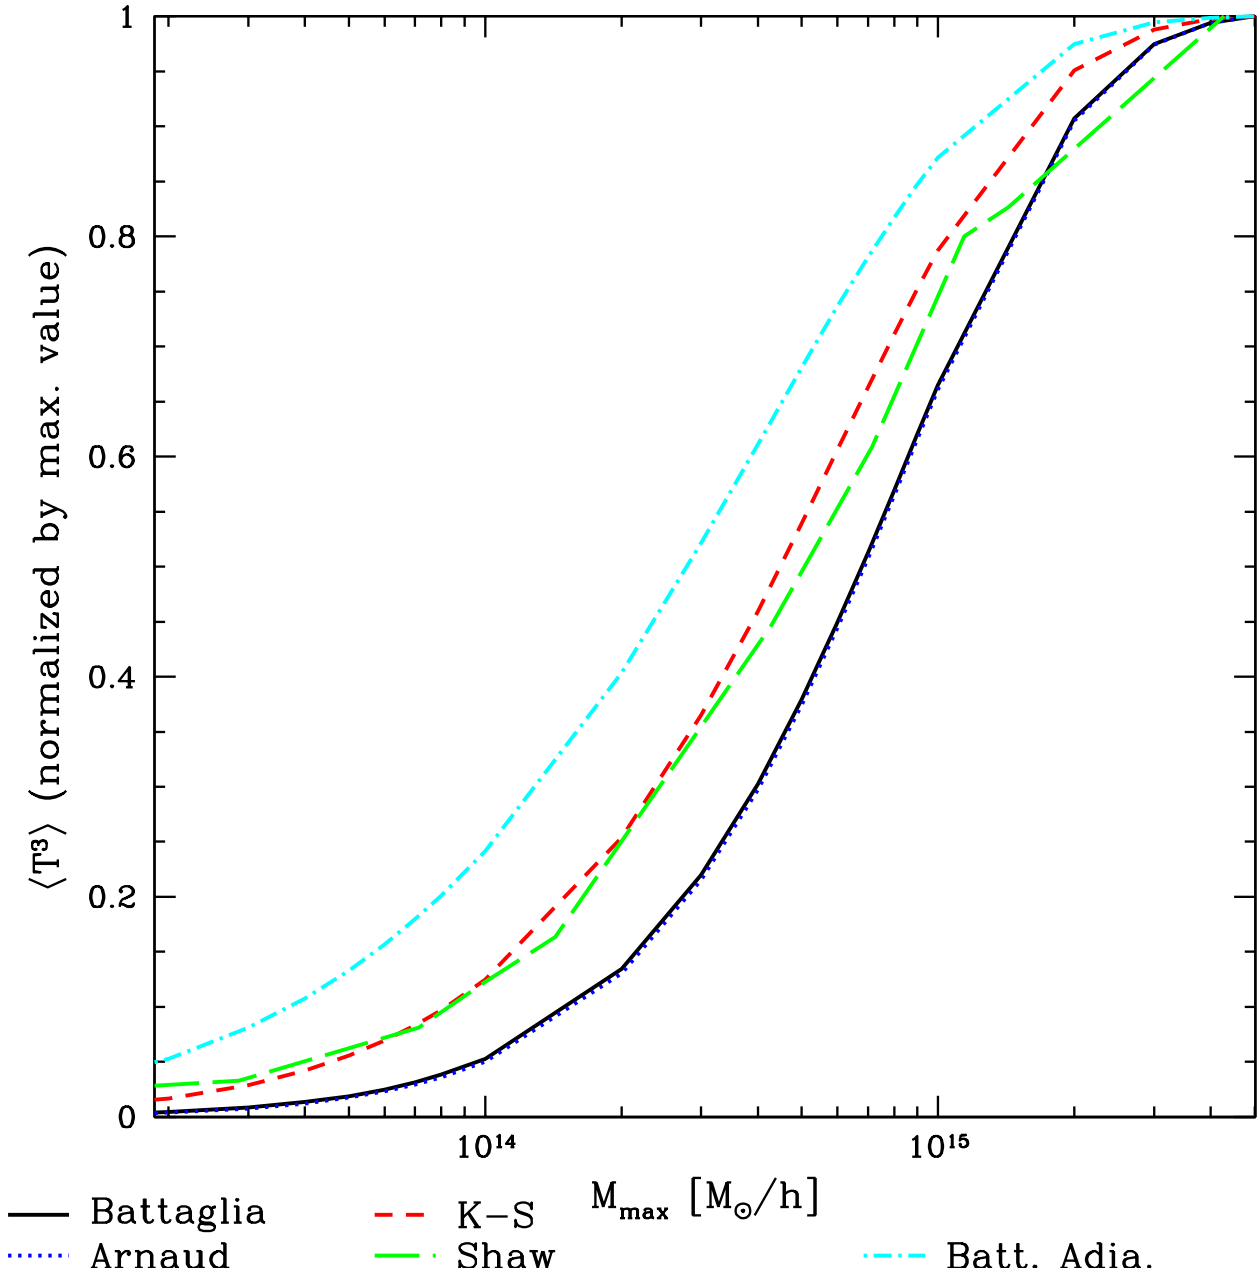

Figure 4.4: Fraction of the tSZ skewness contributed by clusters with virial mass $M < M_{\rm max}$. Comparison with Fig. 3 indicates that the skewness signal arises from higher-mass clusters than those that comprise the variance.

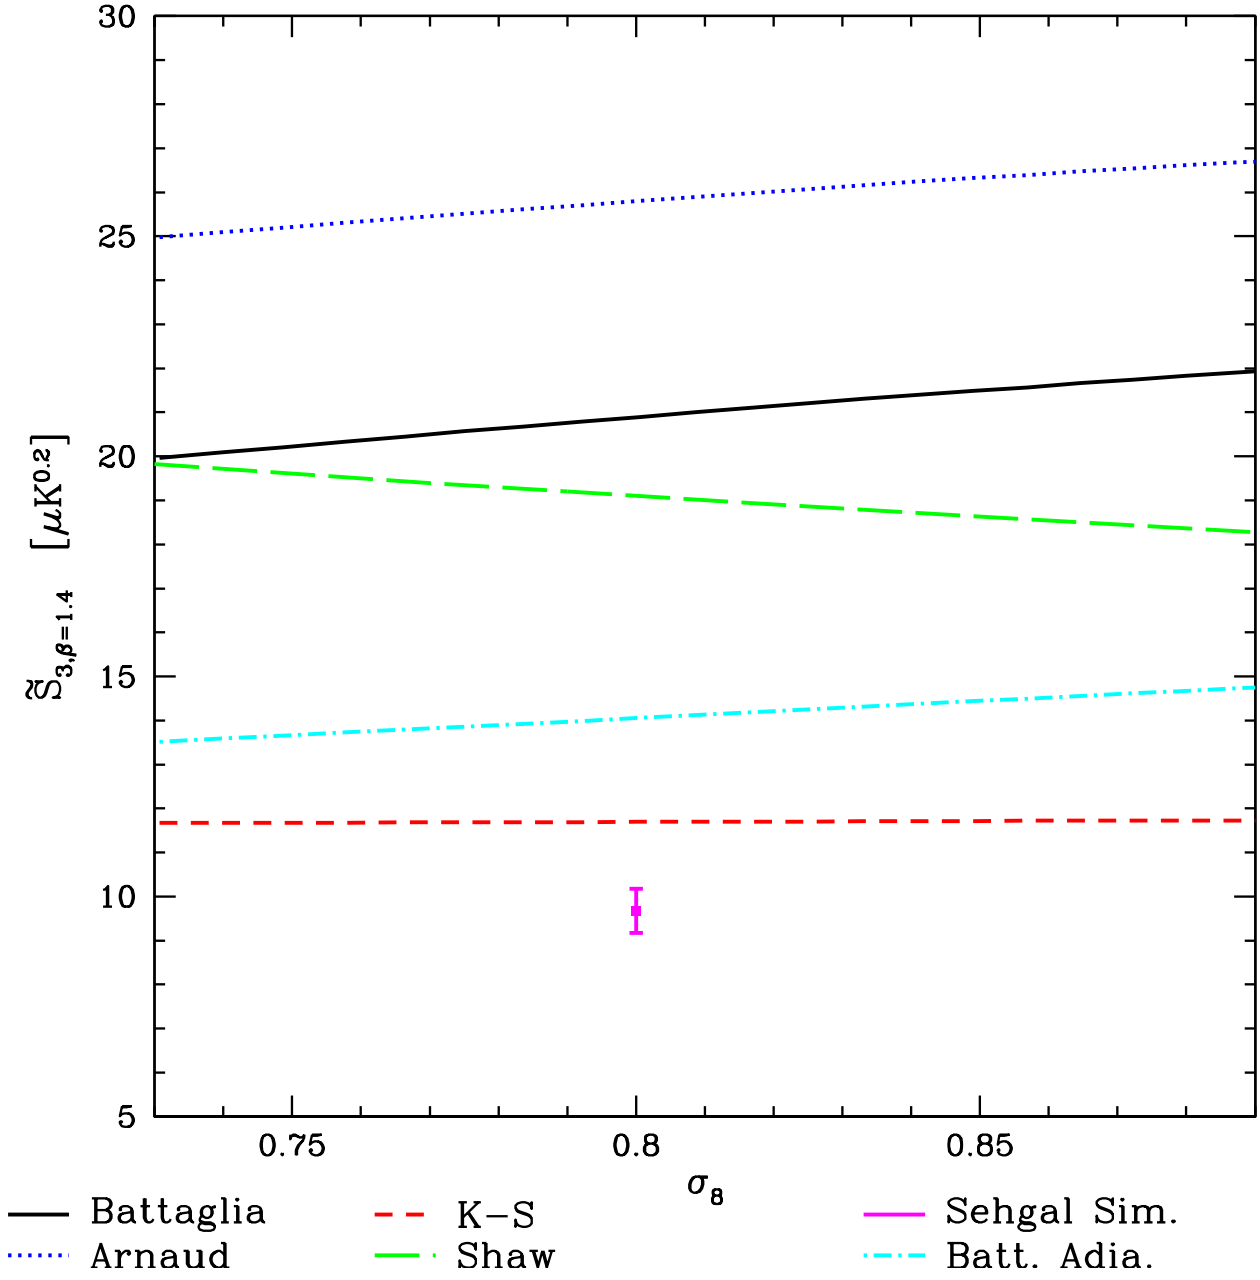

Figure 4.5: The rescaled skewness $\tilde{S}_{3,\beta}$ for $\beta = 1.4$ plotted against σ_8. It is evident that this statistic is nearly independent of σ_8, as expected based on the scalings in Table I. However, it is still dependent on the ICM astrophysics, as represented by the different pressure profiles. Thus, this statistic can be used for determining the correct gas physics model.

4.4 Isolating the Dependence on ICM Astrophysics

As we have described, the cosmological utility of tSZ measurements is limited by their sensitivity to unknown, non-linear astrophysical processes in the ICM, despite their high sensitivity to the underlying cosmology. However, this theoretical uncertainty can be minimized by combining measurements of the tSZ skewness with measurements of the tSZ power spectrum or variance. The use of multiple probes provides an additional statistical handle on both ICM astrophysics and cosmology.

If the tSZ variance and skewness depend differently on the ICM astrophysics and background cosmology, it may be possible to construct a statistic that "cancels" the dependence on one or the other. Here, we focus on a statistic that cancels the dependence on cosmology, but preserves a dependence on the ICM astrophysics. For an alternative approach that attempts to directly cancel the ICM astrophysics dependence, see Appendix A. In particular, we consider the following statistic, which we call the "rescaled skewness":

$$\tilde{S}_{3,\beta} = \frac{|\langle T^3 \rangle|}{\langle T^2 \rangle^{\beta}} \,, \tag{4.6}$$

where the value of β is left to be determined. From the results in Table I, we know that $\langle T^2 \rangle \propto \sigma_8^{6.6-7.9}$ and $\langle T^3 \rangle \propto \sigma_8^{9.7-11.5}$ for these pressure profiles. Thus, if we choose $\beta = 1.4$, then the resulting statistic will cancel the dependence on σ_8, and thus should be nearly independent of the background cosmology, since σ_8 is the dominant cosmological parameter for the tSZ observables. Nonetheless, some dependence on the ICM astrophysics may remain. We investigate this dependence in Fig. 5. It is evident that the rescaled skewness with $\beta = 1.4$ is nearly independent of σ_8, as expected based on the scalings. However, this statistic still shows a dependence on the gas physics model — for $\sigma_8 = 0.817$, the scatter in $\tilde{S}_{3,\beta=1.4}$ between the different pressure profiles is $\approx 35\%$. For the same value of σ_8, the scatter in $\langle T^2 \rangle$ is $\approx 50\%$ between these profiles, while for $\langle T^3 \rangle$ it is $\approx 35\%$. Thus, the scatter in this statistic is essentially identical to that in the skewness. The rescaled skewness with

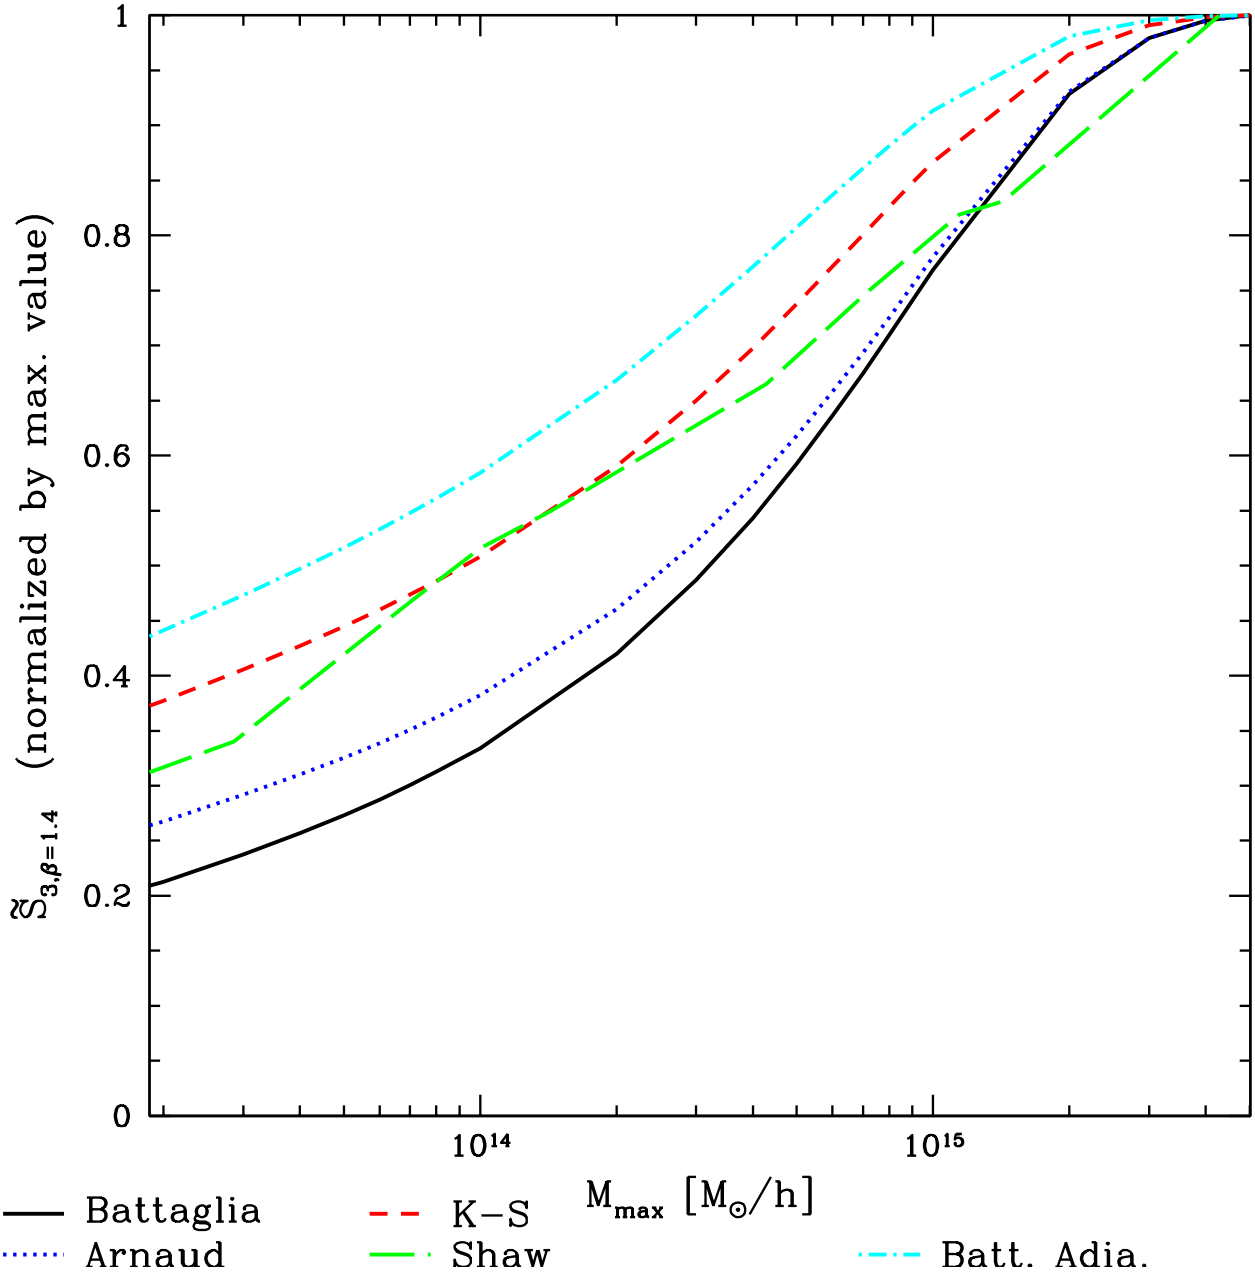

Figure 4.6: Fraction of the rescaled skewness $\tilde{S}_{3,\beta}$ for $\beta = 1.4$ contributed by clusters with virial mass $M < M_{\rm max}$. A significant fraction of the signal comes from low-mass objects — even more so than the variance (Fig. 3). One can thus interpret this statistic as a measure of the gas fraction in low-mass clusters: pressure profiles that yield $f_{gas} \approx \Omega_b/\Omega_m$ in low-mass objects lead to a small value for this statistic (since they give much larger values of $\langle T^2 \rangle$), while pressure profiles that include significant feedback effects — thus lowering the gas fraction in low-mass objects — lead to larger values for this statistic.

$\beta = 1.4$ hence has a very useful property: because it only depends weakly on σ_8, it can be used to determine the correct ICM gas physics model, nearly independent of the background cosmology. Although the value of β could be chosen slightly differently to minimize the small residual dependence on σ_8 for any of the individual pressure profiles, we find that $\beta = 1.4$ is the best model-independent choice available, especially given current levels of observational precision and theoretical uncertainty. After fixing the value of β, one can proceed to measure this statistic from the data (see §IV).

Moreover, the ordering of the results in Fig. 5 suggests a possible interpretation of this statistic. The pressure profiles for which $f_{gas} \approx \Omega_b/\Omega_m$ in all halos — including low-mass groups and clusters — yield low values for $\tilde{S}_{3,\beta=1.4}$, while the profiles with feedback prescriptions that lower f_{gas} in low-mass objects yield high values of this statistic. This can be explained by the fact that the additional contributions from low-mass clusters lead to higher values of $\langle T^2 \rangle$ in the former set of profiles, while $\langle T^3 \rangle$ is not significantly affected, since it is dominated by contributions from more massive objects. Looking at Eq. (4.6), it then follows that the profiles without f_{gas} suppression have lower values for $\tilde{S}_{3,\beta=1.4}$. Thus, the rescaled skewness with $\beta = 1.4$ is a measure of the typical gas fraction in low-mass groups and clusters.

We verify this interpretation by calculating the characteristic mass range responsible for the $\tilde{S}_{3,\beta=1.4}$ signal. Fig. 6 shows the fraction of this signal contributed by clusters with virial mass $M < M_{\rm max}$. As anticipated, it receives significant contributions from very low-mass objects: $\approx 30 - 50\%$ of the amplitude comes from clusters with $M < 6$–7×10^{13} $\mathrm{M}_\odot/h$ (this depends quite strongly on the particular profile used, as seen in the figure). A measurement of this quantity will thus constrain the average amount of gas in these low-mass, high-redshift objects, which have not yet been observed by other techniques. We present a first application of this technique to data from ACT and SPT in the following section.

4.5 Application to ACT and SPT Data

ACT and SPT have measured the tSZ signal in maps of the microwave sky with arcminute angular resolution. In order to apply the rescaled skewness with $\beta = 1.4$ to the data, we rephrase this statistic to match the observational quantities as closely as possible. In particular, we use the SPT measurement of C_{3000} [10], the amplitude of the tSZ power spectrum at $\ell = 3000$, and the ACT measurement of $\langle \tilde{T}^3 \rangle$ [38], the filtered tSZ skewness (defined more precisely below).

We use C_{3000} rather than $\langle T^2 \rangle$ in order to circumvent uncertainties that arise when converting between the tSZ power spectrum and variance. ACT and SPT effectively measure the amplitude of the tSZ power spectrum only at scales of order $\ell = 3000$, and therefore a template for C_ℓ must be used in order to calculate the tSZ variance from this measurement. Such a template relies on an ICM physics model, and since this model is what we are hoping to constrain, we circumvent the calculation of the variance by working directly with C_{3000} instead of $\langle T^2 \rangle$. We use the SPT measurement of C_{3000} due to its higher signal-to-noise ratio (SNR) than the measurement from ACT. Finally, note that SPT measures this quantity at an effective frequency of 152.9 GHz [10], and thus we perform the relevant calculations at this frequency.

ACT recently reported the first detection of a higher-point tSZ observable: the filtered skewness, $\langle \tilde{T}^3 \rangle$ [38]. This quantity is similar to the tSZ skewness that we calculate above, but its value has been modified by an ℓ-space filter applied to the ACT maps, as well as a temperature fluctuation cut-off used to remove outlying pixels in the ACT maps. We explicitly account for these steps in our calculations. First, we Fourier transform the projected temperature decrement profile in Eq. (4.1) and apply the ℓ-space filter used in [38] to each "cluster" of mass M and redshift z in the integrals of Eq. (4.2). We then inverse Fourier transform to obtain the filtered temperature decrement profile for each cluster in real space. Second, we place each cluster in an idealized ACT pixel and compute the observed temperature decrement, accounting carefully for geometric effects that can arise depending

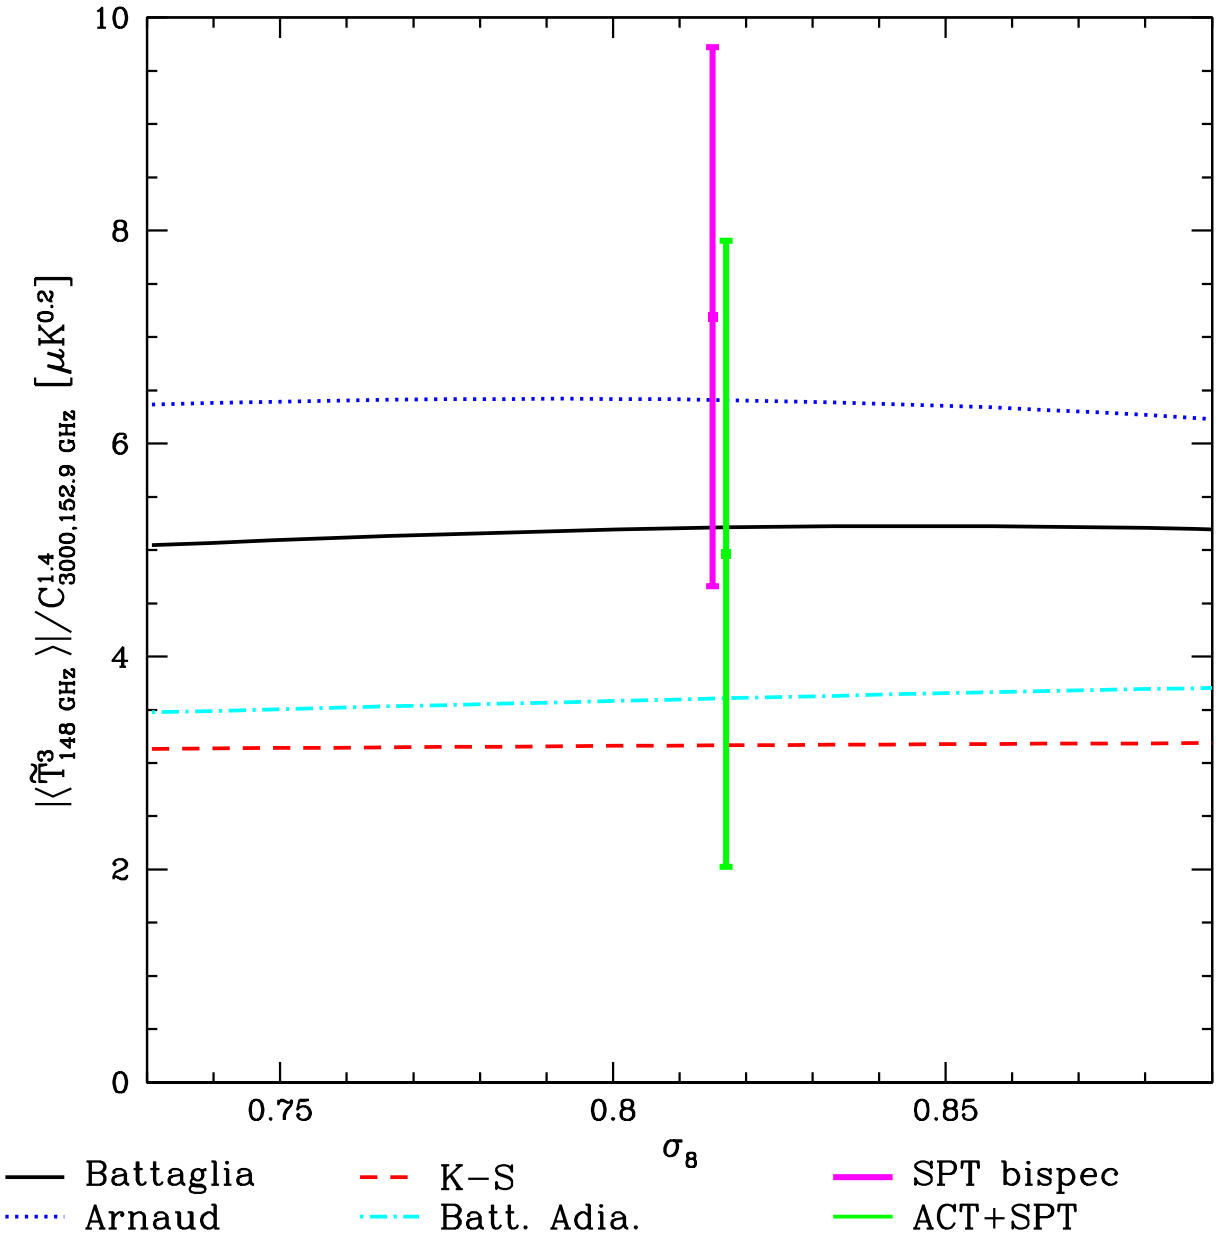

Figure 4.7: Similar to Fig. 5, but calculated in terms of observed tSZ statistics, namely, C_{3000} (the amplitude of the tSZ power spectrum at $\ell = 3000$) and $\langle\tilde{T}^3\rangle$ (the filtered skewness as defined in [38]). The green point shows the constraint using the SPT measurement of C_{3000} [10] and the ACT measurement of $\langle\tilde{T}^3\rangle$ [38]. The magenta point shows an updated constraint using the SPT measurement of the tSZ bispectrum (converted to $\langle\tilde{T}^3\rangle$) [52] instead of the ACT skewness measurement. Note that this figure has been updated from the published version.

on the alignment of the cluster and pixel centers. If the observed temperature decrement exceeds the 12σ cut-off used in [38], then we discard this cluster from the integrals. We thus replicate the data analysis procedure used in the ACT measurement as closely as possible. The net effect is to reduce the value of the tSZ skewness by up to $90-95\%$. The reduction comes almost entirely from the ℓ-space filtering; the cut-off used in the second step only has a small effect. Finally, note that the ACT measurement is at an effective frequency of 148 GHz [9], and thus we perform the relevant calculations at this frequency.

We use these calculations to construct a statistic analogous to $\tilde{S}_{3,\beta=1.4}$, but defined in terms of the observational quantities: $|\langle\tilde{T}^3\rangle|/C_{3000}^{1.4}$. In order for this statistic to work in the manner seen in Fig. 5, we thus require the scalings of C_{3000} and $\langle\tilde{T}^3\rangle$ with σ_8 to be close to those reported in Table I for $\langle T^2\rangle$ and $\langle T^3\rangle$, respectively. Our calculations verify this claim, as can be seen immediately in Fig. 7.

Fig. 7 shows the result from ACT and SPT plotted against the results of our theoretical calculations. Unfortunately, the error bar on the data point is too large to deduce a preference for any particular pressure profile (though the central value is closer to the profiles with significant feedback than those with an unsuppressed value of f_{gas} in all halos). Note that the error bar includes significant contributions due to cosmic variance (i.e., sample variance) resulting from the limited sky coverage of ACT and SPT, as detailed in [38] and [10]. The nearly full-sky results from *Planck* will have far smaller cosmic variance contributions to the error. Moreover, the error is currently dominated by the uncertainty on $\langle\tilde{T}^3\rangle$, which will greatly decrease in the near future — an upcoming SPT measurement should increase the SNR on $\langle\tilde{T}^3\rangle$ by a factor of 3 [34].

Nonetheless, other theoretical issues must still be overcome in order to fully characterize this technique. In particular, the tSZ power spectrum receives non-negligible contributions (up to 15% at $\ell = 3000$) due to deviations about the mean pressure profile found in simulations [29]. The tSZ skewness likely receives similar contributions, which are not accounted for in our halo model-based calculations. Future studies incorporating the results of hydro-

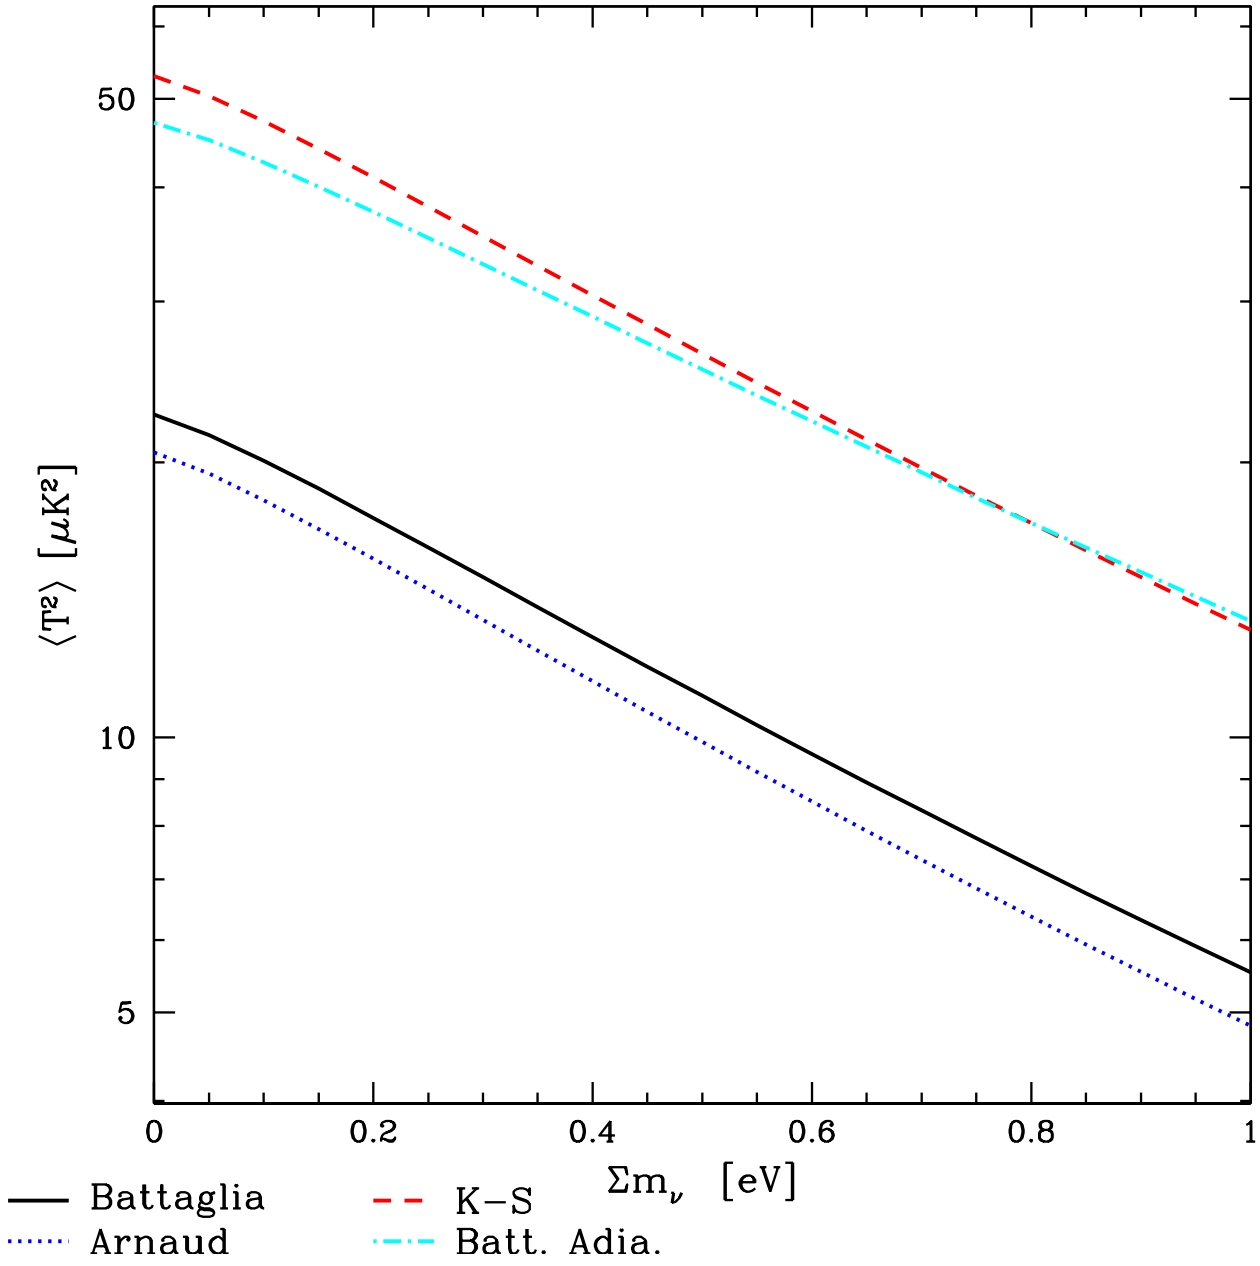

Figure 4.8: The tSZ variance versus the sum of the neutrino masses Σm_ν, with $\Delta^2_{\mathcal{R}} = 2.46 \times 10^{-9}$ (its WMAP5 value). The dependence is not precisely captured by a simple scaling as for σ_8 (note that the axes are log-linear), but the curves are well-fit by quadratic polynomials.

dynamical cosmological simulations will be necessary to fully understand this method. Fig. 7 is simply a proof of concept for this technique, which may soon give interesting constraints on the ICM astrophysics of low-mass groups and clusters. By using the constrained pressure profile to interpret measurements of the power spectrum or skewness, it will be possible to derive stronger constraints on σ_8 from tSZ measurements.

4.6 Future Cosmological Constraints

4.6.1 Extension to other parameters

While we have focused on σ_8, the techniques that we have outlined are in principle sensitive to many cosmological parameters. These include the standard parameters $\Omega_b h^2$, to which the tSZ signal is somewhat sensitive (though this parameter can be very well constrained by the primordial CMB), and Ω_m, to which the tSZ signal is not particularly sensitive (especially in the range $0.15 < \Omega_m < 0.4$) [11]. Here, we instead focus on currently unknown parameters, to which the tSZ moments are sensitive through the halo mass function. Such parameters include the sum of the neutrino masses Σm_ν, the non-Gaussianity parameters (e.g., f_{NL}), and the dark energy equation of state $w(z)$. All of these parameters modify the number of massive clusters in the low-redshift universe: for example, non-zero neutrino masses suppress the number of clusters, while positive f_{NL} provides an enhancement. Thus, these parameters also suppress or enhance the amplitude of the tSZ moments.

As a first example, we compute the tSZ variance and skewness for a fixed WMAP5 background cosmology with massive neutrinos added. To be clear, we fix $\Delta^2_{\mathcal{R}} = 2.46 \times 10^{-9}$ (its WMAP5 value), not $\sigma_8 = 0.817$, as the presence of massive neutrinos will lead to a lower value of σ_8 inferred at the low redshifts from which the tSZ signal originates. We compute the change in the mass function due to the neutrinos following a prescription similar to that of [39], except that we input the suppressed linear theory matter power spectrum to the

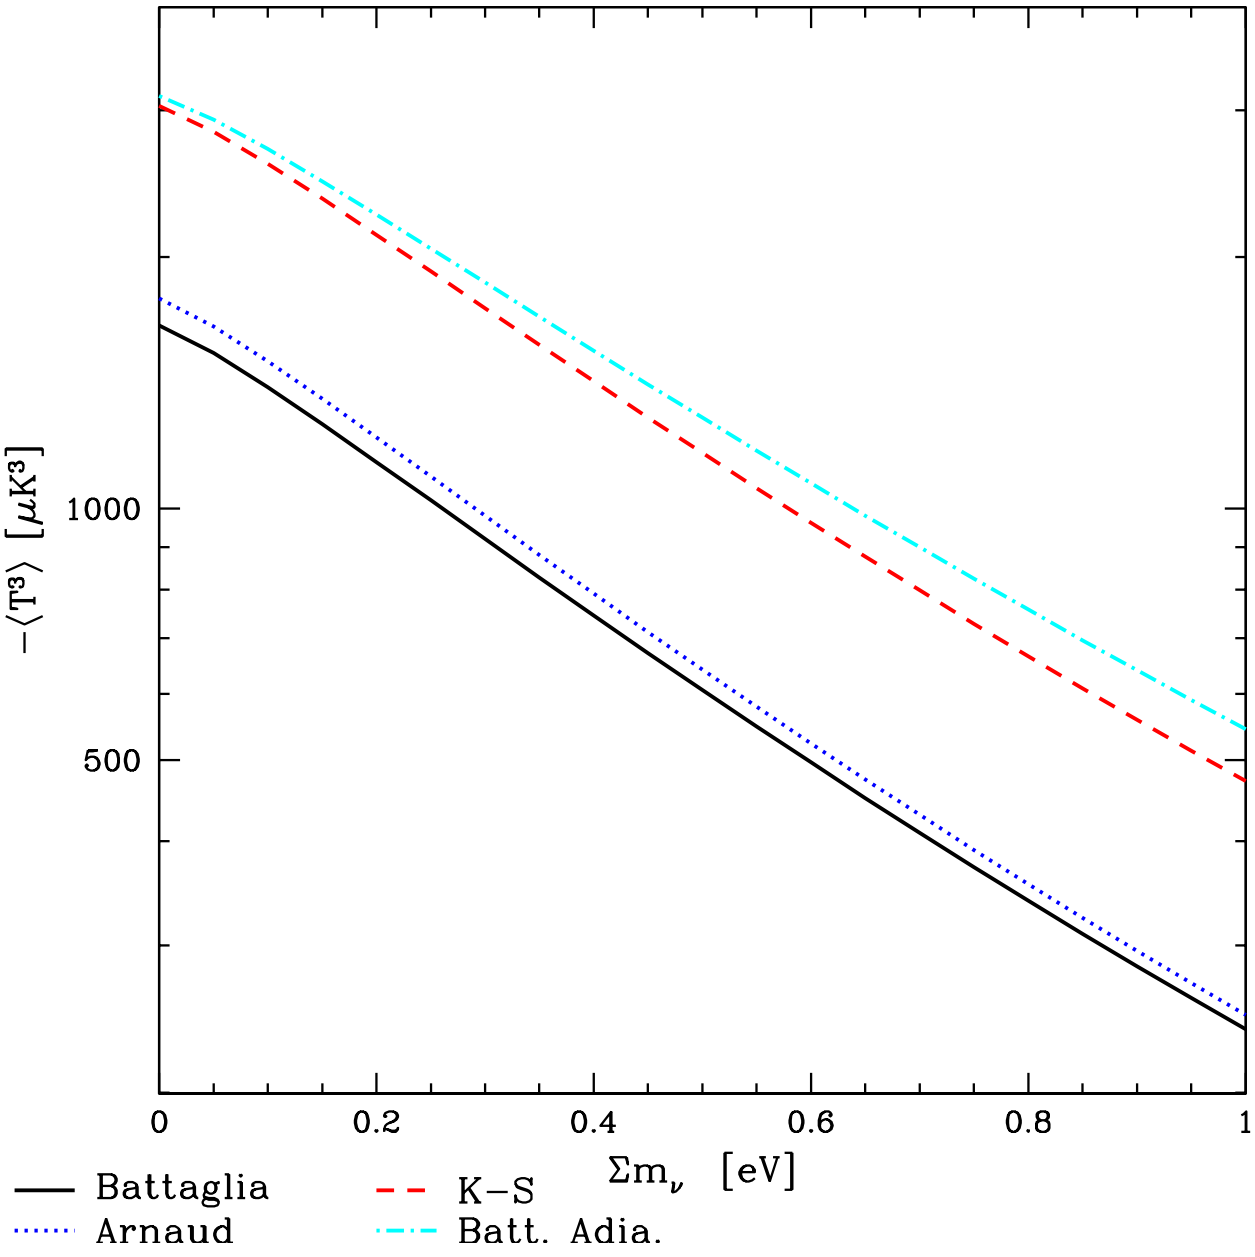

Figure 4.9: The tSZ skewness versus the sum of the neutrino masses Σm_ν, with $\Delta^2_{\mathcal{R}} = 2.46 \times 10^{-9}$ (its WMAP5 value). The dependence is not precisely captured by a simple scaling as for σ_8 (note that the axes are log-linear), but the curves are well-fit by cubic polynomials.

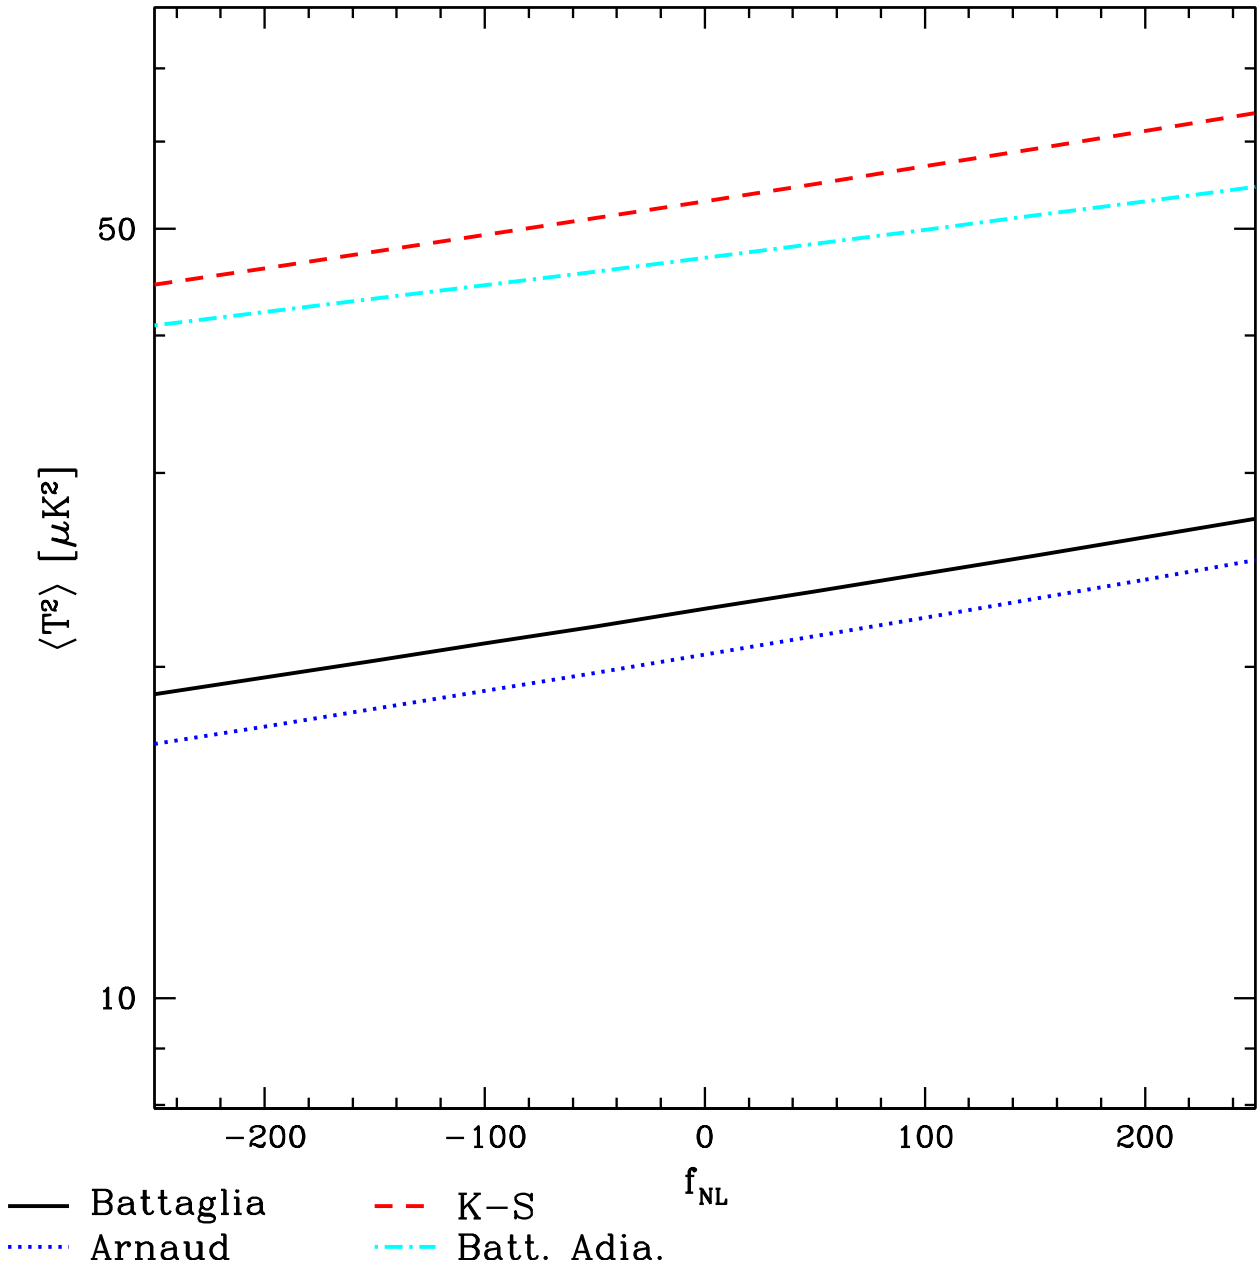

Figure 4.10: The tSZ variance versus f_{NL}, with $\Delta^2_{\mathcal{R}} = 2.46 \times 10^{-9}$ (its WMAP5 value). The dependence is not precisely captured by a simple scaling as for σ_8 (note that the axes are log-linear), but the curves are well-fit by quadratic polynomials.

Tinker mass function rather than that of [40], as was done in [39]. A more precise recipe can be found in [41, 42], but this approach should capture the relevant physical effects.

The results of this analysis are shown in Figs. 8 and 9. As expected, the tSZ variance and skewness both exhibit sensitivity to Σm_ν, although it is currently overwhelmed by the uncertainty in the ICM astrophysics. However, using the method outlined in the previous section to constrain the ICM astrophysics and thus break its degeneracy with cosmology, a measurement of the sum of the neutrino masses with the tSZ effect might be feasible.

As a second example, we compute the tSZ variance and skewness for a fixed WMAP5 background cosmology with a non-zero value of f_{NL} added. Again, to be clear, we fix $\Delta^2_{\mathcal{R}} = 2.46 \times 10^{-9}$ (its WMAP5 value), not $\sigma_8 = 0.817$, as the non-zero value of f_{NL} will change the value of σ_8 inferred at the low redshifts from which the tSZ signal originates. We compute the change in the mass function due to primordial non-Gaussianity using the results of [43]. Although this is technically a "friends-of-friends" (FOF) mass function, we follow [44] in assuming that the *ratio* of the non-Gaussian mass function to the Gaussian mass function is nearly universal (and hence applicable to both FOF and spherical overdensity mass functions), even if the underlying mass function itself is not. While this assumption may not be valid at the percent level, our approach should capture the relevant physical effects of primordial non-Gaussianity on the tSZ moments. In particular, $f_{NL} > 0$ leads to a significant increase in the number of massive halos at late times, while $f_{NL} < 0$ has the opposite effect [45, 43]. Thus, the amplitudes of the tSZ statistics should increase or decrease accordingly, as they are sourced by these massive halos.

The results of this analysis are shown in Figs. 10 and 11. As expected, the tSZ variance and skewness both exhibit sensitivity to f_{NL}, although for the currently allowed range of values ($-10 < f_{NL} < 74$ [46] for the "local" shape) it is overwhelmed by the uncertainty in the ICM astrophysics. However, using the method outlined in the previous section to constrain the ICM astrophysics and thus break its degeneracy with cosmology, a tight constraint on primordial non-Gaussianity with the tSZ effect might be achievable. In addition, it is

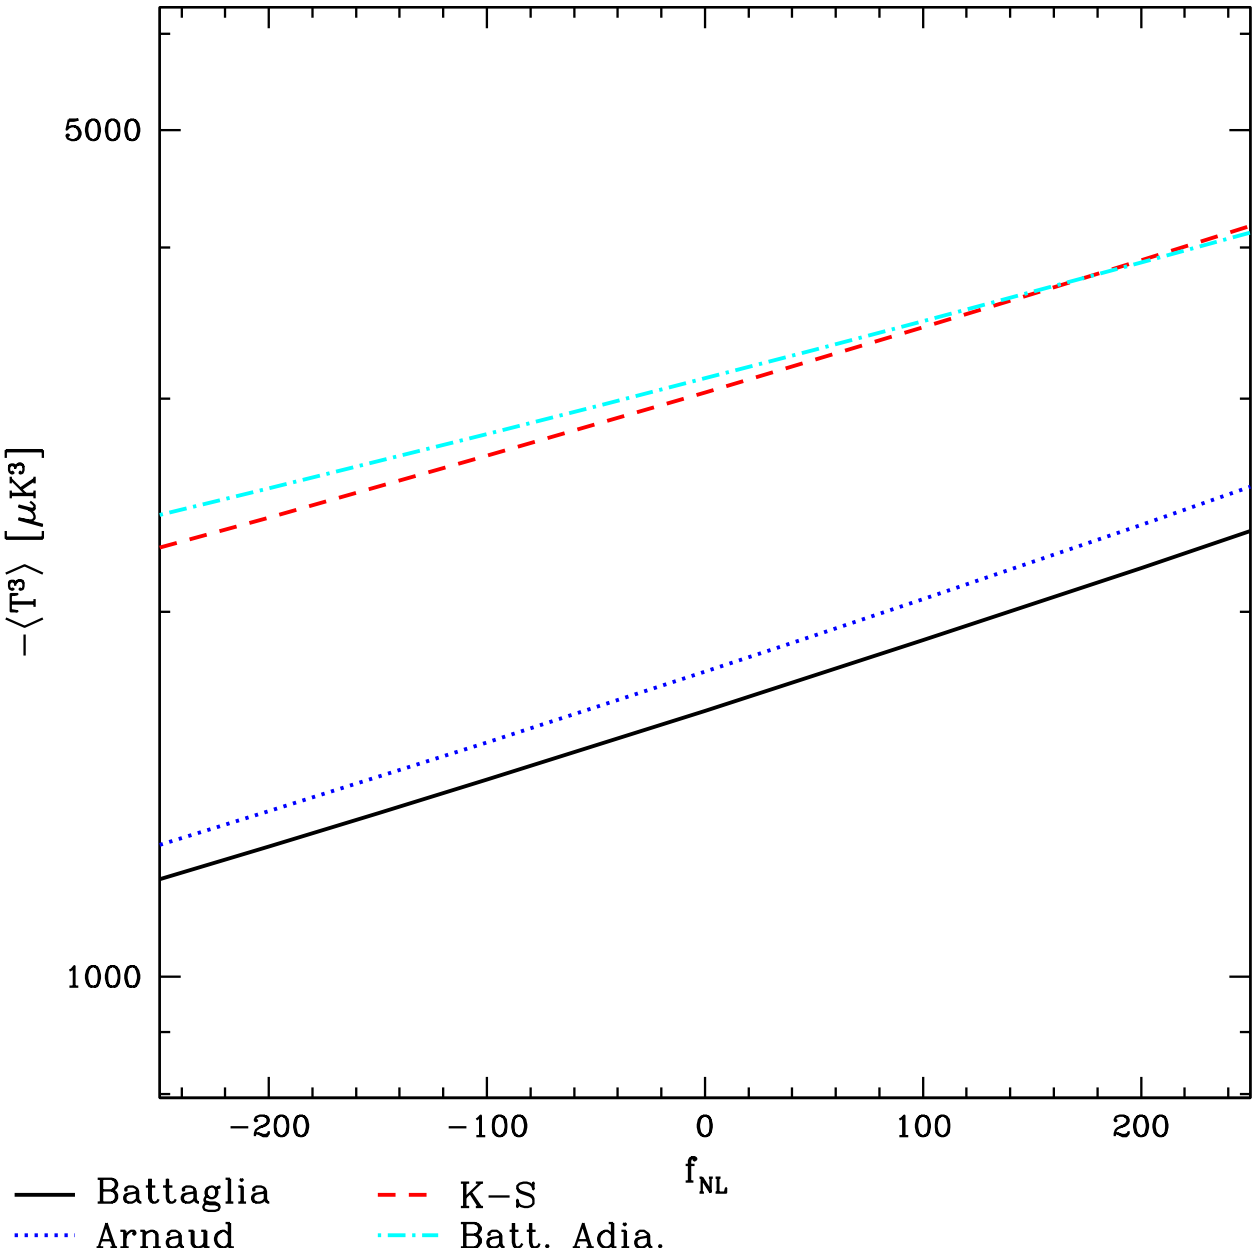

Figure 4.11: The tSZ skewness versus f_{NL}, with $\Delta^2_{\mathcal{R}} = 2.46 \times 10^{-9}$ (its WMAP5 value). The dependence is not precisely captured by a simple scaling as for σ_8 (note that the axes are log-linear), but the curves are well-fit by cubic polynomials.

worth noting that the tSZ moments are sensitive to non-Gaussianity on cluster scales, rather than the large scales probed by the primordial CMB or large-scale halo bias. Thus, the tSZ signal can constrain f_{NL} on much smaller scales than other observables, allowing for complementary constraints that are relevant for models of inflation that predict strongly scale-dependent non-Gaussianity [47].

4.6.2 Estimating the constraints from *Planck*

We briefly estimate the expected constraints on σ_8 and Σm_ν using tSZ measurements from *Planck*. *Planck*'s spectral coverage will allow the separation of the tSZ signal from other CMB components, likely yielding a full-sky tSZ map. However, given issues with bandpass uncertainties and CO contamination, it is difficult to predict the SNR for a *Planck* measurement of $\tilde{S}_{3,\beta=1.4}$. Based on the Sehgal simulation analysis above (which covers an octant of the sky), we anticipate that a cosmic variance-limited full-sky measurement could achieve a SNR ≈ 90 for the tSZ variance or a SNR ≈ 35 for the tSZ skewness. For $\tilde{S}_{3,\beta=1.4}$, we estimate a SNR ≈ 55, which implies that the ICM astrophysics can be constrained to the $\lesssim 5\%$ level. This result will allow for a $\lesssim 1\%$ constraint on σ_8 after subsequently applying the correct model to interpret tSZ measurements. Note that these estimates fully account for both Poisson and cosmic variance error, as they are derived from a cosmological simulation. Technically, we have neglected the fact that the cosmic variance error will scale with σ_8 (see the discussion in [38]), but as long as the true value is not so far from $\sigma_8 = 0.8$ as to conflict with all recent measurements of this parameter, our estimates will be accurate.

Forecasting a Σm_ν constraint is somewhat more difficult because it is highly degenerate with σ_8, and we have not accounted for other possible parameter dependences (e.g., one might expect $\Omega_m h^2$ to enter through its role in the mass function). However, as described in [39], the effect of massive neutrinos is essentially captured by a low-redshift suppression of σ_8 as compared to the value inferred from the primordial CMB at $z = 1100$ (i.e., $\Delta^2_{\mathcal{R}}$). This line of reasoning was also used in the neutrino mass constraint derived by the *Chandra*

Cluster Cosmology Project [48], who found an upper bound of $\Sigma m_\nu < 0.33$ eV at 95% CL. Based on the results of [39], $\Sigma m_\nu \approx 0.1$ eV is roughly equivalent to a $\approx 3\%$ decrease in σ_8 as inferred at $z = 1$ compared to $z = 1100$. The WMAP5 measurement of σ_8 has an uncertainty of $\approx 3\%$ [20], and the equivalent result from *Planck* should be $\lesssim 1\%$ [49]. If the technique that we have outlined works as planned, it could yield a $\lesssim 1\%$ measurement of σ_8 at low redshift. This corresponds to a sensitivity to $\Sigma m_\nu \approx 0.1$–0.2 eV. Given the known lower bound of $\Sigma m_\nu \gtrsim 0.05$ eV, it might even be possible to detect the sum of the neutrino masses through tSZ measurements. Clearly the *Planck* results will have a smaller SNR than the full-sky, cosmic variance-limited assumption made above — and there are important degeneracies between the relevant parameters — but nonetheless these methods show great promise in constraining any parameter that affects the tSZ signal through the mass function.

4.7 Acknowledgments

We thank Nick Battaglia, Eiichiro Komatsu, Marilena LoVerde, Daisuke Nagai, Jeremiah Ostriker, Neelima Sehgal, Laurie Shaw, David Spergel, and Matias Zaldarriaga for useful conversations. We also acknowledge Eiichiro Komatsu's online cosmology routine library, which was a helpful reference. As this work neared completion, we became aware of related results from S. Bhattacharya et al. [34], with whom we subsequently exchanged correspondence. BDS was supported by an NSF Graduate Research Fellowship.

4.8 Addendum

4.8.1 Directly Canceling the Gas Physics

In this appendix, we present a different approach to circumventing the degeneracy between ICM astrophysics and cosmological parameters in tSZ measurements. In particular, we in-

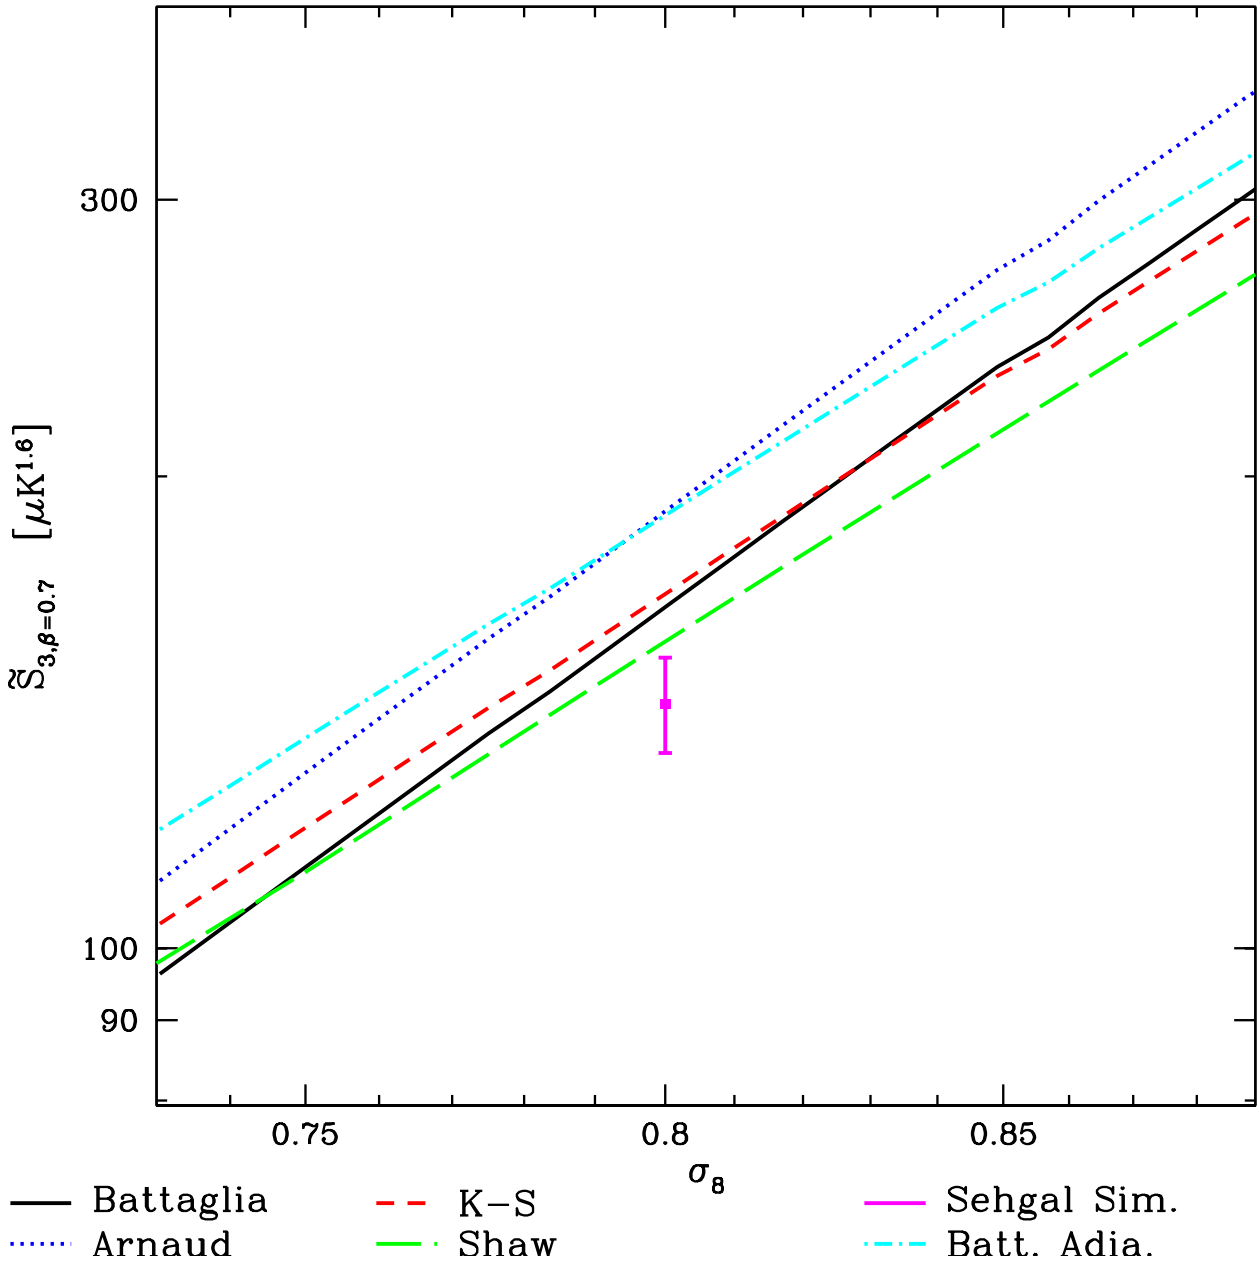

Figure 4.12: The rescaled skewness $\tilde{S}_{3,\beta}$ for $\beta = 0.7$ plotted against σ_8. Clearly there is significantly less scatter for this statistic between the different pressure profiles, but it still scales as σ_8^{5-6}. Thus, this value of β provides a method for constraining cosmological parameters that is fairly independent of the details of the gas physics. The simulation point is somewhat low, but the disagreement is not statistically significant. This point was also not used in determining the optimal value of β.

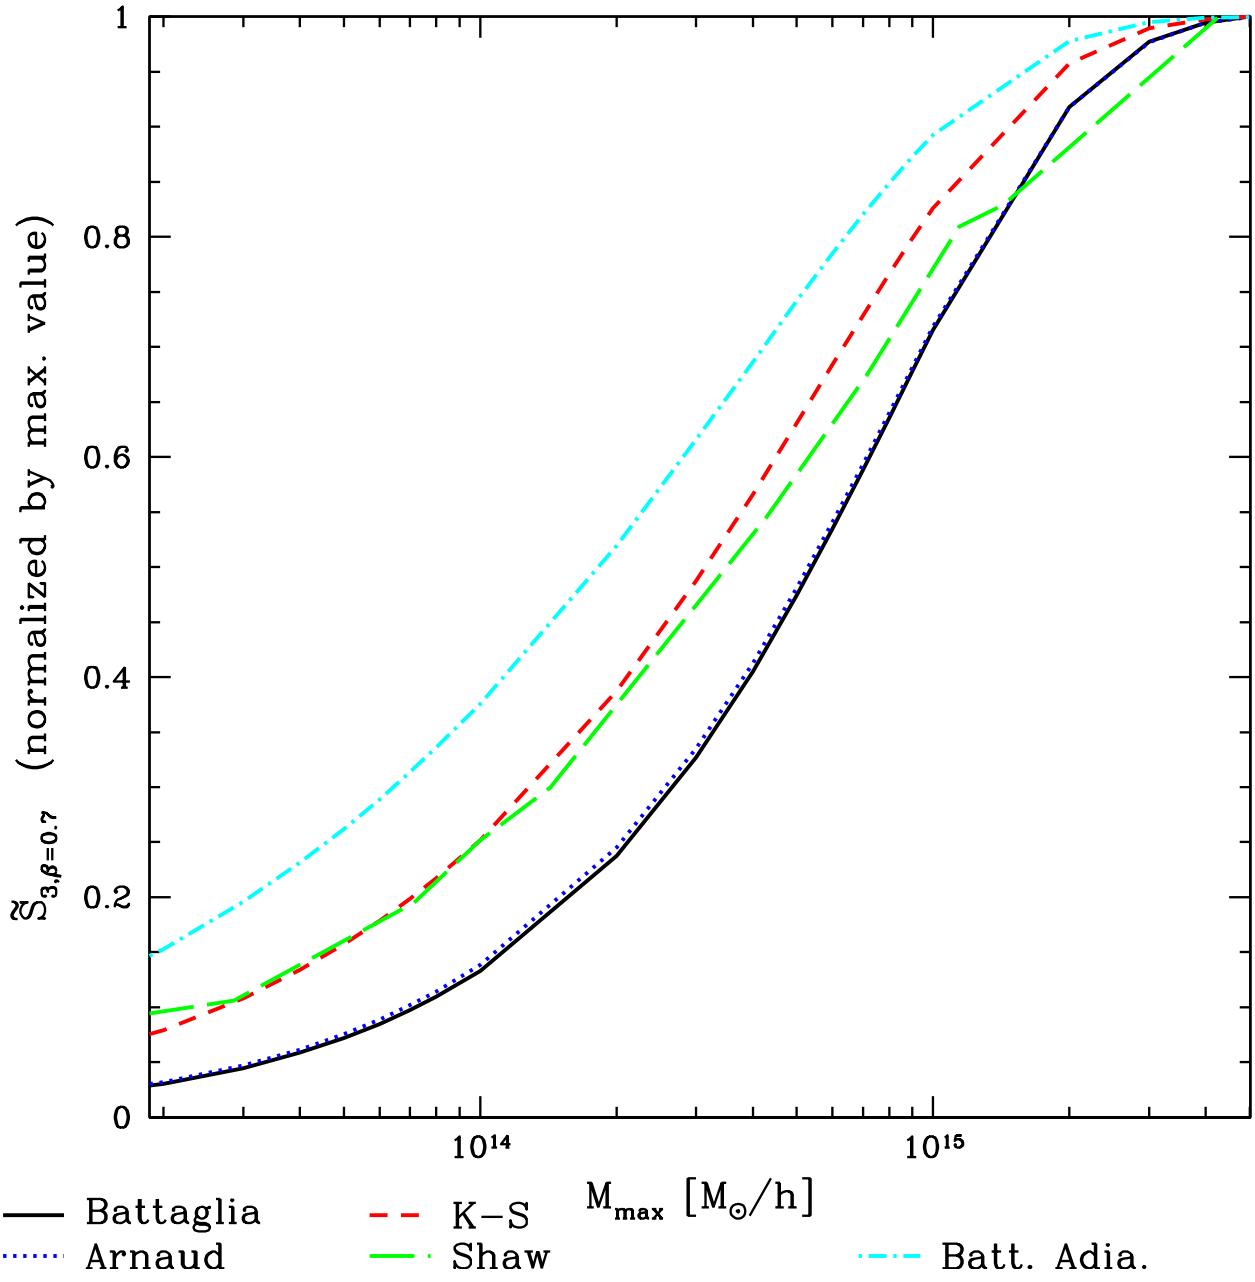

Figure 4.13: Contributions to the rescaled skewness with $\beta = 0.7$ from clusters of various masses, for each of the different pressure profiles. The contributions to this statistic are very similar to those found for the skewness in Fig. 4, as one would expect based on the definition of $\tilde{S}_{3,\beta=0.7}$.

vestigate the rescaled skewness statistic in Eq. (4.6) for other values of the exponent β to see if they might directly cancel the gas physics dependence, but preserve a sensitivity to the underlying cosmology. We heuristically motivate the magnitude of this exponent in the following way. We consider the effect of increasing the amount of energy input through feedback on both the tSZ variance and skewness in a theoretical model for the ICM. Increasing the feedback should substantially reduce the signal from low-mass clusters, and generally affect the pressure profile in the outer regions of clusters. However, high-mass clusters are much less affected by feedback, so their tSZ signal should be essentially unchanged. This fact implies that the variance, which depends more on the signal from low-mass clusters (as shown earlier) and receives contributions at large distances from the cluster center, should be reduced much more by feedback than the skewness, which depends more on the signal from massive clusters and is mostly produced in the inner regions of clusters. In order to construct a rescaled skewness in Eq. (4.6) that is unchanged as the amount of astrophysical feedback is increased, we expect to exponentiate the variance (which is more sensitive to feedback) to a lower power than the skewness, in order to compensate. Thus, we anticipate $\beta < 1$ for a statistic that is somewhat insensitive to the ICM astrophysics, but remains sensitive to the underlying cosmology.

To derive this exponent quantitatively, we find the value of β that minimizes the scatter between different choices of pressure profile for the same values of the cosmological parameters. We focus specifically on σ_8, as the tSZ moments are most sensitive to this parameter. For $0.73 < \sigma_8 < 0.9$ (and all other parameters taking their WMAP5 values), we find that $\beta = 0.7$ minimizes the scatter between the different pressure profiles. As these profiles span a wide range of ICM prescriptions and tSZ power spectra predictions, we expect this result to be fairly robust. We verify this claim by computing the rescaled skewness with $\beta = 0.7$ for the Sehgal simulation. The results are shown in Fig. 12. Although the simulation point is somewhat low, it is only 1σ away from the halo model-based results at $\sigma_8 = 0.8$.

In order to gain more understanding of the rescaled skewness with $\beta = 0.7$, we also investigate the characteristic mass scales that contribute to this statistic. The results are shown in Fig. 13. For a WMAP5 cosmology, we find that $\tilde{S}_{3,\beta=0.7}$ receives $\approx 25 - 50\%$ of its amplitude from clusters with $M < 2$–3×10^{14} $M_\odot/h$. In general, the signal for $\tilde{S}_{3,\beta=0.7}$ is dominated by mass scales similar to those that comprise the skewness, i.e., higher masses than those responsible for the variance. It is interesting to note that this statistic receives contributions from rather different mass scales for the different profiles in Fig. 13, even though the final result is fairly independent of the profile choice (as seen in Fig. 12).

We also note that the same statistic with $\beta = 0.7$ behaves similarly when applied to Σm_ν, as seen in Fig. 14. To be clear, we do not re-fit for the value of β that minimizes the difference between the various pressure profiles in this plot. If we did perform such a fit, the best-fitting exponent to minimize the dispersion is again $\beta = 0.7$, in exact agreement with the best-fitting value for σ_8.

Despite the promising nature of the $\tilde{S}_{3,\beta=0.7}$ statistic, we must emphasize that we lack a strong theoretical motivation for this quantity, and thus one must be cautious in assuming its insensitivity to changes in the gas physics. For example, if one alters the amount of non-thermal pressure support by changing an overall parameter that affects the normalization of all halos' pressure profiles (e.g., the α_0 parameter in the Shaw model [33]), then this statistic will be affected in a nontrivial manner. It is likely that requiring the model to match the observed X-ray data for massive, low-redshift clusters will prevent a drastic change along these lines, but it is still a possibility. In addition, as mentioned earlier, the tSZ skewness is fairly sensitive to the choice of mass function used in these calculations (though the variance is less affected), which will add additional scatter and uncertainty to the relation presented in Fig. 12.

Nonetheless, we are encouraged by the mild sensitivity of $\tilde{S}_{3,\beta=0.7}$ to the complicated astrophysics of the ICM. Moreover, it retains a strong dependence on σ_8. In particular, using the results presented in Table I for $\langle T^2 \rangle$ and $\langle T^3 \rangle$, we find that $\tilde{S}_{3,\beta=0.7} \propto \sigma_8^{5-6}$. For

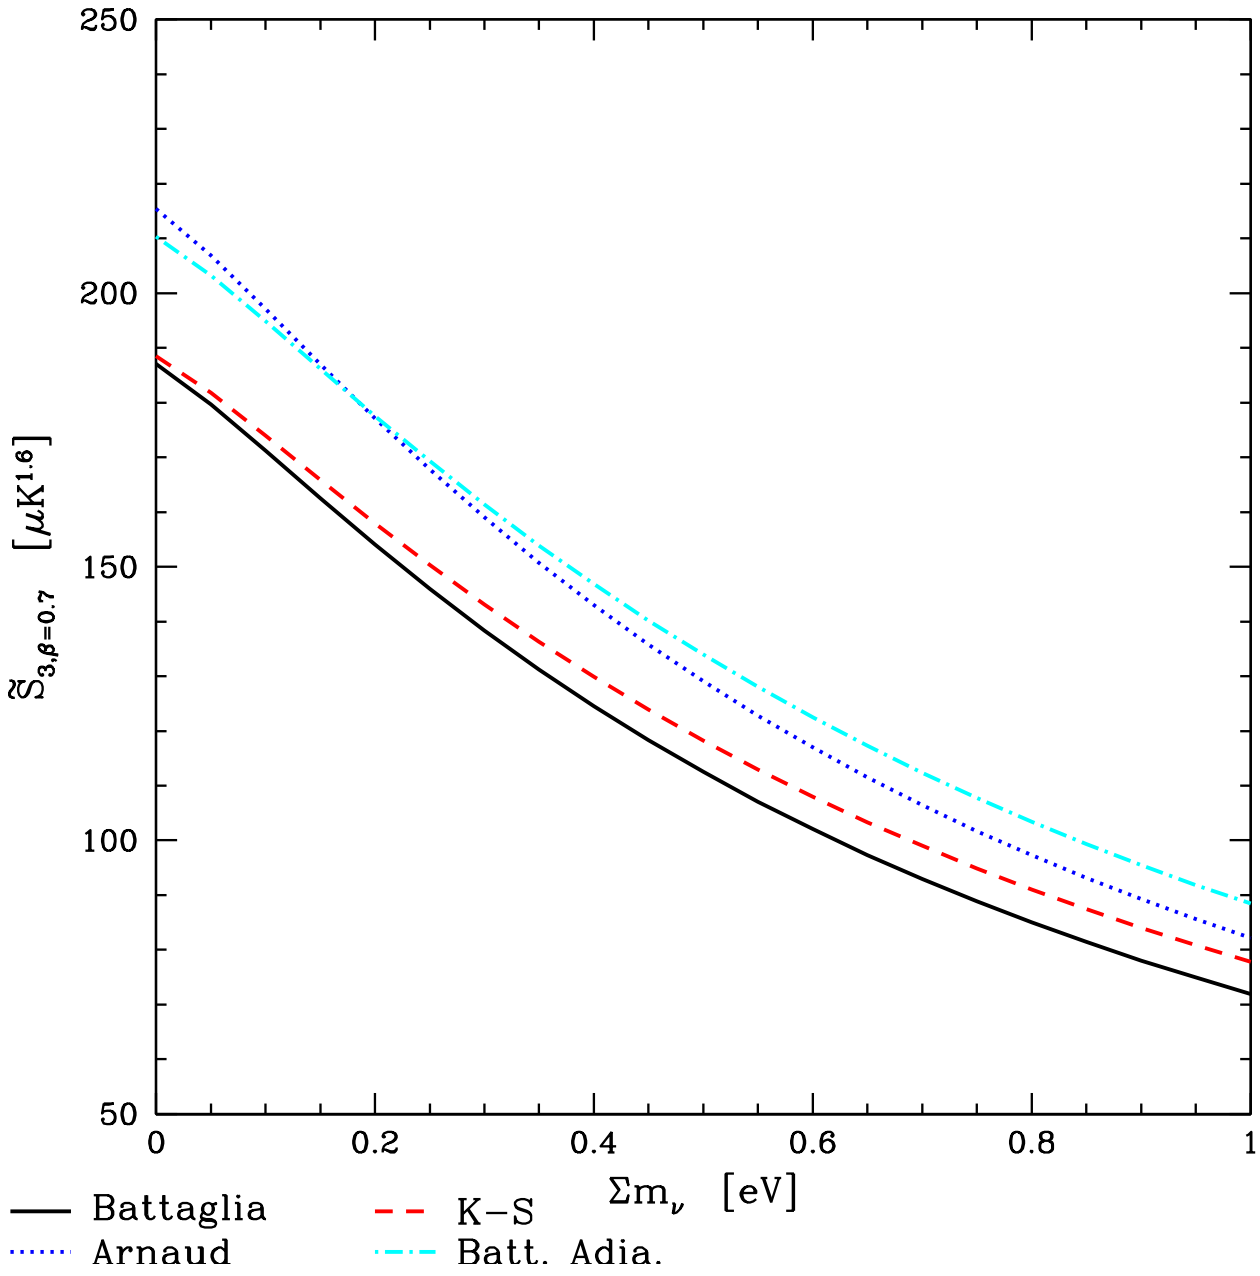

Figure 4.14: Similar to Fig. 12, but now showing the rescaled skewness with $\beta = 0.7$ plotted against the sum of the neutrino masses Σm_ν. We assume $\Delta^2_{\mathcal{R}} = 2.46 \times 10^{-9}$ (its WMAP5 value). There is significantly reduced scatter between the pressure profiles for this statistic, and it retains a leading-order quadratic dependence on Σm_ν.

a fixed value of $\sigma_8 = 0.817$, the scatter in $\tilde{S}_{3,\beta=0.7}$ between the different profiles is only 7%. This scatter is much smaller than that in the variance or skewness: for the same value of $\sigma_8 = 0.817$, the scatter in $\langle T^2 \rangle$ is $\approx 50\%$ between these profiles, while for $\langle T^3 \rangle$ it is $\approx 35\%$. This statistic could thus be used for a high-precision determination of σ_8 from tSZ measurements, with significantly reduced sensitivity to theoretical systematic uncertainties due to unknown physics in the ICM. For a full-sky, CV-limited experiment, we estimate a SNR ≈ 40 for $\tilde{S}_{3,\beta=0.7}$, which implies a $\lesssim 1\%$ constraint on σ_8, in agreement with the estimates presented for $\tilde{S}_{3,\beta=1.4}$ earlier.

However, before applying this technique to real data, its insensitivity to the gas physics prescription must be tested with detailed numerical tSZ simulations, as we lack a good understanding of the theoretical origin of $\tilde{S}_{3,\beta=0.7}$. The rescaled skewness with $\beta = 1.4$, though, is well-motivated theoretically and provides a clear method to constrain the ICM astrophysics in low-mass halos, which will greatly reduce the theoretical systematic uncertainty in tSZ measurements. We defer an application of the $\beta = 0.7$ approach to future work.

4.8.2 Higher-Order tSZ Moments

In the main text of the chapter, we focus on the second (variance) and third (unnormalized skewness) moments of the tSZ signal. In this appendix, we provide additional calculations of the fourth moment (the unnormalized kurtosis) and investigate a "rescaled kurtosis" statistic similar to the rescaled skewness defined in Eq. (4.6). The unnormalized kurtosis is simply given by Eq. (4.2) with $N = 4$. As for the variance and skewness, we compute this quantity for each of the various pressure profiles described in the text, and investigate the dependence on σ_8. We define an amplitude A_4 and scaling α_4 precisely analogous to $A_{2,3}$ and $\alpha_{2,3}$ in Eq. (4.5). The results for each of the profiles are given in Table II.

As for the variance and skewness, the profiles that account for feedback processes (Battaglia and Arnaud) have a steeper scaling with σ_8 than those that do not, because the kurtosis signal for these profiles is dominated to an even greater extent by the most massive,

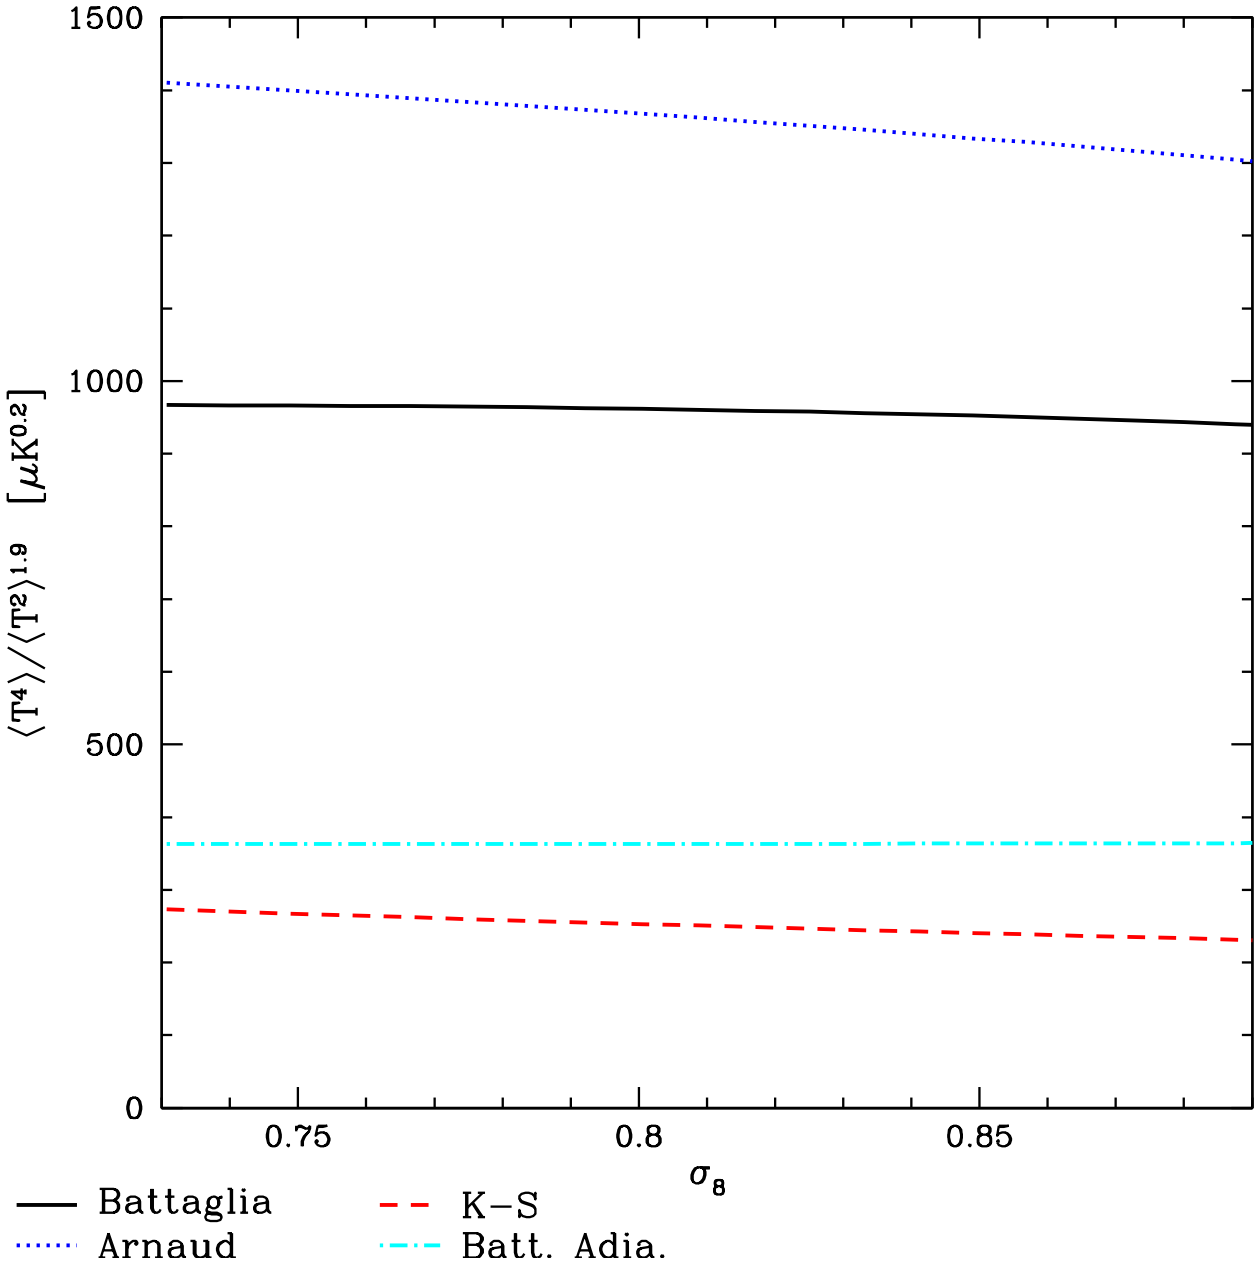

Figure 4.15: The rescaled kurtosis $\langle T^4\rangle/\langle T^2\rangle^{1.9}$ plotted against σ_8. It is evident that this statistic is nearly independent of σ_8, as expected based on the scalings in Tables I and II. However, it is still dependent on the ICM astrophysics, as represented by the different pressure profiles. Thus, this statistic — like the rescaled skewness — can be used for determining the correct gas physics model.

Profile	A_4 [10^5 μK^4]	α_4
Arnaud	4.22	14.7
Battaglia	3.58	14.5
Batt. Adiabatic	5.48	12.5
Komatsu-Seljak	4.70	13.5

Table 4.2: Amplitudes and power-law scalings with σ_8 for the tSZ kurtosis. The first column lists the pressure profile used in the calculation (note that all calculations use the Tinker mass function). The amplitudes are specified at $\sigma_8 = 0.817$, the WMAP5 maximum-likelihood value. All results are computed at $\nu = 150$ GHz.

rare halos. Additionally, note that the scatter between the various profiles is even smaller than that for the skewness (and much smaller than that for the variance).

Furthermore, it is possible to construct a "rescaled kurtosis" statistic similar to the rescaled skewness defined in Eq. (4.6), which should be nearly independent of the background cosmology. In particular, based on the scalings in Tables I and II, one can immediately see that the quantity $\langle T^4 \rangle / \langle T^2 \rangle^{1.9}$ effectively cancels the dependence on σ_8 for each of the profiles. This result is clearly seen in Fig. 15, which presents this statistic as a function of σ_8. Encouragingly, a clear dependence on the ICM astrophysics persists, as for the rescaled skewness in Fig. 5. Note that the same hierarchy is seen between the different profiles in both figures — again, this arises because of the suppression of the tSZ signal in low-mass halos for the profiles that include significant feedback, which decreases their value of the variance, thus increasing their value of the rescaled skewness and kurtosis. In addition, it appears that the rescaled kurtosis (Fig. 15) is somewhat less dependent on σ_8 than the rescaled skewness (Fig. 5), although the difference is not very significant over the feasible range of σ_8. Physically, this is likely due to the fact that the kurtosis is dominated even more than the skewness by very massive halos, for which the different pressure profiles do not differ widely in their predictions. Thus, the various profiles have a similar relative scaling between the kurtosis and the variance, making the rescaled kurtosis nearly independent of σ_8 for the same choice of "rescaling exponent". We verify this interpretation in Fig. 16, which shows the fraction of this signal contributed by clusters with virial mass $M < M_{\rm max}$. Compared to the rescaled skewness (Fig. 6), the rescaled kurtosis is sourced more predominantly by

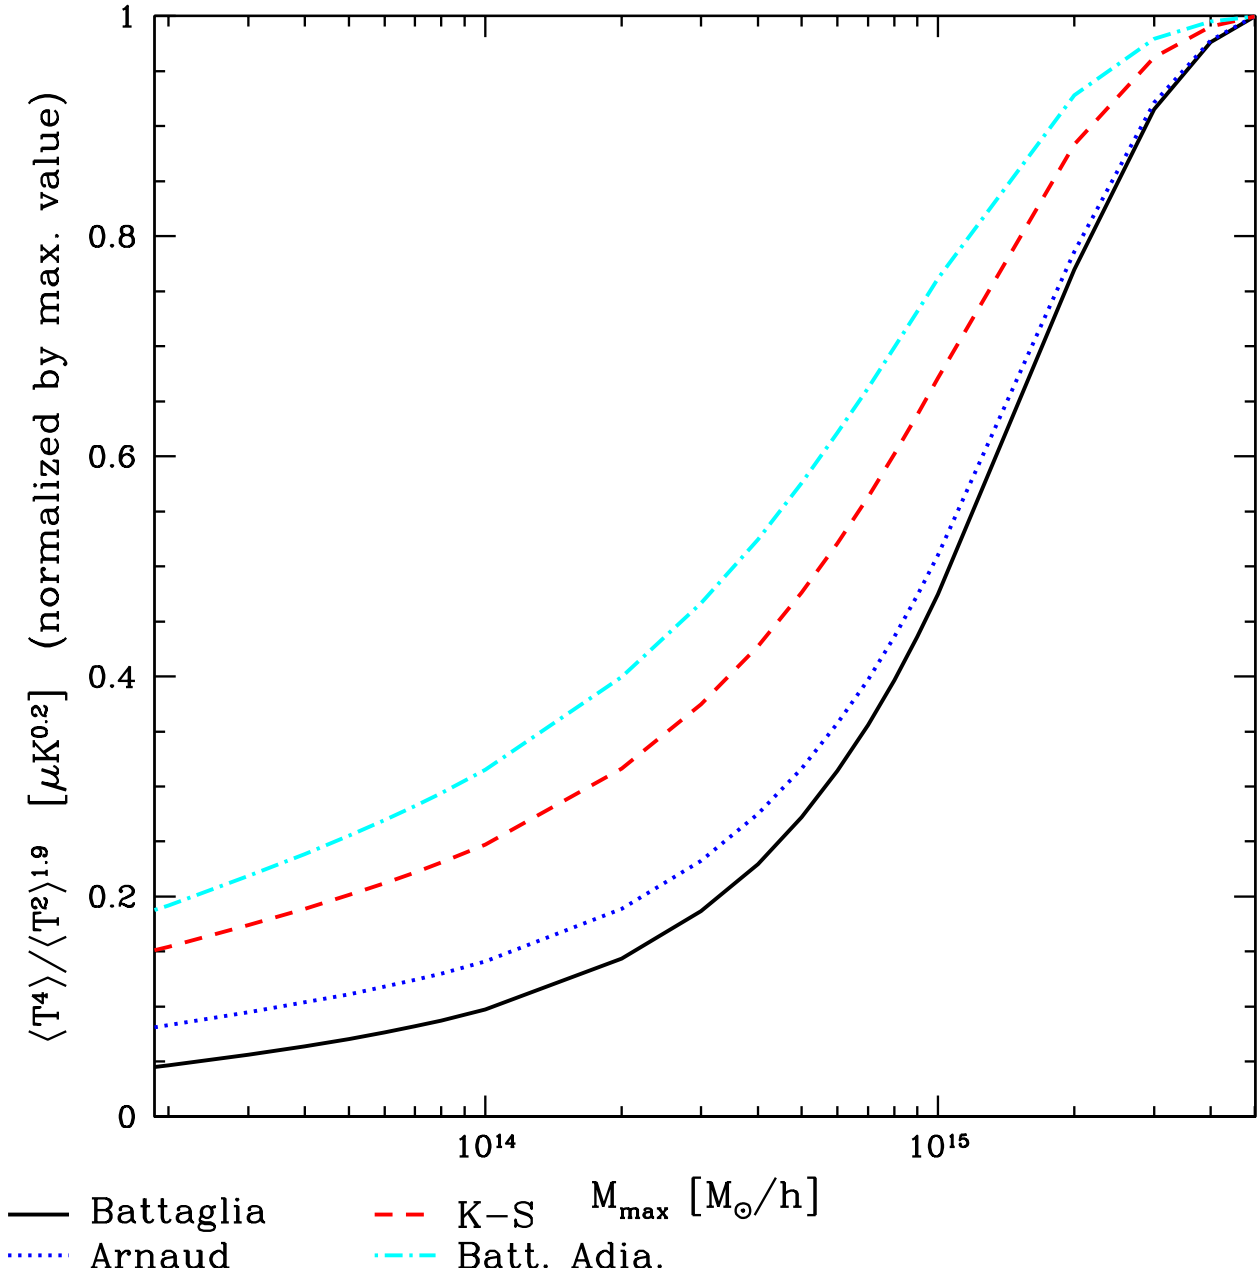

Figure 4.16: Fraction of the rescaled kurtosis $\langle T^4\rangle/\langle T^2\rangle^{1.9}$ contributed by clusters with virial mass $M < M_{\rm max}$. Although a non-trivial fraction of the signal comes from low-mass objects, comparison with Fig. 6 indicates that the rescaled kurtosis is mostly sourced by more massive halos than those responsible for the rescaled skewness.

fairly massive halos. The rescaled kurtosis thus also provides a route to determining the gas physics of the ICM, nearly independent of the background cosmology.

Unfortunately, a measurement of the tSZ kurtosis appears extremely challenging with current data, both because of confusion with other signals (e.g., the kinetic SZ effect, point sources, and so on) and because the cosmic variance for this statistic is quite large, as it is sourced by very massive clusters. Previous work on the kinetic SZ kurtosis allows some rough estimates of this contamination to be made, however [50, 51]. In [50], estimates are given of the normalized (dimensionless) skewness and kurtosis for the tSZ and kSZ effects, as determined from hydrodynamical simulations. Based on their results, it appears that the dimensionless kurtosis due to the tSZ effect is at least a few times larger than that due to the kSZ effect. Thus, assuming that the variance of the kSZ signal is no larger than that of the tSZ signal, this implies that the value of $\langle T^4 \rangle$ that we have calculated for the tSZ effect will be at least a few times larger than the contribution due to the kSZ effect (at 150 GHz). Hence, the kSZ contamination may not be a major problem for a detection of the tSZ kurtosis. However, contamination from point sources will likely remain a significant problem for an experiment with only one or two frequencies. Thus, ACT and SPT are unlikely to make a detection, though *Planck* may do so. Neglecting any potential contaminants and using the Sehgal simulation analysis described earlier, we find that a cosmic variance-limited full-sky experiment could achieve a SNR ≈ 18 for the tSZ kurtosis or a SNR ≈ 22 for the rescaled kurtosis. The results from *Planck* will clearly have a somewhat lower SNR than these estimates, but it may still be worthwhile to investigate these quantities using the forthcoming *Planck* sky maps.

Bibliography

[1] Sunyaev, R. A., & Zeldovich, Y. B. 1970, Ap&SS, 7, 3

[2] Reese, E. D., Mroczkowski, T., Menanteau, F., et al. 2011, arXiv:1108.3343

[3] Plagge, T. J., Marrone, D. P., Abdulla, Z., et al. 2012, arXiv:1203.2175

[4] Lancaster, K., Birkinshaw, M., Gawroński, M. P., et al. 2011, Mon. Not. R. Astron. Soc, 418, 1441

[5] AMI Consortium, Shimwell, T. W., Rodriguez-Gonzalvez, C., et al. 2011, arXiv:1101.5590

[6] Marriage, T. A., Acquaviva, V., Ade, P. A. R., et al. 2011, ApJ, 737, 61

[7] Williamson, R., Benson, B. A., High, F. W., et al. 2011, ApJ, 738, 139

[8] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2011, A&A, 536, A8

[9] Dunkley, J., Hlozek, R., Sievers, J., et al. 2011, ApJ, 739, 52

[10] Reichardt, C. L., Shaw, L., Zahn, O., et al. 2011, arXiv:1111.0932

[11] Komatsu, E., & Seljak, U. 2002, Mon. Not. R. Astron. Soc, 336, 1256

[12] Cooray, A., Hu, W., & Tegmark, M. 2000, ApJ, 540, 1

[13] Rubiño-Martín, J. A., & Sunyaev, R. A. 2003, Mon. Not. R. Astron. Soc, 344, 1155

[14] Holder, G. P., McCarthy, I. G., & Babul, A. 2007, Mon. Not. R. Astron. Soc, 382, 1697

[15] Munshi, D., Joudaki, S., Smidt, J., & Coles, P. 2011, arXiv:1106.0766

[16] Fowler, J. W., Niemack, M. D., Dicker, S. R., et al. 2007, Applied Optics, 46, 3444

[17] Swetz, D. S., Ade, P. A. R., Amiri, M., et al. 2011, ApJS, 194, 41

[18] Carlstrom, J. E., Ade, P. A. R., Aird, K. A., et al. 2011, PASP, 123, 568

[19] Schaffer, K. K., Crawford, T. M., Aird, K. A., et al. 2011, ApJ, 743, 90

[20] Komatsu, E., Dunkley, J., Nolta, M. R., et al. 2009, ApJS, 180, 330

[21] Nozawa, S., Itoh, N., Suda, Y., & Ohhata, Y. 2006, Nuovo Cimento B Serie, 121, 487

[22] Cole, S., & Kaiser, N. 1988, Mon. Not. R. Astron. Soc, 233, 637

[23] Komatsu, E., & Kitayama, T. 1999, ApJ, 526, L1

[24] Tinker, J. L., Robertson, B. E., Kravtsov, A. V., et al. 2010, ApJ, 724, 878

[25] Bryan, G. L., & Norman, M. L. 1998, ApJ, 495, 80

[26] Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, ApJ, 490, 493

[27] Duffy, A. R., Schaye, J., Kay, S. T., & Dalla Vecchia, C. 2008, Mon. Not. R. Astron. Soc, 390, L64

[28] Tinker, J., Kravtsov, A. V., Klypin, A., et al. 2008, ApJ, 688, 709

[29] Battaglia, N., Bond, J. R., Pfrommer, C., & Sievers, J. L. 2011, arXiv:1109.3711

[30] Arnaud, M., Pratt, G. W., Piffaretti, R., et al. 2010, A & A, 517, A92

[31] Battaglia, N., Bond, J. R., Pfrommer, C., Sievers, J. L., & Sijacki, D. 2010, ApJ, 725, 91

[32] Piffaretti, R., & Valdarnini, R. 2008, A&A, 491, 71

[33] Shaw, L. D., Nagai, D., Bhattacharya, S., & Lau, E. T. 2010, ApJ, 725, 1452

[34] Bhattacharya, S., Nagai, D., Shaw, L., Crawford, T., & Holder, G. P. 2012, arXiv:1203.6368

[35] Sun, M., Sehgal, N., Voit, G. M., et al. 2011, ApJ, 727, L49

[36] Sehgal, N., Bode, P., Das, S., et al. 2010, ApJ, 709, 920

[37] Battaglia, N., Bond, J. R., Pfrommer, C., & Sievers, J. L. 2011, arXiv:1109.3709

[38] Wilson, M. J., Sherwin, B. D., Hill, J. C., et al. 2012, arXiv:1203.6633

[39] Ichiki, K., & Takada, M. 2011, arXiv:1108.4688

[40] Bhattacharya, S., Heitmann, K., White, M., et al. 2011, ApJ, 732, 122

[41] Brandbyge, J., Hannestad, S., Haugb olle, T., & Wong, Y. Y. Y. 2010, J. Cosmology Astropart. Phys., 9, 14

[42] Marulli, F., Carbone, C., Viel, M., Moscardini, L., & Cimatti, A. 2011, Mon. Not. R. Astron. Soc, 418, 346

[43] LoVerde, M., & Smith, K. M. 2011, J. Cosmology Astropart. Phys., 8, 3

[44] Wagner, C., Verde, L., & Boubekeur, L. 2010, J. Cosmology Astropart. Phys., 10, 22

[45] Dalal, N., Doré, O., Huterer, D., & Shirokov, A. 2008, Phys. Rev. D, 77, 123514

[46] Komatsu, E., Smith, K. M., Dunkley, J., et al. 2011, ApJS, 192, 18

[47] LoVerde, M., Miller, A., Shandera, S., & Verde, L. 2008, J. Cosmology Astropart. Phys., 4, 14

[48] Vikhlinin, A., Kravtsov, A. V., Burenin, R. A., et al. 2009, ApJ, 692, 1060

[49] Planck Blue Book. http://www.rssd.esa.int/SA/PLANCK/docs/Bluebook-ESA-SCI%282005%291_V2.pdf

[50] Zhang, P., Pen, U.-L., & Wang, B. 2002, ApJ, 577, 555

[51] Castro, P. G. 2004, Phys. Rev. D, 70, 049902

[52] Crawford, T. M., Schaffer, K. K., Bhattacharya, S., et al. 2013, arXiv:1303.3535

Chapter 5

The Atacama Cosmology Telescope: A Measurement of the Thermal Sunyaev-Zel'dovich One-Point PDF

5.1 Abstract

We present a measurement of the one-point probability distribution function (PDF) of the thermal Sunyaev-Zel'dovich (tSZ) signal in filtered Atacama Cosmology Telescope (ACT) 148 GHz temperature maps. The tSZ PDF contains information from all zero-lag moments of the tSZ field, making it a particularly sensitive probe of the amplitude of cosmic density fluctuations σ_8. We use a combination of semi-analytic halo model calculations and numerical simulations to compute the PDF signal and its covariance matrix, accounting carefully for all sources of noise. We remove contamination from nearly all astrophysical sources by measuring only the negative side of the 148 GHz PDF. The non-Gaussian tail induced by the tSZ effect is visible by eye in the ACT PDF. We furthermore demonstrate that the tSZ PDF can be used to probe intracluster gas physics in a manner that potentially breaks the degeneracy with σ_8. Applying these methods to the ACT 148 GHz PDF, we constrain

$\sigma_8 = 0.767 \pm 0.017$ assuming a fixed gas pressure profile model calibrated from cosmological hydrodynamics simulations. If we allow both σ_8 and the global normalization of the gas pressure profile model (relative to the fiducial model) P_0 to vary, we find $\sigma_8 = 0.748^{+0.024}_{-0.020}$ and $P_0 = 1.11 \pm 0.14$. We thus demonstrate the possibility of simultaneously constraining cosmology and intracluster gas physics using tSZ data alone.

5.2 Introduction

In recent years studies of the cosmic microwave background (CMB) temperature anisotropies have progressed beyond measurements of the primordial fluctuations seeded by inhomogeneities in the baryon-photon plasma at $z \approx 1100$, moving on to the study of secondary fluctuations induced by various physical processes at $z \lesssim 10$. This progress has been possible due to significant improvements in resolution and sensitivity, as demonstrated in the current generation of CMB experiments, including the Atacama Cosmology Telescope (ACT)/ACTPol [1, 2, 3, 4], South Pole Telescope (SPT)/SPTPol [5, 6, 7], Planck [8], and many others. Of particular note is the rapid growth in measurements of the Sunyaev-Zel'dovich (SZ) effect, which have been greatly facilitated by the new instrumentation. The SZ effect is a spectral distortion induced in the CMB by the scattering of CMB photons off free electrons along the line-of-sight (LOS) [9, 10]. This encompasses the kinetic SZ (kSZ) effect, due to bulk motions of free electrons along the LOS, and the thermal SZ (tSZ) effect, due to the thermal motions of hot electrons along the LOS, which are predominantly found in the intracluster medium (ICM) of massive galaxy clusters. The kSZ effect was first detected recently in ACT data [11], but this analysis will focus on the tSZ signal. The tSZ effect has now been measured over a wide range of halo masses and redshifts (e.g., [12, 13, 14, 15, 16]), both in direct observations and blind surveys, and has also been studied indirectly through its contribution to the power spectrum and higher-point functions of CMB temperature maps [17, 18, 19, 20, 21].

This paper is focused in particular on a novel indirect, statistical approach to measuring the tSZ signal by using the full one-point probability distribution function (PDF) of the tSZ-induced temperature fluctuations. Indirect measurements such as the tSZ power spectrum or higher-point functions possess some advantages over direct methods (i.e., finding and counting clusters). For example, there is no selection function applied to the map to find clusters, and hence no selection effect-related systematics, such as Malmquist or Eddington bias. In addition, no choice of "aperture" is required within which a cluster mass is defined (though this can be a useful intermediate step in halo model calculations of tSZ statistics), and indeed no measurement of individual cluster masses is performed, which can be an expensive procedure. Likewise, the use of "average" gas pressure profile prescriptions is sensible when computing tSZ statistics, since one is computing a population-level quantity (i.e., there is no need to apply an average/universal profile to individual objects). Lastly, as has been known for many years, tSZ statistics are very sensitive to σ_8, the rms amplitude of cosmic density fluctuations on scales of 8 Mpc/h (e.g., [49, 40, 30, 48]). The tSZ one-point PDF inherits all of these advantages, although it comes with the same disadvantages as other indirect tSZ probes — for example, no redshift information is available, as the entire LOS is integrated over. In addition, like other indirect methods, the tSZ PDF possesses the simultaneous advantage and disadvantage of sensitivity to lower-mass, higher-redshift halos than those found in cluster counts. This is an advantage as the raw increase in statistical power from this sensitivity is large; it is a disadvantage because the ICM gas physics in these systems can be significantly more uncertain than that in massive, low-redshift systems with deep X-ray, optical, and lensing observations (e.g., [42, 29, 40, 48]). However, we will demonstrate that the tSZ PDF allows the opportunity to constrain the gas physics in these systems, while simultaneously constraining cosmology.

Nearly all cosmological constraints derived from measurements of tSZ statistics thus far are based solely on the one-halo term. In effect, these measurements are just indirect ways of counting clusters, including those below the signal-to-noise (SNR) threshold for

direct detection. Clustering information in the two-, three-, and higher-point terms is not measured. The only exceptions are recent hints of the two-halo term in the Planck y-map power spectrum [21] and measurements of the two-halo term in the cross-correlation of tSZ and gravitational lensing maps [31, 32]. This situation thus motivates the consideration of a statistic that fully utilizes the information in the tSZ one-halo term. The obvious candidate for such a statistic is the tSZ one-point PDF.

The tSZ PDF includes the information contained in all zero-lag moments of the tSZ field. It has been shown that higher-point tSZ statistics scale with increasingly high powers of σ_8, with $\langle \Delta T^2 \rangle \propto \sigma_8^{7-8}$, $\langle \Delta T^3 \rangle \propto \sigma_8^{10-12}$, $\langle \Delta T^4 \rangle \propto \sigma_8^{13-15}$, and continuing on in this manner [40, 48]. The underlying reason is that increasingly higher-point tSZ statistics are dominated by contributions from increasingly rarer, more massive clusters that lie deeper in the exponential tail of the mass function. The tSZ PDF includes information from all of these higher moments, which eventually reach a sample variance-dominated limit as one proceeds to consider progressively higher-point statistics. Apart from discarding redshift information, the tSZ PDF is thus likely close to an optimal one-halo-term tSZ statistic for constraining cosmological parameters.

Moreover, by naturally including information from all moments of the tSZ field, the tSZ PDF possesses the ability to break the degeneracy usually present in tSZ statistics between variations in cosmological parameters and variations in the ICM gas physics model. This idea is a natural generalization of the method presented in [40], which used information in the two- and three-point functions alone to explicitly separate the influence of cosmology (i.e., σ_8, which is by far the dominant parameter) from that of ICM physics. The basic idea is quite simple: given the scalings of $\langle \Delta T^2 \rangle$ and $\langle \Delta T^3 \rangle$ with σ_8, one can trivially construct a statistic which cancels the effect of σ_8, and it turns out that this statistic (the "rescaled skewness") still possesses a sensitivity to the ICM physics model (see Fig. 7 of [40]). The tSZ PDF implicitly contains all of this information, as well as the dependence of all higher tSZ moments on the cosmological and astrophysical parameters. The different scaling of these

moments with σ_8 and the ICM physics thus allows the PDF to naturally break the degeneracy between the two (see Fig. 5.5 in Section 5.4.3). The only question is whether a given experiment possesses sufficient statistical power to break the degeneracy. This comes down to the sensitivity, angular resolution, and sky coverage of the experiment. The sensitivity and resolution determine whether the non-Gaussianity in the tSZ field can be measured: if the noise is too large or if the beam is so large that many clusters fall within its resolution, the observed tSZ field will be Gaussianized. The sky coverage determines whether sample variance will generate overwhelming errors in the tSZ PDF. In this paper, we measure the tSZ PDF in ACT 148 GHz data in a ≈ 324 square degree region along the celestial equator. We find that the data is at the level to begin to break the cosmology–ICM degeneracy, but not strongly. It is possible that extant data from SPT and Planck may however be sufficient to do so.

Despite the advantages highlighted above, the tSZ PDF also possesses some downsides. For example, this observable discards most of the angular information contained in the full polyspectra of the tSZ field, which can potentially be used to constrain ICM pressure profile models, or potentially to measure clustering terms in tSZ statistics (e.g., the two-halo term in the tSZ power spectrum [72, 30, 21]). In addition, a precise measurement of the tSZ PDF requires all non-tSZ contributions in the temperature map to be very well-understood, including the CMB, instrumental and atmospheric noise, and other astrophysical signals. We largely circumvent this issue by working with only the $\Delta T < 0$ side of the 148 GHz ACT PDF, in which the non-Gaussian contributions are dominated by the tSZ signal. Theoretically, the tSZ PDF requires essentially the same modeling inputs as other tSZ statistics, in particular an understanding of the ICM gas pressure profile over a fairly wide range of masses and redshifts. For the ACT Equatorial sensitivity and resolution, we find that the PDF is mostly dominated by systems at preferentially lower redshifts and higher masses than those that dominate the tSZ power spectrum at $\ell = 3000$, implying that the modeling uncertainties should not be overwhelming, similar to the situation with the tSZ skewness and kurtosis [40, 48]. A

final potential difficulty arises from the currently poorly constrained correlation between the cosmic infrared background (CIB) and the tSZ field [73, 74]. The CIB is comprised of the cumulative infrared emission from billions of unresolved dusty star-forming galaxies, some of which may be spatially co-located in or near the galaxy groups and clusters sourcing the tSZ signal. Fortunately, since the tSZ PDF is a one-point statistic, it is unaffected by the clustering term in the CIB–tSZ correlation; the only possible issue comes from CIB sources "filling in" the CMB temperature decrements induced by the tSZ effect at 148 GHz. We investigate this issue using the ACT 218 GHz maps as detailed in Section 5.6.

To our knowledge, this paper represents the first work to use the tSZ PDF as a cosmological probe. Similar approaches have been investigated for many years in the weak lensing literature (e.g., [62]), and our method is quite similar to the measurement of "peak counts" in weak lensing maps (e.g. [63]), as well as the traditional $P(D)$ analysis used in radio point source studies for decades [64, 65]. To our knowledge, the first theoretical calculation of the tSZ PDF was performed in [66] using the "peak-patch" picture of cosmic structure formation, a somewhat different approach than we adopt here (see Section 5.4). In subsequent years, a number of simulation groups measured the PDF of the tSZ field in cosmological hydrodynamics simulations [75, 70, 67, 68, 69, 71]. These groups noted the significant non-Gaussianity in the tSZ PDF and suggested this as a way to distinguish the tSZ fluctuations from the (Gaussian) primary CMB field in a single-frequency experiment. Other analyses at the time attempted to compute the tSZ PDF semi-analytically using an Edgeworth expansion in combination with halo model calculations of tSZ moments [39]; unfortunately, this expansion is numerically unstable when applied to the tSZ PDF (with terms beyond the skewness) due to the oscillating sign of each term and the large magnitude of the higher tSZ moments [40]. The Edgeworth expansion likely does not converge when applied to the tSZ PDF. Thus, instead of working in terms of an expansion around a Gaussian or using perturbation theory, in this paper we introduce a method to compute the tSZ PDF directly using the halo model, similar to approaches already used to calculate the tSZ

power spectrum [49], tSZ bispectrum [48], and real-space tSZ moments [40]. We then use these calculations to interpret the tSZ PDF measured with ACT data, yielding constraints on cosmological parameters and the physics of the ICM. The data set used here is nearly identical to that used in our earlier ACT skewness analysis [19]; the improved precision of our results compared to that work is entirely due to the use of a more powerful tSZ statistic.

The rest of this paper is organized as follows. Section 5.3 describes the ACT 148 GHz data used in the analysis and presents its PDF. Section 5.4 describes our halo model-based theoretical approach to the tSZ PDF and verifies its accuracy using numerical simulations. We also describe how to account for noise in the PDF and outline our models for the various non-tSZ contributions to the ACT PDF. In addition, we investigate the sensitivity of the tSZ PDF to cosmological and astrophysical parameters, and calculate the contributions to the signal from different mass and redshift ranges. Section 5.5 presents the simulation pipeline that we use to generate Monte Carlo realizations of the data, which are needed for covariance matrix estimation and tests of the likelihood. Section 5.6 describes the likelihood function that we use to interpret the data and presents constraints on cosmological and astrophysical parameters using our tSZ PDF measurements. We discuss the results and conclude in Section 5.7.

We assume a flat ΛCDM cosmology throughout. Unless stated otherwise, parameters are set to the WMAP9+eCMB+BAO+H_0 maximum-likelihood values [76]. In particular, $\sigma_8 = 0.817$ is our fiducial value for the rms amplitude of linear density fluctuations on 8 Mpc/h scales at $z = 0$. All masses are quoted in units of $M_\odot/h$, where $h \equiv H_0/(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}})$ and H_0 is the Hubble parameter today. All distances and wavenumbers are in comoving units of Mpc/h.

5.3 Data

The Atacama Cosmology Telescope [1, 2, 3] is a six-meter telescope located on Cerro Toco in the Atacama Desert of northern Chile. It observed the sky between 2007 and 2010 at 148, 218, and 277 GHz using three 1024-element arrays of superconducting bolometers. The ACT data characterization and mapmaking techniques are described in [3]. Measurements of the CMB temperature power spectrum using the full three seasons of ACT data on both the celestial equator and southern strip regions are described in [22]. These measurements are interpreted using the likelihood described in [23], yielding constraints on cosmological and secondary parameters presented in [17]. In addition to the primordial CMB, ACT data also contain the imprint of massive galaxy clusters through the tSZ effect. A sample of high-significance clusters discovered in the ACT data are reported in [12], with additional optical and X-ray characterization described in [24]. These clusters are also used to constrain cosmological parameters in [12].

This analysis is a direct extension of the results presented in [19], in which the tSZ signal in the ACT Equatorial map (from all groups and clusters in the field) is measured through its contribution to the skewness of the temperature distribution in the map. Here we interpret effectively the same data with a more powerful technique, the one-point PDF of the temperature distribution, which includes not only the skewness but also all higher-order (zero-lag) moments. The maps used in this analysis are described in [22]; they are nearly identical to those used in [19], which were simply preliminary versions of the same data. We work with a co-added map made from two seasons of observation at 148 GHz in the equatorial field in 2009 and 2010. The map is comprised of six $3^\circ \times 18^\circ$ patches of sky. The map noise level in CMB temperature units is $\approx 18\,\mu$K arcmin. In contrast to [19], we do not use any 218 GHz data in this work (apart from a brief check to ensure CIB emission is not affecting the tSZ results in Section 5.6). By restricting our analysis to negative temperature values in the 148 GHz histogram, we remove essentially any contaminating IR emission (this contamination motivated the use of the 218 GHz data in [19]). The maps used in this analysis

are calibrated to WMAP using the methods described in [25, 22]. To prevent ringing effects arising when we Fourier transform the maps, we apodize them with a mask that smoothly increases from zero to unity over 0.1° from the edge of the maps.

The processing of the 148 GHz map in this analysis is very similar to that in [19], with a few minor exceptions. First, in order to upweight tSZ signal in the data over the noise from the atmosphere and CMB on large scales and the instrumental noise on small scales, we filter the maps in Fourier space using the filter shown in Fig. 1 of [19]. We consistently apply the same Fourier-space filter throughout this work to data, simulations, and theoretical calculations, unless explicitly stated. Note that in contrast to [19], we do not remove signal in the $-100 < \ell_{\rm dec} < 100$ stripe along the Fourier axis corresponding to declination, as the scan noise in the maps is no longer problematic enough to require this step. After filtering, we remove the edges of the maps in order to prevent any edge effects that could occur as a result of Fourier transforming, despite the apodization described above.

Second, as in [19], we use template subtraction [26] to remove high-SNR point sources from the map (primarily radio sources). A full description of the method is provided in [19]. In addition, we mask a circular region of radius 5 arcmin around the subtracted sources after applying the Fourier-space filter described in the previous paragraph, in order to reduce any ringing effects that might arise. We verify that no non-Gaussian structure is produced in the histogram of simulated maps after point source subtraction, filtering, and masking. Note that in contrast to [19], we do not apply a mask to remove all pixels with values more than 12σ from the map mean, as our aim is to study the negative tail of the pixel PDF. However, the histogram of the data used in this analysis contains no such pixels, so this criterion is rendered irrelevant. In total, 13.3% of the 148 GHz map is masked in the processing procedure, but the masked pixels are random with respect to the tSZ field and therefore should not alter the signal we are concerned with in this analysis.

Fig. 5.1 shows the PDF (i.e., histogram) of the pixel temperature values in the filtered, processed ACT Equatorial 148 GHz temperature map. We use bins of width 10 μK, extending

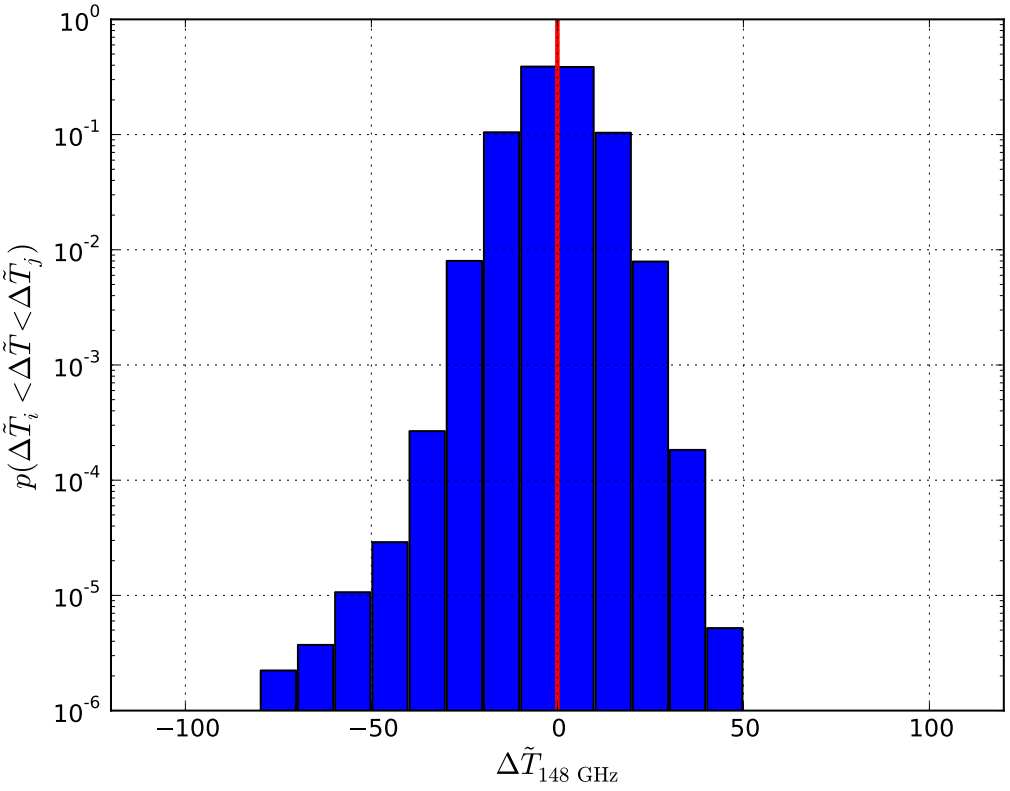

Figure 5.1: PDF of pixel temperature values in the ACT 148 GHz Equatorial CMB map after filtering and point source masking. The solid red vertical line denotes $\Delta\tilde{T} = 0$. The tSZ effect is responsible for the significant non-Gaussian tail on the negative side of the PDF ($\Delta\tilde{T} < 0$). This figure is effectively a re-binned version of Fig. 2 from [19], although the positive side of the PDF is slightly different as we do not apply the 218 GHz CIB mask constructed in that work to the 148 GHz map here (see text for details). The mask is unnecessary because we do not consider the positive side of the PDF ($\Delta\tilde{T} > 0$) in this analysis, since it contains essentially no tSZ signal.

to $\pm 120\,\mu$K, well beyond where the negative tail of the distribution cuts off. The number of bins is primarily motivated by the difficulty in computing the covariance matrix of the PDF for a very large number of bins (see Sections 5.4.4 and 5.5). Fig. 5.1 is equivalent to Fig. 2 in [19], but with a coarser binning and slightly more pixels in the $\Delta\tilde{T} > 0$ tail of the distribution (since we have not applied the additional 218 GHz IR mask from [19]). As we do not consider the $\Delta\tilde{T} > 0$ side of the PDF, this excess is irrelevant in our analysis. The key feature in Fig. 5.1 is the non-Gaussian tail at $\Delta\tilde{T} < 0$ sourced by the tSZ effect.

5.4 Theory

5.4.1 Thermal SZ Effect

The tSZ effect arises from the inverse Compton scattering of CMB photons off of hot electrons along the LOS, e.g., the hot, dilute plasma found in the ICM of galaxy clusters [10]. This interaction produces a frequency-dependent shift in the observed CMB temperature in the direction of the electrons responsible for the scattering. Neglecting relativistic corrections (e.g., [27]), which are relevant only for the most massive clusters in the universe ($\gtrsim 10^{15}\ M_{\odot}/h$), the tSZ-induced change in the observed CMB temperature ΔT at angular position $\vec{\theta}$ on the sky with respect to the center of a cluster of mass M at redshift z is given by [10]:

$$\begin{aligned} \frac{\Delta T(\vec{\theta}, M, z)}{T_{\rm CMB}} &= g_\nu y(\vec{\theta}, M, z) \\ &= g_\nu \frac{\sigma_T}{m_e c^2} \int_{\rm LOS} P_e\left(\sqrt{l^2 + d_A^2 |\vec{\theta}|^2}, M, z\right) dl\,, \end{aligned} \tag{5.1}$$

where $g_\nu = x\coth(x/2) - 4$ is the tSZ spectral function with $x \equiv h\nu/k_B T_{\rm CMB}$, $y(\vec{\theta}, M, z)$ is the Compton-y parameter, σ_T is the Thomson scattering cross-section, m_e is the electron mass, and $P_e(\vec{r}, M, z)$ is the ICM electron pressure at (three-dimensional) separation $\vec{r}$ from the cluster center. Our theoretical calculations assume that the pressure profile is spherically symmetric, i.e., $P_e(\vec{r}, M, z) = P_e(r, M, z)$ where $r = |\vec{r}|$. The integral in Eq. (5.1) is computed along the LOS such that $r^2 = l^2 + d_A(z)^2\theta^2$, where $d_A(z)$ is the angular diameter distance to redshift z and $\theta \equiv |\vec{\theta}|$. We assume that the electron pressure $P_e(\vec{r}, M, z)$ is related to the thermal gas pressure via $P_{th} = P_e(5X_H + 3)/2(X_H + 1) = 1.932 P_e$, where $X_H = 0.76$ is the primordial hydrogen mass fraction.

Our definition of the virial mass M, our convention for the concentration-mass relation (from [28]), and our fiducial ICM pressure profile model (from [29]) are identical to those in [30, 31], and we refer the reader to those works for further details. The pressure profile

model is a parametrized fit to stacked profiles extracted from the "AGN feedback" simulations described in [33]. The model fully describes the pressure profile of the ICM gas as a function of cluster mass (M_{200c}) and redshift. The pressure profile model and the halo mass function are the only necessary inputs for our semi-analytic calculations of the tSZ PDF (see Section 5.4.2).

The "AGN feedback" simulations from which the pressure profile is extracted include virial shock heating, radiative cooling, and sub-grid prescriptions for star formation and feedback from supernovae and active galactic nuclei (AGN). Non-thermal pressure support due to bulk motion and turbulence in the ICM outskirts are captured by the smoothed particle hydrodynamics (GADGET-2) method used for the simulations. The pressure profile extracted from these simulations has been found to agree with many recent X-ray and tSZ measurements [34, 35, 36, 37, 21, 38, 31]. Nonetheless, the precise nature of the pressure-mass relation remains a significant source of uncertainty in current cosmological constraints from tSZ measurements (e.g., [12, 14, 15, 19, 20, 31]). In this work, we parametrize the uncertainty by allowing the overall normalization P_0 of the pressure-mass relation to vary. This approach is equivalent to (and can be approximately mapped onto) a constraint on the so-called "hydrostatic mass bias", a free parameter that sets the overall normalization of the "universal pressure profile" from [34], which is derived using hydrostatic X-ray masses. Our fiducial "AGN feedback" model corresponds approximately to a hydrostatic mass bias $\approx 15\%$ for the massive, low-redshift population of clusters studied in [34], though this value varies with cluster-centric radius (see Fig. 2 of [33]). Note that the pressure-mass normalization (or hydrostatic mass bias) is a quantity averaged over all clusters in the population under study; it is expected to be a function of cluster mass and redshift, and likely to possess scatter about its mean, but to lowest order we aim to constrain its mean value over the cluster population. Note that the pressure profile fitting function that we use from [29] does include mass and redshift dependences in the amplitude of the pressure profile; in this work, we constrain the overall normalization relative to this fitting function. Our approach

is identical to that used in the forecasts of [30]: the parameter P_0 is defined *relative* to the fiducial value found in [29], i.e., $P_0^{\rm this\,work} = P_0^{\rm Battaglia}/18.1$. Thus, $P_0 = 1$ corresponds to the fiducial ICM pressure profile model, whereas $P_0 < 1$ (> 1) implies a typical thermal pressure less (greater) than that expected in the fiducial model, corresponding to a larger (smaller) value of the hydrostatic mass bias than that already present in the fiducial model. We demonstrate below that the tSZ PDF possesses sufficient power to simultaneously constrain P_0 and cosmological parameters.

In all of the relevant calculations for the ACT analysis, an ℓ-space Wiener filter is used to increase the tSZ signal-to-noise in the data. The filter is identical to that used in [19]. In addition, we use temperature maps in which the beam has not been deconvolved in order to prevent pixel-to-pixel noise correlations in the final maps as much as possible. We include these effects in ℓ-space in the semi-analytic calculations by Fourier transforming $y(\theta, M, z)$ for each cluster, multiplying by the filter and beam functions, and then Fourier transforming back to real space. The pixel window function could also be included in the same manner, but it is extremely close to unity over the ℓ-range on which the Wiener filter is non-zero, so we simply set it to unity in our calculations. From this point onward, it should be understood unless stated otherwise that $y(\theta, M, z)$ refers to the filter- and beam-convolved (noiseless) y-profile for each cluster of mass M and redshift z.

5.4.2 Noiseless tSZ-Only PDF

The tSZ PDF encapsulates all of the statistical information in the zero-lag moments of the tSZ field. The most straightforward and accurate way to compute the tSZ PDF is using cosmological hydrodynamics simulations, an approach which was pursued by several groups roughly a decade ago [75, 70, 67, 68, 69, 71]. In order to constrain cosmology using the tSZ PDF, however, a less computationally taxing (semi-)analytic method is required. Previous works have attempted to analytically compute the tSZ PDF using the Edgeworth expansion or other approximation schemes based on the assumption of weak non-Gaussianity (e.g., [39]).

While such approaches may be valid after smoothing the tSZ field with a large beam and/or convolving it with experimental noise (either of which can Gaussianize the signal), the tSZ field itself is indeed highly non-Gaussian [40], and thus such schemes are unlikely to accurately describe its PDF. We instead compute the tSZ PDF using a halo model-based approach. This method requires some assumptions, as we describe below, but does not assume that the tSZ field is weakly non-Gaussian.

The value of the tSZ PDF integrated over a bin $b_i \equiv (y_i, y_{i+1})$ corresponds to the expected fraction of the sky subtended by Compton-y values within that bin. Thus, a straightforward way to compute the signal is to calculate the relevant fraction of sky contributed by each cluster of mass M and redshift z, and then sum over all the clusters in the universe. Let $f(y, M, z)dy$ be the sky area subtended by the y-profile of a cluster of mass M and redshift z in the range $(y, y + dy)$. Define $\theta(y, M, z)$ to be the inverse of $y(\theta, M, z)$ in Eq. (5.1). Then $f(y, M, z)dy$ is given by

$$f(y, M, z)dy = 2\pi\theta(y, M, z)\frac{d\theta}{dy}|_{y,M,z}dy\,. \tag{5.2}$$

The noiseless tSZ PDF $\langle P_i\rangle$ in bin b_i is then given by summing up the contributions to this bin from all clusters in the universe:

$$\begin{aligned}\langle P_i\rangle_{\rm noiseless} &= \int dz\frac{d^2V}{dzd\Omega}\int dM\frac{dn}{dM}\int_{y_i}^{y_{i+1}} dy f(y, M, z)\\ &= \int dz\frac{d^2V}{dzd\Omega}\int dM\frac{dn}{dM}\pi\left(\theta^2(y_i, M, z) - \theta^2(y_{i+1}, M, z)\right)\,,\end{aligned} \tag{5.3}$$

where $d^2V/dzd\Omega$ is the comoving volume element per steradian and $dn(M, z)/dM$ is the halo mass function (the number of halos at redshift z per unit comoving volume and per unit mass), for which we use the prescription of Tinker et al [41]. Note that the PDF is normalized such that $\Sigma_i\langle P_i\rangle = 1$. For the noiseless, tSZ-only case, the PDF is sharply peaked at $y = 0$, since most of the sky is not subtended by massive galaxy clusters; for the noisy

case (see Section 5.4.3), this peak is smoothed out by the presence of CMB fluctuations, instrumental noise, and other contributions to the microwave sky.

The primary assumption in our halo model-based approach is the neglect of any effects due to overlapping clusters along the LOS. This approximation breaks down if one naïvely extends the computation down to arbitrarily low masses or assigns an arbitrarily large outer boundary for each cluster. The validity of the approximation is encoded in the naïve sky fraction subtended by all clusters in the calculation, assuming no overlaps:

$$F_{\rm clust} = \int dz \frac{d^2V}{dz d\Omega} \int dM \frac{dn}{dM} \pi \theta^2_{\rm out}(M,z)\,, \tag{5.4}$$

where $\theta_{\rm out}(M,z) \equiv r_{\rm out}(M,z)/d_A(z)$ is the outer angular boundary of a cluster of mass M at redshift z. We take the outer boundary to be twice the virial radius (as defined in [30, 31]), which is roughly the location of the virial shock seen in simulations of cluster formation [42]. The results of our analysis of the ACT tSZ PDF are not sensitive to this choice, because the y-values in the outer regions of clusters are very small, and thus they only contribute to the noise-dominated regime of the measured PDF (see below). The approximation of neglecting LOS overlaps is valid as long as $F_{\rm clust} \ll 1$. Numerical testing using the mass function of [41] indicates that this approximation breaks down when integrating the mass function down to a mass scale $\sim 10^{14} M_\odot/h$. Naïvely integrating the mass function below this mass scale yields $F_{\rm clust} \approx 1$, indicating the breakdown of the non-overlapping assumption. Thus, we take the mass limits of the integrals in all expressions in this work to be $2 \times 10^{14} M_\odot/h < M < 5 \times 10^{15} M_\odot/h$, which give $F_{\rm clust} \approx 0.3$ (we verify with simulations in Section 5.5 that overlap effects remain negligible for these mass limits). The redshift integration limits are $0.005 < z < 3$ (the lower limit is chosen to prevent divergences at $z = 0$). We have checked that all computations are converged with these limits; note that with the noise and beams appropriate for ACT, clusters at or below the lower mass limit are simply absorbed in the noise-dominated region of the PDF (see Section 5.4.3). We will further verify these

approximations by comparing to simulations. However, it is important to note that due to this approximation, a direct comparison of our semi-analytic results to a tSZ-only map extracted from a hydrodynamical simulation will be expected to disagree somewhat in the low-y regime, since the semi-analytic calculation does not include all of the low-y signal by construction. We verify below that the ACT noise and beams are large enough that the measured PDF is not sensitive to Compton-y values sourced by clusters below our lower mass cutoff.

We validate the halo model-based theory described above by comparing to the tSZ PDF directly extracted from cosmological hydrodynamics simulations. We use the simulations from which our fiducial pressure profile model was extracted [33] in order to facilitate a like-for-like comparison. We extract 390 Compton-y maps from the simulations, each of area $(4.09^\circ)^2$. The maps are direct LOS projections of randomly rotated and translated simulation volumes, with an upper redshift cut at $z = 1$ in order to decrease correlations due to common high-redshift objects in the extracted maps. We apply the same redshift cut to the theoretical calculations in order to be consistent (only for this particular exercise), and in addition we consider a lower mass cutoff in Eq. (5.3) of $2 \times 10^{13} M_\odot/h$ (along with the fiducial $2 \times 10^{14} M_\odot/h$). For this calculation, we use the same values for the cosmological parameters as those used in the simulations (see [33]). In addition, note that no filtering or beam convolution is applied to either the simulation maps or the theoretical calculations for this comparison (i.e., for this calculation only, $y(\theta, M, z)$ refers to the physical, non-filter- or beam-convolved Compton-y profile). The results are shown in Fig. 5.2. Overall, the agreement is excellent. At the largest $|\Delta T|$-values, the slight discrepancy is due to an excess of massive clusters in the simulations of [33] compared to the Tinker mass function (see the appendix of [29]). At the smallest $|\Delta T|$-values, discrepancies emerge due to the breakdown of our assumption that clusters do not overlap along the LOS — this is responsible for the unphysical value of $p > 1$ in the lowest bin for the $M_{\rm min} = 2 \times 10^{13} M_\odot/h$ calculation. One can also see that the lowest bins receive some contributions from halos with masses between

$2 \times 10^{13} M_\odot/h$ and $2 \times 10^{14} M_\odot/h$; however, we will demonstrate in the following section that the tSZ signal from all such objects is subsumed into the noise in the ACT data. It is important to note that the results shown in Fig. 5.2 have not been filtered or convolved with the noise, and thus the ΔT values are not directly comparable to those shown in the rest of the plots in this paper (which are labeled $\Delta\tilde{T}$ for clarity). The comparison is simply a test of the halo model framework.

In particular, the general agreement seen in Fig. 5.2 provides evidence supporting three assumptions made in the halo model approach. First, the comparison verifies that neglecting scatter in the $P_e(r, M, z)$ relation does not invalidate our semi-analytic theory; this scatter is obviously present in the hydrodynamical simulations, which agree well with our results computed using only the mean $P_e(r, M, z)$ relation from [29]. Second, the comparison indicates that tSZ signal from non-virialized regions is negligible for our analysis — this signal is present in the maps from the simulations (which are direct LOS integrations), but neglected in the halo model approach. The halo model calculations, when extended down to low enough masses, appear to account for essentially all of the signal in the simulations. Finally, the comparison verifies that the halo model's neglect of effects due to overlapping clusters along the LOS is safe for a minimum mass cutoff satisfying the $F_{\rm clust} \ll 1$ criterion described above. Overall, the comparison indicates that the semi-analytic halo model theory is an accurate description of the tSZ PDF in this regime.

5.4.3 Noisy PDF

The noiseless tSZ PDF computed in the previous section is instructive for gaining intuition and comparing to hydrodynamical simulations, but is clearly insufficient for a comparison to actual observations of the microwave sky — a treatment of the non-tSZ sky components is required. For the purposes of the tSZ PDF, all other sky components are a source of noise; we must therefore convolve the tSZ PDF with the various PDFs describing these sources of

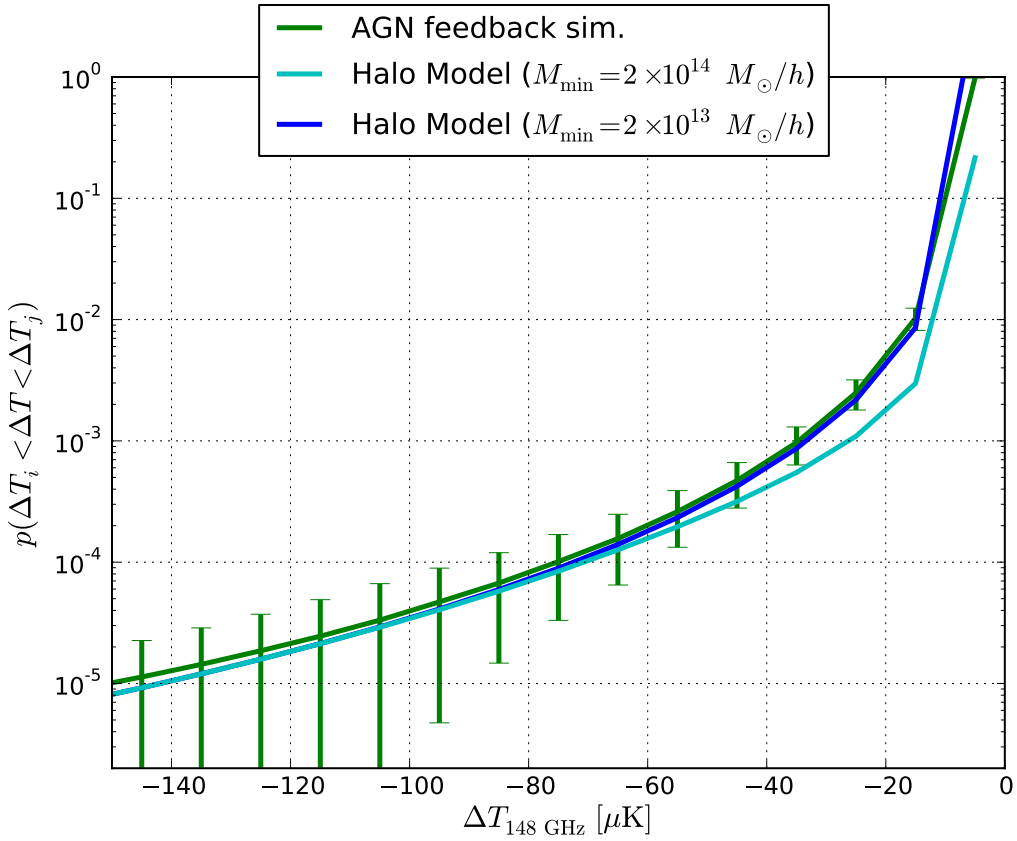

Figure 5.2: Comparison of the noiseless halo model computation of the tSZ PDF using Eq. (5.3) to the PDF directly measured from the cosmological hydrodynamics simulations of [33]. The error bars are computed from the scatter amongst the 390 maps extracted from the simulation. Note that no filtering or noise convolution has been performed for these calculations, and thus the ΔT values cannot be directly compared to those shown in the other plots in the paper. The comparison demonstrates that the halo model approach works extremely well, except for discrepancies at low-$|\Delta T|$ values arising from the breakdown of the halo model assumptions (see the text for discussion).

noise. We will describe the general case and then specify the details of each component's contribution to the noise.

Let $\rho_i(y)$ be the probability of observing a signal in bin $b_i \equiv (y_i, y_{i+1})$ given an input (physical) signal y:

$$\rho_i(y) = \int_{y_i}^{y_{i+1}} dy' N(y - y') \,, \tag{5.5}$$

where $N(y - y')$ is the noise kernel. The contribution of a single cluster to bin b_i after including noise is then given by:

$$\begin{aligned} g_i(M, z) &= \int dy f(y, M, z) \rho_i(y) \\ &= \int d\theta \; 2\pi\theta \; \rho_i \left(y(\theta, M, z) \right) \,, \end{aligned} \tag{5.6}$$

where $\rho_i(y)$ is defined by Eq. (5.5), $f(y, M, z)$ is defined by Eq. (5.2), and the second line is a computationally simpler way of expressing the same quantity. Note that the results in the noiseless case discussed in Section 5.4.2 can be recovered by taking $\rho_i(y) = \Theta(y - y_i) - \Theta(y - y_{i+1})$, where $\Theta(x)$ is the Heaviside step function. Finally, adding the contribution from regions of the sky with zero intrinsic y-signal that fluctuate (due to noise) into bin b_i, we obtain the final expression for the one-point PDF of the noisy Compton-y field:

$$\langle P_i \rangle = (1 - F_{\rm clust}) \rho_i(0) + \int dz \frac{d^2V}{dz d\Omega} \int dM \frac{dn}{dM} \int dy \, g_i(M, z) \,, \tag{5.7}$$

where the first term represents the contributions from noise fluctuations in intrinsically zero-y pixels and the second term represents the "real" tSZ contributions after noise convolution. Eq. (5.7) has the convenient property that adding or subtracting low-mass clusters in the calculation of $F_{\rm clust}$ or the integral in the second term does not change the result, as this procedure simply shifts these contributions between either of the two terms. Similarly, slight changes to the definition of the outer boundary of a cluster do not matter, as these regions are low-y and thus subsumed into the noise-dominated region of the PDF (the same

statement holds for changes in the upper redshift integration limit). These statements are contingent upon the approximations described in Section 5.4.2, in particular that cluster overlaps along the LOS can be neglected. More precisely, we must be in the regime where $F_{\rm clust} \ll 1$. As mentioned earlier, we find that the approximations hold for our ACT analysis calculations with mass limits given by $\{2 \times 10^{14} M_\odot/h, 5 \times 10^{15} M_\odot/h\}$ and $r_{\rm out} = 2r_{\rm vir}$, which give $F_{\rm clust} \approx 0.3$.

We now detail the construction of the noise kernel $N(y-y')$. In the case of homogeneous, uncorrelated, Gaussian noise, the noise kernel is

$$N_G(y-y') = \frac{1}{\sqrt{2\pi\sigma_N^2}} e^{-(y-y')^2/(2\sigma_N^2)}\,, \tag{5.8}$$

where σ_N is the single-pixel RMS noise. The Gaussian approximation is reasonable for computing forecasts, but in the ACT data analysis a more detailed treatment of the noise is required. In particular, for ACT we must account for the non-Gaussian properties of both the instrumental and atmospheric noise contributions, as well as the nearly Gaussian noise contribution from the primordial CMB fluctuations (although this component is significantly reduced by the ℓ-space filter used in the analysis). There are also noise contributions from unresolved point sources and the CIB sourced by dusty star-forming galaxies — however, we largely mitigate these sources of noise by considering only the $\Delta\tilde{T} < 0$ region of the ACT 148 GHz PDF, as described in Section 5.3. Our procedure for constructing the noise kernel $N(y-y')$ appropriate for the ACT analysis is as follows:

- Instrumental and atmospheric noise: this component dominates the total noise in our filtered map. Instead of assuming it to be Gaussian, we measure the instrumental and atmospheric noise directly using null maps constructed from the differences between single-season "split" maps. The null maps are processed in exactly the same manner as the 148 GHz data map, but contain no cosmological or astrophysical signal. Thus they provide an accurate characterization of the instrumental and atmospheric noise. We

fit a cubic spline to this noise PDF, which we will refer to as $N_{\rm inst}(x)$. Its contribution to the RMS of the filtered map is $\sigma_{\rm inst} = 5.78\,\mu$K. Note that although it is symmetric about $\Delta\tilde{T} = 0$ (and thus does not affect our earlier tSZ skewness results [19]), the noise PDF has non-Gaussian tails that must be accounted for in this analysis.

- CMB: the CMB dominates the remainder of the noise budget. As discussed below, its effects will effectively be marginalized over in the analysis, but to set the fiducial level we compute the CMB power spectrum C_ℓ for our fiducial WMAP9 cosmology using CAMB[1] and then obtain the corresponding variance via

$$\sigma^2_{\rm CMB} = \sum_\ell \frac{2\ell+1}{4\pi} C_\ell F_\ell^2 b_\ell^2, \tag{5.9}$$

  where F_ℓ is the filter introduced in Section 5.3 and b_ℓ is the ACT 148 GHz beam [44]. This computation yields $\sigma_{\rm CMB} = 5.22\,\mu$K. We treat the CMB PDF $N_{\rm CMB}(x)$ as a Gaussian with this variance.

- Other contaminants: other sources of power in the microwave sky also contribute to the PDF in the ACT 148 GHz map. The most important of these components are the CIB and emission from radio and infrared point sources. Although our source masking procedure greatly reduces the IR and radio contributions to the PDF, there is still an unresolved component (including Poisson and clustered terms) in the map. The CIB and source contributions are restricted to the $\Delta T > 0$ region of the PDF in the absence of ℓ-space filtering, but leak slightly into the $\Delta\tilde{T} < 0$ region after we filter out the zero-mode of the map. Excluding the $\Delta\tilde{T} > 0$ region removes nearly all of the CIB and point source contributions in our analysis. We also discard the first bin in the $\Delta\tilde{T} < 0$ region to further ensure that no CIB or point source emission affects our results (see Section 5.6). However, these sources still contribute to the noise with which the entire PDF must be convolved; we model their contribution as

[1] `http://camb.info`

a Gaussian with variance determined from the best-fit foreground parameters for the ACT Equatorial map in [17]. In particular, we sum the best-fit results from [17] for all of the 148 GHz power spectra sourced by foregrounds except tSZ (which is not a foreground for our analysis). We then compute the resulting variance in our filtered map $\sigma^2_{\rm fg}$ using Eq. (5.9). We find $\sigma_{\rm fg} = 2.33\,\mu$K. We treat the foregrounds' PDF $N_{\rm fg}(x)$ as a Gaussian with this variance. Since the CMB is also Gaussian, we can convolve the foreground and CMB PDFs trivially to obtain a combined Gaussian PDF with variance $\sigma^2_{\rm fid} = \sigma^2_{\rm CMB} + \sigma^2_{\rm fg}$. In the PDF analysis presented in Section 5.6, we will marginalize over this variance, thus effectively desensitizing our analysis to any two-point information in the data. We use $\sigma_{\rm fid}$ to set the fiducial level of this CMB+foreground variance.

This procedure accounts for the Gaussian contributions due to CIB, point sources, kSZ, and Galactic dust. The lowest-order non-Gaussian contribution is the three-point function due to the CIB — we test for any impact in our results due to unmodeled non-Gaussian CIB contributions in Section 5.6 using the full CIB PDF measured from the simulations of [45], and find no detectable contamination. Higher-order non-Gaussian contributions exist due to the kSZ signal, which has zero skewness but non-zero kurtosis. However, its amplitude is far smaller than the non-Gaussian moments of the tSZ signal [40] (and CIB), and its kurtosis is undetectable at ACT noise levels. We test for its influence using the simulations of [45] and again find no detectable contamination. Finally, gravitational lensing of the CMB in principle induces a non-zero kurtosis, but its amplitude is small enough to be undetectable in a full-sky, cosmic-variance limited experiment [46]. We thus neglect it in our calculations.

Combining these results, we construct the total noise kernel, $N(y-y')$, by convolving all of the components with one another:

$$N(y-y') = \int dy'' N_{\rm inst}(y-y'-y'') \int dy''' N_{\rm CMB}(y''-y''') N_{\rm fg}(y''') \,. \tag{5.10}$$

This result for $N(y - y')$ can be directly applied to Eq. (5.5) to compute the probability of observing a signal in bin b_i given an input signal y, which completes the calculation of the noisy tSZ PDF in Eq. (5.7).

Fig. 5.3 shows the noisy tSZ PDF computed for our fiducial cosmology (WMAP9) and pressure profile ($P_0 = 1$, see Section 5.4.1, with the contributions from the tSZ signal (i.e., the second term in Eq. (5.7)) and the noise (i.e., the first term in Eq. (5.7)) shown individually. The lowest-$|\Delta\tilde{T}|$ bins near the peak of the PDF are dominated by the noise, as expected, while there is a clear non-Gaussian tail extending to negative $\Delta\tilde{T}$ values sourced by the tSZ effect. As a validation of our approach, Fig. 5.4 shows the same quantities for two variations in the lower mass cutoff used in the integrals of Eq. (5.7). We claimed above that the contributions of low-mass objects would simply be shifted between the two terms in Eq. (5.7), rendering the calculation insensitive to the exact value of the lower mass cutoff $M_{\rm min}$, as long as the criterion $F_{\rm clust} \ll 1$ is satisfied. Fig. 5.4 verifies this claim, showing that the total noisy tSZ PDF is unchanged whether we decrease the cutoff to $M_{\rm min} = 1.5 \times 10^{14} M_\odot/h$ (left panel of Fig. 5.4) or increase the cutoff to $M_{\rm min} = 1.5 \times 10^{14} M_\odot/h$ (right panel of Fig. 5.4). Note that we cannot use values of $M_{\rm min}$ significantly less than shown in the figure, since otherwise the $F_{\rm clust}$ criterion is not satisfied and the assumption of no cluster overlaps along the LOS breaks down. By comparing Figs. 5.3 and 5.4, it is clear that the total PDF remains unchanged — the contributions to the low-$|\Delta\tilde{T}|$ bins from the tSZ term and the noise term simply increase or decrease relative to one another. As argued above, this is simply due to the shifting of signal from low-mass clusters (whose signal is well below the noise in the map) from one term to the other. Quantitatively, the mean of the ratio between the fiducial PDF in Fig. 5.3 and the PDFs shown in the left and right panels of Fig. 5.4 over the twelve $\Delta\tilde{T}$ bins is 1.003 and 0.993, respectively. We take this as strong evidence that our theoretical calculations are under control and properly converged.

Fig. 5.5 presents the dependence of the noisy tSZ PDF on the most relevant cosmological and astrophysical parameters in our model, showing the ratio with respect to the fiducial

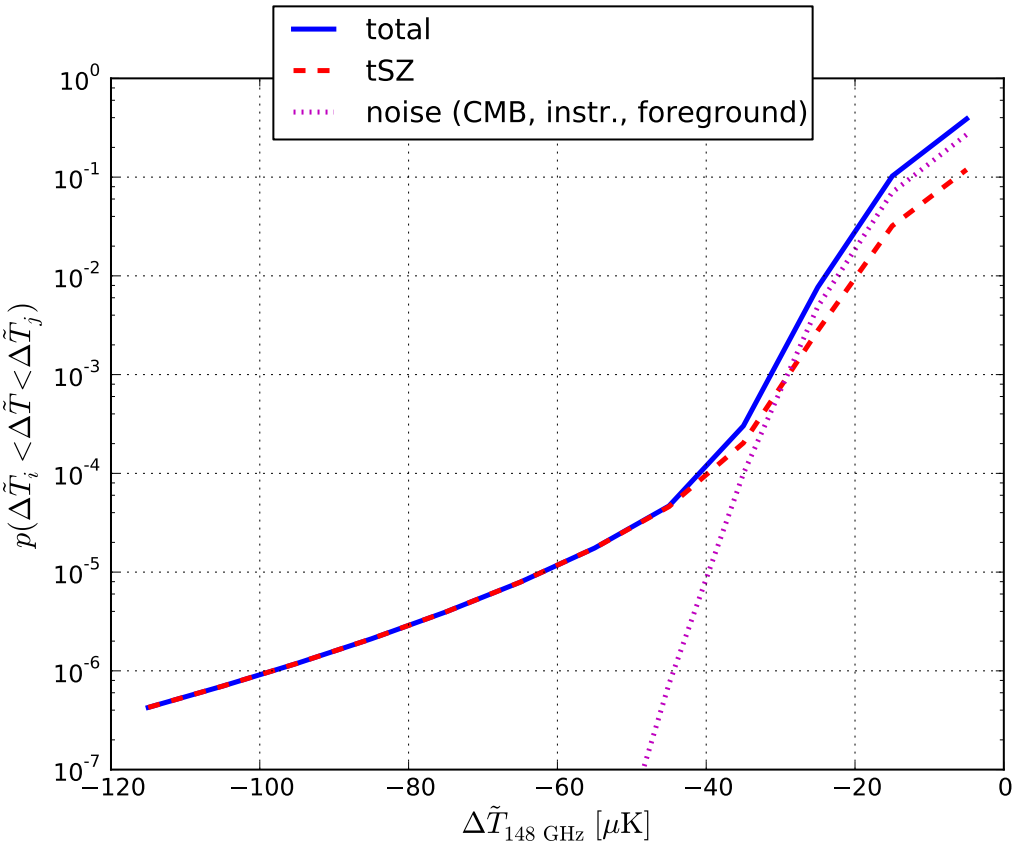

Figure 5.3: The noisy tSZ PDF computed for our fiducial cosmology and pressure profile model using the semi-analytic halo model framework described in Section 5.4.3. The dashed red curve shows the tSZ contribution given by the second term in Eq. (5.7), the dotted magenta curve shows the noise contribution given by the first term in Eq. (5.7), and the solid blue line shows the total signal.

case as each parameter is increased or decreased by 5%. As is well-known, nearly all tSZ observables are very sensitive to σ_8 (e.g., [49, 40, 48]), and we find that the tSZ PDF is no exception. Quoting simple power-law scalings, in the most negative $\Delta\tilde{T}$ bin shown in Fig. 5.5, the PDF scales as $\sim \sigma_8^{16}$, with the dependence remaining quite steep ($\sim \sigma_8^{10-15}$) in all bins until the noise contributions become important. These scalings compare favorably to those for the tSZ power spectrum ($\sim \sigma_8^{7-8}$) [49, 43, 42, 40] or bispectrum ($\sim \sigma_8^{10-12}$) [47, 48, 40].

The signal also depends non-trivially on Ω_m, the matter density, and P_0, the normalization of the pressure-mass relation $P_e(r, M, z)$. It is more sensitive to the latter, with the PDF in the most negative $\Delta\tilde{T}$ bin shown in Fig. 5.5 scaling as $\sim P_0^{4.4}$, while the same bin scales

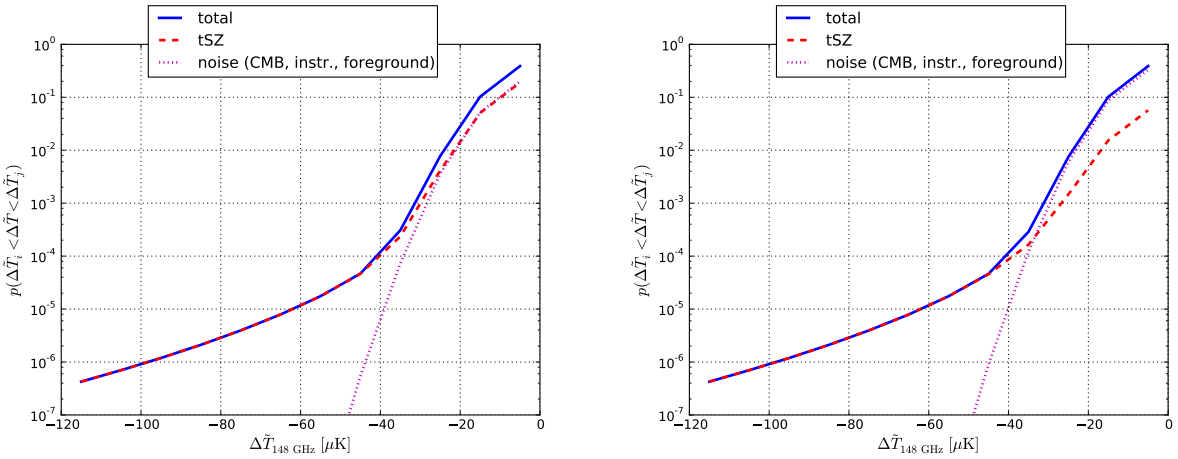

Figure 5.4: The noisy tSZ PDF computed for our fiducial cosmology and pressure profile model using the semi-analytic halo model framework described in Section 5.4.3, but with the lower mass cutoff in the integrals set to $1.5 \times 10^{14} M_\odot/h$ (left panel) and $3 \times 10^{14} M_\odot/h$ (right panel), slightly below and above the fiducial cutoff of $2 \times 10^{14} M_\odot/h$ (Fig. 5.3). Comparing the two panels and Fig. 5.3, it is evident that the total predicted tSZ PDF remains unchanged as the lower mass cutoff is varied — the contributions from low-mass clusters are simply shifted between the tSZ and noise terms in Eq. (5.7), as seen in the dashed red and dotted magenta curves, respectively. The non-Gaussian, tSZ-dominated tail of the PDF is dominated by clusters well above these mass cutoffs, and is thus insensitive to the exact cutoff choice as well.

with Ω_m as $\sim \Omega_m^{2.5}$. Due to the computational cost, we do not consider freeing additional cosmological or astrophysical parameters in our model. Simple Fisher matrix estimates indicate that σ_8 and Ω_m are the most relevant cosmological parameters, while the PDF's dependence on P_0 is comparable to that on the mean outer logarithmic slope of the pressure profile (see [30] for similar calculations). For simplicity, we only allow P_0 to vary. Future analyses with higher SNR may be able to simultaneously constrain several of these parameters, but a more efficient computational implementation will be necessary. A key point to note in Fig. 5.5 is that the ratio of the scaling with σ_8 and P_0 changes across the bins. In other words, these parameters are not completely degenerate in the tSZ PDF. This result suggests that a high-precision measurement of the tSZ PDF can simultaneously constrain cosmology (σ_8) and the pressure-mass relation (P_0), breaking the degeneracy between these quantities that is currently the limiting systematic in cluster cosmology analyses (e.g. [12, 31, 15, 14]). Finally, note that although we have quoted simple scalings for the PDF values in some bins,

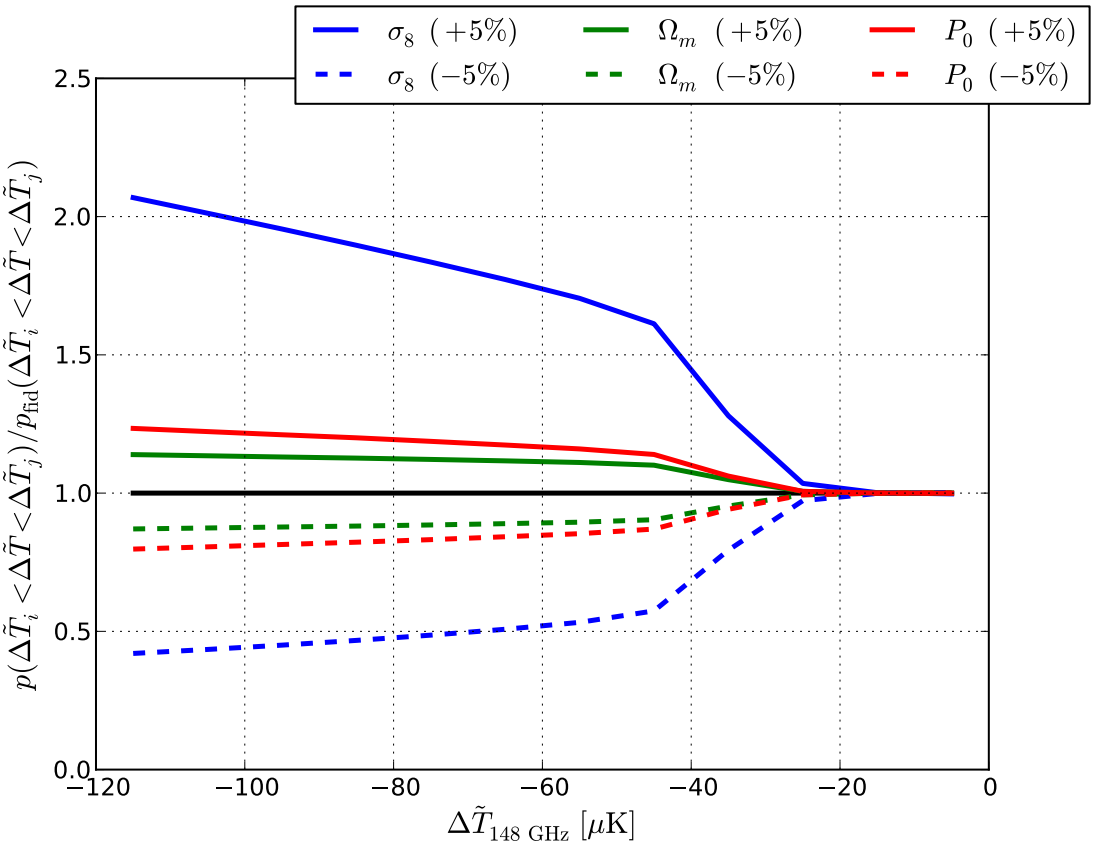

Figure 5.5: Dependence of the tSZ PDF on cosmological and astrophysical parameters. The plot shows the ratio with respect to the fiducial tSZ PDF (see Fig. 5.3) for ±5% variations for each parameter in the model $\{\sigma_8, \Omega_m, P_0\}$. The sensitivity to σ_8 is quite pronounced in the tSZ-dominated tail, with the values in these bins scaling as $\sim \sigma_8^{10-16}$. As expected, the bins in the noise-dominated regime are almost completely insensitive to variations in the cosmological or astrophysical parameters.

we do not use these scalings in our analysis; we evaluate the tSZ PDF directly using Eq. (5.7) at each sampled point in parameter space.

Fig. 5.6 shows the contributions to the noisy tSZ PDF in our fiducial calculation from various mass and redshift ranges. In the noise-dominated bins at low $|\Delta\tilde{T}|$ values, there is effectively no information about the mass and redshift contributions. Proceeding into the tSZ-dominated tail, the plots demonstrate the characteristic mass and redshift scales contributing to the signal. As expected, the outermost bins in the tail are sourced by the most massive clusters in the universe ($M \gtrsim 10^{15}\, M_\odot/h$), and hence are dominated by low-redshift contributions ($z \gtrsim 0.5$), since such clusters are extremely rare at high redshift. The first

tSZ-dominated bins beyond the noise ($|\Delta\tilde{T}| \approx 40$–$50\,\mu$K) are sourced by moderately massive clusters ($2 \times 10^{14}\,M_\odot/h \lesssim M \lesssim 9 \times 10^{14}\,M_\odot/h$), although still predominantly at fairly low redshifts. These statements are entirely dependent on the noise level and angular resolution of the experiment under consideration; an experiment with much lower noise and/or higher resolution than ACT could probe the PDF sourced by lower-mass, higher-redshift objects. The results here are specific to the ACT Equatorial 148 GHz map.

These results indicate that the theoretical modeling uncertainty for the objects dominating the tSZ PDF signal in the ACT map should not be extreme — there are many observational constraints on massive, low-redshift clusters, and theoretical considerations indicate that the thermodynamics of such objects should be dominated by gravitation rather than non-linear energy input from active galactic nuclei, turbulence, and other mechanisms (e.g., [42, 50]). Nonetheless, a substantial fraction of the signal arises from objects below the threshold for direct detection in blind mm-wave cluster surveys: for the ACT Stripe 82 cluster sample, the (rough) 90% completeness threshold for clusters detected at $S/N > 5$ (for which optical confirmation is 100% complete for $z < 1.4$) is $M_{500c} \approx 5 \times 10^{14}\,M_\odot/h_{70} = 3.5 \times 10^{14}\,M_\odot/h$ over $0.2 \lesssim z \lesssim 1.0$ (low-redshift clusters are difficult to detect due to confusion with primordial CMB anisotropies). Converting to the virial mass definition used here, this limit corresponds to roughly $M \approx 6$–$6.5 \times 10^{14}\,M_\odot/h$, depending on the redshift considered. However, the incompleteness increases dramatically for $z < 0.2$, so these numbers should only be taken as an approximate guide (see Section 3.6 and Fig. 11 of [12] for full details). This suggests that any additional statistical constraining power for the tSZ PDF compared to the number counts arises from the inclusion of objects that are potentially missing in the latter, in particular $z \lesssim 0.2$ groups and clusters over a wide range of masses (which are numerous) or low-mass, high-redshift objects (which are quite rare).

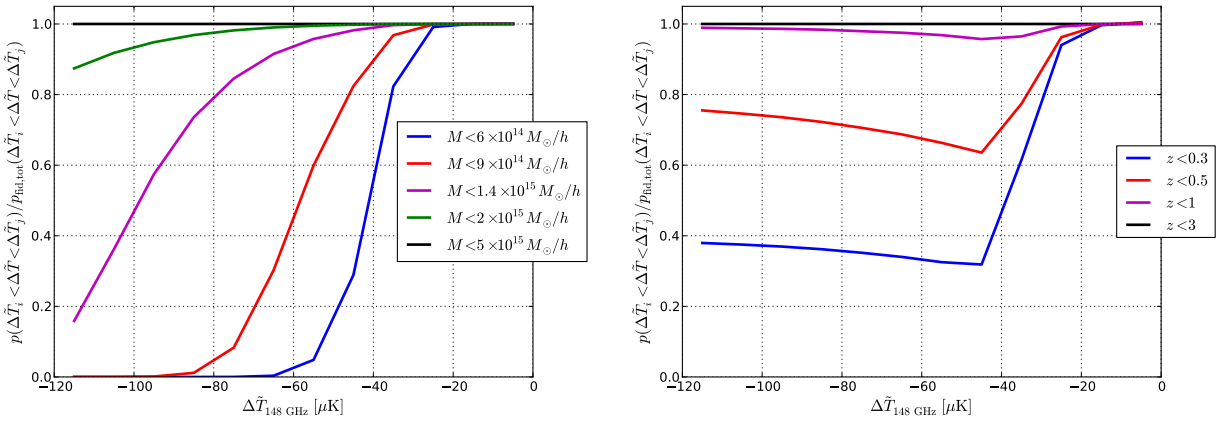

Figure 5.6: Mass and redshift contributions to the noisy tSZ PDF. The curves show the fraction of the total signal in each bin that is sourced by clusters at a mass or redshift below the specified value. In the low-$|\Delta\tilde{T}|$ bins, the total signal is dominated by noise (see Fig. 5.3), and thus there is essentially no information about the mass or redshift contributions. Below $\Delta\tilde{T} \approx -40\,\mu$K, the tSZ signal dominates over the noise, and robust inferences can be made about the mass and redshift contributions, as shown. Note that these results are specific to the noise levels and angular resolution of the filtered ACT Equatorial map, and must be recomputed for different experimental scenarios. See the text for further discussion.

5.4.4 Covariance Matrix

In the limit of uncorrelated, homogeneous noise, the covariance matrix of the noisy tSZ PDF can be computed using the semi-analytic halo model framework described in the previous section. The covariance matrix receives contributions from the Poisson statistics of the finite map pixelization, the Poisson statistics of the clusters, and the cosmic variance of the underlying density realization (often called the "halo sample variance" (HSV) term in other contexts) [51, 52, 53]. The HSV term is only relevant for low $|\Delta\tilde{T}|$-values, where less massive clusters become important in the PDF; this is analogous to the situation for the tSZ power spectrum covariance matrix, where the HSV term becomes relevant at high-ℓ where less massive halos dominate the power spectrum (see Fig. 1 in [54]). The pixel Poisson term only contributes to the diagonal elements of the covariance matrix. The cluster Poisson term, which arises from the finite sampling of the density field, dominates in the large-$|\Delta\tilde{T}|$ tail of the PDF. Note that the off-diagonal components of the tSZ PDF covariance matrix are

non-trivial, as a single cluster contributing to multiple bins in the PDF produces obvious bin-to-bin correlations. Using the notation defined in the previous section, the covariance matrix of the tSZ PDF, $\mathrm{Cov}_{ij} \equiv \langle P_i P_j \rangle - \langle P_i \rangle \langle P_j \rangle$, is given by:

$$\begin{aligned}
\langle P_i P_j \rangle - \langle P_i \rangle \langle P_j \rangle \;=\; & \frac{\langle P_i \rangle}{N_{\rm pix}} \delta_{ij} \; + \qquad \text{(pixel Poisson term)} \\
& \frac{1}{4\pi f_{\rm sky}} \int dz \frac{d^2V}{dz d\Omega} \int dM \frac{dn}{dM} \left(g_i(M,z) - \rho_i(0) \pi \theta_{out}^2(M,z) \right) \times \\
& \left(g_j(M,z) - \rho_j(0) \pi \theta_{out}^2(M,z) \right) + \qquad \text{(cluster Poisson term)} \\
& \int dz \frac{H(z)}{c\chi^2} \left(\frac{d^2V}{dz d\Omega} \right)^2 \times \\
& \left[\int dM_1 \frac{dn}{dM_1} b(M_1,z) \left(g_i(M_1,z) - \rho_i(0) \pi \theta_{out}^2(M_1,z) \right) \right] \times \\
& \left[\int dM_2 \frac{dn}{dM_2} b(M_2,z) \left(g_j(M_2,z) - \rho_j(0) \pi \theta_{out}^2(M_2,z) \right) \right] \times \\
& \int \frac{\ell d\ell}{2\pi} P_{\rm lin} \left(\frac{\ell}{\chi}, z \right) \left| \tilde{W}(\ell \Theta_s) \right|^2 , \qquad \text{(cluster HSV term)} \qquad (5.11)
\end{aligned}$$

where $N_{\rm pix} \approx 4 \times 10^6$ is the number of pixels in the map, $f_{\rm sky}$ is the total observed sky fraction, $b(M,z)$ is the linear halo bias, $P_{\rm lin}(k,z)$ is the linear matter power spectrum, and $\tilde{W}(\vec{\ell})$ is the Fourier transform of the survey window function, which is defined such that $\int d^2\theta W(\vec{\theta}) = 1$. We have used the Limber approximation [55] to obtain the HSV term and have assumed a circular survey geometry for simplicity, with $4\pi f_{\rm sky} = \pi \Theta_s^2$. We use the bias prescription of [56] and compute the linear matter power spectrum using CAMB.

Note that the covariance matrix in Eq. (5.11) depends on the cosmological and astrophysical parameters in our model. The correct approach would be to compute the PDF covariance matrix simultaneously with the PDF signal at each point in parameter space (e.g., [57]). However, due to the computational expense of this method (which is extreme when simulations are required — see the next section), we choose to make the approximation that the covariance matrix depends only on σ_8. Based on the results shown in Fig. 5.5 for the PDF signal, this approximation should be quite accurate, since the dependence of the covariance matrix on these parameters should scale roughly as the square of the PDF signal's depen-

dence. In other words, σ_8 should be by far the most relevant parameter. Regardless, given the computational expense required for the simulations described in the next section, we cannot go beyond this approximation at present.

For the current analysis, we choose to neglect the HSV term based on calculations using our fiducial model that indicate it contributes $\lesssim 10\%$ to the total variance in the bins over which the ACT PDF is sensitive to the tSZ term. The HSV term becomes more important when probing to lower masses, and in the future it should likely be included. This approximation also serves to improve the computational speed with which we can calculate the covariance matrix, given the extra integrals needed for the HSV term.

In addition, it is likely that the SNR of the tSZ PDF can be improved by masking some massive, low-redshift clusters, which contribute significant sample variance to the total error (through the cluster Poisson term in Eq. (5.11)). Similar effects arise in the covariance matrix of the tSZ power spectrum [60, 30] and bispectrum [48, 20]. However, the masking of clusters also introduces additional uncertainty into the measurement of tSZ statistics, as it requires knowledge of the Y–M relation (i.e., pressure–mass relation), leading to an uncertainty in the masking threshold. Finding the optimal masking scenario requires balancing these competing considerations (see [20]). Due to computational constraints, we do not attempt to find an optimal masking scenario in this analysis, and leave all tSZ signal unmasked in the ACT map.

While it is extremely useful to have a semi-analytic prescription for the covariance matrix, in practice the underlying approximation of uncorrelated noise is not valid for our analysis of the ACT 148 GHz PDF. Even if there were no instrumental or atmospheric noise, the CMB itself is a source of noise in our analysis, and it obviously has a non-trivial angular correlation function. The ℓ-space filtering applied to the maps in order to upweight the tSZ signal also has the byproduct of correlating the noise across pixels in real space. Thus, we require realistic Monte Carlo simulations of the ACT data in order to accurately estimate the covariance matrix.

5.5 Simulations

We construct realistic simulations of our data set in order to account for the effects of correlated, inhomogeneous noise in the covariance matrix and in order to robustly estimate any biases in our likelihood analysis in Section 5.6. Although knowledge of the one-point PDF of the component signals (regardless of their spatial correlations) suffices to compute the total one-point PDF of the filtered ACT 148 GHz map, as described in Section 5.4.3, this knowledge is not sufficient to compute the covariance matrix of the PDF in the presence of correlated noise. It is straightforward to show that in this case estimation of the covariance matrix requires knowledge of the contribution of each pair of pixels in the map to the covariance between bin i and bin j (for each component signal in the map). Assuming homogeneity, this would depend only on the distance between the two pixels for each pair; however, the noise in our map is not perfectly homogeneous, invalidating this possible simplification. Thus, even for the noise alone (including CMB, instrumental, and atmospheric noise), we would require simulations. We choose to take the additional step of simply simulating all relevant components in the map. Note that computing the one-point PDF itself depends only on the average properties of a single pixel in the map — there is no sum over pairs of pixels — and hence a simple convolution of the component PDFs suffices to compute the total PDF, as outline in Section 5.4.3. The computation of the covariance matrix in the presence of correlated, inhomogeneous noise necessitates simulations, which also allow tests of some assumptions in the semi-analytic calculations of the PDF, including the approximation of no overlapping clusters along the LOS.

The simulation pipeline is comprised of tools produced for earlier ACT studies, in particular the CMB temperature power spectrum and CMB lensing analyses [58, 59], as well as new tSZ tools constructed specifically for this analysis. The steps in the pipeline are as follows:

1. We generate a random CMB temperature field seeded by the angular power spectrum of our fiducial WMAP9 cosmology; this temperature field is then gravitationally lensed (which is not necessary for this analysis as argued in Section 5.4.3, but was already built into the pipeline). The details of these steps can be found in [58].

2. We generate a Gaussian random field seeded by the best-fit power spectrum of all non-tSZ foreground components as determined in [17], following the same method with which the random CMB temperature field is generated. These foreground components are described in Section 5.4.3.

3. We generate a map of Poisson-distributed tSZ clusters drawn from the Tinker mass function [41] integrated over the same mass and redshift limits used in our semi-analytic calculations. The Compton-y profile of these clusters is computed using our fiducial pressure profile model from [29], which is discretized on a grid of ACT-sized pixels (roughly $0.25\,\mathrm{arcmin}^2$).

4. We sum the lensed CMB map, foreground map, and tSZ cluster map.

5. We add in randomly-seeded realistic ACT noise using the method described in Section 4 of [58]. The seed for the noise is generated from difference maps of subsets of the ACT Equatorial data, and hence accounts for all non-trivial correlation properties and inhomogeneity in the noise. The noisy simulated maps are then processed with the appropriate ACT 148 GHz beam [44].

6. We filter and process the maps using exactly the same procedure as used on the ACT 148 GHz data map (see Section 5.3).

We generate 476 simulated ACT Equatorial maps using this procedure, corresponding to ≈ 155000 square degrees of simulated sky. We then compute the mean PDF and the PDF covariance matrix using the 476 simulated maps. Furthermore, since the covariance matrix of the PDF is a function of σ_8 (see the previous section), we repeat this process for 20 values

of σ_8 between 0.7 and 0.9. The simulations account for all relevant signals in the $\Delta\tilde{T} < 0$ region of the 148 GHz PDF, with the possible exception of higher-order moments of the unresolved point source, CIB, or kSZ fields.

We verify that the tSZ cluster map is insensitive to the details of the pixelization by increasing the resolution by a factor of two (i.e., factor of four in area) and re-running the third step in the pipeline listed above. We also verify that edge effects do not introduce any bias in the results. The simulated maps allow us to directly quantify the effect of overlapping clusters along the LOS: we find that for our fiducial mass and redshift limits, about 10% of the sky area populated by clusters (which is about 30% of the full sky) consists of LOS overlaps of multiple objects. This test verifies that the assumptions in our semi-analytic calculations are valid. We demonstrate this consistency directly in Fig. 5.7, which displays the tSZ PDF for our fiducial cosmology and pressure profile computed using the semi-analytic method in Section 5.4.3 and using the simulations described above. The results are in nearly perfect agreement.

In order to obtain the covariance matrix of the tSZ PDF for use in the likelihood analysis in the next section, we combine the smooth semi-analytic results found using Eq. (5.11) with the covariance matrices obtained from the simulations. Attempts to use the simulated covariance matrices directly led to ill-behaved likelihood functions due a lack of convergence in far off-diagonal elements of the covariance matrix. Although 476 simulations should suffice to determine the covariance matrix of 12 (correlated) Gaussian random variables, the values of the PDF in the tSZ-dominated bins are not perfectly Gaussian-distributed, especially in the highest-$|\Delta\tilde{T}|$ bins, which are sourced by rare, massive clusters. Moreover, the convergence of the simulation-estimated covariance matrix depends on the value of σ_8, since many fewer clusters are distributed for low σ_8 values compared to high σ_8 values, thus requiring many more simulations to achieve convergence.

Motivated by these issues, in practice we combine the smooth semi-analytic results computed using Eq. (5.11) with corrections computed from the simulations. In particular, the

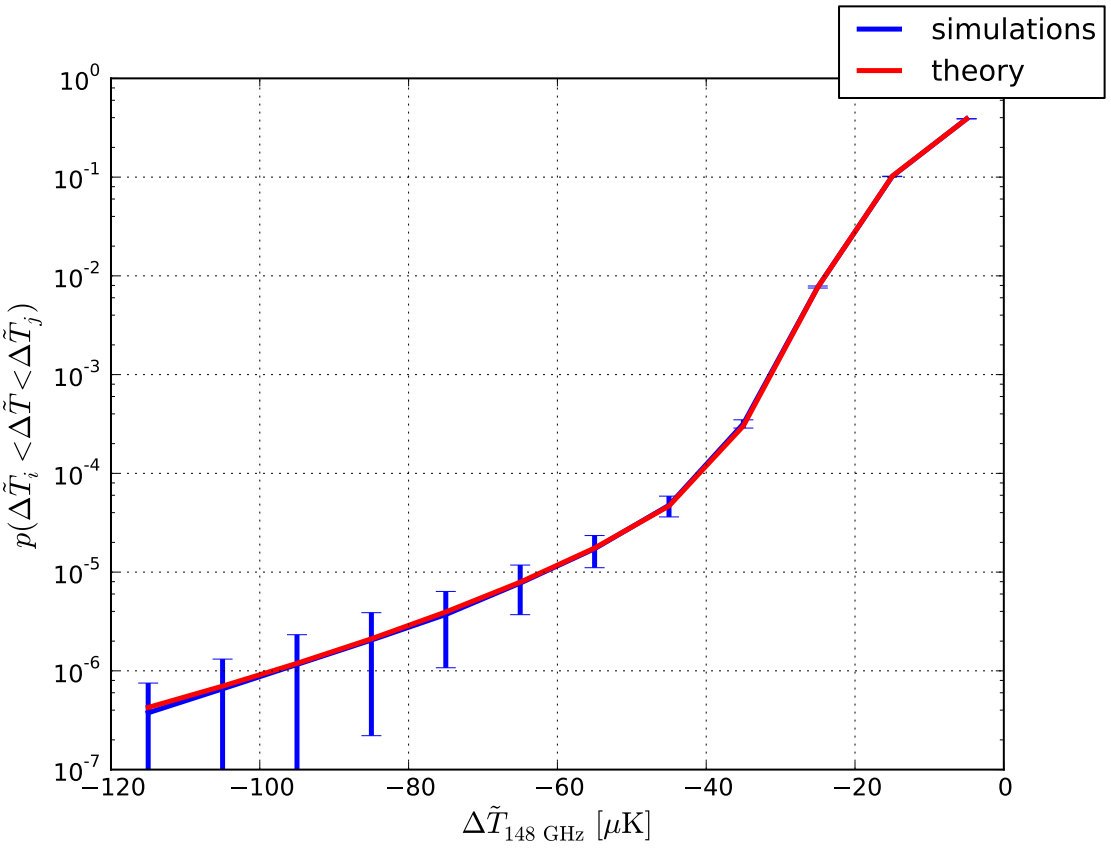

Figure 5.7: Comparison of the semi-analytic halo model computation of the noisy tSZ PDF using Eq. (5.7) to the PDF directly measured from the simulated maps described in Section 5.5. The error bars are computed from the scatter amongst the 476 simulated maps. The two curves are nearly indistinguishable. This validates the halo model approximation that clusters do not overlap along the LOS (see Section 5.4.2).

diagonal elements of the simulated covariance matrices are well-converged for all values of σ_8, as are all of the low-$|\Delta\tilde{T}|$ elements (on- and off-diagonal), since those bins are highly populated in all realizations. Our hybrid method for computing the covariance matrix is thus as follows (note that in the following "bin i" refers to the ith bin in the tSZ PDF, starting at $\Delta\tilde{T} = 0$ for bin 0 and counting to progressively more negative bins). For bins 0–2, we compute the ratio of the simulated covariance matrix to the semi-analytic covariance matrix and fit it with a linear function of σ_8 for each of the six independent elements in this sub-matrix. We perform the same procedure for the diagonal elements of the sub-matrix spanned by bins 3–11. For the off-diagonal elements of this sub-matrix, we compute no correction

factor since the simulation results and the semi-analytic results are in reasonable agreement (within the convergence limitations of the simulated results). Finally, for the off-diagonal elements linking the two sub-matrices (e.g., Cov_{03}, Cov_{13}, Cov_{04}, etc.), we compute a correction given by the square root of the product of the corresponding on-diagonal elements' corrections. We then apply the correction factors computed using the linear fits to each of the semi-analytic covariance matrices as a function of σ_8. In general, the correction factors computed in this procedure are most significant in the low-$|\Delta T|$ bins, where the noise treatment in the semi-analytic calculations is most inaccurate. In the tSZ-dominated tail, the semi-analytic and simulation results are in quite good agreement (i.e., within a factor of 1–2). The output of this pipeline is an estimate of the covariance matrix that is a smooth function of σ_8 and simultaneously includes the effects of correlated, inhomogeneous noise captured in the simulations. Fig. 5.8 presents the covariance matrix computed for our fiducial cosmology and pressure profile. In particular, bin-to-bin correlations induced by correlated noise and by clusters that contribute tSZ signal to multiple bins are clearly visible.

5.6 Interpretation

5.6.1 Likelihood Function

We use the measurement of the tSZ PDF in the ACT Equatorial 148 GHz data to constrain the cosmological and astrophysical parameters in our model. Given that σ_8 is the most relevant cosmological parameter by a significant margin (see Fig. 5.5), we only allow it to vary amongst the ΛCDM parameters. We will further consider scenarios in which the overall normalization of the pressure-mass relation, P_0, is either fixed or is free to vary. Finally, we marginalize all results over a nuisance parameter corresponding to the variance in the PDF, σ^2_{nuis}. The fiducial value of this nuisance parameter is $\sigma^2_{\mathrm{nuis}} = \sigma^2_{\mathrm{fid}}$ given in Section 5.4.3. By marginalizing over it, we desensitize our results to any unknown Gaussian component in the map, including the contribution of any residual non-tSZ foregrounds or CMB. We place a

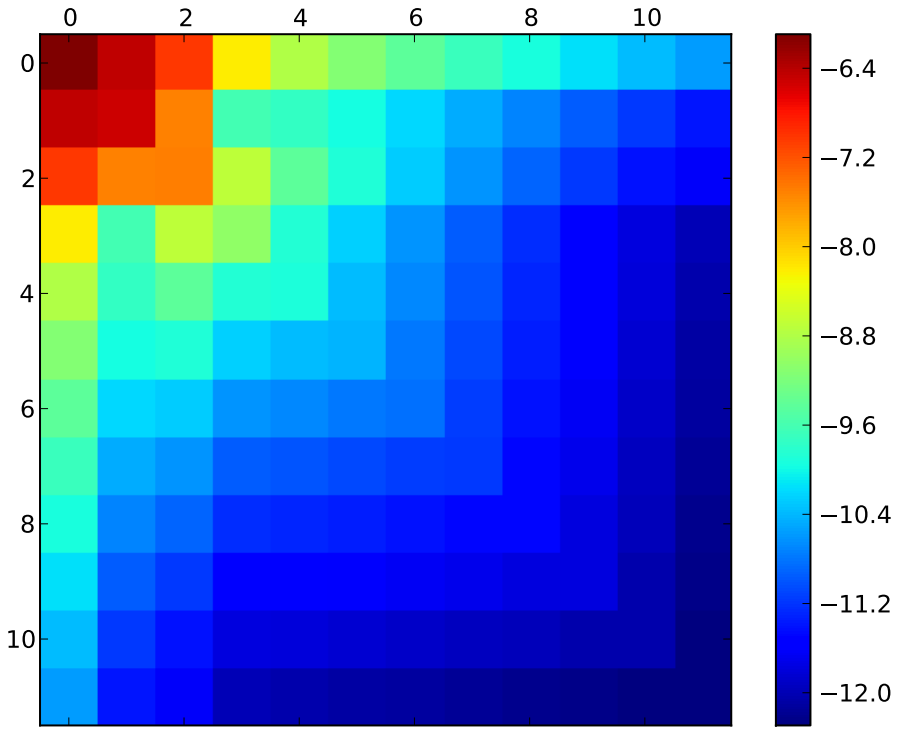

Figure 5.8: Covariance matrix of the noisy tSZ PDF computed for our fiducial cosmology and pressure profile model using the semi-analytic halo model framework described in Section 5.4.4 with corrections computed from the Monte Carlo simulations described in Section 5.5. The axes are labeled by bin number, with the top left corner corresponding to the lowest $|\Delta\tilde{T}|$ bin shown in Fig. 5.3 (i.e. the noise-dominated regime) and the bottom right corner corresponding to the highest $|\Delta\tilde{T}|$ bin shown in Fig. 5.3 (i.e. the tSZ-dominated regime). The specific quantity plotted is $\log_{10}(|\mathrm{Cov}|)$, with values labeled by the colors (the absolute value is necessary because some of the off-diagonal bins representing correlations between the noise-dominated region and the tSZ-dominated region are negative). Note the significant amount of bin-to-bin correlation, especially amongst the tSZ-dominated bins in the tail of the PDF.

Gaussian prior of width $0.1\,\mu$K on $\sigma_{\rm nuis}$, since its fiducial value is computed using the CMB power spectrum (which is well-known) and the best-fit non-tSZ foreground power spectra from [17] which are directly measured in the Equatorial map. The primary motivation for $\sigma^2_{\rm nuis}$ is to capture any small amount of residual CIB or kSZ that could leak into the $\Delta\tilde{T} < 0$ PDF. In order to further prevent this leakage, we discard the lowest-$|\Delta\tilde{T}|$ bin in the PDF in the likelihood analysis, which corresponds to roughly the 1σ fluctuations in the map (the RMS of the filtered map is $8.28\,\mu$K, and the bins are $10\,\mu$K wide). We are thus left with 11 bins, most of which are dominated by tSZ signal.

We construct the following likelihood function:

$$\begin{aligned}\mathcal{L}(\sigma_8, P_0, \sigma^2_{\rm nuis}) &= \frac{1}{\sqrt{(2\pi)^{N_b}\det(\mathrm{Cov}(\sigma_8))}} \times \\ &\exp\left(-\frac{1}{2}(P_i(\sigma_8, P_0, \sigma^2_{\rm nuis}) - \hat{P}_i)(\mathrm{Cov}^{-1})_{ij}(\sigma_8)(P_j(\sigma_8, P_0, \sigma^2_{\rm nuis}) - \hat{P}_j)\right) \end{aligned} \quad (5.12)$$

where $N_b = 11$ is the number of bins in the measurement, Cov_{ij} is the covariance matrix described in the previous section (computed as a function of σ_8 only), P_i is the PDF value in bin i predicted by the semi-analytic theory described in Section 5.4.3, and $\hat{P}_i$ is the measured PDF value in bin i. As noted above, we marginalize all results over $\sigma^2_{\rm nuis}$. The likelihood function in Eq. (5.12) obviously rests on the assumption of Gaussianity, which breaks down in the bins far in the tSZ-dominated tail of the PDF, which are rarely populated (in fact, the outermost four bins in the ACT PDF are empty). An approach based on Poisson statistics may be more appropriate for these bins, but this raises questions about properly accounting for bin-to-bin correlations, which Fig. 5.8 indicates are strong. The only reference we are aware of in the literature that presents a likelihood for the PDF in general (not the tSZ PDF in particular) is [61]. However, they assume that pixels, and hence bins, are completely uncorrelated, which is a highly inaccurate approximation for our purpose. In the absence of a more suitable likelihood, we use Eq. (5.12) and quantify any resulting bias in our constraints using simulations.

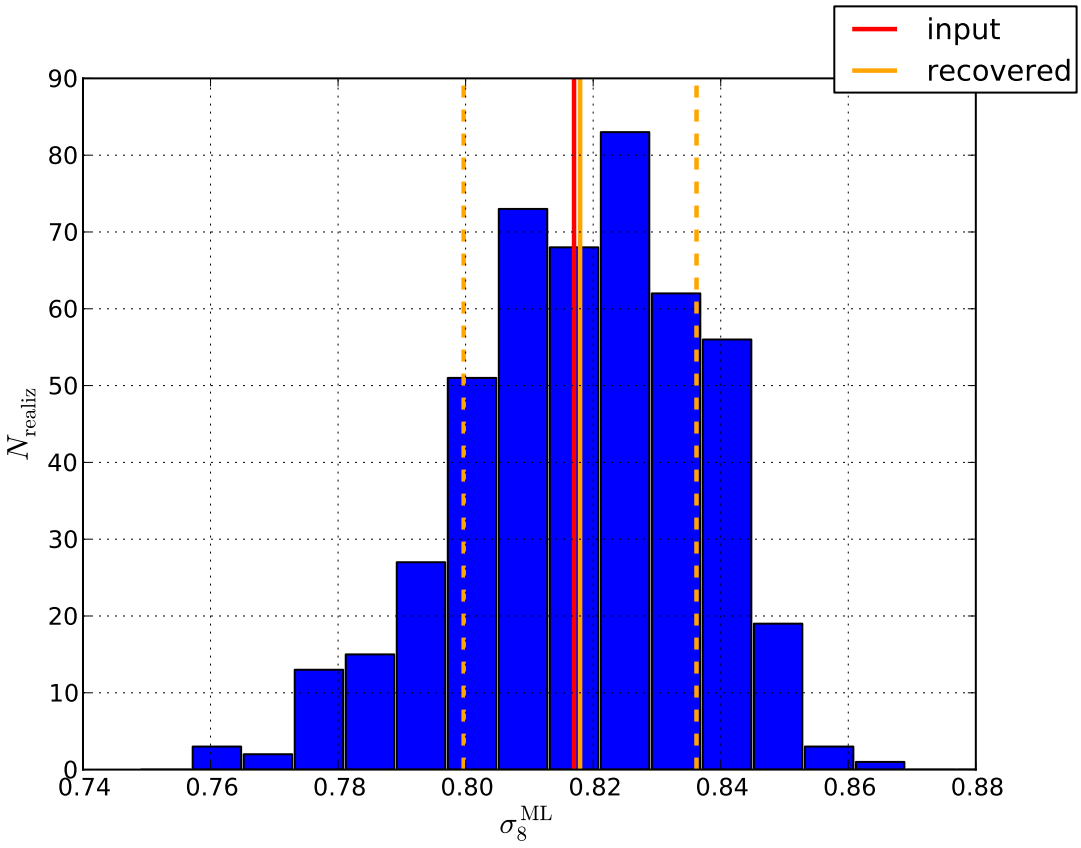

Figure 5.9: Histogram of ML σ_8 values recovered from 476 Monte Carlo simulations processed through the likelihood in Eq. (5.12) with P_0 fixed to unity and the nuisance parameter $\sigma^2_{\rm nuis}$ marginalized over. The input value, $\sigma_8 = 0.817$ (indicated by the red vertical line), is recovered in an unbiased fashion: the mean recovered value is $\langle\sigma_8^{\rm ML}\rangle = 0.818 \pm 0.018$ (indicated by the orange vertical lines).

We first consider the case in which P_0 is fixed to its fiducial value, i.e., unity. We place no prior on σ_8, and treat $\sigma^2_{\rm nuis}$ as described above, marginalizing over it in our results. We process 476 Monte Carlo simulations described in Section 5.5 for our fiducial model ($\sigma_8 = 0.817$, $\sigma^2_{\rm nuis} = \sigma^2_{\rm fid}$, $P_0 = 1$) through the likelihood. We marginalize the likelihood over $\sigma^2_{\rm nuis}$ for each simulation and then find the maximum-likelihood (ML) value for σ_8. The resulting histogram of ML σ_8 values is shown in Fig. 5.9. The mean recovered value is $\langle\sigma_8^{\rm ML}\rangle = 0.818 \pm 0.018$, which agrees extremely well with the input value $\sigma_8 = 0.817$. This result indicates that the P_0-fixed likelihood for σ_8 is unbiased.

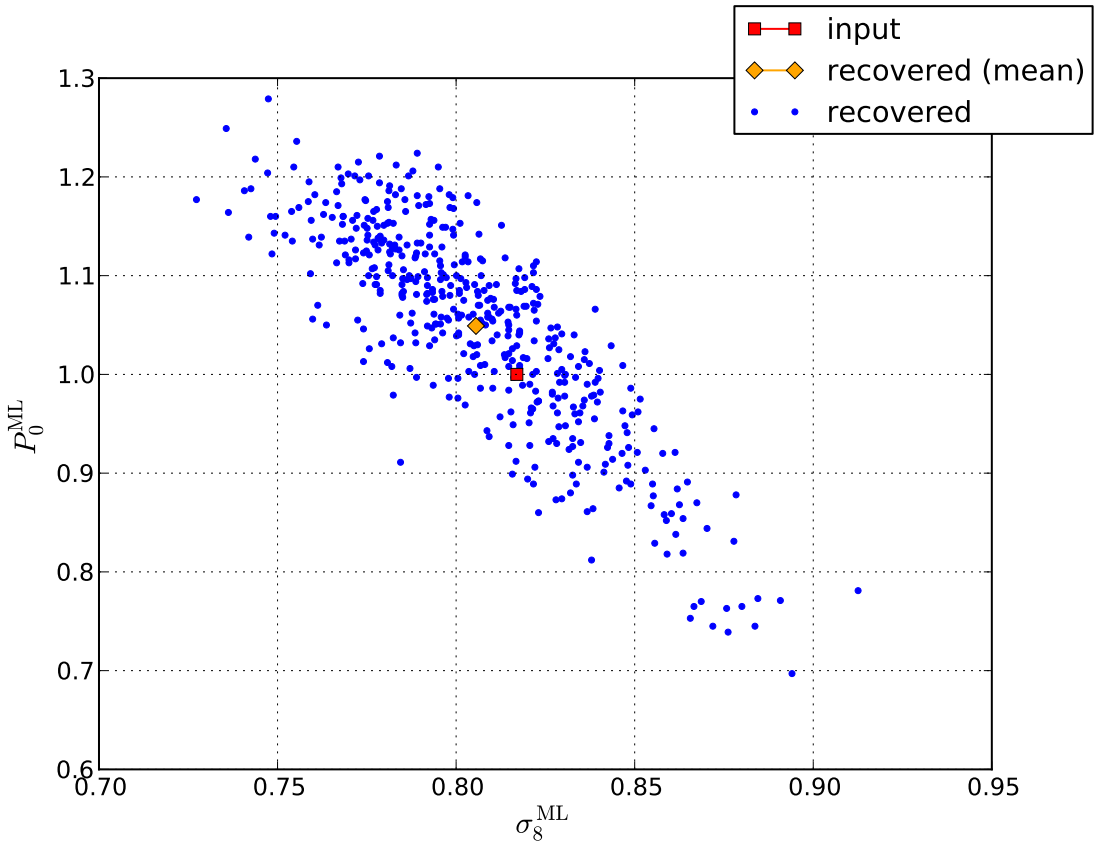

Figure 5.10: Scatter plot of ML recovered values of σ_8 and P_0 using 476 Monte Carlo simulations processed through the likelihood in Eq. (5.12) with both parameters allowed to vary (the nuisance parameter $\sigma^2_{\rm nuis}$ can also vary and is marginalized over). The approximate degeneracy between σ_8 and P_0 seen here can be inferred from the scalings shown in Fig. 5.5. The input values, $\sigma_8 = 0.817$ and $P_0 = 1$, are represented by the red square; the mean recovered ML result is represented by the orange diamond. A bias can clearly be seen; this is further investigated in the 1D marginalized results in Fig. 5.11.

We next consider the case in which both P_0 and σ_8 are allowed to vary. We place the same prior on $\sigma^2_{\rm nuis}$ as before and marginalize over it to obtain a 2D likelihood for P_0 and σ_8. The parameter dependences shown in Fig. 5.5 suggest that the effects of σ_8 and P_0 on the tSZ PDF should not be completely degenerate. However, a full breaking of the degeneracy would require a high-SNR measurement of the PDF over a wide range of $\Delta\tilde{T}$ values, so that the different changes in the shape of the PDF produced by the two parameters could be distinguished. Testing the 2D likelihood for P_0 and σ_8 using our Monte Carlo simulations, we find that the degeneracy cannot be strongly broken at the current ACT sensitivity and resolution. Thus, in order to allow us to quote meaningful marginalized 1D constraints on these parameters, we place a weak prior on P_0 motivated by the many recent observational studies that agree with our fiducial pressure profile (the "AGN feedback" fit from [29]) [34, 35, 36, 38, 31]. In particular, [31] used a measurement of the tSZ – CMB lensing cross-power spectrum to constrain P_0 in the context of a fixed WMAP9 background cosmology, finding that $P_0 = 1.10 \pm 0.22$. We thus place a Gaussian prior of width 0.2 on P_0, centered on the fiducial value $P_0 = 1$. (Note that this prior thus also encompasses the constraint on P_0 inferred in [31] using the Planck collaboration's measurement of the tSZ auto-power spectrum [21].)

We process the 476 Monte Carlo simulations of our fiducial model through the full (σ_8, P_0, T_0) likelihood. We marginalize over T_0 and compute the ML point in the (σ_8, P_0) 2D parameter space. The results are plotted in Fig. 5.10. A bias in the recovered values is evident. This bias is further detailed in Fig. 5.11, which shows the histograms of recovered σ_8 and P_0 values obtained after marginalizing the 2D likelihood over the other parameter. The mean recovered values are $\langle\sigma_8^{\rm ML}\rangle = 0.806 \pm 0.030$ and $\langle P_0^{\rm ML}\rangle = 1.05 \pm 0.10$, which disagree with the input values of $\sigma_8 = 0.817$ and $P_0 = 1$, respectively. The cause of this bias is currently unclear. It may be a symptom of the failure of the Gaussian approximation underlying our likelihood — in the outermost bins in the tSZ-dominated tail of the PDF, this approximation must break down as rare clusters dominate the signal. However, the error

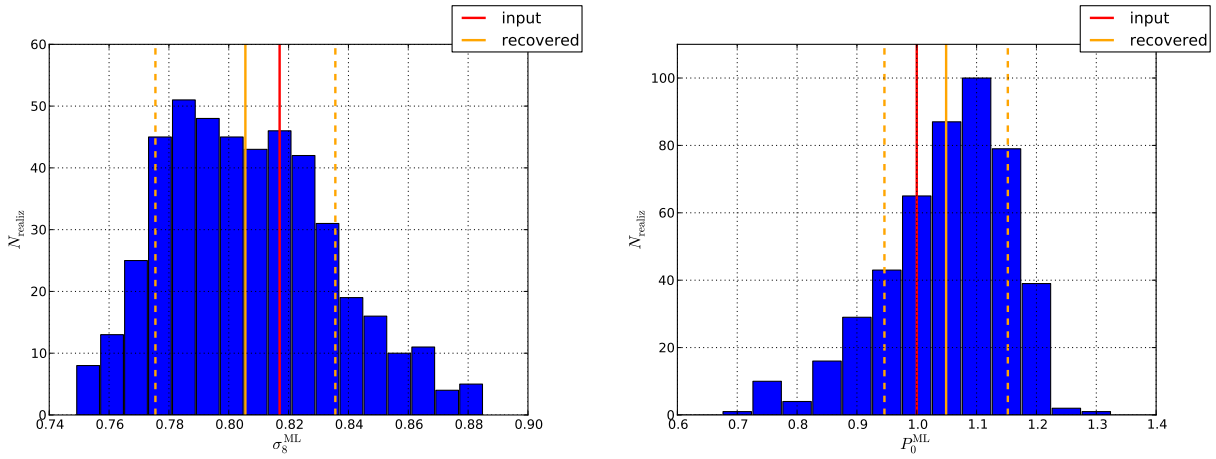

Figure 5.11: Histogram of ML σ_8 values (left panel) and P_0 values (right panel) recovered from 476 Monte Carlo simulations processed through the likelihood in Eq. (5.12) after marginalizing over all other parameters. The input values, $\sigma_8 = 0.817$ and $P_0 = 1$, are indicated by the red vertical lines in each panel. There is evidence of a bias in the likelihood: the mean recovered values are $\langle\sigma_8^{\rm ML}\rangle = 0.806 \pm 0.030$ and $\langle P_0^{\rm ML}\rangle = 1.05 \pm 0.10$ (indicated by the orange lines in each panel). The cause of this bias is presently unclear (see text for discussion), and we treat it as a systematic which must be corrected for in the ACT data analysis.

bars are quite large on the outermost bins, which should downweight their relative influence in the likelihood. An alternative possibility is that it is a poor approximation to treat the covariance matrix as a function of σ_8 only (see Section 5.4.4), and its dependence on P_0 must be modeled as well. For now, we treat the biases seen in Figs. 5.10 and 5.11 as systematics requiring corrections to be imposed on the constraints derived using the ACT data in the following section.

5.6.2 Constraints

We apply the likelihood described in the previous section to the measured PDF of the filtered, processed ACT Equatorial 148 GHz maps described in Section 5.3. All results are marginalized over $\sigma^2_{\rm nuis}$, effectively removing all sensitivity to the two-point statistics (i.e., variance) of the data. In all constraints quoted in the following, we report the best-fit value as

the mean of the marginalized likelihood, while the lower and upper error bounds correspond to the 16% and 84% points in the marginalized cumulative distribution, respectively.

We first set $P_0 = 1$ and let only σ_8 vary, analogous to the simulated results shown in Fig. 5.9. The resulting 1D likelihood for σ_8 is shown in Fig. 5.12. We find

$$\sigma_8 = 0.767 \pm 0.017\,. \tag{5.13}$$

This result can be directly compared to the primary constraint derived from the tSZ skewness measurement from the same data set in [19], which found $\sigma_8 = 0.79 \pm 0.03$ under the same set of assumptions (i.e., a fixed gas pressure profile model and all other cosmological parameters held fixed). The error bar on the PDF-derived result is roughly a factor of 2 smaller than that from the skewness alone. The additional statistical power contained in the higher moments of the tSZ field beyond the skewness is responsible for this decrease in the error (again, note that we are not using two-point level information in this analysis, i.e., the tSZ power spectrum/variance). The central value of σ_8 is also somewhat lower than that found in [19]. The most plausible explanation for this result is the empty bins in the tail of the ACT PDF (see Figs. 5.1 and 5.15) — the higher moments beyond the skewness are progressively more dominated by the bins far in the tail, and hence their influence likely pulls σ_8 down somewhat from the skewness-derived result.

At this level of precision on σ_8, it is essential to consider uncertainties in the gas pressure profile model as well. Moreover, as we have highlighted earlier, the tSZ PDF does possess the promise of breaking the degeneracy between σ_8 and ICM astrophysics. Thus we consider the likelihood with both σ_8 and P_0 allowed to vary (in addition to $\sigma^2_{\rm nuis}$). In the previous section we found evidence of a bias in this likelihood using simulations, and we will correct our final constraints on σ_8 and P_0 using these results below.

Fig. 5.13 shows the 2D likelihood for σ_8 and P_0, with no bias corrections applied. Although the parameters are fairly degenerate, the ACT PDF does weakly break the degener-

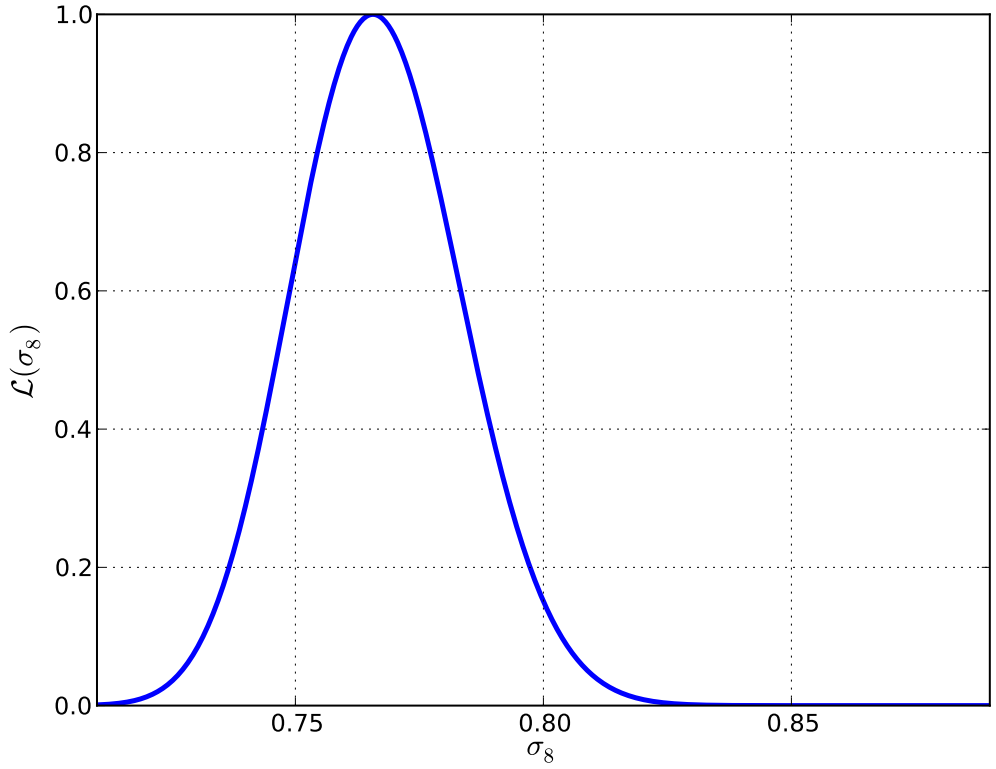

Figure 5.12: Likelihood of σ_8 computed from the ACT 148 GHz PDF using the likelihood function described in Section 5.6.1 with P_0 fixed to unity (i.e., only σ_8 and $\sigma^2_{\rm nuis}$ are free to vary, and the latter is marginalized over).

acy, using tSZ data alone. Fig. 5.14 presents the corresponding 1D marginalized likelihood for each parameter, also with no bias corrections applied. We find

$$\begin{aligned} \sigma_8 &= 0.738^{+0.024}_{-0.020} \\ P_0 &= 1.17 \pm 0.14\,. \end{aligned} \tag{5.14}$$

Note that the error on P_0 is substantially smaller than the width of the prior, indicating that the measured ACT PDF is able to constrain both cosmology and ICM astrophysics simultaneously. We have also verified that with no prior on P_0 we obtain results consistent with those in Eq. (5.14) at the $\approx 1\sigma$ level.The prior on P_0 downweights unrealistically large values of P_0 that otherwise the likelihood tends to prefer.

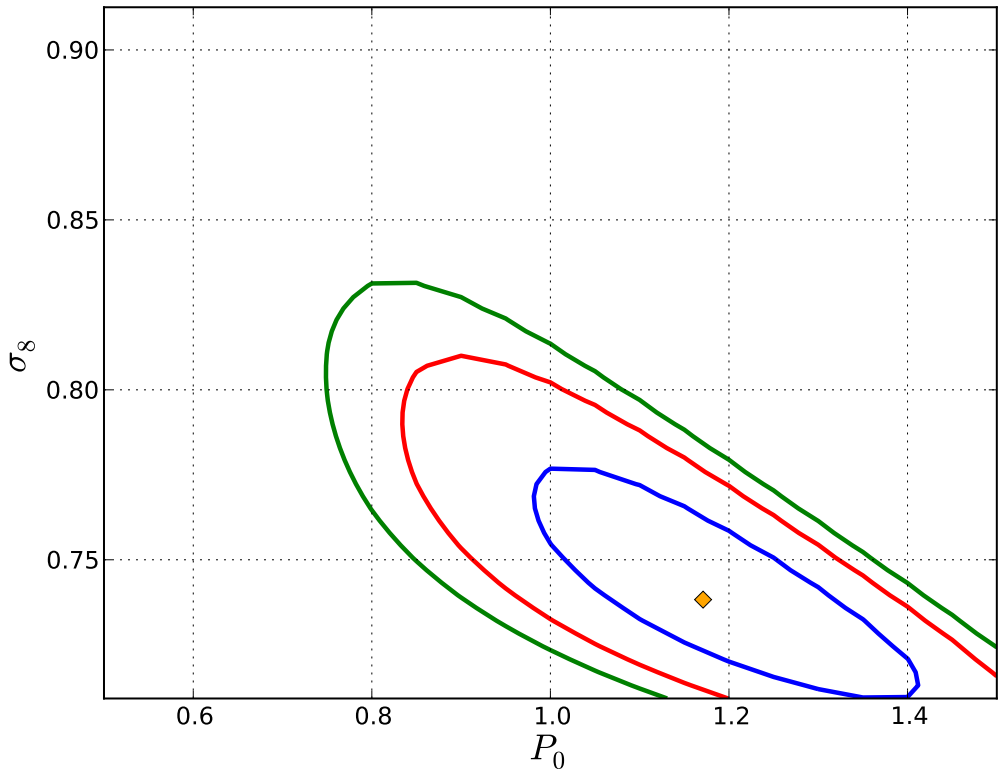

Figure 5.13: Likelihood contours for σ_8 and P_0 computed from the ACT 148 GHz PDF using the likelihood function described in Section 5.6.1 with both parameters (and $\sigma^2_{\rm nuis}$) allowed to vary. The blue, red, and green curves represent the 68%, 95%, and 99% confidence level contours, respectively. The orange diamond represents the ML point in the 2D parameter space. These results are not corrected for the biases found in Section 5.6.1. Note that the PDF measurement is indeed able to weakly break the degeneracy between σ_8 and P_0, using tSZ data alone.

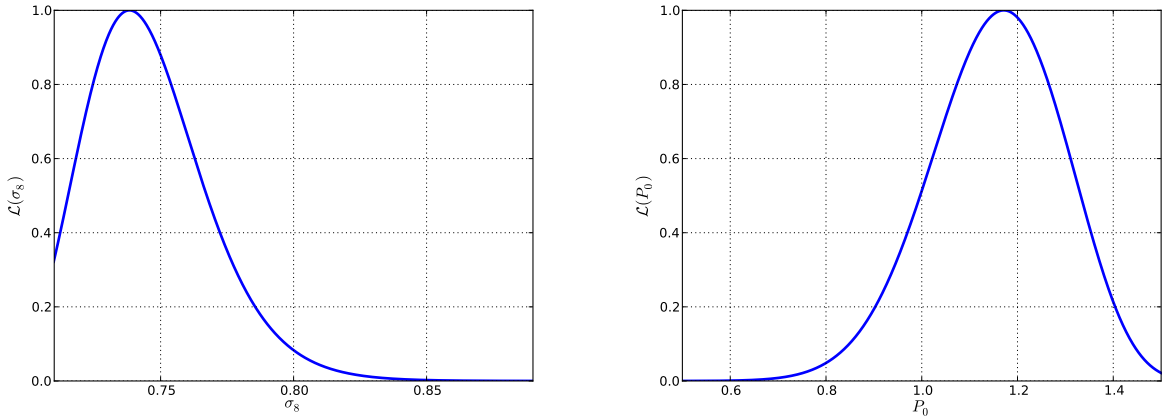

Figure 5.14: One-dimensional likelihood profiles for σ_8 (left panel) and P_0 (right panel) after marginalizing the 2D likelihood shown in Fig. 5.13 over all other parameters. These results are not corrected for the biases found in Section 5.6.1.

Fig. 5.15 presents a comparison of the ACT PDF to the ML PDF models found for both the σ_8-only likelihood and the (σ_8, P_0) likelihood. We leave out the lowest bin in the PDF in these plots as it was not fit in the likelihood analysis. The ML model is quite similar in the two cases, which is not surprising given the degeneracy between σ_8 and P_0. The empty bins in the tail of the ACT PDF appear to have a noticeable effect in driving the overall best-fit result to a lower amplitude in both cases. This is investigated further in Fig. 5.16, which shows the difference between the ACT PDF and the ML model for both likelihood scenarios divided by the square root of the diagonal elements of the covariance matrix of the ML model. We emphasize that the diagonal-only nature of this plot obscures the significant bin-to-bin correlations in the PDF, and thus no strong inferences should be derived. Nonetheless, it is interesting that the best-fit model lies slightly high over most of the range of the PDF, with the outer bins apparently responsible for the pull toward lower amplitudes. Note that none of the results shown in Figs. 5.15 and 5.16 are computed for bias-corrected values of σ_8 and P_0.

Finally, to complete our analysis we compute corrections to the constraints given in Eq. (5.14) in order to account for the systematic biases found in Section 5.6.1. This bias is perhaps also hinted at by the substantial shift in the σ_8 result in Eq. (5.14) from the P_0-

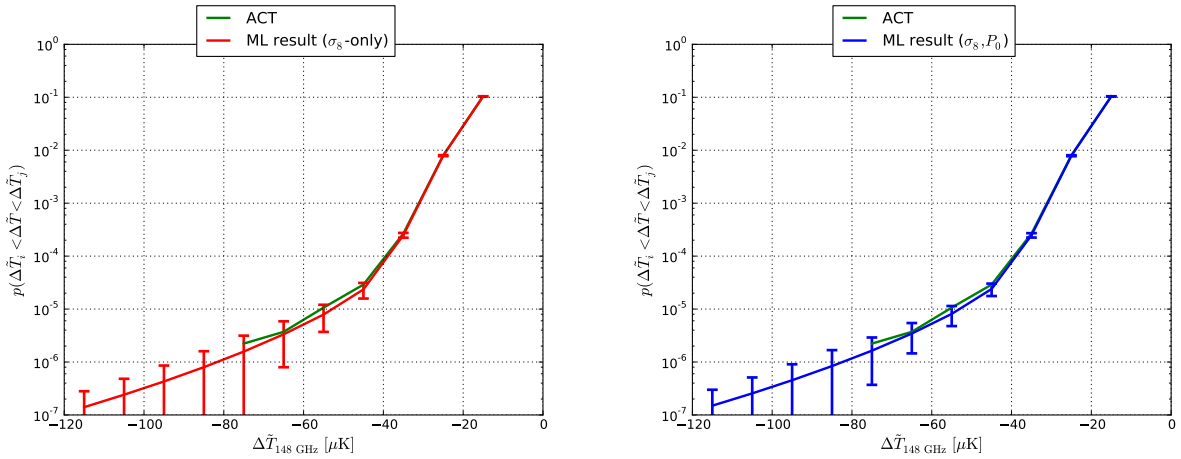

Figure 5.15: Comparison of the measured ACT PDF to the ML PDF model for the σ_8-only likelihood (left panel) and (σ_8, P_0) likelihood (right panel).

fixed result presented in Eq. (5.13). We correct for the bias on each parameter assuming it is multiplicative, i.e., σ_8 is corrected by a factor of 0.817/0.806 and P_0 by a factor of 1.0/1.05. The results are statistically identical if we compute the correction assuming an additive bias. This procedure yields our final constraints on σ_8 and P_0:

$$\begin{aligned} \sigma_8 &= 0.748^{+0.024}_{-0.020} \\ P_0 &= 1.11 \pm 0.14\,. \end{aligned} \tag{5.15}$$

5.7 Discussion and Outlook

In this paper we have presented the first application of the one-point PDF to statistical measurements of the tSZ effect. After introducing a theoretical framework in which to compute the noisy tSZ PDF and its covariance matrix, we measure the signal in the ACT 148 GHz Equatorial data and use it to constrain σ_8 and P_0. Directly comparing the fixed-P_0 result to our earlier tSZ skewness analysis [19] indicates that the PDF decreases the error bar on σ_8 by a factor of ≈ 2. Furthermore, we demonstrate that the tSZ PDF can break the degeneracy between σ_8 and P_0 using tSZ data alone, although the ACT data can only

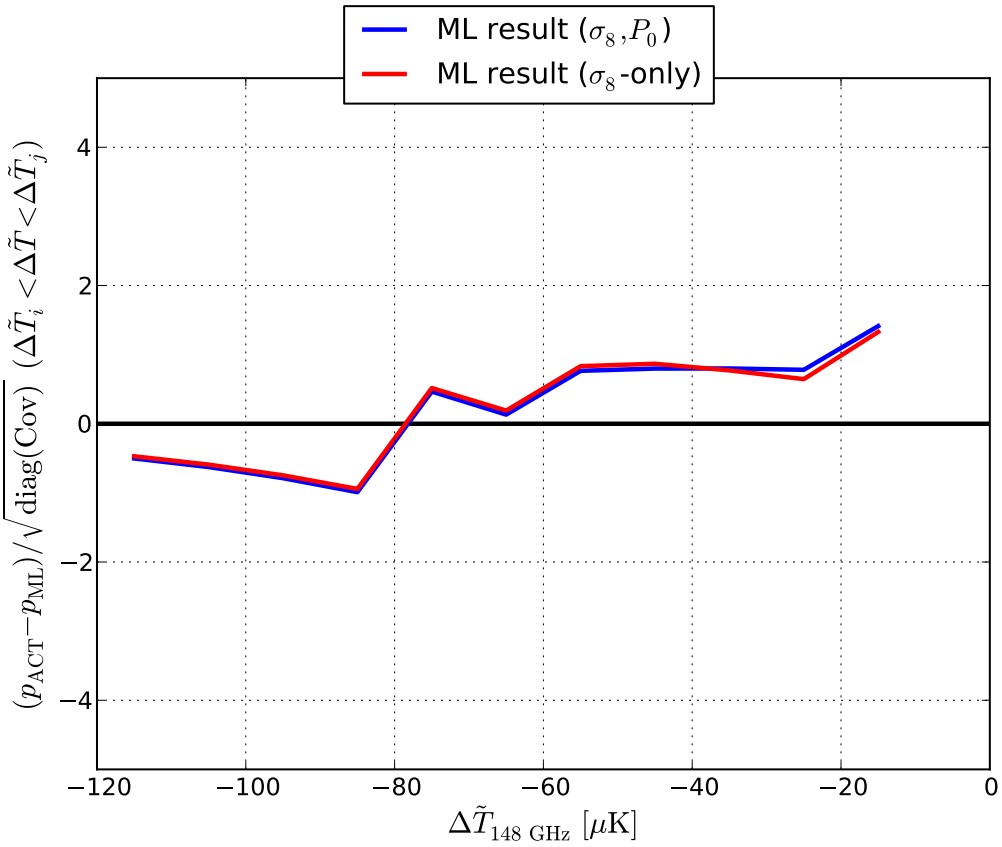

Figure 5.16: Difference between the measured ACT PDF and the ML model (for both likelihood scenarios presented in Fig. 5.15) divided by the square root of the diagonal elements of the covariance matrix of the ML model. Note that this plot obscures the significant bin-to-bin correlations in the PDF (see Fig. 5.8), and thus extreme caution must be used in any interpretation. However, the plot clearly indicates that the empty outermost bins in the tail of the ACT PDF are responsible for pulling the best-fit amplitude of the ML models down from the value that would likely be preferred by the populated bins.

do so weakly. Extant data from SPT and Planck may allow a stronger breaking of this degeneracy. A key requirement for the use of the PDF statistic is a detailed understanding of the noise properties in the map, preferably in such a way as to allow efficient simulation of many noise realizations to understand its inhomogeneity and spatial correlations. The increased sensitivity of ACTPol and SPTPol should allow for improved constraints using the tSZ PDF in the near future, although Planck may have a sample variance advantage given its full-sky coverage. We leave a detailed study of the trade-off in the PDF between sensitivity/resolution and sky coverage for future work.

There are a number of clear extensions to this approach that can be implemented in the near future. Of particular interest is the 2D joint tSZ–weak lensing PDF, $p(y, \kappa)$, which should strongly break the degeneracy between cosmological and ICM parameters present in tSZ statistics. Effectively, the weak lensing information provides information on the mass scale contributing to a given Compton-y signal. This allows simultaneous determination of the pressure-mass relation and cosmological parameters at a level closer to the full forecasted precision of the tSZ statistics. Another extension of interest is the measurement of the tSZ PDF after applying different multipole-space filters, as opposed to the single-filter approach used in this work. By measuring the PDF as a function of filter scale, it should be possible to recover much of the angular information lost by using only the zero-lag moments of the tSZ field rather than the full polyspectra. This angular information then provides constraints on the structure of the ICM pressure profile. We leave these considerations for future work, with an eye to higher-SNR measurements of the tSZ PDF very soon.

5.8 Acknowledgments

This work was supported by the U.S. NSF through awards AST-0408698, PHY-0355328, AST-0707731 and PIRE-0507768, as well as by Princeton Univ., the Univ. of Pennsylvania, FONDAP, Basal, Centre AIUC, RCUK Fellowship (JD), NASA grant NNX08AH30G (SD,

AH, TM), NSERC PGSD (ADH), NSF AST-0546035 and AST-0807790 (AK), NSF PFC grant PHY-0114422 (ES), KICP Fellowship (ES), SLAC no. DE-AC3-76SF00515 (NS), ERC grant 259505 (JD), BCCP (SD), the NSF GRFP (BDS, BLS), and NASA Theory Grant NNX12AG72G and NSF AST-1311756 (JCH, DNS). We thank B. Berger, R. Escribano, T. Evans, D. Faber, P. Gallardo, A. Gomez, M. Gordon, D. Holtz, M. McLaren, W. Page, R. Plimpton, D. Sanchez, O. Stryzak, M. Uehara, and Astro-Norte for assistance with ACT. ACT operates in the Parque Astronómico Atacama in northern Chile under the auspices of Programa de Astronomía, a program of the Comisión Nacional de Investigación Científica y Tecnológica de Chile (CONICYT).

Bibliography

[1] Fowler, J. W., Niemack, M. D., Dicker, S. R., et al. 2007, Applied Optics, 46, 3444

[2] Swetz, D. S., Ade, P. A. R., Amiri, M., et al. 2011, ApJS, 194, 41

[3] Dünner, R., Hasselfield, M., Marriage, T. A., et al. 2013, ApJ, 762, 10

[4] Niemack, M. D., Ade, P. A. R., Aguirre, J., et al. 2010, Proc. SPIE, 7741

[5] Carlstrom, J. E., Ade, P. A. R., Aird, K. A., et al. 2011, PASP, 123, 568

[6] Schaffer, K. K., Crawford, T. M., Aird, K. A., et al. 2011, ApJ, 743, 90

[7] Austermann, J. E., Aird, K. A., Beall, J. A., et al. 2012, Proc. SPIE, 8452

[8] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5062

[9] Zel'dovich, Y. B., & Sunyaev, R. A. 1969, Ap&SS, 4, 301

[10] Sunyaev, R. A., & Zeldovich, Y. B. 1970, Ap&SS, 7, 3

[11] Hand, N., Addison, G. E., Aubourg, E., et al. 2012, Physical Review Letters, 109, 041101

[12] Hasselfield, M., Hilton, M., Marriage, T. A., et al. 2013, J. Cosmology Astropart. Phys., 7, 8

[13] Gralla, M. B., Crichton, D., Marriage, T. A., et al. 2013, arXiv:1310.8281

[14] Reichardt, C. L., Stalder, B., Bleem, L. E., et al. 2013, ApJ, 763, 127

[15] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5080

[16] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, A&A, 557, A52

[17] Sievers, J. L., Hlozek, R. A., Nolta, M. R., et al. 2013, J. Cosmology Astropart. Phys., 10, 60

[18] Story, K. T., Reichardt, C. L., Hou, Z., et al. 2013, ApJ, 779, 86

[19] Wilson, M. J., Sherwin, B. D., Hill, J. C., et al. 2012, Phys. Rev. D, 86, 122005

[20] Crawford, T. M., Schaffer, K. K., Bhattacharya, S., et al. 2014, ApJ, 784, 143

[21] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, arXiv:1303.5081

[22] Das, S., Louis, T., Nolta, M. R., et al. 2014, J. Cosmology Astropart. Phys., 4, 14

[23] Dunkley, J., Calabrese, E., Sievers, J., et al. 2013, J. Cosmology Astropart. Phys., 7, 25

[24] Menanteau, F., Sifón, C., Barrientos, L. F., et al. 2013, ApJ, 765, 67

[25] Hajian, A., Acquaviva, V., Ade, P. A. R., et al. 2011, ApJ, 740, 86

[26] Marriage, T. A., Baptiste Juin, J., Lin, Y.-T., et al. 2011, ApJ, 731, 100

[27] Nozawa, S., Itoh, N., Suda, Y., & Ohhata, Y. 2006, Nuovo Cimento B Serie, 121, 487

[28] Duffy, A. R., Schaye, J., Kay, S. T., & Dalla Vecchia, C. 2008, Mon. Not. R. Astron. Soc, 390, L64

[29] Battaglia, N., Bond, J. R., Pfrommer, C., & Sievers, J. L. 2012, ApJ, 758, 75

[30] Hill, J. C., & Pajer, E. 2013, Phys. Rev. D, 88, 063526

[31] Hill, J. C., & Spergel, D. N. 2014, J. Cosmology Astropart. Phys., 2, 30

[32] Van Waerbeke, L., Hinshaw, G., & Murray, N. 2014, Phys. Rev. D, 89, 023508

[33] Battaglia, N., Bond, J. R., Pfrommer, C., Sievers, J. L., & Sijacki, D. 2010, ApJ, 725, 91

[34] Arnaud, M., Pratt, G. W., Piffaretti, R., et al. 2010, A & A, 517, A92

[35] Sun, M., Sehgal, N., Voit, G. M., et al. 2011, ApJ, 727, L49

[36] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, A&A, 550, A131

[37] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, A&A, 554, A140

[38] Hajian, A., Battaglia, N., Spergel, D. N., et al. 2013, J. Cosmology Astropart. Phys., 11, 64

[39] Cooray, A. 2000, Phys. Rev. D, 62, 103506

[40] Hill, J. C., & Sherwin, B. D. 2013, Phys. Rev. D, 87, 023527

[41] Tinker, J., Kravtsov, A. V., Klypin, A., et al. 2008, ApJ, 688, 709

[42] Shaw, L. D., Nagai, D., Bhattacharya, S., & Lau, E. T. 2010, ApJ, 725, 1452

[43] Trac, H., Bode, P., & Ostriker, J. P. 2011, ApJ, 727, 94

[44] Hasselfield, M., Moodley, K., Bond, J. R., et al. 2013, ApJS, 209, 17

[45] Sehgal, N., Bode, P., Das, S., et al. 2010, ApJ, 709, 920

[46] Kesden, M., Cooray, A., & Kamionkowski, M. 2002, Phys. Rev. D, 66, 083007

[47] Holder, G. P., McCarthy, I. G., & Babul, A. 2007, Mon. Not. R. Astron. Soc, 382, 1697

[48] Bhattacharya, S., Nagai, D., Shaw, L., Crawford, T., & Holder, G. P. 2012, ApJ, 760, 5

[49] Komatsu, E., & Seljak, U. 2002, Mon. Not. R. Astron. Soc, 336, 1256

[50] Battaglia, N., Bond, J. R., Pfrommer, C., & Sievers, J. L. 2012, ApJ, 758, 74

[51] Hu, W., & Kravtsov, A. V. 2003, ApJ, 584, 702

[52] Takada, M., & Bridle, S. 2007, New Journal of Physics, 9, 446

[53] Takada, M., & Spergel, D. N. 2013, arXiv:1307.4399

[54] Zhang, P., & Sheth, R. K. 2007, ApJ, 671, 14

[55] Limber, D. N. 1954, ApJ, 119, 655

[56] Tinker, J. L., Robertson, B. E., Kravtsov, A. V., et al. 2010, ApJ, 724, 878

[57] Eifler, T., Schneider, P., & Hartlap, J. 2009, A&A, 502, 721

[58] Das, S., Marriage, T. A., Ade, P. A. R., et al. 2011, ApJ, 729, 62

[59] Das, S., Sherwin, B. D., Aguirre, P., et al. 2011, Physical Review Letters, 107, 021301

[60] Shaw, L. D., Zahn, O., Holder, G. P., & Doré, O. 2009, ApJ, 702, 368

[61] Huffenberger, K. M., & Seljak, U. 2005, New A, 10, 491

[62] Jain, B., & Van Waerbeke, L. 2000, ApJ, 530, L1

[63] Kratochvil, J. M., Haiman, Z., & May, M. 2010, Phys. Rev. D, 81, 043519

[64] Scheuer, P. A. G. 1957, Proceedings of the Cambridge Philosophical Society, 53, 764

[65] Condon, J. J. 1974, ApJ, 188, 279

[66] Bond, J. R., & Myers, S. T. 1996, ApJS, 103, 63

[67] Springel, V., White, M., & Hernquist, L. 2001, ApJ, 549, 681

[68] Springel, V., White, M., & Hernquist, L. 2001, ApJ, 562, 1086

[69] Refregier, A., & Teyssier, R. 2002, Phys. Rev. D, 66, 043002

[70] Seljak, U., Burwell, J., & Pen, U.-L. 2001, Phys. Rev. D, 63, 063001

[71] Zhang, P., Pen, U.-L., & Wang, B. 2002, ApJ, 577, 555

[72] Komatsu, E., & Kitayama, T. 1999, ApJ, 526, L1

[73] Addison, G. E., Dunkley, J., & Spergel, D. N. 2012, Mon. Not. R. Astron. Soc, 427, 1741

[74] Mesinger, A., McQuinn, M., & Spergel, D. N. 2012, Mon. Not. R. Astron. Soc, 422, 1403

[75] da Silva, A. C., Barbosa, D., Liddle, A. R., & Thomas, P. A. 2000, Mon. Not. R. Astron. Soc, 317, 37

[76] Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, ApJS, 208, 19

Appendix A

The Atacama Cosmology Telescope: A Measurement of the Thermal Sunyaev-Zel'dovich Effect Using the Skewness of the CMB Temperature Distribution

A.1 Abstract

We present a detection of the unnormalized skewness $\langle \tilde{T}^3(\hat{\mathbf{n}}) \rangle$ induced by the thermal Sunyaev-Zel'dovich (tSZ) effect in filtered Atacama Cosmology Telescope (ACT) 148 GHz cosmic microwave background temperature maps. Contamination due to infrared and radio sources is minimized by template subtraction of resolved sources and by constructing a mask using outlying values in the 218 GHz (tSZ-null) ACT maps. We measure $\langle \tilde{T}^3(\hat{\mathbf{n}}) \rangle = -31 \pm 6\ \mu\mathrm{K}^3$ (Gaussian statistics assumed) or $\pm 14\ \mu\mathrm{K}^3$ (including non-Gaussian corrections) in the filtered ACT data, a 5σ detection. We show that the skewness is a sen-

sitive probe of σ_8, and use analytic calculations and tSZ simulations to obtain cosmological constraints from this measurement. From this signal alone we infer a value of $\sigma_8 = 0.79^{+0.03}_{-0.03}$ (68% C.L.) $^{+0.06}_{-0.06}$ (95% C.L.). Our results demonstrate that measurements of non-Gaussianity can be a useful method for characterizing the tSZ effect and extracting the underlying cosmological information.

A.2 Introduction

Current observations of the cosmic microwave background (CMB) anisotropies on arcminute scales using experiments such as the Atacama Cosmology Telescope [ACT; 1, 3, 2] and the South Pole Telescope [SPT; 4, 5] probe not only the primordial microwave background fluctuations sourced 13.7 billion years ago, but also measure secondary anisotropies caused by more recent and less distant physical processes. Such secondary anisotropies are induced by infrared (IR) dusty galaxies and radio sources, gravitational lensing, and the Sunyaev-Zel'dovich (SZ) effect. The SZ effect [6, 7] arises due to the inverse Compton scattering of CMB photons off high energy electrons located predominantly in hot gas in galaxy clusters (the intra-cluster medium, or ICM). This scattering modifies the spectrum of CMB photons in the direction of a cluster in a way that depends on both the thermal energy contained in the ICM (the thermal SZ effect) as well as the peculiar velocity of the cluster with respect to the CMB rest frame (the kinetic SZ effect). The kinetic SZ effect simply increases or decreases the amplitude of the CMB spectrum in the direction of a cluster, but the thermal SZ (tSZ) effect modifies the CMB spectrum in a frequency-dependent manner. The tSZ effect is characterized by a decrease (increase) in the observed CMB temperature at frequencies below (above) 218 GHz in the direction of a galaxy cluster due to inverse Compton scattering. The thermal effect is generally at least an order of magnitude larger than the kinetic effect for a typical massive cluster at 148 GHz. Measurements of the tSZ signal, which is proportional to the integrated ICM pressure along the line of sight, can be used to observe the high

redshift universe, constrain cosmological parameters – in particular σ_8, the variance of matter fluctuations on scales of 8 Mpc/h – and probe baryonic physics in the ICM.

The tSZ signal has so far primarily been studied either by directly resolving individual clusters in arcminute-scale CMB maps [8, 9, 10, 11, 12] or by measuring it statistically through its presence in the small-scale CMB power spectrum [13, 14]. However, it is by no means obvious that the power spectrum is the best way to characterize the statistical properties of the tSZ field. Indeed, measuring the tSZ signal in the power spectrum is challenging because there are many other sources of CMB power on arcminute scales: primordial CMB fluctuations, CMB lensing, instrumental noise, dusty star-forming IR galaxies, and radio sources. In order to disentangle these contributions to the power spectrum and isolate the amplitude of the tSZ signal, a sophisticated multifrequency analysis is required, which involves modeling the power spectrum contribution of each of these components in at least two frequency bands.

In this appendix we instead measure the tSZ signal using the unnormalized skewness of the filtered temperature fluctuation $\left\langle \tilde{T}^3(\hat{\mathbf{n}}) \right\rangle$. This quantity has the significant advantage that, unlike measurements of the tSZ effect through the power spectrum, its measurement does not require the subtraction of Gaussian contributions, because it is only sensitive to non-Gaussian signals with non-zero skewness. The primordial CMB (which is assumed to be Gaussian on these scales) and instrumental noise (which is Gaussian) hence do not contribute to it. In addition, CMB lensing and the kinetic SZ effect do not induce skewness (as they are equally likely to produce positive and negative fluctuations), and so do not contribute either. The primary contributions to this quantity are thus only the tSZ effect and point sources. These signals have a different frequency dependence. Furthermore, the tSZ signal contributes negative skewness, whereas radio and IR point sources contribute positive skewness. These characteristics allow the tSZ signal to be effectively isolated and studied, as first pointed out in [15].

Measurements of the skewness also possess significant advantages from an astrophysical perspective. A consistent problem plaguing studies of the tSZ power spectrum has been theoretical uncertainty in the ICM electron pressure profile [16, 17, 18], especially in the low-mass, high-redshift groups and clusters that contribute much of the signal. As discussed in the following section in detail, the tSZ skewness signal is dominated by characteristically higher-mass, lower-redshift clusters than those that source the power spectrum signal. The ICM astrophysics for these objects is better constrained by X-ray observations and they are less sensitive to energy input from non-gravitational sources [18, 19]. Thus, the theoretical systematic uncertainty in modeling the tSZ skewness is correspondingly lower as well. In addition, at 148 GHz, dusty star-forming galaxies are less prevalent in massive, low-redshift clusters (which contribute more to the skewness) than in high-redshift groups and clusters (which contribute more to the tSZ power spectrum) [20]. Thus, we expect the correlation between tSZ signal and dusty galaxy emission, which can complicate analyses of the tSZ effect, to be smaller for a measurement of the skewness.

Moreover, the tSZ skewness scales with a higher power of σ_8 than the tSZ power spectrum amplitude. This result is precisely what one would expect if the signal were dominated by higher-mass, rarer objects, as the high-mass tail of the mass function is particularly sensitive to a change in σ_8. This provides the prospect of tight constraints on cosmological parameters from the skewness that are competitive with constraints from the power spectrum.

In this appendix, we first explain the usefulness of the skewness as a cosmological probe by theoretically deriving its scaling with σ_8 as well as the characteristic masses of the objects sourcing the signal. Subsequent sections of the appendix describe how we measured this skewness in the ACT data. We describe how the ACT temperature maps are processed in order to make a reliable measurement of the unnormalized skewness due to the tSZ effect, and discuss how contamination from IR dusty galaxies and radio point sources is minimized. We report the measurement results and discuss how the errors are calculated. Finally, we discuss the cosmological constraints and associated uncertainties derived from this measurement.

We assume a flat ΛCDM cosmology throughout, with parameters set to their WMAP5 values [21] unless otherwise specified. All masses are quoted in units of $M_\odot/h$, where $h \equiv H_0/(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}})$ and H_0 is the Hubble parameter today.

A.3 Skewness of the tSZ Effect

In this section, we investigate the Nth moments of the pixel probability density function, $\langle T^N \rangle \equiv \langle T(\hat{\mathbf{n}})^N \rangle$, focusing on the specific case of the unnormalized skewness $\langle T^3 \rangle$. We show that the unnormalized skewness $\langle T^3 \rangle$ has a steeper scaling with σ_8 than the power spectrum amplitude and is dominated by characteristically higher-mass, lower-redshift clusters, for which the ICM astrophysics is better constrained and modeled. As explained earlier, these characteristics make tSZ skewness measurements a useful cosmological probe.

In order to calculate the Nth moment of the tSZ field, we assume the distribution of clusters on the sky can be adequately described by a Poisson distribution (and that contributions due to clustering and overlapping sources are negligible [22]). The Nth moment is then given by

$$\langle T^N \rangle = \int dz \frac{dV}{dz} \int dM \frac{dn(M,z)}{dM} \int d^2\boldsymbol{\theta}\, T(\boldsymbol{\theta}; M, z)^N \,, \tag{A.1}$$

where $T(\boldsymbol{\theta}; M, z)$ is the tSZ temperature decrement at position $\boldsymbol{\theta}$ on the sky with respect to the center of a cluster of mass M at redshift z:

$$\begin{aligned} T(\boldsymbol{\theta}; M, z) &= g(\nu) T_{\mathrm{CMB}} \frac{\sigma_T}{m_e c^2} \\ &\times \int P_e\left(\sqrt{l^2 + d_A^2(z)|\boldsymbol{\theta}|^2}; M, z\right) dl \,, \end{aligned} \tag{A.2}$$

where $g(\nu)$ is the spectral function of the tSZ effect, $d_A(z)$ is the angular diameter distance to redshift z, and the integral is taken over the electron pressure profile $P_e(\boldsymbol{r}; M, z)$ along the line of sight.

For a given cosmology, Eqs. (A.1) and (A.2) show that there are two ingredients needed to calculate the $N^{\rm th}$ tSZ moment (in addition to the comoving volume per steradian dV/dz, which can be calculated easily): (1) the halo mass function $dn(M,z)/dM$ and (2) the electron pressure profile $P_e(\boldsymbol{r};M,z)$ for halos of mass M at redshift z. We use the halo mass function of [23] with the redshift-dependent parameters given in their Eqs. (5)–(8). While uncertainties in tSZ calculations due to the mass function are often neglected, they may be more important for the skewness than the power spectrum, as the skewness is more sensitive to the high-mass exponential tail of the mass function. We estimate the uncertainty arising from the mass function by performing alternate calculations with the mass function of [24], which predicts more massive clusters at low redshift than [23] for the same cosmology. As an example, the predicted skewness calculated using the pressure profile of [16] with the mass function of [23] is $\approx 35\%$ lower than the equivalent result using the mass function of [24]. However, the derived scalings of the variance and skewness with σ_8 computed using [24] are identical to those found below using [23]. Thus, the scalings calculated below are robust to uncertainties in the mass function, and we use them later to interpret our skewness measurement. However, we rely on cosmological simulations to obtain predicted values of the tSZ skewness. We do not consider alternate mass functions any further in our analytic calculations.

We consider three different pressure profiles from [16, 17, 25] in order to evaluate the theoretical uncertainty in the scaling of the tSZ skewness with σ_8. These profiles differ in how they are derived and in the ICM physics they assume. They thus provide a measure of the scatter in the scalings of the variance and skewness with σ_8 due to uncertainties in the gas physics.

Finally, in order to make a faithful comparison between the theory and data, we convolve Eq. (A.2) with the Fourier-space filter described in the subsequent data analysis sections of this appendix. In addition, we account for the 12σ pixel fluctuation cutoff used in the data analysis (see below) by placing each "cluster" of mass M and redshift z in the integrals of Eq. (A.1) in an idealized ACT pixel and computing the observed temperature decrement,

accounting carefully for geometric effects that can arise depending on the alignment of the cluster and pixel centers. If the calculated temperature decrement exceeds the 12σ cutoff, then we do not include this cluster in the integrals. These steps cannot be neglected, as the filter and cutoff reduce the predicted tSZ skewness amplitude by up to 95% compared to the pure theoretical value [26]. Most of this reduction is due to the filter, which modestly suppresses the temperature decrement profile of a typical cluster; this suppression strongly affects the skewness because it is a cubic statistic.

The analytic theory described above determines the scaling of the $N^{\rm th}$ tSZ moment with σ_8. In particular, we compute Eq. (A.1) with $N = 2$, $N = 3$, and $N = 6$ for each of the chosen pressure profiles while varying σ_8. The scalings of the variance ($N = 2$), the skewness ($N = 3$), and the sixth moment ($N = 6$, which we require for error calculation) with σ_8 are well-described by power laws for each of these profiles: $\left\langle \tilde{T}^{2,3,6} \right\rangle \propto \sigma_8^{\alpha_{2,3,6}}$. For the profile of [16], we find $\alpha_2 = 7.8$, $\alpha_3 = 11.1$, and $\alpha_6 = 16.7$; for the profile of [17], we find $\alpha_2 = 8.0$, $\alpha_3 = 11.2$, and $\alpha_6 = 15.9$; and for the profile of [25], we find $\alpha_2 = 7.6$, $\alpha_3 = 10.7$, and $\alpha_6 = 18.0$. Note that the scaling of the variance matches the scaling of the tSZ power spectrum amplitude that has been found by a number of other studies, as expected (e.g., [25, 27]). The scaling of the unnormalized skewness is similar to that found by [28], who obtained $\alpha_3 = 10.25$. Also, note that the skewness scaling is modified slightly from its pure theoretical value [26] due to the Fourier-space filter and pixel fluctuation cutoff mentioned above. The overall conclusion is that the skewness scales with a higher power of σ_8 than the variance (or power spectrum). We use this scaling to derive a constraint on σ_8 from our measurement of the skewness below.

In addition, we compare the characteristic mass scale responsible for the tSZ skewness and tSZ variance (or power spectrum) signals. Analytic calculations show that the tSZ power spectrum amplitude typically receives $\approx 50\%$ of its value from halos with $M < 2$–3×10^{14} $\mathrm{M}_\odot/h$, while the tSZ skewness receives only $\approx 20\%$ of its amplitude from these less massive objects. This indicates that the clusters responsible for the tSZ skewness signal are

better theoretically modeled than those responsible for much of the tSZ power spectrum, both because massive clusters have been observed more thoroughly, and because more massive clusters are dominated by gravitational heating and are less sensitive to non-linear energy input from active galactic nuclei, turbulence, and other mechanisms [18, 19]. We verify this claim when interpreting the skewness measurement below, finding that the systematic theoretical uncertainty (as derived from simulations) is slightly smaller than the statistical error from the measurement, though still non-negligible.

A.4 Map Processing

A.4.1 Filtering the Maps

The Atacama Cosmology Telescope [1, 3, 2] is a 6m telescope in the Atacama Desert of Chile, which operated at 148, 218, and 277 GHz using three 1024-element arrays of superconducting bolometers. The maps used in this analysis were made over three years of observation in the equatorial region during 2008–2010 at 148 GHz, and consist of six $3^\circ \times 18^\circ$ patches of sky at a noise level of ≈ 21 μK arcmin. In our source mask construction we also use maps of the same area made in 2008 at a frequency of 218 GHz. The maps were calibrated as in [29]. We apodize the maps by multiplying them with a mask that smoothly increases from zero to unity over 0.1° from the edge of the maps.

Although atmospheric noise is removed in the map-making process, we implement an additional filter in Fourier space to remove signal at multipoles below $\ell = 500$ (ℓ is the magnitude of the Fourier variable conjugate to sky angle). In addition, we remove a stripe for which $-100 < \ell_{\rm dec} < 100$ along the Fourier axis corresponding to declination to avoid contamination by scan noise. Furthermore, to increase the tSZ signal-to-noise, we apply a Wiener filter which downweights scales at which the tSZ signal is subdominant. This (non-optimal) filter is constructed by dividing the best-fit tSZ power spectrum from [13] by the total average power spectrum measured in the data maps, i.e. $C_\ell^{tSZ}/C_\ell^{tot}$. For multipoles

above $\ell = 6 \times 10^3$, the tSZ signal is completely dominated by detector noise and point sources, and hence we remove all power above this multipole in the temperature maps. The final Fourier-space filter, shown in Fig. 1, is normalized such that its maximum value is unity. As it is constructed using the binned power spectrum of the data, it is not perfectly smooth; however, we apply the same filter consistently to data, simulations, and analytic theory, and thus any details of the filter do not bias the interpretation of our result. After filtering, the edges of the maps are cut off to reduce any edge effects that might occur upon Fourier transforming despite apodization. Simulations verify that no additional skewness is introduced by edge effects into a trimmed map.

A.4.2 Removing Point Sources

In order to obtain a skewness signal due only to the tSZ effect, any contamination of the signal by point sources must be minimized. These objects consist of IR dusty galaxies and radio sources. We use two approaches to eliminate the point source contribution: template subtraction and masking using the 218 GHz channel.

In the template subtraction method [30], which we use to remove resolved point sources (mainly bright radio sources), sources with a signal-to-noise (S/N) greater than five are first identified in a match-filtered map. A template with the shape of the ACT beam is then scaled to the appropriate peak brightness of each source, and this profile is subtracted from the raw data. The process is iterated following the CLEAN algorithm [31] until no more sources can be identified. We verify that this procedure does not introduce skewness into the maps (e.g., through oversubtraction) by checking that similar results are obtained using a different procedure for reducing source contamination, in which we mask and in-paint pixels which contain bright sources with S/N > 5 [32].

We take a second step to suppress the lower-flux, unresolved point sources (mainly dusty galaxies) that remain undetected by the template subtraction algorithm. At 218 GHz, dusty galaxies are significantly brighter than at 148 GHz and the tSZ effect is negligible. We

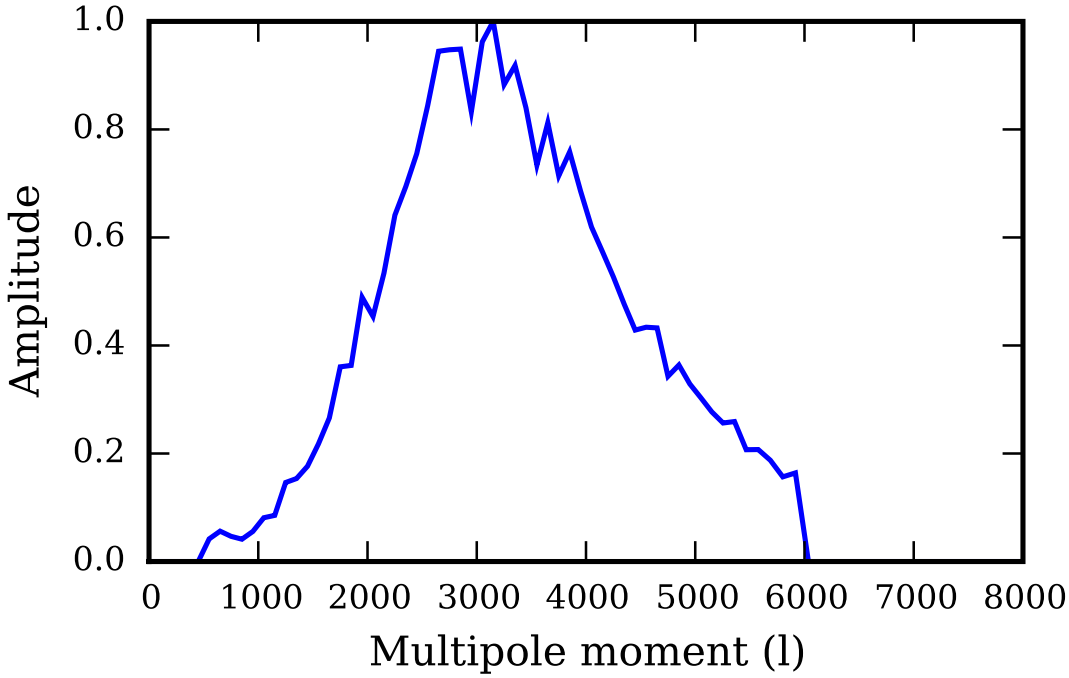

Figure A.1: The Wiener filter applied to the ACT temperature maps before calculating the unnormalized skewness. This filter upweights scales on which the tSZ signal is large compared to other sources of anisotropy.

construct a dusty galaxy mask by setting all pixels (which are approximately 0.25 arcmin^2) to zero that have a temperature in the 218 GHz maps larger than a specified cutoff value. This cutoff is chosen to be 3.2 times the standard deviation of the pixel values in the filtered 148 GHz map (3.2σ). This procedure ensures regions with high flux from dusty galaxies are masked. We also set to zero all pixels for which the temperature is lower than the negative of this cutoff, so that the masking procedure does not introduce spurious skewness into the lensed CMB distribution, which is assumed to have zero intrinsic skewness. The mask is then applied to the 148 GHz map to reduce the point source contribution. Simulations ([33] for IR sources) verify that the masking procedure does not introduce spurious skewness into the 148 GHz maps.

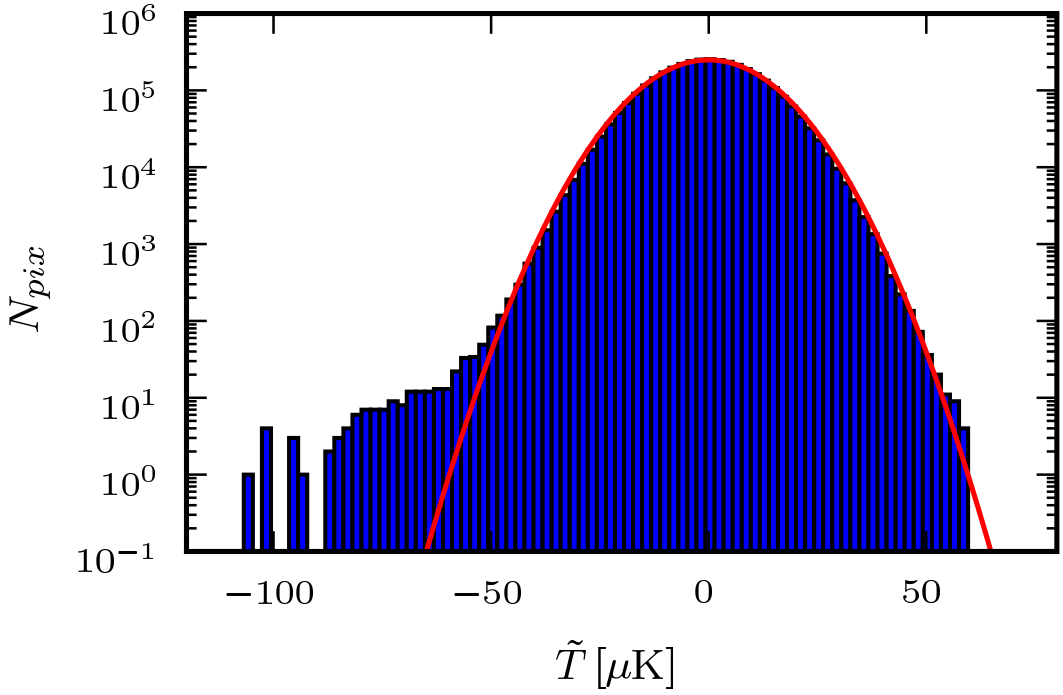

Figure A.2: Histogram of the pixel temperature values in the filtered, masked ACT CMB temperature maps. A Gaussian curve is overlaid in red.

Finally, all pixels more than twelve standard deviations (12σ) from the mean are also removed from the 148 GHz maps. Due to the ringing around very positive or negative pixels caused by the Wiener filter, the surrounding eight arcminutes of these points are also masked. This additional step slightly increases the S/N of the skewness measurement by reducing the dependence on large outliers, enhances the information content of low moments by truncating the tail of the pixel probability density function, and ensures that any anomalous outlying points from possibly mis-subtracted bright radio sources do not contribute to the skewness signal. Overall, 14.5% of the 148 GHz map is removed by the masking procedure, though the removed points are random with respect to the tSZ field and should not change the signal.

A.5 Results

A.5.1 Evaluating the Skewness

We compute the unnormalized skewness of the filtered and processed 148 GHz maps by simply cubing and averaging the pixel values in real space. The result is $\langle \tilde{T}^3 \rangle = -31 \pm 6\ \mu\mathrm{K}^3$, a 5σ deviation from the null result expected for a signal without any non-Gaussian components. The skewness of the CMB temperature distribution in our filtered, processed maps is visible in the pixel value histogram shown in Fig. 2 (along with a Gaussian curve overlaid for comparison). It is evident that the Gaussian CMB has been recovered on the positive side by point source masking, with the apparent truncation beyond 50 μK due to the minute probability of such temperatures in the Gaussian distribution. The likelihood corresponding to our measurement of the skewness is shown in Fig. 3.

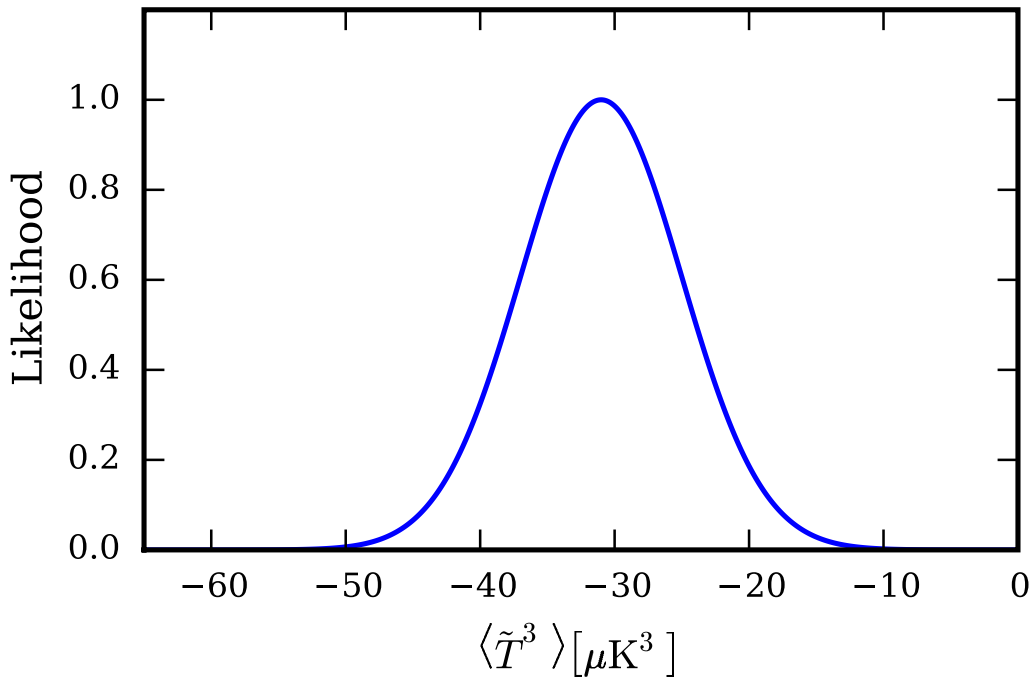

Figure A.3: Likelihood of the skewness measurement described in the text (with Gaussian statistics assumed).

The "Gaussian statistics assumed" error on the skewness includes only Gaussian sources of noise. We calculate this error by using map simulations that consist of Gaussian random fields with the same power spectrum as that observed in the data, including beam effects. These simulations contain Gaussian contributions from IR, SZ, and radio sources, the primordial lensed CMB, and detector noise. This estimate thus does not include the error resulting from non-Gaussian corrections, which (after source subtraction) are due to the non-Gaussian tSZ signal. Though CMB lensing is also a non-Gaussian effect it does not contribute to the error on the skewness, as the connected part of the six-point function is zero to lowest order in the lensing potential, and the connected part of the three-point function is also negligibly small (see [34] and references therein).

We calculate errors that include non-Gaussian corrections by constructing more realistic simulations. To construct such simulations, we add maps with simulated tSZ signal from [35], which assume $\sigma_8 = 0.8$, to realizations of a Gaussian random field which has a spectrum such that the power spectrum of the combined map matches that observed in the ACT temperature data. Given the simulated sky area, we obtain 39 statistically independent simulated maps, each of size 148 deg^2. By applying an identical procedure to the simulations as to the data, measuring the scatter amongst the patches, and appropriately scaling the error to match the 237 deg^2 of unmasked sky in the processed map, we obtain a full error (including non-Gaussian corrections) on the unnormalized skewness of 14 μK^3. While this error is a robust estimate it should be noted that the "error on the error" is not insignificant due to the moderate simulated volume available. The scatter of skewness values measured from each of the simulated maps is consistent with a Gaussian distribution. The estimate for the full error is coincidentally the same as the standard error, 14 μK^3, estimated from the six patches into which the data are divided. The full error is used below in deriving cosmological constraints from the skewness measurement.

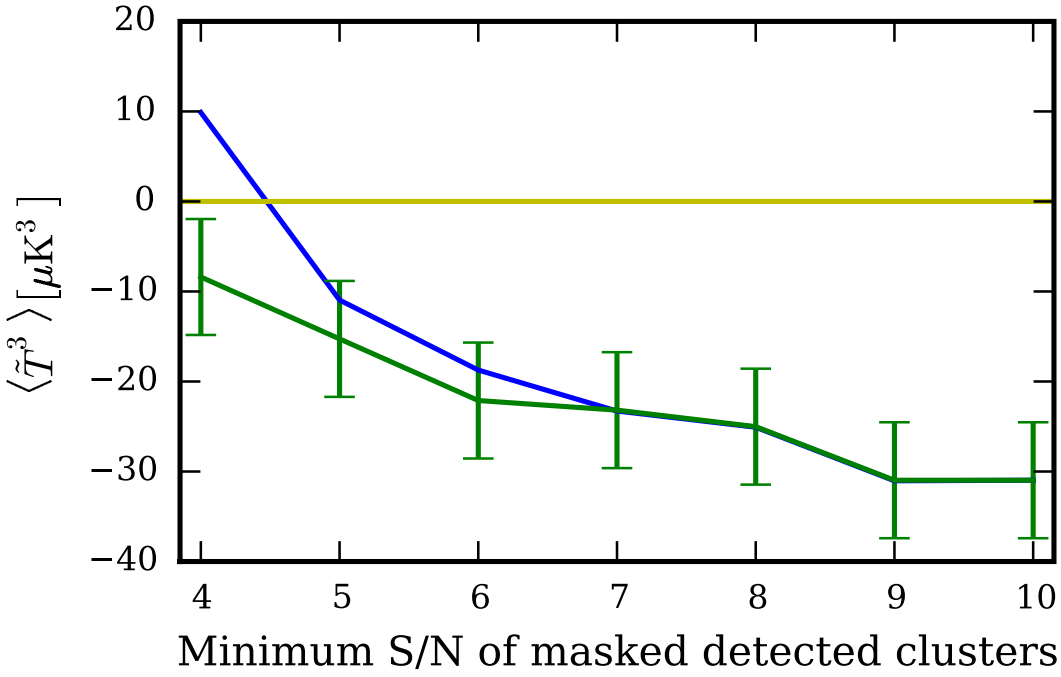

Figure A.4: Plot of the skewness signal as a function of the minimum S/N of the clusters that are masked (this indicates how many known clusters are left in the data, unmasked). The blue line is calculated using the full cluster candidate catalog obtained via matched filtering, while the green line uses a catalog containing only optically-confirmed clusters [38]. Both lines have identical errors, but we only plot them for the green line for clarity. Confirmed clusters source approximately two-thirds of the signal, which provides strong evidence that it is due to the tSZ effect. Note that one expects a positive bias of $\approx 4\ \mu\mathrm{K}^3$ for the S/N = 4 point of the blue line due to impurities in the full candidate catalog masking the tail of the Gaussian distribution.

A.5.2 The Origin of the Signal

Is the skewness dominated by massive clusters with large tSZ decrements – as suggested by theoretical considerations described earlier – or by more numerous, less massive clusters? To investigate this question, we mask clusters in our data which were found in the 148 GHz maps using a matched filter as in [36]. All clusters detected above a threshold significance value are masked; we vary this threshold and measure the remaining skewness in order to determine the origin of the signal.

Fig. 4 shows a plot of the signal against the cluster detection significance cutoff. We include calculations using both the full cluster candidate catalog obtained via matched filtering as well a catalog containing only clusters confirmed optically using the methodology of [37] on the SDSS Stripe 82 [38]. The SDSS Stripe 82 imaging data cover $\approx 80\%$ of the total map area, and thus some skewness signal will necessarily arise from objects not accounted for in this catalog. The results for these two catalogs agree when masking clusters with S/N ≥ 7, but differ slightly when masking lower S/N clusters. This effect is likely due to the small shortfall in optical follow-up area as well as a small number of false detections (i.e., impurity) in the candidate clusters that have not yet been optically followed up.

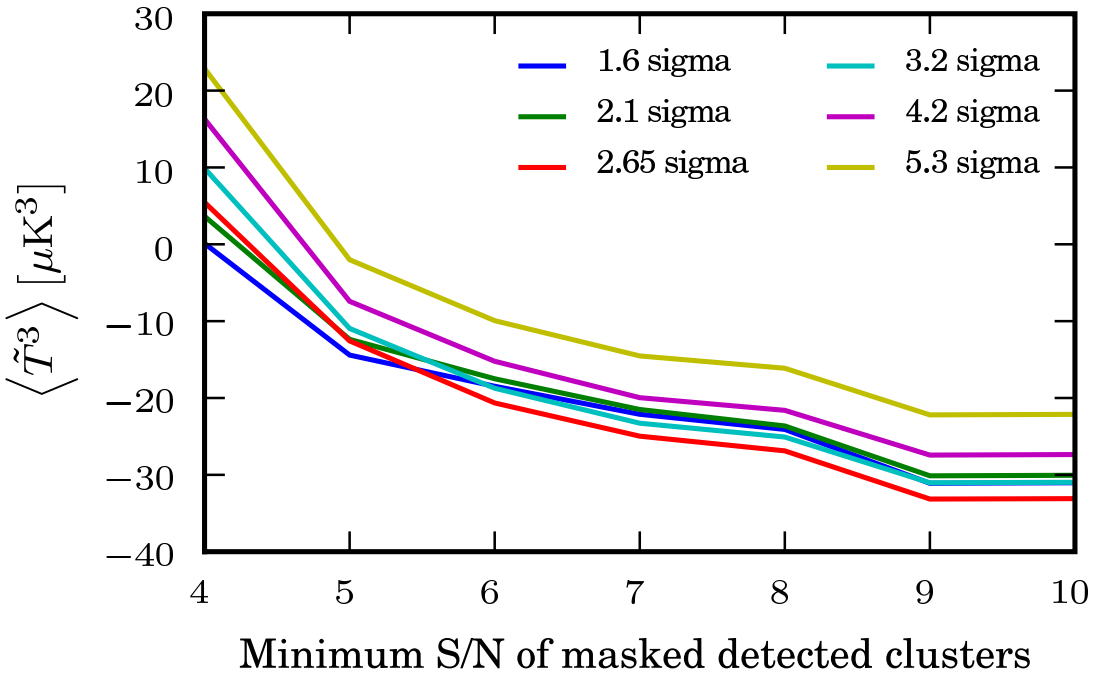

Figure A.5: A test for IR source contamination: similar to the blue line in Fig. 4, but with a range of values of the cutoff used to construct an IR source mask in the 218 GHz band. Any cutoff below $\approx 3.2\sigma$ gives similarly negative results and thus appears sufficient for point source removal, where $\sigma = 10.3$ μK is the standard deviation of the 148 GHz maps. For comparison, the standard deviation of the 218 GHz maps is ≈ 2.2 times larger. The percentages of the map which are removed for the masking levels shown, from the least to the most strict cut, are 0.7, 2.5, 8.4, 14.5, 23.7, and 36.6%.

Using either catalog, Fig. 4 implies that just under half of the tSZ skewness is obtained from clusters that lie below a 5σ cluster detection significance, while the remainder comes from the brightest and most massive clusters. The results of [36] suggest that clusters detected at 5σ significance are roughly characterized by a mass $M_{500} = 5 \times 10^{14} M_{\odot}/h$, where M_{500} is the mass enclosed within a radius such that the mean enclosed density is 500 times the critical density at the cluster redshift. This value corresponds to a virial mass of roughly $M = 9 \times 10^{14} M_{\odot}/h$, which was also found to be the mass detection threshold for high-significance ACT clusters in [37]. Fig. 4 thus demonstrates that roughly half of the tSZ skewness signal is due to massive clusters with $M \gtrsim 10^{15} M_{\odot}/h$. The theoretical calculations described earlier give similar results for the fraction of the signal coming from clusters above and below this mass scale, which is significantly higher than the characteristic mass scale responsible for the tSZ power spectrum signal.

Finally, the positive value in the full candidate catalog line shown in Fig. 4 when masking clusters above S/N = 4 is consistent with zero. When masking at this level (with the candidate catalog which contains some impurities), we slightly cut into the negative pixel values in the Gaussian component of Fig. 2, leading to a small spurious positive skewness. For the points we plot, we calculate that this bias is only non-negligible for the S/N = 4 cut, where it is $\approx 4\ \mu\mathrm{K}^3$. This bias effectively explains the small positive offset seen in Fig. 4. However, we discuss positive skewness due to any possible residual point source contamination below. Overall, the dependence of the measured skewness on cluster masking shown in Fig. 4 provides strong evidence that it is caused by the tSZ effect.

A.5.3 Testing for Systematic Infrared Source Contamination

Despite our efforts to remove point sources, a small residual point source contamination of the signal could remain, leading to an underestimate of the amplitude of the tSZ skewness. To investigate this systematic error source, we vary the level at which point sources are masked in the 218 GHz maps (the original level is 3.2 times the standard deviation of the pixel values

in the filtered 148 GHz map (3.2σ), as described above). The results of this test are shown in Fig. 5, which uses the full catalog of cluster candidates as described in Fig. 4, since the optically-confirmed catalog does not yet cover the entire ACT map. Note that masking at 3.2σ results in a skewness measurement which agrees with its apparent asymptotic limit as the IR contamination is reduced, within the expected fluctuations due to masking. While a slightly more negative skewness value can be measured for some masking levels stricter than the 3.2σ level chosen in the analysis, fluctuations upon changing the unmasked area of the sky are expected, so that it can not be rigorously inferred that IR contamination is reduced between 3.2σ and 2.65σ. Fig. 5 suggests that masking at the 3.2σ level sufficiently removes any contamination by IR sources, and stricter cuts will reduce the map area and increase statistical errors unnecessarily.

However, to further estimate the residual point source contamination in the 148 GHz maps, we process simulations of IR sources from [33] (with source amplitudes scaled down by 1.7 to match recent observations, as in [39]) with the same masking procedure as that applied to the data (described in §A.4.2), creating a mask in a simulated 218 GHz map, and applying it to a simulated IR source signal at 148 GHz. We find a residual signal of $\left\langle \tilde{T}^3 \right\rangle = 3.9 \pm 0.1\ \mu\mathrm{K}^3$. We treat this result as a bias in deriving cosmological constraints from the tSZ skewness in the following section.

We also investigate a linear combination of the 148 and 218 GHz maps that should have minimal IR source levels, namely, an appropriately scaled 218 GHz map subtracted from a 148 GHz map. Assuming that the spatial distribution of the point sources is not affected by the difference in observation frequency between 148 and 218 GHz and a single spectral index can be applied to all sources, a simple factor of ≈ 3.2 [13] relates a point source's signal in the two different frequency bands. We find that the appropriate linear combination (subtracting 1/3.2 times the 220 GHz map from the 148 GHz map) produces a signal in agreement with that resulting from the previously described masking procedure, although

the additional noise present in the 218 GHz maps slightly reduces the significance of the detection.

A.6 Cosmological Interpretation

To obtain cosmological information from the measured amplitude of the unnormalized skewness, we compare our results with two different sets of tSZ simulations [33, 35]. Both sets of simulations are run with $\sigma_8 = 0.8$, but differ in their treatment of the ICM. The simulation of [35] is a fully hydrodynamic cosmological simulation that includes sub-grid prescriptions for feedback from active galactic nuclei, star formation, and radiative cooling. The simulation also captures non-thermal pressure support due to turbulence and other effects, which significantly alters the ICM pressure profile. The simulation of [33] is a large dark matter-only N-body simulation that is post-processed to include gas according to a polytropic prescription. This simulation also accounts for non-thermal pressure support (though with a smaller amount than [35]), and matches the low-redshift X-ray data presented in [17].

We perform the same filtering and masking as that applied to the data in order to analyze the simulation maps. For both simulations, the filtering reduces the signal by $\approx 95\%$ compared to the unfiltered value. For the simulations of [35], we measure $\left\langle \tilde{T}^3 \right\rangle^S = -37\ \mu\mathrm{K}^3$, with negligible errors (the superscript S indicates a simulated value). However, this value is complicated by the fact that these simulations only include halos below $z = 1$. An analytic estimate for the skewness contribution due to halos with $z > 1$ from Eq. (A.1) gives a 6% correction, which yields $\left\langle \tilde{T}^3 \right\rangle^S = -39\ \mu\mathrm{K}^3$. For the simulations of [33], we measure $\left\langle \tilde{T}^3 \right\rangle^S = -50\ \mu\mathrm{K}^3$, with errors also negligible for the purposes of cosmological constraints.

We combine these simulation results with our calculated scalings of the skewness and the sixth moment with σ_8 to construct a likelihood:

$$\mathcal{L}(\sigma_8) = \exp\left(-\frac{\left(\left\langle \tilde{T}^3 \right\rangle^D - \left\langle \tilde{T}^3 \right\rangle^{\mathrm{th}}(\sigma_8) \right)^2}{2\sigma^2_{\mathrm{th}}(\sigma_8)} \right) \tag{A.3}$$

where $\left\langle \tilde{T}^3 \right\rangle^D$ is our measured skewness value and the theoretically expected skewness as a function of σ_8 is given by

$$\left\langle \tilde{T}^3 \right\rangle^{\rm th}(\sigma_8) = \left\langle \tilde{T}^3 \right\rangle^S \left(\frac{\sigma_8}{0.8}\right)^{\alpha_3} . \tag{A.4}$$

The likelihood in Eq. (A.3) explicitly accounts for the fact that $\sigma^2_{\rm th}$, the variance of the skewness, depends on σ_8 — a larger value of σ_8 leads to a larger expected variance in the tSZ skewness signal. In particular, the variance of the tSZ skewness is described by a sixth moment, so it scales as $\sigma_8^{\alpha_6}$. As determined above, the Gaussian and non-Gaussian errors on the skewness are 6 μK^3 and $\sqrt{14^2 - 6^2}$ μK^3 = 12.6 μK^3, respectively. We approximate the dependence of the full error on σ_8 by assuming that only the non-Gaussian component scales with σ_8; while this is not exact, as some of the Gaussian error should also scale with σ_8, small differences in the size or scaling of the Gaussian error component cause negligible changes in our constraints on σ_8.

Finally, although we have argued previously that IR source contamination is essentially negligible, we explicitly correct for the residual bias as calculated in the previous section. Thus, we replace $\left\langle \tilde{T}^3 \right\rangle^D = -31$ μK^3 with $\left\langle \tilde{T}^3 \right\rangle^D_{corr} = -31 - 3.9$ μK^3 in Eq. (A.3). (Note that this bias correction only shifts the central value derived for σ_8 below by roughly one-fifth of the 1σ confidence interval.) Moreover, in order to be as conservative as possible, we also model the effect of residual point sources by including an additional IR contamination error (with the same value as the residual IR source contamination, 3.9 μK^3) in our expression for the variance of the skewness:

$$\sigma^2_{\rm th}(\sigma_8) = 6^2 \ \mu{\rm K}^6 + 12.6^2 \left(\frac{\sigma_8}{0.8}\right)^{\alpha_6} \ \mu{\rm K}^6 + 3.9^2 \ \mu{\rm K}^6. \tag{A.5}$$

Using the likelihood in Eq. (A.3), we obtain confidence intervals and derive a constraint on σ_8. Our likelihood and hence our constraints depend in principle on which simulation we use to calculate $\left\langle \tilde{T}^3 \right\rangle^S$, as well as on the values we choose for α_3 and α_6. Using the simulations of [35] and the scalings determined above for the profile from [16], we find $\sigma_8 = 0.79^{+0.03}_{-0.03}$

(68% C.L.) $^{+0.06}_{-0.06}$ (95% C.L.). In Table I, we compare the constraints on σ_8 obtained from the use of different scalings and simulated skewness values; the constraints are insensitive to both the pressure profile used to derive the scaling laws and the choice of simulation used to compute the skewness.

For comparison, the final release from the *Chandra* Cluster Cosmology Project found $\sigma_8 = 0.803 \pm 0.0105$, assuming $\Omega_m = 0.25$ (there is a strong degeneracy between σ_8 and Ω_m for X-ray cluster measurements that probe the mass function) [40]. Perhaps more directly comparable, recent studies of the tSZ power spectrum have found $\sigma_8 = 0.77 \pm 0.04$ (statistical error only) [13] and $\sigma_8 = 0.807 \pm 0.016$ (statistical error and approximately estimated systematic error due to theoretical uncertainty) [41]. Our results are also comparable to recent constraints using number counts of SZ-detected clusters from ACT and SPT, which found $\sigma_8 = 0.851 \pm 0.115$ (fully marginalizing over uncertainties in the mass-SZ flux scaling relation) [42] and $\sigma_8 = 0.807 \pm 0.027$ (marginalizing over uncertainties in an X-ray-based mass-SZ flux scaling relation) [43], respectively. Although more than half of the tSZ skewness signal that we measure is sourced by detected clusters (i.e., the same objects used in the number counts analyses), our method also utilizes cosmological information from clusters that lie below the individual detection threshold, which gives it additional statistical power. Finally, note that we have fixed all other cosmological parameters in this analysis, as σ_8 is by far the dominant parameter for the tSZ skewness [44]. However, marginalizing over other parameters will slightly increase our errors.

To evaluate the theoretical systematic uncertainty in the amplitude of the filtered skewness due to unknown ICM astrophysics, we test the effect of different gas prescriptions by analyzing simulations from [35] with all forms of feedback, radiative cooling and star formation switched off, leading to an adiabatic ICM gas model. For these adiabatic simulations we find $\left\langle \tilde{T}^3 \right\rangle^S = -56\ \mu\mathrm{K}^3$ (after applying the 6% correction mentioned earlier), which for the skewness we measure in our data would imply $\sigma_8 = 0.77^{+0.02}_{-0.02}$ (68% C.L.)$^{+0.05}_{-0.05}$ (95% C.L.). Turning off feedback and all sub-grid physics is a rather extreme case, so the systematic the-

oretical uncertainty for a typical simulation with some form of feedback should be slightly

	Battaglia $\left\langle \tilde{T}^3 \right\rangle^S$	Sehgal $\left\langle \tilde{T}^3 \right\rangle^S$
Battaglia α_3, α_6	$0.79^{+0.03\ +0.06}_{-0.03\ -0.06}$	$0.77^{+0.03\ +0.05}_{-0.02\ -0.05}$
Arnaud α_3, α_6	$0.79^{+0.03\ +0.06}_{-0.03\ -0.06}$	$0.77^{+0.02\ +0.05}_{-0.02\ -0.05}$
K-S α_3, α_6	$0.79^{+0.03\ +0.07}_{-0.03\ -0.06}$	$0.77^{+0.03\ +0.06}_{-0.03\ -0.05}$

Table A.1: Constraints on σ_8 derived from our skewness measurement using two different simulations and three different scalings of the skewness and its variance with σ_8. The top row lists the simulations used to calculate the expected skewness for $\sigma_8 = 0.8$ [35, 33]; the left column lists the pressure profiles used to calculate the scaling of the skewness and its variance with σ_8 [16, 17, 25]. The errors on σ_8 shown are the 68% and 95% confidence levels.

smaller than the statistical error from the measurement, though still non-negligible. This contrasts with measurements of σ_8 via the tSZ power spectrum, for which the theoretical systematic uncertainty is comparable to or greater than the statistical uncertainty [13, 41]. As highlighted earlier, this difference can be traced to the dependence of the power spectrum amplitude on the ICM astrophysics within low-mass, high-redshift clusters. The skewness, on the other hand, is dominated by more massive, lower-redshift clusters that are less affected by uncertain non-gravitational feedback mechanisms and are more precisely constrained by observations. Nonetheless, as the statistical uncertainty decreases on future measurements of the tSZ skewness, the theoretical systematic error will quickly become comparable, and thus additional study of the ICM electron pressure profile will be very useful.

A.7 Conclusions

As the thermal Sunyaev-Zel'dovich field is highly non-Gaussian, measurements of non-Gaussian signatures such as the skewness can provide cosmological constraints that are competitive with power spectrum measurements. We have presented a first measurement of the unnormalized skewness $\left\langle \tilde{T}^3(\hat{\mathbf{n}}) \right\rangle$ in ACT CMB maps filtered for high signal to noise. As this is a purely non-Gaussian signature, primordial CMB and instrumental noise cannot be confused with or bias the signal, unlike measurements of the tSZ power spectrum. We measure the skewness at 5σ significance: $\left\langle \tilde{T}^3(\hat{\mathbf{n}}) \right\rangle = -31 \pm 6\ \mu\mathrm{K}^3$ (Gaussian statistics as-

sumed). Including non-Gaussian corrections increases the error to $\pm$ 14 μK^3. Using analytic calculations and simulations to translate this measurement into constraints on cosmological parameters, we find $\sigma_8 = 0.79^{+0.03}_{-0.03}$ (68% C.L.) $^{+0.06}_{-0.06}$ (95% C.L.), with a slightly smaller but non-negligible systematic error due to theoretical uncertainty in the ICM astrophysics. This detection represents the first realization of a new, independent method to measure σ_8 based on the tSZ skewness, which has different systematic errors than several other common methods. With larger maps and lower noise, tSZ skewness measurements promise significantly tighter cosmological constraints in the near future.

A.8 Acknowledgments

As this manuscript was being prepared, we learned of related theoretical work by the authors of [44], and we acknowledge very helpful discussions with the members of this collaboration. This work was supported by the U.S. NSF through awards AST-0408698, PHY-0355328, AST-0707731 and PIRE-0507768, as well as by Princeton Univ., the Univ. of Pennsylvania, FONDAP, Basal, Centre AIUC, RCUK Fellowship (JD), NASA grant NNX08AH30G (SD, AH, TM), NSERC PGSD (ADH), NSF AST-0546035 and AST-0807790 (AK), NSF PFC grant PHY-0114422 (ES), KICP Fellowship (ES), SLAC no. DE-AC3-76SF00515 (NS), ERC grant 259505 (JD), BCCP (SD), and the NSF GRFP (BDS, BLS). We thank B. Berger, R. Escribano, T. Evans, D. Faber, P. Gallardo, A. Gomez, M. Gordon, D. Holtz, M. McLaren, W. Page, R. Plimpton, D. Sanchez, O. Stryzak, M. Uehara, and Astro-Norte for assistance with ACT. ACT operates in the Parque Astronómico Atacama in northern Chile under the auspices of Programa de Astronomía, a program of the Comisión Nacional de Investigación Científica y Tecnológica de Chile (CONICYT).

Bibliography

[1] Fowler, J. W., Niemack, M. D., Dicker, S. R., et al. 2007, Applied Optics, 46, 3444

[2] Dunner, R., Hasselfield, M., Marriage, T. M., et al. in prep.

[3] Swetz, D. S., Ade, P. A. R., Amiri, M., et al. 2011, ApJS, 194, 41

[4] Carlstrom, J. E., Ade, P. A. R., Aird, K. A., et al. 2011, PASP, 123, 568

[5] Schaffer, K. K., Crawford, T. M., Aird, K. A., et al. 2011, ApJ, 743, 90

[6] Zel'dovich, Y. B., & Sunyaev, R. A. 1969, Ap&SS, 4, 301

[7] Sunyaev, R. A., & Zel'dovich, Y. B. 1970, Ap&SS, 7, 3

[8] Staniszewski, Z., Ade, P. A. R., Aird, K. A., et al. 2009, ApJ, 701, 32

[9] Vanderlinde, K., Crawford, T. M., de Haan, T., et al. 2010, ApJ, 722, 1180

[10] Hincks, A. D., Acquaviva, V., Ade, P. A. R., et al. 2010, ApJS, 191, 423

[11] Marriage, T. A., Acquaviva, V., Ade, P. A. R., et al. 2011, ApJ, 737, 61

[12] Sehgal, N., Trac, H., Acquaviva, V., et al. 2011, ApJ, 732, 44

[13] Dunkley, J., Hlozek, R., Sievers, J., et al. 2011, ApJ, 739, 52

[14] Keisler, R., Reichardt, C. L., Aird, K. A., et al. 2011, ApJ, 743, 28

[15] Rubiño-Martín, J. A., & Sunyaev, R. A. 2003, Mon. Not. R. Astron. Soc, 344, 1155

[16] Battaglia, N., Bond, J. R., Pfrommer, C., & Sievers, J. L. 2011, arXiv:1109.3711

[17] Arnaud, M., Pratt, G. W., Piffaretti, R., et al. 2010, A&A, 517, A92

[18] Shaw, L. D., Nagai, D., Bhattacharya, S., & Lau, E. T. 2010, ApJ, 725, 1452

[19] Battaglia, N., Bond, J. R., Pfrommer, C., & Sievers, J. L. 2011, arXiv:1109.3709

[20] Hajian, A., Viero, M. P., Addison, G., et al. 2012, ApJ, 744, 40

[21] Komatsu, E., Dunkley, J., Nolta, M. R., et al. 2009, ApJS, 180, 330

[22] Komatsu, E., & Kitayama, T. 1999, ApJ, 526, L1

[23] Tinker, J., Kravtsov, A. V., Klypin, A., et al. 2008, ApJ, 688, 709

[24] Sheth, R. K., & Tormen, G. 2002, Mon. Not. R. Astron. Soc, 329, 61

[25] Komatsu, E., & Seljak, U. 2002, Mon. Not. R. Astron. Soc, 336, 1256

[26] Hill, J. C., & Sherwin, B. D. 2012, arXiv:1205.5794

[27] Trac, H., Bode, P., & Ostriker, J. P. 2011, ApJ, 727, 94

[28] Holder, G. P., McCarthy, I. G., & Babul, A. 2007, Mon. Not. R. Astron. Soc, 382, 1697

[29] Hajian, A., Acquaviva, V., Ade, P. A. R., et al. 2011, ApJ, 740, 86

[30] Marriage, T. A., Baptiste Juin, J., Lin, Y.-T., et al. 2011, ApJ, 731, 100

[31] Högbom, J. A. 1974, A&AS, 15, 417

[32] Bucher, M., & Louis, T. 2011, arXiv:1109.0286

[33] Sehgal, N., Bode, P., Das, S., et al. 2010, ApJ, 709, 920

[34] Lewis, A., & Challinor, A. 2006, Physics Reports, 429, 1

[35] Battaglia, N., Bond, J. R., Pfrommer, C., Sievers, J. L., & Sijacki, D. 2010, ApJ, 725, 91

[36] Marriage, T. A., Acquaviva, V., Ade, P. A. R., et al. 2011, ApJ, 737, 61

[37] Menanteau, F., González, J., Juin, J.-B., et al. 2010, ApJ, 723, 1523

[38] Menanteau, F., et al. in prep.

[39] van Engelen, A., Keisler, R., Zahn, O., et al. 2012, arXiv:1202.0546

[40] Vikhlinin, A., Kravtsov, A. V., Burenin, R. A., et al. 2009, ApJ, 692, 1060

[41] Reichardt, C. L., Shaw, L., Zahn, O., et al. 2011, arXiv:1111.0932

[42] Sehgal, N., Trac, H., Acquaviva, V., et al. 2011, ApJ, 732, 44

[43] Reichardt, C. L., Stalder, B., Bleem, L. E., et al. 2012, arXiv:1203.5775

[44] Bhattacharya, S., Nagai, D., Shaw, L., Crawford, T., & Holder, G. P. 2012, arXiv:1203.6368