To the Horizon and Beyond:

Weak Lensing of the CMB
and
Binary Inspirals into Horizonless Objects

Thesis by

Michael Kesden

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California

2005
(Submitted June 13, 2005)

This thesis is dedicated to my mother, whose inquiring mind challenges the scientific method everyday, my father, whose deep admiration for scientific discovery began with the *Microbe Hunters* many years ago, and the rest of my family who continue to sustain me throughout my life.

Acknowledgements

First and foremost, I would like to acknowledge the advice and support given to me throughout my graduate career by my advisor, Prof. Marc Kamionkowski . He has a tremendous intuition for identifying exciting research problems, and a gift for explaining the importance of his discoveries to any audience. The research presented in this thesis could not have been completed without substantial assistance from my collaborators, Prof. Asantha Cooray and Dr. Jonathan Gair, who generously shared their considerable expertise in all areas of theoretical astrophysics. Also worthy of recognition are the past and present members of the TAPIR research group, including but not limited to Eric Agol, Andrew Benson, Roger Blandford, Latham Boyle, Avery Broderick, Ben Collins, Adrienne Erickcek, Alison Farmer, Steve Furlanetto, Peter Goldreich, Dan Grin, Lee Lindblom, Geoffrey Lovelace, Ilya Mandel, Milos Milosavljevic, Peng Oh, Margaret Pan, Arielle Phillips, Sterl Phinney, Frans Pretorius, Jonathan Pritchard, Anthony Pullen, Yasser Rathore, Mike Santos, Re'em Sari, Naoki Seto, Kris Sigurdson, Tristan Smith, Sherry Suyu, Kip Thorne, and Nevin Weinberg, who created such a stimulating environment to work in over the past five years. Finally, I would like to thank the faculties and staff of the California Institute of Technology and Princeton University for providing me a world-class education in theoretical physics, and the National Science Foundation and NASA for generously supporting my graduate research.

Abstract

This thesis examines two predictions of general relativity: weak lensing and gravitational waves. The cosmic microwave background (CMB) is gravitationally lensed by the large-scale structure between the observer and the last-scattering surface. This weak lensing induces non-Gaussian correlations that can be used to construct estimators for the deflection field. The error and bias of these estimators are derived and used to analyze the viability of lensing reconstruction for future CMB experiments.

Weak lensing also affects the one-point probability distribution function of the CMB. The skewness and kurtosis induced by lensing and the Sunayev-Zel'dovich (SZ) effect are calculated as functions of the angular smoothing scale of the map. While these functions offer the advantage of easy computability, only the skewness from lensing-SZ correlations can potentially be detected, even in the limit of the largest amplitude fluctuations allowed by observation.

Lensing estimators are also essential to constrain inflation, the favored explanation for large-scale isotropy and the origin of primordial perturbations. $B-$mode polarization is considered to be a "smoking-gun" signature of inflation, and lensing estimators can be used to recover primordial $B-$modes from lensing-induced contamination. The ability of future CMB experiments to constrain inflation is assessed as functions of survey size and instrumental sensitivity.

A final application of lensing estimators is to constrain a possible cutoff in primordial density perturbations on near-horizon scales. The paucity of independent modes on such scales limits the statistical certainty of such a constraint. Measurements of the deflection field can be used to constrain at the 3σ level the existence of a cutoff large enough to account for current CMB observations.

A final chapter of this thesis considers an independent topic: the gravitational-wave (GW) signature of a binary inspiral into a horizonless object. If the supermassive objects at galactic centers lack the horizons of traditional black holes, inspiraling objects could emit GWs after passing within their surfaces. The GWs produced by such an inspiral are calculated, revealing distinctive features potentially observable by future GW observatories.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Those who have taken upon them to lay down the law of nature as a thing already searched out and understood ... have therein done philosophy and the sciences great injury.

-Francis Bacon, *Novum Organum* (1620)

Science is a celebration of the unknown. As long as curiosity endures and the law of nature remains even partially hidden, scientists will seek to advance the boundary between what is known and what is unknown. Research takes place on this horizon, but rarely are the horizons more explicit than in the research that comprises this thesis. As indicated by the title, this thesis consists of research in two almost entirely distinct areas: the weak lensing of the cosmic microwave background (CMB) and the gravitational waves (GWs) produced during a binary inspiral into a horizonless object. The first topic involves one of the few avenues we have to explore scales at and beyond the *particle horizon*, the boundary separating us from portions of the Universe from which we have been causally disconnected since the end of the inflationary epoch (Liddle & Lyth, 2000). The primordial CMB consists of photons that last scattered during recombination when protons and electrons first combined to form neutral hydrogen. At a mean redshift of $z \simeq 1,100$, the last-scattering surface reaches 97% of the way to the particle horizon and the lowest multipole moments of the CMB are sensitive to fluctuautions on even larger scales. In fact, if inflation lasts only several e-foldings longer than strictly necessary to solve the flatness and horizon problems, the CMB dipole and quadrupole may offer clues to the preinflationary era itself (Turner, 1991). This primordial CMB is gravitationally lensed by large-scale structure (LSS) extending all the way back to the last-scattering surface, and thus, the lowest multipole moments of the weak-lensing deflection field probe scales the size of the particle horizon and larger, just like the primordial CMB. We thus see that attempts to reconstruct the weak-lensing deflection field offer us the opportunity to explore beyond a horizon in a way few other observables do.

The second area considered in this thesis – the GWs produced during a binary inspiral into a horizonless object – provides us a chance to probe another horizon predicted by general relativity. The surfaces of black holes are *event horizons*, boundaries of the causal past of future null infinity (Misner et al., 1973). No null geodesics exist which cross this boundary, implying that not even photons can escape from the black hole's interior into the exterior, asymptotically flat spacetime. Substantial observational evidence suggests that supermassive compact objects exist at the centers of most galaxies. These objects should accrete stars from the surrounding galactic cusps that will spiral inwards as energy and angular momentum are lost to gravitational radiation. If the objects truly are supermassive black holes (SMBHs), they should possess event horizons as predicted by general relativity. No further gravitational radiation should reach the outside world once the accreted star passes within the event horizon and the black-hole perturbations are damped away. However, observation certainly allows room for the object to be something significantly less compact than a black hole, such as a supermassive boson star (SMBS). In this case, no horizon will exist and one would generically expect to receive GWs long after the horizon would have curtailed emission from a comparably massive black hole. By calculating the GW signal as the inspiraling star plunges beyond the radius at which one would expect to encounter the event horizon, we hope to develop an important observational test of the existence and nature of the boundary of this previously unexplored region.

Each of the following chapters of this thesis consists of a previously published paper in a peer-reviewed research journal, either *Physical Review Letters* or *Physical Review D*. The texts of these articles have not been changed; only the typesetting and bibliographies have been altered to conform with Caltech Ph.D. thesis regulations. Each of the chapters has its own brief introduction motivating the problem at hand, but I would like to take the remainder of this chapter to place the research in the context of other recent research in the field. Although a comprehensive review of all of the theoretical and experimental research on the CMB and GW-driven binary inspirals is well beyond the scope of this thesis, I hope that a summary of the most relevant research in each field will help to provide an appreciation for the significance of my research and help to elucidate future avenues of investigation.

1.1 Weak Lensing of the CMB

1.1.1 Origins of the CMB

CMB physics became an experimental science in 1965, when Penzias & Wilson (1965) concluded that the isotropic, unpolarized 3.5 K excess antenna temperature they had observed was cosmological in origin. In a companion paper, Dicke et al. (1965) analyzed the profound implications of this revolutionary discovery. The previously observed Hubble expansion combined with the Friedmann equations of general relativity had already suggested that the Universe had been hotter and denser in the past. The existence of a microwave background, however, would imply that this expansion could be traced all the way back to a time when pair creation and Thomson scattering kept the photons and electrons in thermal equilibrium. If the observed 25% relic abundance of helium had indeed been produced through big-bang nucleosynthesis (BBN) as suggested by others (Alpher et al., 1948), then the baryon density at the well-defined temperature where BBN occurs could be approximately calculated. When combined with a measurement of the baryon density today and the well known scaling of this density with the scale factor, $\rho(t) \propto a^{-3}(t)$, an estimate for the temperature of the CMB could be obtained. The actual measurement of the 3.5 ± 1 K microwave background by Penzias and Wilson provided cosmologists with a crucial constraint on the expansion history of the universe, and even allowed them to consider the possibility of a massless scalar field or other exotic dark component (Dicke et al., 1965).

The detection of the isotropic CMB motivated a great deal of theoretical work into the nature of possible CMB anisotropies. On the largest scales, Sachs & Wolfe (1967) explored how linear perturbations in the gravitational potential might lead to anisotropies in the CMB as photons propagated from the last-scattering surface to the observer. Though their ignorance of cosmological parameters and the spectrum of primordial perturbations caused them to set an erroneous lower bound of 1% on the amplitude of CMB temperature fluctuations, the physics of their analysis was sound and the integrated Sachs-Wolfe (ISW) effect is recognized today as the dominant contribution to the CMB power spectrum on large scales. In Kesden et al. (2003) – Chapter 4 of this thesis – I explore the ISW effect and its correlation with weak lensing in the presence of a large-scale cutoff. On smaller scales, Silk (1968) recognized that perturbations in the CMB below a certain critical size would be damped by radiative diffusion before the epoch of galaxy formation. The existence of galaxy clusters and structure on Mpc scales today thus implies fluctuations in the CMB on scales of $10''$ and greater. Sunyaev & Zel'dovich (1970) recognized that the connection between CMB fluctuations and adiabatic density perturbations after recombination was not as simple as proposed by Silk. Silk assumed that if the density fluctuations $\delta\rho/\rho \simeq 1$ observed today grew linearly with the scale factor $a = 1/(1+z)$ as in an Einstein-de Sitter universe, then recombination at redshift $z \simeq 1000$ would imply CMB temperature fluctuations of order $\delta T/T \simeq 10^{-3}$. Sunyaev and Zel'dovich showed that continuity in the velocity perturbations of the photon-baryon fluid during recombination implied an amplitude jump of the adiabatic perturbations by a factor of order 10^2. The observed CMB temperature fluctuations would be further suppressed by this factor, leading to $\delta T/T \simeq 10^{-5}$ in agreement with present-day observations. Sunyaev & Zel'dovich (1970) also noted that the Jeans length was growing prior to recombination, implying that perturbations on a given scale would transition from growing modes to standing waves as they slipped below the Jeans length. Conditions for joining these solutions through the transition required that sound waves on all scales began to oscillate with the same initial phase. However, the smaller-scale perturbations would become standing waves sooner and thus, would have more time to oscillate prior to recombination. This led to distinct peaks in the CMB power spectrum for those scales on which perturbations underwent an integral number of oscillations. This argument led Sunyaev and Zel'dovich to predict the acoustic-peak structure of the CMB, decades before it was spectacularly confirmed by Boomerang, Maxima, DASI, and other experiments (de Bernardis et al., 2000; Hanany et al., 2000; Halverson et al., 2002). The prediction of acoustic peaks was further supported by Peebles & Yu (1970), who conducted a direct numerical integration of the collision equation for the photon distribution function through the era of recombination. This analysis went beyond the treatment of photons and baryons as a viscous fluid, which was strictly valid only at early times when the photon mean free path was much shorter than the wavelength of the perturbations. Peebles & Yu (1970) again found that the physics of recombination imposed two length scales into the CMB: the sound horizon at recombination and the scale below which dissipation of the perturbations through radiative diffusion became appreciable. Although the scales and amplitudes of the CMB peaks would remain model-dependent quantities until their detection, the physics behind their existence was firmly established by the research of this period.

Peebles & Yu (1970) remarked that while they had analyzed the evolution of density perturbations through recombination, "It is well to bear in mind that in this calculation the initial density fluctuations are invoked in an ad hoc manner because we do not have a believable theory of how they may have originated." Such a theory would not be developed until 1981, when Guth (1981) proposed a period of exponential inflation in the early Universe to solve two long-standing problems of the standard hot big-bang cosmology. The horizon problem related to the lack of a causal mechanism to account for the observed isotropy of the CMB on scales larger than the particle horizon at recombination. Assuming that the dynamics of the expanding Universe was dominated purely by its observed radiation and matter content, the present-day Universe consisted of 10^{83} causally disconnected regions, back when the mean temperature was $T_0 = 10^{17}$ GeV. Immense fine-tuning would be required to explain the homogeneity of these regions in the absence of a causal process. The second problem is known as the flatness problem and arises from the instability of the flat $\Omega = 1$ cosmological solution. In order for the present-day density to be in the range $0.01 < \Omega < 10$ – conservative even in 1981 – the mean density must have differed from the critical density by less than a part in 10^{55} back when $T_0 = 10^{17}$ GeV. Guth proposed to solve both of these problems by positing a period of inflation during which the Universe was vacuum-energy dominated and the Hubble constant did not vary with scale factor. During this time, the physical size of causally connected regions, along with the radius of curvature, would increase by over 65 e-foldings so that even today, such scales would be well beyond the Hubble radius. A flat, homogenous Universe is thus seen to be an unavoidable consequence of inflation.

While the importance of resolving the flatness and horizon problems cannot be overstated, inflation's true predictive power came from its ability to generically produce the scale-invariant spectrum of initial density perturbations long sought by Peebles and others. These perturbations arose from quantum fluctuations in the scalar field driving inflation. Albrecht & Steinhardt (1982) and Linde (1982a) initially introduced the scalar field as a modification to Guth's inflationary scenario that could provide a "graceful exit" from the inflationary phase to the usual radiation-dominated Friedmann-Robertson-Walker Universe. Guth presumed that inflation occurred as the Universe cooled in a metastable state below the temperature of a first-order phase transition. Albecht, Steinhardt, and Linde proposed "new inflation", in which spontaneous symmetry breaking (SSB) introduces a nonzero global minimum in the scalar-field potential below a given temperature. Inflation occurs as the scalar field slowly rolls towards this new minimum, and naturally ends once the minimum is reached. Rapid fluctuation of the scalar field near the minimum transfers the vacuum energy to a thermal distribution of particles and radiation, allowing the hot big bang to proceed as previously considered. A group of authors (Guth & Pi, 1982; Hawking, 1982; Linde, 1982b; Starobinsky, 1982; Bardeen et al., 1983) recognized that the existence of an event horizon during inflation naturally implied spatially varying zero-point fluctuations in the scalar field. Fluctuations with a given comving wavenumber k are produced when that scale exits the horizon during inflation, and only reenter the observable Universe long after inflation has ceased. The near scale-invariance of these perturbations follows from the approximate time-translation symmetry of the Universe during inflation, and the inability of causal processes to operate once the modes have left the horizon. The incorporation of inflation into the hot big-bang scenario finally led to a "standard model" of cosmology with definite predictions ripe for testing in the crucible of experiment.

The first positive detection of the CMB anisotropies occurred in 1991, when the Cosmic Background Explorer (COBE) observed a 30 ± 5 μK rms sky variation when smoothed with a 10° Gaussian beam (Smoot et al., 1992). This discovery led to a flurry of observational effort to characterize the CMB anisotropies on all scales, culminating with the all-sky temperature map currently being produced by the WMAP collaboration (Bennett et al., 2003). These maps have been combined with other observations of large-scale structure to estimate the cosmological parameters with unprecedented precision (Balbi et al., 2000; Lange et al., 2001; Spergel et al., 2003). Measurement of the CMB temperature and polarization fields with ever increasing resolution has become a thriving industry, with a bevy of ground-based, balloon, and satellite experiments in all stages of development. On the theoretical front, characterizing the subtly different implications of various flavors of inflation has become a cottage industry in some quarters. Kamionkowski & Kosowsky (1999) is an excellent review of the implications of inflation and other recent developments in particle physics for CMB observables. Other efforts have focused on accurately predicting the secondary anisotropies imposed on the CMB since recombination, and developing new data analysis techniques for handling the extremely large data sets being produced by current and future CMB experiments. Hu & Dodelson (2002) provided an excellent recent review of progress in these areas. The next subsection will consider the weak gravitational lensing of the CMB by large-scale structure, the secondary anisotropy most relevant to the research presented in this thesis.

1.1.2 Lensing by Large-Scale Structure

Gravitational lensing – the deflection of photon trajectories by density perturbations – has been a fundamental prediction of general relativity since it was first developed by Einstein in 1915. The measurement of the weak-lensing deflection of starlight by A. S. Eddington during his famous 1919 expedition to observe a solar eclipse in West Africa was instrumental in convincing the scientific community of the validity of relativity theory. Weak lensing of background galaxies by foreground large-scale structure (LSS) has been considered theoretically since the 1960's. Sachs (1961) proved several important theorems concerning the propagation of photons along the null geodesics of a vacuum spacetime. Kristian & Sachs (1966) recognized that these theorems implied a new "distortion effect", namely that "in any anisotropic model, all distant objects in a particular direction on the celestial sphere may appear distorted, with a definite preferential direction for their longest dimension." Gunn (1967) studied this distortion effect – now called weak lensing – in a variety of inhomogenous cosmologies and calculated the correlation function for the shear field responsible for the lensing.

Weak lensing by LSS remained theoretical until Lynds & Petrosian (1989) convincingly showed that the luminous arcs they had observed behind a foreground cluster were in fact, lensed galaxies. This initiated an outpouring of further observational and theoretical work. Using the net tangential alignment of 22-26 B mag blue galaxies, Tyson et al. (1990) actually reconstructed the lens structure of the cluster A1689. Miralda-Escudé (1991) constructed a formalism for calculating the observable galaxy ellipticity correlation functions for arbitrary distortion correlations and intrinsic ellipticities. He then applied this formalism to calculate these correlations for Poisson-distributed lenses and distributions given by an Einstein-de Sitter cold dark matter model. Kaiser (1992) conducted a similar analysis of the relation between correlations of the lensing shear field and observed ellipticities, but he emphasized the advantages of performing this analysis in Fourier space. The deflection field $\delta\theta_i$ of a source at comving distance w is a weighted integral along the line-of-sight of the tangential gradient of the Newtonian potential $\phi(\mathbf{w}, \eta)$,

$$\delta\theta_i = \frac{-2}{w} \int_0^w dw'(w - w')\phi_{,i}(\mathbf{w}', 1 - w') \, . \tag{1.1}$$

This implies that Fourier modes of the projected density field will induce deflections parallel to their wave number. Kaiser recognized that appropriately weighted Fourier modes of the observed galaxy ellipticities could serve as estimators for the *same* Fourier modes of the projected density field. This is very similar to the argument used in Kesden et al. (2003), Chaper 2 of this thesis, to conclude that appropriately weighted Fourier modes of quadratic combinations of CMB observables could serve as estimators for Fourier modes of the deflection field. Kaiser then combined estimators for individual Fourier modes to produce estimators for the dark matter power spectrum, and he even addressed issues of bias like those considered in Kesden et al. (2003). In many ways, Kaiser (1992) provided the theoretical underpinning for much of the later work in CMB-based lensing reconstruction.

Even before the first observation of the large-scale CMB anisotropies by COBE in 1991, it was realized that the CMB would be gravitationally lensed by foreground LSS in much the same way high-redshift galaxies were. Blanchard & Schneider (1987) recognized that at each position x on the sky, gravitational bending will result in an apparent displacement field $\lambda(x)$. This displacement field will remap the CMB so that the apparent temperature one observes in a direction x is what one would have observed in a direction $x + \lambda(x)$ in the absence of lensing. Lensing conserves specific intensity, and hot and cold patches of the CMB are equally likely to be lensed, provided the displacement field $\lambda(x)$ is uncorrelated with the primordial anisotropies at the last-scattering surface. In this case, Blanchard and Schneider arrived at the important conclusion that "the variance of the fluctuating part of the background is conserved by any arbitrary physical deviation field." Kesden et al. (2002a), included as Chapter 3 of this thesis, is devoted to an analysis of the one-point probability distribution function (PDF) in the presence of lensing and shows that not just the variance, but all higher moments of the distribution are unchanged by weak lensing.

The form of the CMB power spectrum was largely unconstrained before COBE, and Blanchard & Schneider (1987) chose to model it as a simple Gaussian,

$$P(k) = \left| \frac{\delta \hat{T}}{T}(k) \right|^2 \propto \exp(-\theta_0^2 k^2/2)\,. \tag{1.2}$$

While this form failed to capture the rich structure of the acoustic peaks, it allowed Blanchard and Schneider to qualitatively predict the shifting of power from large to small scales calculated in more accurate modern analyses.

This prediction of enhanced CMB power on small scales was not shared by other contemporary authors who believed that lensing would lead to an overall reduction in CMB power below scales of several arcminutes. Kashlinsky (1988) considered the possibility of multiple gravitational lensing by compact objects like discrete galaxies. At the time, optimistic assumptions like primordial isocurvature perturbations, an Einstein-de Sitter ($\Omega = 1$) universe, and completed structure formation before redshift $z_i \gtrsim 5$ made the possibility of a large optical depth to multiple lensing more viable than it is today. Such a scenario would supposedly lead to the phase mixing of photons from different primordial patches and a corresponding decrease in the amplitude of observed fluctuations. Linder (1988) also claimed that the effective observed temperature in a direction ϕ would be a weighted average of the primordial temperatures in directions $\phi + \theta$,

$$T_{\rm eff}(\phi) = \int d^2\theta\, p(\theta) T(\phi + \theta)\,, \tag{1.3}$$

where $p(\theta)$ was the probability that the lensing deflection field assumed a value of θ in direction ϕ. While such an ensemble average would indeed provide a mechanism for isotropization of anisotropies as claimed by Linder, the Universe consists of a single realization of the deflection field and therefore, no such average should be reflected in observed quantities.

A consensus of the effects of lensing on the CMB was finally reached during the renewed theoretical investigations in the aftermath of COBE. Seljak (1996) gave a definitive treatment of the effect of lensing on the CMB temperature power spectrum. Using Limber's equation in Fourier space to relate three-dimensional density perturbations to projected quantities, Seljak derived an expression $\sigma(\theta)$ for the lensing-induced dispersion in the angular separartion of photons separated by an angle θ. He then showed that the observed power spectrum $\tilde{C}_l$ could be expressed as a convolution of the primordial power spectrum C_l,

$$\tilde{C}_l = \int_0^\infty C_{l'} \frac{dl'}{\sqrt{2\pi}\epsilon l'} e^{-(l-l')^2/2(\epsilon l')^2} \,, \tag{1.4}$$

where $\epsilon \equiv \sigma(\theta)/\theta$. This lensed power spectrum was evaluated numerically for both Einstein-de Sitter and ΛCDM cosmologies, and found to have little effect on scales above a degree. The above convolution did, however, lead to broadening of the acoustic peaks for $l \gtrsim 1000$. Seljak suggested that this effect might need to be taken into account in future high-precision efforts at parameter estimation. Although Seljak (1996) showed that the effect of lensing on the CMB power spectrum was not as important as suggested by previous authors, lensing does have significant implications for higher-order correlation functions as discussed in the next subsection.

1.1.3 Non-Gaussianity

An essential model-independent prediction of single-field inflation is that primordial perturbations should be described by Gaussian statistics (Bardeen et al., 1983). The perturbations remain small and evolve linearly from the end of inflation through the recombination era, implying that this Gaussianity should be preserved in the primordial CMB. The statistics of Gaussian fields are determined *entirely* by their power spectrum; all higher-order correlations can be expressed in terms of the power spectrum. However, it has long been understood that the dynamics of gravitational instability are nonlinear, and that once the dispersion $\sigma \equiv < \delta^2 >^{1/2}$ of density perturbations δ exceeds unity, the results of linear perturbation theory are no longer valid. Juszkiewicz et al. (1995) studied the evolution of the perturbations in the weakly nonlinear regime in the context of structure formation, and derived a consistent power series expansion in σ for the one-point PDF of δ. The coefficients in this expansion are polynomials of cumulants, those portions of the central moments that vanish for a Gaussian distribution. Although the equations governing the growth of structure differ from those pertaining to weak lensing, the mathematical form of the expansion in the weakly non-Gaussian regime is the same and proves essential to my treatment of the one-point PDF of the lensed CMB in Chapter 3. Explicit estimators $\bar{\xi}_N$ for the cumulants of a PDF are defined for a finite number of samples N_T in Hui & Gaztanaga (1999). Traditionally, such estimators are constructed from nonlinear combinations of estimators for the central moments of a distribution. This introduces what Hui and Gaztanaga refer to as "nonlinear estimation bias", even if the initial estimators are themselves unbiased. This issue is fundamental to the bias associated with power-spectrum estimators derived from estimators for individual modes as discussed in Chapter 2. It is also responsible for the confusion pertaining to the variance associated with different estimators for the skewness and kurtosis as described in the final section of Chapter 3.

Nonlinearity occurs not just in the growth of large-scale structure, but also in the remapping of the observed CMB anisotropies brought about by weak lensing. Bernardeau (1997) realized that the weak-lensing deflection field would induce non-Gaussian correlations in the observed CMB map, and that these correlations would lead to a nonvanishing connected part of the four-point correlation function. In the flat-sky limit, i.e., on scales for which Fourier transforms can be used in place of decomposition into spherical harmonics, Bernardeau derived expressions for the real-space four-point function in terms of the primordial CMB power spectrum and the perturbation power spectrum $P(k)$ that gives rise to weak lensing. He then considered the dependence of the four-point function on the geometry of the configuration, and calculated the amplitude of the correlation in a flat $\Omega_m = 0.3$, $\Omega_\Lambda = 0.7$ cosmology. He concluded that this lensing-induced non-Gaussianity should be observable in full-sky CMB maps, but that further theoretical work was needed to determine the optimal means of searching for this non-Gaussianity. Bernardeau (1998) proposed one such means of searching for non-Gaussianity: the statistics of the local curvature matrix

$$c_{ij} \equiv \frac{d^2\delta_T}{dx_i dx_j} \tag{1.5}$$

of temperature fluctuations δ_T. Defining an ellipticity e from this curvature matrix, Bernardeau examined how lensing affected the probability distribution function of e for different values of the smoothing scale and power spectrum normalization σ_8. He found that lensing could be more readily detected in simulated maps using the local ellipticity than the simple collapsed four-point function, where much of the effect cancels due to averaging over the different configurations. This cancellation is central to the conclusions of Chapter 3, where I examine whether smoothing on an appropriate scale could interfere with the cancellation enough to make the collapsed four-point function detectable.

Analysis of higher-order correlations in harmonic space began with Spergel & Goldberg (1999), where the bispectrum

$$\langle a_{l_1 m_1} a_{l_2 m_2} a_{l_3 m_3} \rangle \equiv \begin{pmatrix} l_1 & l_2 & l_3 \\ m_1 & m_2 & m_3 \end{pmatrix} B_{l_1 l_2 l_3} \tag{1.6}$$

was considered in a cosmological context and an estimator for this quantity was constructed from all-sky maps. As described previously, the growth of LSS is nonlinear, and the resulting non-Gaussianity leaves an imprint in the CMB through the Rees-Sciama effect. The contribution of this effect to the bispectrum was determined to be undetectable in a ΛCDM universe, even by an experiment limited only by cosmic variance. In a companion paper (Goldberg & Spergel, 1999), cross-correlations between the weak-lensing field and the integrated Sachs-Wolfe or Sunyaev-Zel'dovich (SZ) effects were found to induce similar contributions to the bispectrum. For most choices of cosmological parameters, the lensing-SZ correlations were found to dominate, and in Chapter 3 we show that this effect could even lead to observable non-Gaussianity in the one-point function for high values of σ_8. A definitive treatment of the CMB bispectrum was provided by Hu (2000a), who performed a detailed calculation of the

CMB power spectra and bispectra for both temperature and polarization. The flat-sky approximation in which the maps are decomposed into Fourier modes instead of spherical harmonics was found to introduce a 10% error into the calculation. Surprisingly, this error was not limited to the largest scales where one would naturally expect the flat-sky approximation to break down, but rather extended to $l \simeq 1000$. Hu attributed this peculiar result to the mode-coupling nature of weak lensing, whereby poorly approximated large-scale modes of the lensing field provide the dominant contribution to small-scale CMB power through remapping.

This shifting of power from large to small scales was explored in a more general context by Zaldarriaga (2000). If the primordial CMB was described by a uniform gradient on a sufficiently small scale, possibly due to the effects of Silk damping, Fourier modes of the observed temperature map below this scale would be directly proportional to the corresponding modes in the deflection field $\delta\theta$,

$$T(\mathbf{l}) = \tilde{T}_x \delta\theta_x(\mathbf{l}) + \tilde{T}_y \delta\theta_y(\mathbf{l}) \,, \tag{1.7}$$

where $\tilde{T}_x$ and $\tilde{T}_y$ are partial derivatives of the temperature map in the x and y directions, respectively. Zaldarriaga proposed to test this hypothesis by searching for a correlation between small-scale power and the large-scale gradients that source this power as described by Eq. (1.7). The statistics of temperature gradients in the observed map was also considered by Villumsen (2001) in the hope that non-Gaussianity might be detectable in the one-point PDF. Villumsen found that the kurtosis of the temperature gradient was linearly proportional to the variance of the convergence κ, defined as

$$\kappa \equiv -1/2 \nabla \cdot \delta\theta \tag{1.8}$$

and further claimed that this kurtosis was a maximum-likelihood estimator for σ_κ^2. Increasing knowledge of the detailed effect of lensing on the CMB motivated authors to consider not just a statistical detection of lensing, but a complete reconstruction of the lensing field – a topic addressed in the following subsection.

1.1.4 Lensing Reconstruction

A method for reconstructing the weak-lensing field from CMB observables was first introduced in a *Letter* (Seljak & Zaldarriaga, 1999a), with details of the derivation and performance of this method included in the companion paper (Zaldarriaga & Seljak, 1999). Seljak and Zaldarriaga defined three quadratic combinations of derivatives $T_i \equiv \partial T/\partial x_i$ of the CMB temperature field, and then to linear order, expressed these combinations in terms of the unlensed quantities and symmetric shear tensor $\Phi_{ab} \equiv \partial\delta\theta_a/\partial\theta_b$,

$$\begin{aligned}
\mathcal{S} &\equiv T_x^2 + T_y^2 \\
&= (1+\Phi_{xx}+\Phi_{yy})\tilde{\mathcal{S}} + (\Phi_{xx}-\Phi_{yy})\tilde{\mathcal{Q}} + 2\Phi_{xy}\tilde{\mathcal{U}} \\
\mathcal{Q} &\equiv T_x^2 - T_y^2 \\
&= (1+\Phi_{xx}+\Phi_{yy})\tilde{\mathcal{Q}} + (\Phi_{xx}-\Phi_{yy})\tilde{\mathcal{S}} \\
\mathcal{U} &\equiv 2T_xT_y \\
&= (1+\Phi_{xx}+\Phi_{yy})\tilde{\mathcal{U}} + 2\Phi_{xy}\tilde{\mathcal{S}}\,.
\end{aligned} \tag{1.9}$$

The isotropy of the primordial CMB implies that, averaged over different Gaussian realizations,

$$\begin{aligned}
\langle\tilde{\mathcal{S}}\rangle_{\rm CMB} &= \sigma_{\mathcal{S}} \\
\langle\tilde{\mathcal{Q}}\rangle_{\rm CMB} &= 0 \\
\langle\tilde{\mathcal{U}}\rangle_{\rm CMB} &= 0\,.
\end{aligned} \tag{1.10}$$

Defining the convergence κ and shear fields γ_1 and γ_2 from the shear tensor,

$$\begin{aligned}\kappa &\equiv -1/2(\Phi_{xx}+\Phi_{yy})\\ \gamma_1 &\equiv -1/2(\Phi_{xx}-\Phi_{yy})\\ \gamma_2 &\equiv -\Phi_{xy}\,,\end{aligned} \tag{1.11}$$

Seljak and Zaldarriaga determined that

$$\begin{aligned}\langle\tilde{\mathcal{S}}\rangle_{\rm CMB} &= (1-2\kappa)\sigma_{\mathcal{S}}\\ \langle\tilde{\mathcal{Q}}\rangle_{\rm CMB} &= -2\gamma_1\sigma_{\mathcal{S}}\\ \langle\tilde{\mathcal{U}}\rangle_{\rm CMB} &= -2\gamma_2\sigma_{\mathcal{S}}\,.\end{aligned} \tag{1.12}$$

Because the shear tensor is derived from a scalar deflection potential, κ, γ_1, and γ_2 are not independent quantities, but obey the Fourier-space relations

$$\begin{aligned}\gamma_1(\mathbf{l}) &= \kappa(\mathbf{l})\cos(2\phi_l)\\ \gamma_2(\mathbf{l}) &= \kappa(\mathbf{l})\sin(2\phi_l)\,.\end{aligned} \tag{1.13}$$

This implies that both $\mathcal{S}$ and an appropriate combination of $\mathcal{Q}$ and $\mathcal{U}$ can serve as estimators for the convergence field κ. Seljak and Zaldarriaga calculated the covariance matrix for these estimators on scales both above and below the correlation length ξ, below which power in the primordial CMB is largely suppressed by Silk damping. On scales larger than ξ, cosmic variance in the primordial CMB contributes to the variance of the estimator in addition to instrumental noise. In the small-scale limit, it should be possible in principle to express temperature fluctuations in terms of products of the large-scale gradients and the deflection field as in Eq. (1.7). There should, accordingly, be no cosmic-variance contribution to the reconstruction; a perfect reconstruction would be possible for a no-noise experiment. Seljak and Zaldarriaga tested the performance of their estimators by attempting to reconstruct the deflection field of a simulated galaxy cluster, then discussed the more realistic case of power-spectrum estimation for all-sky experiments like WMAP and Planck. Power-spectrum estimators consist of products of estimators for the individual modes, which are themselves quadratic in the CMB temperature field. Seljak and Zaldarriaga recognized that the noise properties of these estimators would therefore depend on the fourth-order statistics of the CMB, yet they assumed that for this purpose, the CMB could be considered to be Gaussian. This is inconsistent with the very premise of exploiting non-Gaussianity to reconstruct the weak-lensing field. An accurate calculation of the noise associated with power spectrum estimation that properly accounts for all non-Gaussianity to a consistent order in the weak-lensing field is provided in Chapter 2 of this thesis. Despite this approximation, Seljak and Zaldarriaga estimate the WMAP will detect lensing with a signal-to-noise ratio of 3, while for Planck, the figure will be between 15-25. This optimistic appraisal of lensing reconstruction motivated other authors to try to refine the technique further.

Significant advances were made in a series of papers by Wayne Hu. A new quadratic estimator for lensing reconstruction was introduced in Hu (2001a), with extensive derivations of its validity given in the companion paper (Hu, 2001b). Hu approached lensing reconstruction from the perspective of maximizing the information contained in the non-Gaussianity of the CMB. Hu conducted a thorough analysis of the lensing trispectrum in the all-sky limit, as he recognized that the power spectrum of a quadratic estimator was just a partially collapsed measure of the more general four-point function:

$$\langle \Theta_{l_1 m_1} \dots \Theta_{l_4 m_4} \rangle = \sum_{LM} \begin{pmatrix} l_1 & l_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \begin{pmatrix} l_3 & l_4 & L \\ m_3 & m_4 & M \end{pmatrix} (-1)^M Q^{l_1 l_2}_{l_3 l_4}(L) \,, \quad (1.14)$$

where Θ_{lm} are the harmonic coefficients of the temperature field and $Q^{l_1 l_2}_{l_3 l_4}(L)$ can be decomposed into an unconnected Gaussian component and the trispectrum $T^{l_1 l_2}_{l_3 l_4}(L)$.

The symmetry of the trispectrum implied that it could be expressed as a sum of a fundamental pairing $P^{l_1 l_2}_{l_3 l_4}(L)$ and summations over different configurations of that pairing:

$$T^{l_1 l_2}_{l_3 l_4}(L) = P^{l_1 l_2}_{l_3 l_4}(L) + (2L+1)\sum_{L'}\left[(-1)^{l_2+l_3}\begin{Bmatrix} l_1 & l_2 & L \\ l_4 & l_3 & L' \end{Bmatrix} P^{l_1 l_3}_{l_2 l_4}(L') \right.$$
$$\left. +(-1)^{L+L'}\begin{Bmatrix} l_1 & l_2 & L \\ l_3 & l_4 & L' \end{Bmatrix} P^{l_1 l_4}_{l_3 l_2}(L')\right] . \tag{1.15}$$

Hu defined a fully general quadratic estimator in harmonic space:

$$x^{ab}_{LM} \equiv (-1)^M \sum_{l_1 m_1}\sum_{l_2 m_2} x^{ab}_{l_1 l_2}(L)\Theta_{l_1 m_1}\Theta_{l_2 m_2}\sqrt{\frac{2L+1}{4\pi}}\begin{pmatrix} l_1 & l_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \tag{1.16}$$

where $x^{ab}_{l_1 l_2}(L)$ is an arbitrary filtering function, and showed that with appropriate normalization, this statistic could serve as an estimator for the lensing potential. He then chose the weights $x^{ab}_{l_1 l_2}(L)$ to maximize the signal-to-noise (S/N) of this estimator in the limit – valid for $L \lesssim$ several hundred – that the first term of Eq. (1.15) dominated the remaining terms. In Chapter 2, I consider these neglected terms and how they might affect the choice of filtering function. Hu found that his choice of weights boosted the total signal-to-noise ratio for lending detection for Planck to $(S/N)^2 \approx 4050$, compared to $(S/N)^2 \approx 135$ for the gradient-gradient statistic studied by Zaldarriaga & Seljak (1999).

This improvement in signal-to-noise ratio becomes even more pronounced when CMB polarization observables are used to construct additional estimators. The most general polarization field can be described by the Stokes' parameters Q, U, and V (Rybicki & Lightman, 1979), but as the CMB is not expected to be circularly polarized, only Q and U are nonvanishing. Seljak & Zaldarriaga (1997) and Kamionkowski et al. (1997a) independently recognized that the polarization field could be decomposed in a coordinate-independent way into $E-$ and $B-$modes that could be expressed in

Fourier space as linear combinations of these Stokes' parameters:

$$\begin{aligned} E(\mathbf{l}) &= Q(\mathbf{l})\cos(2\phi_{\mathbf{l}}) + U(\mathbf{l})\sin(2\phi_{\mathbf{l}}) \\ B(\mathbf{l}) &= -Q(\mathbf{l})\sin(2\phi_{\mathbf{l}}) + U(\mathbf{l})\cos(2\phi_{\mathbf{l}})\,. \end{aligned} \tag{1.17}$$

Along with the temperature field itself, there are $N = 3$ different CMB observables, and thus, $N(N + 1)/2 = 6$ possible quadratic estimators. Hu & Okamoto (2002) showed that all six of these quadratic statistics are valid estimators for the weak-lensing deflection field, and calculated their covariance matrix in the flat-sky limit. This analysis is extended to all-sky harmonic-space estimators by Okamoto & Hu (2003). Hu and Okamoto found that the polarization-based estimators surpassed the performance of their quadratic temperature statistic for experiments with better than 26μK-arcmin noise and $4'$ resolution. For cosmic-variance-limited experiments, the minimum variance combination of the six estimators provided over an order-of-magnitude improvement over the signal-to-noise ratio possible with temperature alone. They also considered possible applications of lensing reconstruction, including the subtraction of lensing-induced $B-$modes to set constraints on the primordial $B-$modes produced by inflationary gravitational waves. I examine the details of this subtraction at length in Chapter 4.

The quadratic estimators devised by Wayne Hu show great theoretical promise, and recent efforts have sought to verify this promise by testing their effectiveness on simulated data from experiments faced with astrophysical foregrounds, limited sky coverage, and other practical considerations. Amblard et al. (2004) examined the performance of quadratic estimators by creating different Monte-Carlo realizations of the primordial CMB and deflection field, and then determining the lensed map explicitly. They then reconstructed the convergence power spectrum from the lensed map and compared it to the input power spectrum. For a no-noise experiment, the reconstructed power spectrum was biased upwards by $\sim 70\%$ after subtracting of the Gaussian noise determind by Hu. This bias was reduced to $\sim 25\%$ once the second-order noise calculated by Cooray & Kesden (2003) and Kesden et al. (2003) was taken into account, implying that future low-noise, high-resolution experiments will need to consider these terms in their analyses. Other complications discussed by Amblard et al. (2004) include a non-Gaussian deflection field and contamination from the kinetic Sunyaev-Zel'dovich (kSZ) effect. Nonlinear structure formation at low redshifts leads to non-Gaussianity in the deflection field, so Amblard et al. (2004) replaced their Gaussian samples of structure at redshifts less than $z = 2$ with the output of numerical simulations. The resulting non-Gaussian deflection field leads to an $\sim$ 50% bias in the estimated convergence power spectrum, and an $\sim$ 10% shift of power from small to large scales. The kSZ effect is a thermal distortion in the CMB induced by the bulk motion of clusters. Unlike the thermal SZ effect, these distortions have a Planck spectrum and therefore, cannot be removed even in principle by a multi-frequency survey. Amblard et al. (2004) took advantage of the fact that the *same* portions of the sky should be contaminated by both the kinetic and thermal SZ effects, and therefore, a mask determined by the distinguishable thermal SZ effect would eliminate the kinetic effect as well. By masking only 1.4% of the pixels, Amblard et al. (2004) were able to reduce the kSZ-induced bias from over a factor of 2 to $\sim$ 70%. After studying the performance of quadratic estimators for arbitrary noise and resolution, Amblard et al. (2004) applied these results to forecast how well upcoming CMB experiments such as APEX/SZ, ACT, SPT, and Planck could actually reconstruct the convergence power spectrum. SPT was found to be the most effective experiment at lensing reconstruction, owing to its superior combination of low noise, high resolution, and large survey area.

A greater understanding of the difficulties associated with quadratic-based lensing reconstruction is valuable, but a practical approach for moving past these difficulties is even more so. Hirata and Seljak developed maximum-likelihood estimators (MLEs) for the lensing potential using both CMB temperature (Hirata & Seljak, 2003a) and polarization (Hirata & Seljak, 2003b). For measured CMB temperature fluctuations $\hat{\Theta}(\mathbf{n}_i)$ at N points $\{\mathbf{n}_1, \ldots, \mathbf{n}_N\}$, these estimators select the lensing potential field Φ that maximizes the probability density $P(\hat{\Theta}|\Phi)$. This approach is guaranteed to be asymptotically efficient in the limit of large data sets, and retains its validity when second-order contributions in Φ become significant. Unfortunately, it also requires one to solve N coupled nonlinear equations which rapidly become computationally untractable, particularly for polarization. Hirata and Seljak resorted to a series of approximations to make the computational problem more manageable, and arrived at an iterative lensing reconstruction estimator that appeared well-behaved as higher-order terms were taken into account. While this estimator yielded only a $10 - 20\%$ improvement in the RMS error for temperature alone, for high-resolution polarization experiments, they claimed over an order-of-magnitude improvement in the mean squared error compared to quadratic techniques. Some form of approximate maximum likelihood technique will therefore likely be employed for lensing-reconstruction efforts with future high-sensitivity polarization experiments. Having given a review of the theory behind lensing reconstruction, I will discuss two possible applications of this technique in the following subsections.

1.1.5 Detecting Tensor Perturbations

As discussed in subsection 1.1.1, inflation has become an important addition to the standard big-bang paradigm because of its ability both to solve the flatness and horizon problems and to generically predict a scale-invariant spectrum of primordial perturbations. Quantum fluctuations in the scalar field responsible for inflation served as the seeds of structure formation in the early Universe, but it had long been realized that quantum fluctuations in the spacetime metric were "frozen in" during inflation as well. These primordial gravitational waves decayed in amplitude with the scale factor once they re-entered the horizon, but they may have left a detectable signature in the CMB through the Sachs-Wolfe effect. Abbott & Wise (1984) calculated the power spectrum of CMB temperature perturbations generated in this process and showed that it scaled as the fourth power of the energy scale of inflation. Starobinskii (1985) compared the relative contributions of adiabatic density (scalar) perturbations and gravitational wave (tensor) perturbations to primordial CMB anisotropies. He found that the tensor-to-scalar ratio,

$$\delta \equiv \left(\left\langle \left(\frac{\Delta T}{T} \right)^2_{lm} \right\rangle_g \Big/ \left\langle \left(\frac{\Delta T}{T} \right)^2_{lm} \right\rangle_{ad} \right)^{1/2} , \tag{1.18}$$

scaled as $n^{1/2}$ for inflationary potentials $V(\Phi) \propto \Phi^n$, with $\delta \approx 0.5$ in an $n = 4$ model. The contributions of both scalar and tensor modes to CMB polarization was considered by Crittenden et al. (1993), where for a tensor-to-scalar ratio of unity, the tensor contribution was shown to dominate for $l \lesssim 40$. Crittenden et al. (1993) are not optimistic about the prospects of detecting CMB polarization, as they predicted that even on large angular scales ($\gtrsim 1^\circ$), the net polarization would be less than 1% of $\Delta T/T$. Early reionization could boost the net polarized signal to as much as 9%, but in this case the tensor contribution becomes subdominant and could only be distinguished with sensitivities better than 1%. It appeared that even indirect detection of primordial gravitational waves might present an insurmountable observational challenge.

Seljak & Zaldarriaga (1997) and Kamionkowski et al. (1997a) gave renewed hope to efforts aimed at detecting primordial tensor modes by noting that the CMB polarization could be uniquely decomposed into $E-$ and $B-$modes as in Eq. (1.17). Only tensor perturbations source B-mode polarization, so in principle, a nonzero detection of this polarization would be an unambiguous signature of primordial gravitational waves in the early Universe. An explicit calculation of the power spectra and real-space correlation functions of the CMB polarization in the small-scale limit was given by Seljak (1997), along with an analysis of the instrumental sensitivity, resolution, and sky coverage necessary to reconstruct these power spectra in an Einstein-de Sitter CDM Universe. The formalism underlying studies of CMB polarization in the all-sky limit was presented by Kamionkowski et al. (1997b), including an explicit construction of estimators for the power spectra and a calculation of the covariance matrix of these estimators for Gaussian pixel noise. The connection between CMB polarization and inflationary observables was explored more fully in Kamionkowski

& Kosowsky (1998). Explicit relations were derived between the inflaton potential $V(\phi)$, its derivative $V' \equiv \partial V/\partial\phi$, and the amplitudes and spectral indices of scalar and tensor perturbations,

$$\begin{aligned} \mathcal{S} &\equiv 6C_2^{\mathrm{TT,scalar}} = 33.2\left[\frac{V^3}{(V')^2}\right] \\ \mathcal{T} &\equiv 6C_2^{\mathrm{TT,scalar}} = 9.2V \\ 1-n_s &= \frac{1}{8\pi}\left(\frac{V'}{V}\right)^2 - \frac{1}{4\pi}\left(\frac{V'}{V}\right)' \\ n_t &= -\frac{1}{8\pi}\left(\frac{V'}{V}\right)^2 . \end{aligned} \tag{1.19}$$

These relations imply a critical consistency test, $\mathcal{T}/\mathcal{S} \simeq -7n_t$, which could in principle be tested by observations of the CMB temperature and polarization. Kamionkowski & Kosowsky (1998) examined the predicted tensor amplitude $\mathcal{T}$ for various models of $V(\phi)$, and found that while some models led to unobservably small tensor perturbations, others generated perturbations that could be detected for experimental sensitivities of a few $\mu\mathrm{K}\sqrt{\mathrm{sec}}$. Once the tensor perturbations have been detected, their spectral index n_t could be measured from the B-polarization power spectrum and the consistency test could definitively exclude a wide class of inflationary models. Jaffe et al. (2000) extended this analysis to experiments with arbitrary sensitivity, resolution, and sky coverage and found somewhat surprisingly that for a fixed-time experiment, tighter constraints could be placed on $\mathcal{T}$ by sacrificing sky coverage to integrate more deeply on a $2-5^\circ$ patch. Ground-based efforts to detect CMB polarization in surveys of limited size are currently underway, and E-mode polarization has already been observed in this way by the Degree Angular Scale Interferometer (DASI) (Kovac et al., 2002). Improving detector technology will continue to lower the upper bounds on B-mode polarization, and a true detection may occur in the near future if the energy scale of inflation is associated with grand unification $\sim 10^{16}$ GeV.

Unfortunately, real experiments are faced with a variety of astrophysical foregrounds in addition to instrumental noise and limited sky coverage. Tucci et al. (2004) found that Galactic foregrounds (synchrotron emission for $\nu \lesssim 70$ GHz, dust emission at higher frequencies) dominate primordial B-mode polarization at all frequencies for even the most optimistic assumptions about the tensor-to-scalar ratio. However, with a template for foreground emission $\tilde{I}_{\nu_t}(\hat{\mathbf{n}})$ at a frequency ν_t and a position-dependent spectral index $\tilde{\beta}(\hat{\mathbf{n}})$, Tucci et al. (2004) claimed that linear subtraction could reduce these foregrounds by almost two orders-of-magnitude. More troubling is the residual polarization produced by extragalactic radio sources, which then become dominant on subdegree scales ($l \gtrsim 100$). Subtraction of these point sources involves a fundamental tradeoff; an increasingly larger fraction of the total pixels must be masked to remove sources of lower flux. Removing sources with $S > 150$ mJy led to a lower bound of $\sim 10^{-4}$ on the detectable value of the tensor-to-scalar ratio. Yet another astrophysical foreground to consider is the B-mode polarization induced by weak gravitational lensing. Zaldarriaga & Seljak (1998) first studied the effect of lensing on CMB polarization on small scales, and discovered that lensing converts a fraction of E-mode polarization to B-modes and vice versa. They expressed the observed polarization power spectra in terms of the unlensed quantities

$$\begin{aligned}
C_{El} &= C_{\tilde{E}l} + \frac{1}{2}[\mathcal{W}_{1l}^{l'} + \mathcal{W}_{2l}^{l'}]C_{\tilde{E}l'} + \frac{1}{2}[\mathcal{W}_{1l}^{l'} - \mathcal{W}_{2l}^{l'}]C_{\tilde{B}l'} \\
C_{Bl} &= C_{\tilde{B}l} + \frac{1}{2}[\mathcal{W}_{1l}^{l'} - \mathcal{W}_{2l}^{l'}]C_{\tilde{E}l'} + \frac{1}{2}[\mathcal{W}_{1l}^{l'} + \mathcal{W}_{2l}^{l'}]C_{\tilde{B}l'}
\end{aligned} \tag{1.20}$$

and the window functions $\mathcal{W}_{1l}^{l'}$ and $\mathcal{W}_{2l}^{l'}$ which are linear in the power spectrum of the lensing potential. This fractional conversion could cause the primordial E-modes to masquerade as a signal of tensor perturbations if not properly subtracted. The incomplete subtraction of lensing-induced B-modes using quadratic lensing estimators is the subject of Kesden et al. (2002b), Chapter 4 of this thesis. A similar calculation with similar results was performed independently by Knox & Song (2002). Continued improvements in lensing reconstruction should lead to improved subtraction of lensing-induced B-modes, and ultimately tighter constraints on inflation.

1.1.6 Low CMB Quadrupole

Lensing reconstruction has a variety of potential applications to cosmology beyond its use in probing primordial B-mode polarization. One question on which information from lensing reconstruction could possibly be decisive is the existence of a cutoff in the primordial density perturbation power spectrum on horizon scales. The CMB temperature quadrupole first measured by COBE (Hinshaw et al., 1995) was found to be unusually low for a model with a scale-invariant power spectrum, and subsequent WMAP observations showed that the $l = 3$ and $l = 4$ moments were below predictions as well (Bennett et al., 2003). The statistical significance of this discovery is a matter of considerable debate. The WMAP collaboration concluded in Spergel et al. (2003) that only 0.7% of realizations of a scale-invariant power spectrum yielded smaller values of the large-angle statistic

$$S \equiv \int_{-1}^{1/2} [C(\theta)]^2 d(\cos\theta)\,, \tag{1.21}$$

where $C(\theta)$ is the real-space correlation function. Tegmark et al. (2003) re-analyzed the CMB foregrounds observed by WMAP, resulting in a different Galactic cut and somewhat larger low l multipoles than determined by the WMAP collaboration. Gaztanaga et al. (2003) also re-examined the WMAP analysis of large-scale anisotropies, and found that real-space statistics like the correlations function and its moments were more informative than spherical harmonics for probing the largest scales. They found that WMAP observations were fully consistent with a scale-invariant power spectrum, as fully 32% of realizations yielded fits with higher χ^2. This result was highly sensitive to the choice of Galactic cut, in keeping with the discrepancies between Spergel et al. (2003) and Tegmark et al. (2003). One way of interpreting the low CMB multipoles is as a large-scale cutoff k_c in the density perturbation power spectrum $P(k)$. Several authors (Bridle et al., 2003; Contaldi et al., 2003; Cline et al., 2003) studied the WMAP and 2 Degree Field Galaxy Redshift Survey (2DFGRS) data and found that models with cutoffs of $k_c \simeq 3 - 5$ Mpc^{-1} were preferred, though the absence of a cutoff was allowed at the 2σ level. Cline et al. (2003) found that no values of k_c fit the data particularly well, as I also concluded in Kesden et al. (2003), Chapter 5 of this thesis. Cutoffs in $P(k)$ large enough to account for the WMAP results tend to bleed into higher multipole moments producing excessive suppression above $l = 4$. Efstathiou (2003a) compared the frequentist and Bayesian approaches employed by different authors in their statistical analyses, and found that estimates like those of Spergel et al. (2003) were gross overestimations. He suggested that 1 in 10 or 1 in 20 was a more reasonable estimate of the unlikeliness of the WMAP results.

If the low CMB multipoles observed by WMAP are truly significant, what kinds of new physics might they imply? One curiosity first noted by Tegmark et al. (2003) was that the planes of the CMB quadrupole and octupole appear to be aligned. Schwarz et al. (2004) further asserted that several of these planes were aligned at greater than 99% confidence levels with the ecliptic, the equinoxes, and the supergalactic plane. This would imply that the CMB quadrupole and octupole are not cosmological in origin, but instead, result from some local source or unanticipated systematic error. The statistical significance of these results has not been confirmed by other authors, but unknown systematics should always be considered, particularly in light of the previously determined sensitivity to the choice of Galactic cut. Other explanations for the unusually low CMB multipoles tend to invoke new physics on cosmological scales. DeDeo et al. (2003) proposed that the low CMB multipoles result from a suppression of the late-time integrated Sachs-Wolfe (ISW) effect expected in scale-invariant ΛCDM models. Gravitational potential wells will evolve with redshift for $\Omega_m \neq 1$, implying that the depths of these wells will change between the times that CMB photons enter and exit them. This will impart a gravitational redshift to these photons and corresponding perturbation power to the CMB. In the ΛCDM concordance model, the cosmological constant becomes dominant at late times, driving $\Omega_m < 1$ and causing large-scale potential fluctuations to decay away. This induces enhanced power in the CMB quadrupole and octupole through the ISW effect. However, if the late-time acceleration is driven not by a cosmological constant, but some form of dark energy with sound speed $c_s < 1$, this dark energy will cluster and inhibit the decay of large-scale potential wells. The late-time ISW effect and CMB multipoles will be accordingly suppressed. Other proposals to reduce the low-l CMB multipoles require a suppression of large-scale power in the early Universe. As noted earlier in this Introduction, Turner (1991) calculated the minimum number of e-foldings of inflation necessary to solve the flatness and horizon problems. If inflation lasted only ~ 10 e-foldings longer than strictly necessary, the CMB quadrupole and octupole could be sensitive to density perturbations in the pre-inflationary Universe. In this scenario, the low CMB multipoles would be indicative of an extremely homogeneous Universe prior to inflation. Primordial perturbations could also be suppressed during inflation itself, either by modifying the inflaton potential (Contaldi et al., 2003) or introducing a second scalar field (Cline et al., 2003; Feng & Zhang, 2003). Both of these theories introduce a characteristic cutoff scale k_c that violates the scale-invariance predicted by a single, slowly rolling field. Efstathiou (2003b) associated this scale with the curvature scale, and showed that the CMB power spectrum could be made consistent with nonzero curvature by fixing the physical baryon and cold dark matter densities $\omega_b \equiv \Omega_b h^2$, $\omega_c \equiv \Omega_c h^2$ and allowing the Hubble constant h to vary. The low CMB multipoles were suppressed while preserving the locations and heights of the acoustic peaks. Clearly, there is no shortage of theoretical explanations for the small CMB quadrupole and octupole.

How might one distinguish between these explanantions? One hope is that these theories might have other observable consequences besides the suppression of large-scale power. DeDeo et al. (2003) claimed that a variable sound speed for the dark energy would lead to possibly observable oscillations in the perturbation power spectrum on smaller scales as well ($k > 10^{-2} h\,\mathrm{Mpc}^{-1}$). Gurzadyan et al. (2005) argued that spatial curvature could account for both the low CMB multipoles and the degree-scale ellipticity of anisotropies observed in BOOMERanG CMB maps. However, many theories do not provide readily testable supplemental predictions. On the largest scales, the measured error in the CMB power spectrum C_l is already dominated by cosmic variance. There are only $2l+1$ spherical harmonic coefficients a_{lm} at each l, implying that the estimator

$$\hat{C}_l \equiv \frac{1}{2l+1} \sum_{m=-l}^{l} |a_{lm}|^2 \tag{1.22}$$

will have a variance $2C_l/(2l+1)$ even in the absence of foregrounds and instrumental noise. Only by somehow gaining access to additional coefficients a_{lm} on large scales can the noise be driven below the limit imposed by cosmic variance. Kamionkowski & Loeb (1997) proposed one method to do this: by examining the CMB polarization induced by the thermal SZ effect. The magnitude and orientation of this polarization is directly sensitive to the CMB temperature quadrupole seen by the clusters responsible for the SZ effect. For clusters at significant redshift, this quadrupole will be appreciably different from that observed on Earth, providing access to perturbations not visible in the primordial CMB. Kesden et al. (2003), Chapter 5 of this thesis, proposes a new way of observing power on large scales. The lowest multipole moments of the lensing potential – reconstructed via quadratic estimators as described above – can provide a partially independent probe of the same large-scale modes responsible for the low CMB quadrupole. By examining the cross-correlation of the lensing potential with CMB temperature perturbations, I show that lensing reconstruction can distinguish a $k_c \simeq 5.0 \times 10^{-4}\,\mathrm{Mpc}^{-1}$ cutoff from a scale-invariant spectrum at $> 3\sigma$. This approach offers one of the only ways of laying to rest the controversy over the CMB quadrupole.

1.2 Binary Inspirals into Horizonless Objects

In the remainder of this Introduction, I provide an overview of previous work relating to the gravitational-wave (GW) signature of nonstandard binary inspirals, the subject of Kesden et al. (2005), Chapter 6 of this thesis.

1.2.1 Black-Hole Inspirals

Interest in the gravitational radiation produced by accelerating point particles dates back to Einstein (1918), where it was proven in linearized theory that the GWs produced by weakly gravitating, nonrelativistic systems was quadrupolar to lowest order. For a very long time thereafter, there was a lack of consensus in the scientific community as to whether gravitational radiation was a real physical phenomenon or simply an artifact of the theory. It was not immediately apparent how to construct a stress-energy tensor for the gravitational field, and different authors obtained qualitatively different results depending on their choice of coordinate system. Peters & Mathews (1963) finally determined the power radiated in GWs by two point masses moving in an elliptical Keplerian orbit. In the weak-field limit, the spacetime metric can be linearized about the Lorentz metric $\delta_{\mu\nu}$,

$$g_{\mu\nu} = \delta_{\mu\nu} + \kappa h_{\mu\nu} \qquad (|\kappa h_{\mu\nu}| \ll 1)\,, \tag{1.23}$$

where $\kappa^2 \equiv 32\pi G$. The Einstein equation for this metric implies that the trace-reversed metric perturbation

$$\bar{h}_{\mu\nu} \equiv h_{\mu\nu} - 1/2\delta_{\mu\nu}h_{\lambda\lambda} \tag{1.24}$$

obeys the wave equation

$$\Box\bar{h}_{\mu\nu} = -1/2\kappa T_{\mu\nu}\,, \tag{1.25}$$

where $T_{\mu\nu}$ is the stress-energy tensor of the GW source. Peters & Mathews (1963) then expanded both $\bar{h}_{\mu\nu}$ and $T_{\mu\nu}$ into multipoles and kept only the dominant "magnetic quadrupole" type of radiation. Defining the mass quadrupole tensor of the system as

$$Q_{ij} \equiv \sum_{\alpha} m_\alpha x_{\alpha i} x_{\alpha j}\,, \tag{1.26}$$

where the sum α is taken over all particles, the total power radiated is

$$P = \frac{G}{5c^5}\left(\frac{d^3Q_{ij}}{dt^3}\frac{d^3Q_{ij}}{dt^3} - \frac{1}{3}\frac{d^3Q_{ii}}{dt^3}\frac{d^3Q_{jj}}{dt^3}\right)\,. \tag{1.27}$$

This formalism was used to determine the average power emitted over an orbital period by point masses m_1 and m_2 in a Keplerian orbit with semi-major axis a and ellipticity e,

$$\langle P\rangle = \frac{32}{5}\frac{G^4}{c^5}\frac{m_1^2 m_2^2(m_1+m_2)}{a^5(1-e^2)^{7/2}}\left(1+\frac{73}{24}e^2+\frac{37}{96}e^4\right) , \tag{1.28}$$

and to calculate the power radiated into each harmonic of the fundamental frequency $\omega_0 = [G(m_1+m_2)/a^3]^{1/2}$. This analysis is essentially the same approach I use in Chapter 6 to calculate the energy and angular momentum losses, although I generalize to non-Keplerian orbits in the strong-field limit. Peters (1964) generalized his previous results for arbitrary metric perturbations $\bar{h}_{\mu\nu}$, and determined the angular momentum loss as well. A definitive treatment of the multipole expansion of the gravitational field was given by Thorne (1980), where the energy, angular momentum, and linear momentum fluxes were expanded in terms of the time derivatives of the source multipole moments.

The multipole expansions described above are valid in the limit of linearized gravity as in Eq. (1.23), but a proper treatment in the strong-field limit must include perturbations about a background black-hole spacetime. These perturbations were first investigated by Regge & Wheeler (1957), where arbitrary perturbations of the Schwarzschild spacetime were decomposed into tensor spherical harmonics. Only perturbations that maintain the vacuum were allowed, leading to the second-order differential equation

$$\delta R_{\mu\nu} = -\delta\Gamma^{\beta}_{\mu\nu;\beta} + \delta\Gamma^{\beta}_{\mu\beta;\nu} , \tag{1.29}$$

where the perturbed Christoffel symbol is

$$\delta\Gamma^{\alpha}_{\beta\gamma} = \frac{1}{2}g^{\alpha\nu}(h_{\beta\nu;\gamma} + h_{\gamma\nu;\beta} - h_{\beta\gamma;\nu}) . \tag{1.30}$$

The eigenmodes of this differential equation with imaginary frequencies are forbidden by the boundary conditions at infinity and the black-hole horizon, implying that the Schwarzschild black hole is stable against arbitrary small perturbations. Zerilli (1970) considered how these perturbations might be sourced by a point mass traveling on a geodesic of the background Schwarzschild spacetime, and obtained radial differential equations for the amplitudes of the outgoing GWs. Teukolsky (1973) generalized this program to the spacetimes of rotating black holes discovered by Kerr (1963). To do so, he made use of the Newman-Penrose formalism developed by Newman & Penrose (1962), in which all tensor quantities are projected along a tetrad of null vectors l_μ, m_μ, $\bar{m}_\mu$, and n_μ. This projection reduces various identities involving the Riemann tensor to a large number of scalar equations, many of which vanish for background metrics of Petrov Type D which include both Schwarzschild and Kerr. Teukolsky (1973) derived a single master equation involving Newman-Penrose quantities whose solutions provided a complete orthogonal basis for arbitrary black hole perturbations. Techniques for solving this equation were discussed in Teukolsky & Press (1974), along with its application to the scattering of GWs incident on a black hole. A class of transformations derived by Sasaki & Nakamura (1982) reformulated the Teukolsky equation with short-range potential and source function, making its solution more numerically feasible. Although this Teukolsky-Sasaki-Nakamura (TSN) formalism is the rigorously correct way to calculate GWs in a strong-field spacetime, Babak et al. (2005) showed empirically that the more simple, multipole-moment expansion was sufficiently accurate for qualitative studies of binary inspirals, and I will use this approach in Chapter 6.

The approval of the Laser Interferometer Gravitational-wave Observatory (LIGO) in the early 1990's motivated extensive theoretical effort to calculate accurate gravitational waveforms. These waveforms could be used as templates with which detector output could be cross-correlated in a matched-filtering scheme. This approach was outlined in Poisson (1993), where an attempt was made to calculate such waveforms using the TSN formalism described above. Restricting his attention to circular inspirals in the extreme mass-ratio limit $\mu \ll M$, Poisson determined corrections to the quadrupole-approximation energy flux and waveform to order $(v/c)^3$. He also recognized that throughout the binary inspiral – even including the last few orbits prior to the plunge into the event horizon – the radius of the inspiraling particle is decaying adiabatically. This is equivalent to the statement that the orbital time scale is always much shorter than the inspiral time scale. In the approximation that the ratio of these two time scales is infinitely small, the inspiraling particle can be assumed to travel along a geodesic and the waveforms calculated by Poisson can be stitched together over the entire duration of the inspiral. By comparing detector output to pre-determined templates over the entire course of the inspiral, enough signal-to-noise can be accumulated to accurately estimate all relevant source parameters such as the particle masses, orbital eccentricity and inclination, etc. This program for calculating waveform templates was pursued by many authors following Poisson (1993), with an important generalization to the orbits of Kerr black holes being investigated by Hughes (2000). The geodesics of Kerr black holes with spin a are most generally described by three constants of the motion: the energy E, z-component of the angular momentm L_z, and Carter constant

$$Q = p_\theta^2 + \cos^2\theta[a^2(1 - E^2) + L_z^2 \csc^2\theta]\,. \tag{1.31}$$

Unlike changes in the energy and angular momentum, the change in an inspiraling particle's Carter constant cannot be determined from the losses to gravitational radiation, even in the adiabatic limit. The Carter constant is defined from a Killing tensor of the Kerr spacetime, $Q \equiv Q_{\mu\nu}p^\mu p^\nu$. The quadratic dependence of the Carter constant on the inspiraling particle's four-momentum implies that its rate of change is determined by the detailed time dependence of the true radiation-reaction force $f^\nu_{\rm RR}$ over each orbital period,

$$\dot{Q} = 2Q_{\mu\nu}p^\mu f^\nu_{\rm RR}\,. \tag{1.32}$$

Hughes (2000) showed that since circular orbits remain circular as GWs are emitted, the flux in Carter constant could be expressed in terms of the energy and angular momentum fluxes, $\dot{Q} = \dot{Q}(\dot{E}, \dot{L}_z)$. This constraint allowed Hughes to determine how the constants $\{E, L_z, Q\}$ evolve during a binary inspiral from the emitted GWs calculated in the TSN formalism. In Hughes (2001), the final few years of the binary inspirals of particles on inclined, circular orbits are investigated using this formalism, and the waveforms and evolution of radius and inclination angle are determined. The alternative generalization – to orbits in the equatorial plane of a Kerr black hole, but with nonzero eccentricity – was considered by Glampedakis & Kennefick (2002). In this case, the Carter constant identically vanishes, so the time evolution of the semilatus recturm p and eccentricity e can again be determined from the energy and angular momentum flux in the adiabatic limit. Glampedakis & Kennefick (2002) found that contrary to earlier expectations, inspiraling test particles can maintain appreciable eccentricity all the way down to the plunge into the event horizon, and that on the final orbits, the eccentricity will actually increase prior to plunge. Calculating the binary inspiral in the most general eccentric, inclined case is still an open research question. Such a calculation will most likely require an accurate determination of the true radiation-reaction force, which can in principle be derived from the nonsingular component of the local gravitational field at the position of the inspiraling particle. Poisson (2004) provides a current review of efforts to calculate the radiation reaction in the self-force formalism.

1.2.2 Alternative Compact Objects

Although substantial progress has been made in understanding binary inspirals into black holes, much less attention has been given to the more general case of inspirals into some arbitrary compact object. Current observations are fully consistent with objects much less compact than black holes residing at galactic centers. Such objects are generically known as boson stars if they are composed of fundamental particles with integer spins (bosons). Typical boson stars consist of scalar particles supported against gravity by a variety of model-dependent mechanisms. Extremely light boson stars can be supported against collapse purely by the Heisenberg uncertainty principle, while more massive boson stars rely on net rotation or a repulsive self-interaction. Although no massive scalar fields have yet been discovered, they are an essential ingredient of the Standard Model of particle physics and its most popular extensions (Torres et al., 2000). Within the Standard Model, the Higgs boson plays a vital role in generating masses for the Z^0 and $W^{\pm}$ gauge vector bosons. All particle masses implicitly depend on the Higgs mass through loop corrections; in particular, the observed top quark mass $m_t = 173.8 \pm 5$ GeV places an indirect limit on the Higgs mass of $m_h = 104^{+93}_{-49}$ GeV (Djouadi et al., 1998). Extensions of the Standard Model naturally introduce additional massive scalar fields. Supersymmetry, designed to solve the hierarchy problem of particle masses, posits the existence of a bosonic partner for each fermion in the Standard Model. Should one of these many bosons be stable, its expected mass would lie near the supersymmetry-breaking scale $m_{\rm susy} \sim 1$ TeV. Massive scalars also arise in string effective supergravity, where the dilaton mass is expected to be $m_\varphi \simeq 10^{-3}\ (m_{\rm susy}/\ {\rm TeV})^2$ eV. Yet another massive scalar found in particle theory is the axion, the Goldstone boson of Peccei-Quinn symmetry breaking in QCD. Though nonrelativistic because of its nonthermal production, the axion is anticipated to be extremely light, $m_\sigma \gtrsim 10^{-11}$ eV (Torres et al., 2000). Clearly, an abundance of boson candidates are suggested by high-energy physics, providing ample motivation to consider the possibility of a boson star in an astrophysical context.

Despite the twenty-three orders of magnitude spanned by the masses of the bosons proposed above, different self-interactions can lead to boson stars of ($\sim 10^6 M_\odot$) for all these candidates. Torres et al. (2000) classified all models as either mini-boson stars, boson stars, or soliton stars, depending upon their means of support against gravitational collapse. The simplest case is that of mini-boson stars, where the massive scalar field is free and thus, the size is determined by the Compton wavelength (Friedberg et al., 1987a). Mini-boson stars can attain masses up to those at which the Compton wavelength equals the Schwarzschild radius. This relation sets an upper bound on the mini-boson star mass of $M \sim m_{\rm Pl}^2/m$. A boson mass as small as 10^{-16} eV can lead to a mini-boson star of $10^6 M_\odot$. The mini-boson star scenario thus allows extremely light boson masses such as that of the axion to produce compact objects with masses of astrophysical interest. Scalar fields for more generic boson stars are not free; a repulsive interaction potential supplements the Heisenberg uncertainty principle to support the boson star against gravitational collapse. The simple interaction potential $U(\phi) = 1/4\lambda|\phi|^4$ leads to an increase in the boson star radius by a factor $\Lambda^{1/2} \equiv \lambda^{1/2} m_{\rm Pl}/m$ over that of the mini-boson star (Colpi et al., 1986). The

Schwarzschild radius of a compact boson star can increase by a similar factor, implying for the dimensionless coupling $\lambda \sim 1$ an upper bound $M \sim m_{\rm Pl}^3/m^2$ on the boson star mass. Boson stars with masses comparable to Sgr A* would be produced by intermediate mass bosons $m \simeq 1$ MeV. The soliton star – the third and final type of boson star – is characterized by a restricted class of interaction potentials for which bound, stable solutions exist even in the absense of gravity (Lee & Pang, 1992). Without gravity, the potential must have an attractive component for a bound solution to be energetically favored over free particles. Yet the potential must be repulsive at large field values for the vacuum to be stable. The simplest potential that satisfies these criteria is an even sixth-order polynomial with a negative fourth-order term (Friedberg et al., 1987b). Compact soliton stars in this model have masses less than $M \sim m_{\rm Pl}^4/m^3$. This limit is enhanced by a factor $m_{\rm Pl}/m$ over that for a $\lambda|\phi|^4$ potential, requiring a yet more massive boson ($m \simeq 10$TeV) to produce a $10^6 M_\odot$ compact object. As this boson mass is closest to that of the Higgs and the supersymmetry scale, I used this particular model for quantitative comparisons in Chapter 6. This restriction is not constraining, since all boson stars of a given mass should have similar astrophysical properties.

1.2.3 Signatures of New Geometry

How might one distinguish the GWs produced during inspirals into black holes and boson stars? This project was first seriously considered by Ryan (1995), who restricted his attention to the vacuum spacetime outside axisymmetric bodies with reflection symmetry about the equatorial plane. In this case, the spacetime can be fully specified by two infinite series of multipole moments. The use of multipole moments to characterize vacuum metrics was first proposed by Geroch (1970), who proved that an arbitrary static, asymptotically flat, vacuum spacetime could be described by a series of symmetric, trace-free tensors $P, P_a, P_{ab}, \ldots$ at spatial infinity. This approach had the advantages that it was coordinate-independent and offered a ready classical analogy with Newtonian gravitational fields with the same multipole moments. It also appeared to be physically meaningful in that the Schwarzschild solution corresponded to the spacetime with vanishing multipole moments beyond the monopole. Hansen (1974) extended this approach to stationary spacetimes, for which two sets of multipole moments, the mass and angular momentum moments, were required for full generality. He then showed that axisymmetry constrained the tensor structure of these multipole moments, so that all the information could be contained in two series of *scalar* multipole moments. Reflection symmetry about the equatorial plane then required that the odd mass and even angular momentum multipoles must vanish. Hansen calculated the multipole moments of the Kerr spacetime as a concrete example of this formalism. Ryan (1995) applied these results to the calculation of the GW spectrum produced by the binary inspiral of a test particle along a circular, equatorial orbit of such an axisymmetric vacuum spacetime. He derived a power-series expansion in the post-Newtonian parameter $v/c = (\pi GMf/c^3)^{1/3}$ for $\Delta E(f)$, the energy E emitted per logarithmic interval of the frequency f. The coefficients of this expansion are the multipole moments of the spacetime of the central object, so in principle, if the GW spectrum was observed, one could simply read off the multipole moments. The "no hair" theorem implies that all black holes are described by the Kerr geometry, a two-parameter family characterized by the mass and spin. The first two nonvanishing multipole moments uniquely constrain these quantities, so the measurement of higher multipole moments should be able to definitively determine whether the central object is indeed a black hole. Ryan (1997a) applied this formalism to a rapidly spinning boson star with a quartic interaction potential $U(\phi) = 1/4\lambda|\phi|^4$ as described by Colpi et al. (1986). The rotation induced a mass quadrupole moment M_2 in the spinning boson star much larger than that of a Kerr black hole of equivalent mass M and spin S_1. The particular boson star studied by Ryan had $-M_2M/S_1^2 \simeq 24$, in stark contrast to the corresponding value $-M_2M/S_1^2 = 1$ for a Kerr black hole. Such a difference would be manifestly apparent to a GW experiment with sufficient signal-to-noise. The issue of how well a real experiment such as LIGO or the Laser Interferometer Space Antenna (LISA) could actually measure the mulitpole moments was addressed in Ryan (1997b). He found that LIGO would be largely unable to measure the higher multipole moments because of the relatively small number ($\sim 11,000$) of cycles between when the inspiral enters the LIGO band and the final plunge. In contrast, LISA would observe $\sim 4.2 \times 10^5$ cycles over a two-year period and could thus constrain S_1 and M_2 to .13% and 1.5%, respectively for inspirals with $S/N = 10$.

Recent work has focused on the practical difficulties in implementing Ryan's program for determining multipole-moment structure, and on other complementary signals of departures from general relativity. One such difficulty is the lack of model-independent alternatives to black holes with which one might compare the standard black-hole results. Collins & Hughes (2004) developed a formalism for determining the spacetimes of *bumpy* black holes: objects with atypical multipole-moment structures that will reduce to Kerr black holes in a natural limit. This is in contrast to the specific boson-star models employed by Ryan (1997a) and myself in Chapter 6 that are fundamentally distinct objects that cannot be continuously perturbed into black holes. The bumpy black holes proposed by Collins and Hughes are well defined in the strong-field region, making them suitable for calculating the GWs produced during the final orbits before plunge. Another new development is the investigation by Dreyer et al. (2004) of the quasi-normal modes (QNMs) associated with the ringdown of a Kerr black hole immediately following a merger event. The frequency and damping time of *each* mode constrain the mass and spin of the black hole, so measurements of a *spectrum* of different modes from a single merger would overdetermine these parameters. This amounts to a consistency test of the "no hair" theorem, violations of which would indicate deviations from general relativity or that the object was not a true black hole. Dreyer et al. (2004) found that the ringdowns of massive black-hole mergers expected to be observed by LISA should have enough signal-to-noise to provide such a test. Fang & Lovelace (2005) proposed a second effect that might differentiate true black holes from other massive compact objects. The tidal field $\mathcal{E}_{ij}$ of the inspiraling test particle should induce a time-varying mass quadrupole moment

$$\mathcal{I}_{ij} = \frac{32}{45} M^6 \dot{\mathcal{E}}_{ij} \tag{1.33}$$

in the central black hole, which should both emit GWs itself and also exert a force on the inspiraling particle, altering its trajectory. Such effects could, in principle, be identified in sufficiently accurate measurements of the GW spectrum. Alternative compact objects like boson stars should have very different induced mass quadrupole moments, providing a handle by which they could be distinguished from black holes. In Kesden et al. (2005), Chapter 6 of this thesis, I examine yet another such handle: the lack of an event horizon for compact objects like boson stars. An inspiraling test particle interacting purely gravitationally with the compact objects could continue to emit GWs even after falling within the innermost stable orbit. I calculate the waveform produced immediately after the fall, and discuss how it is qualitatively different from any conceivable black-hole waveform. These new tests of black-hole spacetimes, when brought to bear on the signals that will hopefully be detected by LIGO, LISA, and other GW observatories in the near future, offer the hope of greatly advancing our knowledge of general relativity and new physics.

Chapter 2

Lensing Reconstruction with CMB Temperature and Polarization

Reprinted with permission from M. Kesden, A. Cooray, and M. Kamionkowski Phys. Rev. D **67**, 123507 (2003).

Weak gravitational lensing by intervening large-scale structure induces a distinct signature in the cosmic microwave background (CMB) that can be used to reconstruct the weak-lensing displacement map. Estimators for individual Fourier modes of this map can be combined to produce an estimator for the lensing-potenial power spectrum. The naive estimator for this quantity will be biased upwards by the uncertainty associated with reconstructing individual modes; we present an iterative scheme for removing this bias. The variance and covariance of the lensing-potenial power spectrum estimator are calculated and evaluated numerically in a ΛCDM universe for Planck and future polarization-sensitive CMB experiments.

2.1 Introduction

The primordial cosmic microwave background (CMB) was generated when photons first decoupled from the baryonic fluid when the universe was only 400,000 years old. The vast majority of these photons travel unperturbed to the present day, and features of their angular power spectrum such as acoustic peaks and the damping tail (Peebles & Yu, 1970; Sunyaev & Zel'dovich, 1970; Silk, 1968; Hu & Dodelson, 2002) record valuable information about cosmological parameters (Jungman et al., 1996; Bond et al., 1997; Zaldarriaga et al., 1997; Eisenstein et al., 1999). Baryons and dark matter evolve from small inhomogeneities at decoupling into increasingly complicated large-scale structure which can subtly perturb the observed pattern of CMB anisotropies. Assuming that the primordial CMB is Gaussian, non-Gaussian correlations in the observed map can be used to reconstruct the intervening large-scale structure (Seljak & Zaldarriaga, 1999a; Zaldarriaga & Seljak, 1998). In addition to the importance of learning about the large-scale structure itself, reconstruction of the weak-lensing potential generated by structure is essential to constraining tensor perturbations. Weak lensing converts a fraction of the E-mode polarization generated by scalar perturbations at the last-scattering surface into B-mode polarization in the observed map. Only by subtracting this B-mode polarization can one conclusively detect the primordial B-modes which serve as a model-independent signal of tensor perturbations (Kesden et al., 2002b; Knox & Song, 2002).

Understanding lensing reconstruction requires a more detailed discussion of how weak lensing affects the CMB. Weak gravitational lensing deflects the paths of CMB photons as they travel from the last-scattering surface to the observer. This deflection is accomplished by a projected lensing potential which is a weighted line-of-sight integral of the gravitational potential between the observer and the surface of last-scattering. At each point on the sky, lensing remaps the temperature and polarization to that of a nearby point at the last-scattering surface, the deflection angle being the gradient of the aforementioned projected lensing potential. Assuming that this deflection angle is small, the temperature at any point can be expanded in a Taylor series in the gradient of the lensing potential. In Fourier space, this expansion appears as a series of convolutions of individual temperature and projected potential modes. The observed temperature-squared map in Fourier space also appears as a convolution of individual Fourier modes. Subject to an overall normalization dependent on the scale of the Fourier mode, these convolutions cancel in such a manner that each Fourier mode of the temperature-squared map acts as an estimator for the *same* Fourier mode of the projected lensing potential.

Lensing reconstruction as outlined above has been considered previously (Hu, 2001a; Hu & Okamoto, 2002). In these works, two sources of noise were identified, and a filter of the temperature-squared map in Fourier space was chosen to minimize the variance associated with lensing reconstruction subject to these noise sources. The first source is intrinsic signal variance; the observed large-scale structure is one arbitrary member of an ensemble of realizations allowed by theory. The second source of noise, endemic to this method of lensing reconstruction, is a consequence of the nature of the primordial CMB. Like the large-scale structure itself, the pattern of CMB anisotropies at the last-scattering surface is only one of many possible realizations allowed by theory. We do not know *a priori* which of these realizations nature has provided us, and this uncertainty hinders our ability to deconvolve the effects of lensing from true anisotropies at the last-scattering surface. Even if the true pattern of anisotropies at the last-scattering surface was known, the finite amount of power in the CMB at small scales would still constrain lensing reconstruction. Silk damping at the last-scattering surface suppresses CMB power at small scales, while the finite resolution of any real experiment would limit the detection of any signal that is present at small scales. Lensing reconstruction fails below scales at which there is sufficient power, for the same reason that any remapping is indistinguishable, given a uniform background.

Here, we consider a third source of noise neglected in previous studies. The filtered temperature-squared map is an unbiased estimator for the lensing potential in the approximation that a correlation between two given temperature modes is induced only by the single lensing mode whose wavevector is the sum of that of the two temperature modes. In actuality, any combination of two or more lensing modes whose wavevectors sum to this total induce correlations between the two temperature modes. There are many such combinations, but since these correlations add incoherently, we do not expect a systematic bias. Nonetheless, for estimators of each individual lensing mode, we must use our knowledge of other lensing modes to subtract off this unwanted bias. This is an iterative process, and since our knowledge of the lensing map is imperfect, it induces noise in lensing reconstruction. We calculate this additional variance for various estimators constructed from CMB temperature and polarization maps, and show how it compares to the dominant noise sources for an all-sky CMB experiment with a noise-equivalent temperature of 1 μK $\sqrt{\text{sec}}$. Since the lensing-potential power spectrum is a measure of the theoretical uncertainty with which we can predict the value of a given lensing mode, this noise associated with lensing reconstruction causes a systematic overestimation of the lensing-potential power spectrum. This systematic bias must be accounted for in order to compare observations with theoretical predictions.

This paper is organized as follows. In § 2.2, we define the formalism we will use to explore the effects of weak lensing on the CMB. The Taylor expansion of the lensed CMB map in gradients of the lensing potential is given in both real and Fourier space, and the power spectra and trispectra of various components of the CMB temperature map are listed for later use. In § 2.3, we show that the Fourier modes of the temperature-squared map, when properly filtered, can serve as estimators for the Fourier modes of the displacement map with the same wave vector. Using the power spectrum and trispectrum given in the preceding Section, we calculate the variance associated with this estimator including a new component neglected in previous studies. This variance is evaluated numerically using the currently favored ΛCDM cosmological model with baryon density $\Omega_b = 0.05$, cold dark matter density $\Omega_{cdm} = 0.30$, cosmological constant density $\Omega_\Lambda = 0.65$, Hubble parameter h=0.65, and power-spectrum amplitude $\sigma_8 = 0.9$. We then use the displacement estimator for individual Fourier modes to construct an unbiased estimator for the lensing-potential power spectrum in § 2.4, and calculate the variance and covariance associated with this estimator. A few concluding remarks about the implications of our work for future studies are given in § 2.5. The Appendix contains useful formulae related to additional estimators of lensing based on polarization and a combination of temperature and polarization.

2.2 Weak Lensing of the CMB

We consider weak lensing under the flat-sky approximation following Bernardeau (1997); Zaldarriaga (2000); Hu (2000a). As discussed before (Hu, 2000a; Spergel & Goldberg, 1999; Goldberg & Spergel, 1999; Cooray & Hu, 2000; Zaldarriaga & Seljak, 1999; Peiris & Spergel, 2000), weak lensing deflects the path of CMB photons resulting in a remapping of the observed temperature pattern on the sky

$$\begin{aligned}\tilde{\Theta}(\hat{\mathbf{n}}) &= \Theta[\hat{\mathbf{n}} + \nabla\phi(\hat{\mathbf{n}})] \\ &\approx \Theta(\hat{\mathbf{n}}) + \nabla_i\phi(\hat{\mathbf{n}})\nabla^i\Theta(\hat{\mathbf{n}}) + \frac{1}{2}\nabla_i\phi(\hat{\mathbf{n}})\nabla_j\phi(\hat{\mathbf{n}})\nabla^i\nabla^j\Theta(\hat{\mathbf{n}}) + \ldots \qquad (2.1)\end{aligned}$$

where $\Theta(\hat{\mathbf{n}})$ is the unlensed primary component of the CMB in a direction $\hat{\mathbf{n}}$ at the last scattering surface. The observed, gravitationally-lensed temperature map $\tilde{\Theta}(\hat{\mathbf{n}})$ in direction $\hat{\mathbf{n}}$ is that of the unlensed map in direction $\hat{\mathbf{n}} + \nabla\phi(\hat{\mathbf{n}})$ where $\nabla\phi(\hat{\mathbf{n}})$ represents the lensing deflection angle or displacement map. Although a real CMB map will include secondary contributions such as the SZ effect (Sunyaev & Zel'dovich, 1980), we assume that such effects can be distinguished by their frequency dependence (Cooray et al., 2000). They will not be further considered in this paper. A noise component denoted by $\Theta^{\rm n}(\hat{\mathbf{n}})$ due to finite experimental sensitivity must be included as well. Thus, the total observed CMB anisotropy will be $\Theta^{\rm t}(\hat{\mathbf{n}}) = \tilde{\Theta}(\hat{\mathbf{n}}) + \Theta^{\rm n}(\hat{\mathbf{n}})$.

Taking the Fourier transform of the lensed map $\tilde{\Theta}(\hat{\mathbf{n}})$ under the flat-sky approximation, we write

$$\begin{aligned}\tilde{\Theta}(\mathbf{l}) &= \int d\hat{\mathbf{n}}\, \tilde{\Theta}(\hat{\mathbf{n}}) e^{-i\mathbf{l}\cdot\hat{\mathbf{n}}} \\ &= \Theta(\mathbf{l}) - \int \frac{d^2\mathbf{l}'}{(2\pi)^2}\Theta(\mathbf{l}')L(\mathbf{l},\mathbf{l}')\,, \qquad (2.2)\end{aligned}$$

where

$$\begin{aligned}L(\mathbf{l},\mathbf{l}') \equiv\ & \phi(\mathbf{l}-\mathbf{l}')\ [(\mathbf{l}-\mathbf{l}')\cdot\mathbf{l}'] + \frac{1}{2}\int \frac{d^2\mathbf{l}''}{(2\pi)^2}\phi(\mathbf{l}'') \qquad (2.3) \\ & \times\phi(\mathbf{l}-\mathbf{l}'-\mathbf{l}'')\,(\mathbf{l}''\cdot\mathbf{l}')\,[(\mathbf{l}''+\mathbf{l}'-\mathbf{l})\cdot\mathbf{l}'] + \ldots\end{aligned}$$

CMB correlations in Fourier space can be described in terms of a power spectrum and trispectrum as defined in the usual manner

$$\begin{aligned} \left\langle \Theta^i(\mathbf{l}_1)\Theta^i(\mathbf{l}_2) \right\rangle &\equiv (2\pi)^2 \delta_{\rm D}(\mathbf{l}_1+\mathbf{l}_2) C_l^i \,, \\ \left\langle \Theta^i(\mathbf{l}_1)\ldots\Theta^i(\mathbf{l}_4) \right\rangle_c &\equiv (2\pi)^2 \delta_{\rm D}(\mathbf{l}_1+\mathbf{l}_2+\mathbf{l}_3+\mathbf{l}_4) T^i(\mathbf{l}_1,\mathbf{l}_2,\mathbf{l}_3,\mathbf{l}_4) \,, \end{aligned} \tag{2.4}$$

where the angle brackets denote ensemble averages over possible realizations of the primordial CMB, large-scale structure between the observer and the surface of last-scattering, and instrumental noise. The subscript c denotes the connected part of the four-point function, and the superscript i denotes the temperature map being considered ($\Theta^{\rm t}, \tilde{\Theta}$, or $\Theta^{\rm n}$). The lensing-potential power spectrum can be defined analogously as

$$\langle \phi(\mathbf{l})\phi(\mathbf{l}') \rangle_{\rm LSS} = (2\pi)^2 \delta_D(\mathbf{l}+\mathbf{l}') C_l^{\phi\phi} \,, \tag{2.5}$$

where the angle brackets denote an average over all realizations of the large-scale structure. We make the assumption that primordial fluctuations at the last-scattering surface are Gaussian. Gaussian statistics are fully described by a power spectrum; the Gaussian four-point correlator, $\langle \Theta(\mathbf{l}_1)\ldots\Theta(\mathbf{l}_4) \rangle_c$ is zero. The instrumental noise $\Theta^{\rm n}$ is also assumed to be Gaussian, as is the lensing potential ϕ. This second assumption is justified because the dominant contributions to the lensing potential come from intermediate redshifts $1 \lesssim z \lesssim 3$ at which linear theory holds.

Using these definitions, we can calculate the anticipated power spectrum and trispectrum of the observed CMB map. Because the instrumental noise is uncorrelated with the signal, the power spectrum of the observed map is the sum of signal and noise power spectra,

$$C_l^{\Theta\Theta{\rm t}} = \tilde{C}_l^{\Theta\Theta} + C_l^{\Theta\Theta{\rm n}} \,. \tag{2.6}$$

The power spectrum of the noise component is given by:

$$C_l^{\Theta\Theta{\rm n}} = f_{\rm sky} w^{-1} e^{l^2\sigma_b^2} \,, \tag{2.7}$$

where $f_{\rm sky}$ is the fraction of the sky surveyed, w^{-1} is the variance per unit area on the sky, and $\sigma_b = \theta/\sqrt{8\ln 2}$ is the effective beamwidth of the instrument expressed in terms of its full-width half-maximum resolution θ. A CMB experiment that spends a time $t_{\rm pix}$ examining each of $N_{\rm pix}$ pixels with detectors of sensitivity s will have a variance per unit area $w^{-1} = 4\pi(s/T_{\rm CMB})^2/(t_{\rm pix}N_{\rm pix})$ (Knox, 1995). The power spectrum of the lensed CMB can be determined by inserting Eq. (2.2) into Eq. (2.4) as discussed in Hu (2000a)

$$\tilde{C}_l^{\Theta\Theta} = \left[1 - \int \frac{d^2\mathbf{l}_1}{(2\pi)^2} C_{l_1}^{\phi\phi} \left(\mathbf{l}_1 \cdot \mathbf{l}\right)^2\right] C_l^{\Theta\Theta} + \int \frac{d^2\mathbf{l}_1}{(2\pi)^2} C_{|\mathbf{l}-\mathbf{l}_1|}^{\Theta\Theta} C_{l_1}^{\phi\phi} [(\mathbf{l}-\mathbf{l}_1)\cdot\mathbf{l}_1]^2 . \tag{2.8}$$

This result is given to linear order in the lensing-potential power spectrum $C_l^{\phi\phi}$. Lensing neither creates nor destroys power in the CMB, but merely shifts the scales on which it occurs as seen by the fact that

$$\tilde{\sigma}^2 = \int \frac{d^2\mathbf{l}}{(2\pi)^2} \tilde{C}_l^{\Theta\Theta} = \int \frac{d^2\mathbf{l}}{(2\pi)^2} C_l^{\Theta\Theta} = \sigma^2 . \tag{2.9}$$

The observed CMB trispectrum can be calculated in a similar manner; under our assumptions of Gaussian instrumental noise and no secondary anisotropies, the trispectrum of the lensed component $\tilde{\Theta}$ is the sole contribution to the total observed trispectrum

$$\begin{aligned} T^{\rm t}(\mathbf{l}_1,\mathbf{l}_2,\mathbf{l}_3,\mathbf{l}_4) &= \tilde{T}^{\Theta}(\mathbf{l}_1,\mathbf{l}_2,\mathbf{l}_3,\mathbf{l}_4) \\ &= -C_{l_3}^{\Theta\Theta} C_{l_4}^{\Theta\Theta} \Big[C_{|\mathbf{l}_1+\mathbf{l}_3|}^{\phi\phi} [(\mathbf{l}_1+\mathbf{l}_3)\cdot\mathbf{l}_3][(\mathbf{l}_1+\mathbf{l}_3)\cdot\mathbf{l}_4] \\ &\qquad + C_{|\mathbf{l}_2+\mathbf{l}_3|}^{\phi\phi} [(\mathbf{l}_2+\mathbf{l}_3)\cdot\mathbf{l}_3][(\mathbf{l}_2+\mathbf{l}_3)\cdot\mathbf{l}_4] \Big] + \text{Perm.} . \end{aligned} \tag{2.10}$$

The term shown above is manifestly symmetric under the interchange $\mathbf{l}_1 \leftrightarrow \mathbf{l}_2$, while the "+ Perm." represents five additional terms identical in form, but with the replacement of $(\mathbf{l}_1, \mathbf{l}_2)$ and $(\mathbf{l}_3, \mathbf{l}_4)$ with the other five combinations of pairs. The total trispectrum is symmetric under the interchange of any given pair as one would expect. Having established a formalism within which to analyze weak lensing, we now consider the problem of reconstructing the lensing potential from an observed CMB temperature map.

2.3 Lensing-Potential Estimators

In this Section, we examine lensing reconstruction following the approach of Hu & Okamoto (2002), largely adopting their notation as well. The only important difference in notation is that we use $\tilde{\Theta}$ to denote the lensed temperature field and Θ for the unlensed field following Hu (2000a) and most recent papers. Hu & Okamoto (2002) used the opposite convention. For $\mathbf{l} \neq -\mathbf{l}'$ and to linear order in ϕ,

$$\langle \Theta^{\rm t}(\mathbf{l}) \Theta^{\rm t}(\mathbf{l}') \rangle_{\rm CMB} = f_{\Theta\Theta}(\mathbf{l}, \mathbf{l}') \phi(\mathbf{L}) \,, \tag{2.11}$$

where

$$f_{\Theta\Theta}(\mathbf{l}, \mathbf{l}') \equiv C_l^{\Theta\Theta}(\mathbf{L} \cdot \mathbf{l}) + C_{l'}^{\Theta\Theta}(\mathbf{L} \cdot \mathbf{l}') \,, \tag{2.12}$$

and $\mathbf{L} = \mathbf{l} + \mathbf{l}'$. Note that $\langle \quad \rangle_{\rm CMB}$ differs from the unmarked $\langle \quad \rangle$ that first appeared in Eq. (2.4) in that it denotes an ensemble average only over different Gaussian realizations of the primordial CMB and instrument noise; a fixed realization of the large-scale structure is assumed. For the purposes of estimating the large-scale structure actually realized in our observable universe, this is the appropriate average to take to ensure that our estimators are truly unbiased for a typical realization of the primordial CMB. When calculating the noise associated with lensing-potential estimators and again for power spectrum estimation in § 2.4, we will return to the full unmarked ensemble average. Eq. (2.11), an immediate consquence of Eq. (2.2), suggests that a temperature-squared map appropriately filtered in Fourier space can serve as an estimator for the deflection field $\mathbf{d}(\mathbf{L}) \equiv i\mathbf{L}\phi(\mathbf{L})$. Hu and Okamoto define five different estimators for the deflection field constructed from various combinations of the temperature and polarization; we discuss the temperature-squared estimator in this Section and relegate the analogous formulae for polarization estimators to the Appendix. The minimum-variance temperature-squared estimator derived in Hu &

Okamoto (2002) is

$$\mathbf{d}_{\Theta\Theta}(\mathbf{L}) \equiv \frac{i\mathbf{L}A_{\Theta\Theta}(L)}{L^2}\int \frac{d^2\mathbf{l}_1}{(2\pi)^2}\Theta^{\rm t}(\mathbf{l}_1)\Theta^{\rm t}(\mathbf{l}_2)F_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_2)\,, \tag{2.13}$$

where

$$F_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_2) \equiv \frac{f_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_2)}{2C_{l_1}^{\Theta\Theta{\rm t}}C_{l_2}^{\Theta\Theta{\rm t}}}\,, \tag{2.14}$$

$$A_{\Theta\Theta}(L) \equiv L^2\left[\int \frac{d^2\mathbf{l}_1}{(2\pi)^2}f_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_2)F_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_2)\right]^{-1}\,, \tag{2.15}$$

and $\mathbf{l}_2 = \mathbf{L}-\mathbf{l}_1$. Substitution of Eqs. (2.11), (2.12), (2.14), and (2.15) into Eq. (2.13) shows the desired result,

$$\langle \mathbf{d}_{\Theta\Theta}(\mathbf{L})\rangle_{\rm CMB} = \mathbf{d}(\mathbf{L})\,, \tag{2.16}$$

namely that $\mathbf{d}_{\Theta\Theta}(\mathbf{L})$ is indeed an unbiased estimator for the deflection field in Fourier space. We now proceed to calculate the variance of this estimator. At first, we assume a complete knowledge of all lensing modes not examined by this estimator. In that case, we find

$$\begin{aligned}&\langle \mathbf{d}^*_{\Theta\Theta}(\mathbf{L})\cdot\mathbf{d}_{\Theta\Theta}(\mathbf{L}')\rangle_{\rm CMB} - \langle \mathbf{d}^*_{\Theta\Theta}(\mathbf{L})\rangle_{\rm CMB}\cdot\langle \mathbf{d}_{\Theta\Theta}(\mathbf{L}')\rangle_{\rm CMB} = \\ &(\mathbf{L}\cdot\mathbf{L}')\frac{A_{\Theta\Theta}(L)}{L^2}\frac{A_{\Theta\Theta}(L')}{L'^2} \\ &\quad\times\int\frac{d^2\mathbf{l}_1}{(2\pi)^2}\int\frac{d^2\mathbf{l}'_1}{(2\pi)^2}\langle\Theta^{\rm t}(-\mathbf{l}_1)\Theta^{\rm t}(-\mathbf{l}_2)\Theta^{\rm t}(\mathbf{l}_1{}')\Theta^{\rm t}(\mathbf{l}_2{}')\rangle_{\rm CMB}F_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_2)F_{\Theta\Theta}(\mathbf{l}_1{}',\mathbf{l}_2{}') \\ &-\mathbf{d}^*(\mathbf{L})\cdot\mathbf{d}(\mathbf{L}')\,,\end{aligned} \tag{2.17}$$

where $\mathbf{l}_2{}' = \mathbf{L}' - \mathbf{l}_1{}'$. Evaluating the four-point function in the integrand of Eq. (2.17) to second order in the lensing field, we obtain

$$
\begin{aligned}
&\langle \Theta^{\rm t}(-\mathbf{l}_1)\Theta^{\rm t}(-\mathbf{l}_2)\Theta^{\rm t}(\mathbf{l}_1{}')\Theta^{\rm t}(\mathbf{l}_2{}')\rangle_{\rm CMB} = \Bigg[\left(C_{l_1}^{\Theta\Theta} + C_{l_1}^{\Theta\Theta{\rm n}}\right)(2\pi)^2\delta_{\rm D}(\mathbf{L}) \\
&\quad +\phi(-\mathbf{L}) f_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_2) \\
&\quad -\int \frac{d^2\mathbf{l}'}{(2\pi)^2} C_{l'}^{\Theta\Theta}\phi(-\mathbf{l}_1-\mathbf{l}')\phi(-\mathbf{l}_2+\mathbf{l}')\left[\mathbf{l}'\cdot(\mathbf{l}_1+\mathbf{l}')\right]\left[\mathbf{l}'\cdot(\mathbf{l}_2-\mathbf{l}')\right] \\
&\quad -\frac{1}{2}\int \frac{d^2\mathbf{l}'}{(2\pi)^2}\phi(\mathbf{l}')\phi(-\mathbf{L}-\mathbf{l}') \\
&\qquad\qquad\qquad \left\{C_{l_1}^{\Theta\Theta}(\mathbf{l}_1\cdot\mathbf{l}')\left[\mathbf{l}_1\cdot(\mathbf{L}+\mathbf{l}')\right] + C_{l_2}^{\Theta\Theta}(\mathbf{l}_2\cdot\mathbf{l}')\left[\mathbf{l}_2\cdot(\mathbf{L}+\mathbf{l}')\right]\right\}\Bigg] \\
&\times\Bigg[\left(C_{l_1'}^{\Theta\Theta} + C_{l_1'}^{\Theta\Theta{\rm n}}\right)(2\pi)^2\delta_{\rm D}(\mathbf{L}') + \phi(\mathbf{L}')f_{\Theta\Theta}(\mathbf{l}_1{}',\mathbf{l}_2{}') \\
&\quad -\int \frac{d^2\mathbf{l}'}{(2\pi)^2} C_{l'}^{\Theta\Theta}\phi(\mathbf{l}_1{}'-\mathbf{l}')\phi(\mathbf{l}_2{}'+\mathbf{l}')\left[\mathbf{l}'\cdot(\mathbf{l}_1{}'-\mathbf{l}')\right]\left[\mathbf{l}'\cdot(\mathbf{l}_2{}'+\mathbf{l}')\right] \\
&\quad +\frac{1}{2}\int \frac{d^2\mathbf{l}'}{(2\pi)^2}\phi(\mathbf{l}')\phi(\mathbf{L}'-\mathbf{l}') \\
&\qquad\qquad\qquad \left\{C_{l_1'}^{\Theta\Theta}(\mathbf{l}_1{}'\cdot\mathbf{l}')\left[\mathbf{l}_1{}'\cdot(\mathbf{l}'-\mathbf{L}')\right] + C_{l_2'}^{\Theta\Theta}(\mathbf{l}_2{}'\cdot\mathbf{l}')\left[\mathbf{l}_2{}'\cdot(\mathbf{l}'-\mathbf{L}')\right]\right\}\Bigg] \\
&+\text{Perm.}
\end{aligned}
\tag{2.18}
$$

The terms given explicitly in Eq. (2.18) correspond to the correlations between $\Theta^{\rm t}(-\mathbf{l}_1)$ and $\Theta^{\rm t}(-\mathbf{l}_2)$ and those between $\Theta^{\rm t}(\mathbf{l}_1{}')$ and $\Theta^{\rm t}(\mathbf{l}_2{}')$. The "+ Perm." stands for two additional terms, identical in form, arising from the pairings $\langle\Theta^{\rm t}(-\mathbf{l}_1)\Theta^{\rm t}(\mathbf{l}_1{}')\rangle_{\rm CMB}$ $\langle\Theta^{\rm t}(-\mathbf{l}_2)\Theta^{\rm t}(\mathbf{l}_2{}')\rangle_{\rm CMB}$ and $\langle\Theta^{\rm t}(-\mathbf{l}_1)\Theta^{\rm t}(\mathbf{l}_2{}')\rangle_{\rm CMB}$ $\langle\Theta^{\rm t}(-\mathbf{l}_2)\Theta^{\rm t}(\mathbf{l}_1{}')\rangle_{\rm CMB}$. This expression indicates how uncertainty in the CMB at the last-scattering surface propagates into uncertainty in lensing reconstruction for a particular realization $\phi(\mathbf{L})$ of the large-scale structure. To linear order in Eq. (2.18), correlations between the modes $\Theta^{\rm t}(-\mathbf{l}_1)$, $\Theta^{\rm t}(-\mathbf{l}_2)$, $\Theta^{\rm t}(\mathbf{l}_1{}')$, and $\Theta^{\rm t}(\mathbf{l}_2{}')$ are induced by those lensing modes whose wavevectors are the sums of any pair of wavevectors of these four modes. These lensing modes are precisely those forming the diagonals of the quadrilaterals depicted in Fig. 2.1. In practice, we do not know the large-scale structure between us and the last-scattering surface, so we assume a variance given by Eq. (2.5) with a model-dependent power spectrum $C_L^{\phi\phi}$. We must average Eq. (2.17) over different realizations of the large-scale

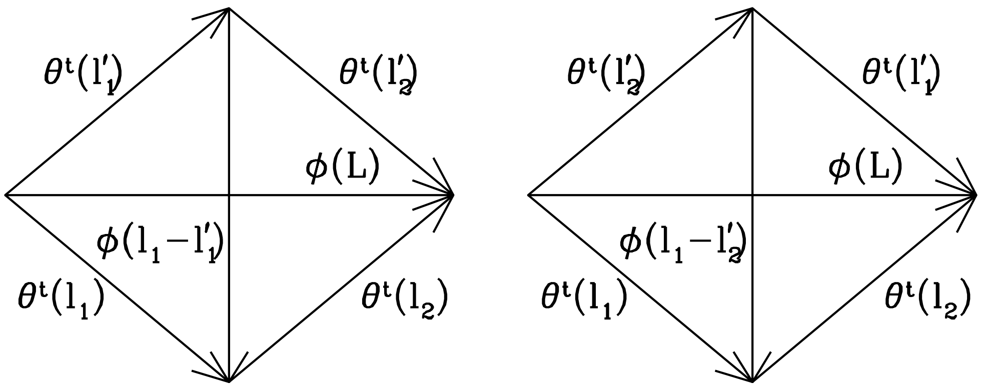

Figure 2.1: The two quadrilaterals consistent with the constraint $\mathbf{L}-\mathbf{L}'=\mathbf{l}_1+\mathbf{l}_2-\mathbf{l}_1{}'-\mathbf{l}_2{}'=0$ for the variance of the estimator $\mathbf{d}_{\Theta\Theta}(\mathbf{L})$. The lensing modes $\phi(\mathbf{L})$, $\phi(\mathbf{l}_1-\mathbf{l}_1{}')$, and $\phi(\mathbf{l}_1-\mathbf{l}_2{}')$ – depicted as diagonals in the above quadrilaterals – induce non-Gaussian couplings between the modes of the observed temperature map represented as sides of these quadrilaterals. They lead to the three groups of linear terms appearing in Eq. (2.18).

structure (denoted by $\langle\quad\rangle_{\rm LSS}$) to obtain the total expected variance of our estimator,

$$\begin{aligned}\Big\langle \langle \mathbf{d}^*_{\Theta\Theta}(\mathbf{L})\cdot\mathbf{d}_{\Theta\Theta}(\mathbf{L}')\rangle_{\rm CMB} - \langle \mathbf{d}^*_{\Theta\Theta}(\mathbf{L})\rangle_{\rm CMB}\cdot\langle \mathbf{d}_{\Theta\Theta}(\mathbf{L}')\rangle_{\rm CMB}\Big\rangle_{\rm LSS} = \\ (\mathbf{L}\cdot\mathbf{L}')\frac{A_{\Theta\Theta}(L)}{L^2}\frac{A_{\Theta\Theta}(L')}{L'^2} \\ \times\int\frac{d^2\mathbf{l}_1}{(2\pi)^2}\int\frac{d^2\mathbf{l}_1'}{(2\pi)^2}\langle\Theta^{\rm t}(-\mathbf{l}_1)\Theta^{\rm t}(-\mathbf{l}_2)\Theta^{\rm t}(\mathbf{l}_1{}')\Theta^{\rm t}(\mathbf{l}_2{}')\rangle F_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_2)F_{\Theta\Theta}(\mathbf{l}_1{}',\mathbf{l}_2{}') \\ -(2\pi)^2\delta_{\rm D}(\mathbf{L}-\mathbf{L}')C_L^{dd},\end{aligned} \tag{2.19}$$

where

$$\langle \mathbf{d}^*(\mathbf{L})\cdot\mathbf{d}(\mathbf{L}')\rangle = (2\pi)^2\delta_{\rm D}(\mathbf{L}-\mathbf{L}')C_L^{dd} = (2\pi)^2\delta_{\rm D}(\mathbf{L}-\mathbf{L}')L^2C_L^{\phi\phi}\,. \tag{2.20}$$

The assumption that the lensing potential is Gaussian imposes the constraint $\mathbf{L}-\mathbf{L}'=\mathbf{l}_1+\mathbf{l}_2-\mathbf{l}_1{}'-\mathbf{l}_2{}'=0$ which closes the quadrilaterals of Fig. 2.1. The average of the four-point correlation function in Eq. (2.19) can be calculated by further averaging

Eq. (2.18) over the large-scale structure. Terms linear in the lensing field vanish when averaged over different realizations of the large-scale structure. Quadratic terms in the lensing field arise as products either of two linear terms or of a zeroth and second-order term. Averages over the product of two linear terms produce the connected part of the four-point correlation function, the trispectrum defined in Eq. (2.4). Averages over the product of a zeroth and second-order term have no connected portion, but instead, furnish an implicit dependence on $C_L^{\phi\phi}$ in the total observed power spectrum of Eq. (2.6). The final result of averaging over the large-scale structure can be expressed in terms of the observed power spectrum and trispectrum,

$$\begin{aligned} \langle \Theta^{\rm t}(-\mathbf{l}_1)\Theta^{\rm t}(-\mathbf{l}_2)\Theta^{\rm t}(\mathbf{l}_1{}')\Theta^{\rm t}(\mathbf{l}_2{}')\rangle \;=\; & (2\pi)^4\Big[C_{l_1}^{\Theta\Theta{\rm t}}C_{l_1'}^{\Theta\Theta{\rm t}}\delta_{\rm D}(\mathbf{L})\delta_{\rm D}(\mathbf{L}') \\ & +C_{l_1}^{\Theta\Theta{\rm t}}C_{l_2}^{\Theta\Theta{\rm t}}\Big\{\delta_{\rm D}(\mathbf{l}_1{}'-\mathbf{l}_1)\delta_{\rm D}(\mathbf{l}_2{}'-\mathbf{l}_2) \\ & \quad +\delta_{\rm D}(\mathbf{l}_2{}'-\mathbf{l}_1)\delta_{\rm D}(\mathbf{l}_1{}'-\mathbf{l}_2)\Big\} \\ & +(2\pi)^{-2}T^{\rm t}(-\mathbf{l}_1,-\mathbf{l}_2,\mathbf{l}_1{}',\mathbf{l}_2{}')\delta_{\rm D}(\mathbf{L}-\mathbf{L}')\Big]\,, \end{aligned} \tag{2.21}$$

where the trispectrum can written in terms of $f_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_2)$ as

$$\begin{aligned} T^{\rm t}(\mathbf{l}_1,\mathbf{l}_2,\mathbf{l}_3,\mathbf{l}_4) \;=\; & C^{\phi\phi}_{|\mathbf{l}_1+\mathbf{l}_2|}f_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_2)f_{\Theta\Theta}(\mathbf{l}_3,\mathbf{l}_4)+C^{\phi\phi}_{|\mathbf{l}_1+\mathbf{l}_3|}f_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_3)f_{\Theta\Theta}(\mathbf{l}_2,\mathbf{l}_4) \\ & +C^{\phi\phi}_{|\mathbf{l}_1+\mathbf{l}_4|}f_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_4)f_{\Theta\Theta}(\mathbf{l}_2,\mathbf{l}_3)\,. \end{aligned} \tag{2.22}$$

This form of the trispectrum is consistent with that of Eq. (2.10) given directly in terms of power spectra. Since $\mathbf{L}\neq 0$, $\delta_{\rm D}(\mathbf{L})=0$ and the first term of Eq. (2.21) vanishes. The remaining two terms containing pairs of delta functions, inserted into Eq. (2.19), yield the dominant contribution to the variance

$$\begin{aligned} &\Big\langle \langle \mathbf{d}^*_{\Theta\Theta}(\mathbf{L})\cdot\mathbf{d}_{\Theta\Theta}(\mathbf{L}')\rangle_{\rm CMB}-\langle\mathbf{d}^*_{\Theta\Theta}(\mathbf{L})\rangle_{\rm CMB}\cdot\langle\mathbf{d}_{\Theta\Theta}(\mathbf{L}')\rangle_{\rm CMB}\Big\rangle_{\rm LSS} = \\ &\qquad (2\pi)^2\delta_{\rm D}(\mathbf{L}-\mathbf{L}')[N^{(0)}_{\Theta\Theta,\Theta\Theta}(L)+\ldots]\,, \end{aligned} \tag{2.23}$$

where $N^{(0)}_{\Theta\Theta,\Theta\Theta}(L) = A_{\Theta\Theta}(L)$. Notice that $N^{(0)}_{\Theta\Theta,\Theta\Theta}(L)$ is zeroth order in the lensing potential ϕ; it depends on the lensing potential power spectrum $C_l^{\phi\phi}$ only implicitly through the total observed power spectrum $C_l^{\Theta\Theta\mathrm{t}}$. The ellipsis represents terms of higher order in $C_l^{\phi\phi}$ that we now proceed to calculate. These terms arise from the trispectrum term of Eq. (2.21) after $\langle\langle \mathbf{d}^*_{\Theta\Theta}(\mathbf{L})\rangle_{\mathrm{CMB}} \cdot \langle \mathbf{d}_{\Theta\Theta}(\mathbf{L}')\rangle_{\mathrm{CMB}}\rangle_{\mathrm{LSS}} = (2\pi)^2\, \delta_{\mathrm{D}}(\mathbf{L}-\mathbf{L}')\; C_L^{dd}$ is removed. Substituting these results into Eq. (2.19), we find that to first order in $C_l^{\phi\phi}$

$$\Big\langle \langle \mathbf{d}^*_{\Theta\Theta}(\mathbf{L}) \cdot \mathbf{d}_{\Theta\Theta}(\mathbf{L}')\rangle_{\mathrm{CMB}} - \langle \mathbf{d}^*_{\Theta\Theta}(\mathbf{L})\rangle_{\mathrm{CMB}} \cdot \langle \mathbf{d}_{\Theta\Theta}(\mathbf{L}')\rangle_{\mathrm{CMB}} \Big\rangle_{\mathrm{LSS}} = (2\pi)^2 \delta_{\mathrm{D}}(\mathbf{L}-\mathbf{L}')[N^{(0)}_{\Theta\Theta,\Theta\Theta}(L) + N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)]\,, \tag{2.24}$$

where $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$ is given by

$$\begin{aligned} N^{(1)}_{\Theta\Theta,\Theta\Theta}(L) = \frac{A^2_{\Theta\Theta}(L)}{L^2} \int \frac{d^2\mathbf{l}_1}{(2\pi)^2} \int \frac{d^2\mathbf{l}'_1}{(2\pi)^2} F_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_2) F_{\Theta\Theta}(\mathbf{l}_1{}',\mathbf{l}_2{}') \\ \times \Big\{ C^{\phi\phi}_{|\mathbf{l}_1-\mathbf{l}_1{}'|} f_{\Theta\Theta}(-\mathbf{l}_1,\mathbf{l}_1{}') f_{\Theta\Theta}(-\mathbf{l}_2,\mathbf{l}_2{}') + C^{\phi\phi}_{|\mathbf{l}_1-\mathbf{l}_2{}'|} f_{\Theta\Theta}(-\mathbf{l}_1,\mathbf{l}_2{}') f_{\Theta\Theta}(-\mathbf{l}_2,\mathbf{l}_1{}') \Big\}\,. \end{aligned} \tag{2.25}$$

The first-order contribution to the noise $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$ involves integrals over the lensing-potential power spectrum, and thus, probes lensing modes with wave vectors different from that of the estimator $\mathbf{d}_{\Theta\Theta}(\mathbf{L})$. It can be interpreted physically as interference from these other modes in the determination of the mode $\mathbf{d}(\mathbf{L})$ being estimated. The filter $F_{\Theta\Theta}(\mathbf{l}_1,\mathbf{l}_2)$ was chosen to optimize the signal-to-noise ratio in the absence of the first-order contribution $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$; it is no longer an optimal filter once this additional noise is taken into account. As long as $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L) \ll N^{(0)}_{\Theta\Theta,\Theta\Theta}(L)$, the noise reduction that can be attained by re-optimizing our filter will not be significant. Formulas analagous to those presented here relevant to the construction of estimators using polarization data are given in the Appendix.

The significance of $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$ for two different experiments is shown in Fig. 2.2 using the currently favored ΛCDM cosmological model with baryon density $\Omega_b = 0.05$, cold dark matter density $\Omega_{cdm} = 0.30$, cosmological constant density $\Omega_\Lambda = 0.65$,

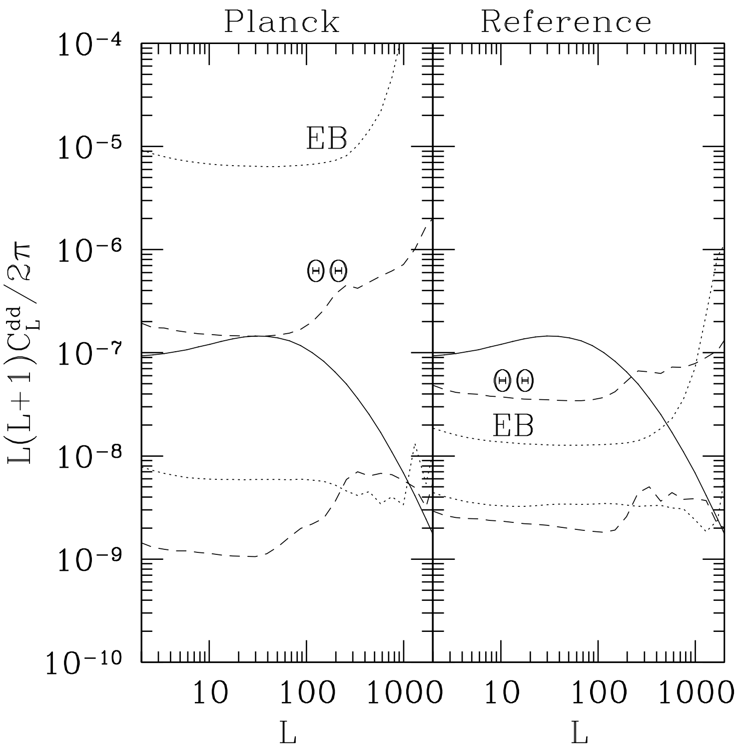

Figure 2.2: Variances with which individual modes $\mathbf{d}(\mathbf{L})$ of the deflection field can be reconstructed by the Planck and reference experiments described in the text. The solid curves are the power spectra C_L^{dd} anticipated for our ΛCDM cosmological model. The upper and lower dashed curves are the zeroth and first-order noise power spectra $N^{(0)}_{\Theta\Theta,\Theta\Theta}(L)$ and $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$, respectively, for the temperature-based estimator, $\mathbf{d}_{\Theta\Theta}(\mathbf{L})$, while the dotted curves are the corresponding noise variances for $\mathbf{d}_{EB}(\mathbf{L})$. A mode $\mathbf{d}(\mathbf{L})$ cannot be reconstructed with signal-to-noise greater than unity when $C_L^{dd} \leqslant N^{(0)}_{\Theta\Theta,\Theta\Theta}(L) + N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$.

Hubble parameter h=0.65, and power-spectrum amplitude $\sigma_8 = 0.9$. The Planck experiment is equivalent to a one-year, full-sky survey with temperature and polarization sensitivities of 12.42 and 26.02 $\mu\mathrm{K}\sqrt{\mathrm{sec}}$, respectively and resolution $\theta = 7.0$ arcminutes as described in § 2.2. The reference experiment has the same resolution, but superior sensitivities of 0.46 and 0.65 $\mu\mathrm{K}\sqrt{\mathrm{sec}}$ for temperature and polarization. These estimates of experimental parameters are identical to those given for the Planck and reference experiments of Hu & Okamoto (2002). The ΘE and EE estimators have noise power spectra intermediate to those of the ΘΘ and EB estimators, while the ΘB estimator has substantially higher noise because the primordial CMB lacks true

B-modes in the absence of inflationary gravitational waves. We see that for Planck, with its comparatively inferior polarization sensitivity, the $\Theta\Theta$ estimator will be best, although it will be unable to detect individual Fourier modes of the deflection field at the 1σ level. The reference experiment, and further experiments with similar sensitivity and even higher resolution, should be able to push 1σ detection of individual $\mathbf{d}(\mathbf{L})$ modes to $L \simeq 1000$ by primarily relying on the EB estimator. In these cases, the secondary noise $N^{(1)}_{EB,EB}(L)$ is only smaller than the dominant noise $N^{(0)}_{EB,EB}(L)$ by a factor of a few, whereas for higher sensitivity (noisier) experiments like Planck, it is smaller by at least an order-of-magnitude. This illustrates an interesting point, apparent from Fig. 2.2, that the zeroth-order noise $N^{(0)}_{\Theta\Theta,\Theta\Theta}(L)$ declines dramatically with decreasing sensitivity until it becomes dominated by cosmic variance while the $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$ is largely unaffected by instrument sensitivity. The reasons for this trend are that instrument noise appears in $N^{(0)}_{\Theta\Theta,\Theta\Theta}(L) = A_{\Theta\Theta}(L)$ through its contribution to the denominator of $F_{\Theta\Theta}(\mathbf{l}_1, \mathbf{l}_2)$ as shown by Eqs. (2.14) and (2.15). Decreasing instrument noise raises the value of $F_{\Theta\Theta}(\mathbf{l}_1, \mathbf{l}_2)$, thereby lowering $N^{(0)}_{\Theta\Theta,\Theta\Theta}(L)$. By contrast, instrument noise is reflected in $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$ through its effects on both $A_{\Theta\Theta}(L)$ and $F_{\Theta\Theta}(\mathbf{l}_1, \mathbf{l}_2)$, as shown in Eq. (2.25). Smaller instrument noise raises $F_{\Theta\Theta}(\mathbf{l}_1, \mathbf{l}_2)$ as before – driving $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$ up in this case – but this is compensated by a decrease in $A_{\Theta\Theta}(L)$ which appears as a pre-factor outside the integrals. These two effects largely cancel each other out, rendering $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$ remarkably insensitive to instrument noise.

2.4 Power Spectrum Estimation

Although complete reconstruction of the deflection field $\mathbf{d}(\mathbf{L})$ is essential for purposes such as B-mode subtraction (Kesden et al., 2002b; Knox & Song, 2002), sometimes estimates of the lensing-potential power spectrum $C_L^{\phi\phi}$ are sufficient. This power spectrum is a model-dependent prediction of theories of large-scale structure formation, and therefore, estimates of the power spectrum from real data could be used to test these theories, as well as the consistency of other determinations of cosmological parameters. Furthermore, since estimates of all the modes $\mathbf{d}(\mathbf{L})$ with $|\mathbf{L}| = L$ can be combined to estimate $C_L^{\phi\phi}$, 1σ detection of the power spectrum can be pushed to much higher L than can that of individual modes. The deflection-field estimator $\mathbf{d}_{\Theta\Theta}(\mathbf{L})$ derived in the preceding Section can be used to construct an estimator for $C_L^{\phi\phi}$. Our first guess for an appropriate lensing-potential power spectrum estimator is

$$D_L \equiv \frac{(2\pi)^2}{AL^2} \frac{1}{2\pi L \Delta L} \int_{a_L} \frac{d^2\mathbf{l}}{(2\pi)^2}\, \mathbf{d}_{\Theta\Theta}(\mathbf{l}) \cdot \mathbf{d}_{\Theta\Theta}(-\mathbf{l})\,, \tag{2.26}$$

where A is the area of the sky surveyed and a_L is an annulus of radius L and width ΔL. We ensemble average our estimator over different realizations of the CMB and large-scale structure using Eq. (2.24) by bringing $\langle\langle \mathbf{d}^*_{\Theta\Theta}(\mathbf{L})\rangle_{\rm CMB} \cdot \langle \mathbf{d}_{\Theta\Theta}(\mathbf{L}')\rangle_{\rm CMB}\rangle_{\rm LSS}$ from the left to the right-hand side. This yields

$$\langle D_L \rangle = \frac{(2\pi)^2}{AL^2} \frac{1}{2\pi L \Delta L} \delta_{\rm D}(0) \int_{a_L} d^2\mathbf{l} \left[l^2 C_l^{\phi\phi} + N^{(0)}_{\Theta\Theta,\Theta\Theta}(l) + N^{(1)}_{\Theta\Theta,\Theta\Theta}(l) \right] . \tag{2.27}$$

The definition of the Dirac delta function,

$$(2\pi)^2 \delta_{\rm D}(\mathbf{l}) \equiv \int_A d\hat{\mathbf{n}}\, e^{i\mathbf{l}\cdot\hat{\mathbf{n}}}\,, \tag{2.28}$$

implies that $\delta_{\rm D}(0) = A/(2\pi)^2$. Furthermore, in the limit that ΔL is small compared to the scales on which $l^2 C_l^{\phi\phi} + N^{(0)}_{\Theta\Theta,\Theta\Theta}(l) + N^{(1)}_{\Theta\Theta,\Theta\Theta}(l)$ is varying, we can evaluate the integrand of Eq. (2.27) at its central value $l = L$ and extract it from the integral. The integral over the annulus a_L cancels the factor $2\pi L \Delta L$ in the denominator, reducing Eq. (2.27) to

$$\langle D_L \rangle = C_L^{\phi\phi} + L^{-2} \left[N^{(0)}_{\Theta\Theta,\Theta\Theta}(L) + N^{(1)}_{\Theta\Theta,\Theta\Theta}(L) \right] . \tag{2.29}$$

D_L is indeed an estimator for the lensing-potential power spectrum $C_L^{\phi\phi}$, albeit a biased one. Note that the *bias* in the power-spectrum estimator D_L is precisely the same as the *variance* shown in Eq. (2.24) with which we were able to determine each individual lensing mode. This is no coincidence; it reflects the fact that there are no grounds *a priori* on which to differentiate the variance with which we can reconstruct individual modes $\mathbf{d}(\mathbf{L})$ from the intrinsic variance C_L^{dd} of the underlying distribution from which they are drawn. To obtain an unbiased estimator to compare with theoretical predictions, we subtract off this unwanted reconstruction variance

$$\widehat{C}_L^{\phi\phi} \equiv D_L - L^{-2}\left[N^{(0)}_{\Theta\Theta,\Theta\Theta}(L) + N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)\right] . \tag{2.30}$$

Since $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$ as defined in Eq. (2.25) itself depends on $C_L^{\phi\phi}$, this subtraction and evaluation must be performed iteratively until a self-consistent solution is obtained. The variance of our estimator $\widehat{C}_L^{\phi\phi}$ can be calculated in the usual manner

$$\sigma^2_{\widehat{C}_L^{\phi\phi}} \equiv \langle(\widehat{C}_L^{\phi\phi})^2\rangle - \langle\widehat{C}_L^{\phi\phi}\rangle^2 = \langle(D_L)^2\rangle - \langle D_L\rangle^2 . \tag{2.31}$$

Evaluating this expression requires us to calculate

$$\begin{aligned}\langle(D_L)^2\rangle = {} & \frac{(2\pi)^4}{A^2L^4}\frac{1}{(2\pi L\Delta L)^2} \times \\ & \int_{a_L}\frac{d^2\mathbf{l}_1{}'}{(2\pi)^2}\int_{a_L}\frac{d^2\mathbf{l}_2{}'}{(2\pi)^2}\left\langle[\mathbf{d}_{\Theta\Theta}(\mathbf{l}_1{}')\cdot\mathbf{d}_{\Theta\Theta}(-\mathbf{l}_1{}')]\,[\mathbf{d}_{\Theta\Theta}(\mathbf{l}_2{}')\cdot\mathbf{d}_{\Theta\Theta}(-\mathbf{l}_2{}')]\right\rangle . \end{aligned} \tag{2.32}$$

Since $\mathbf{d}_{\Theta\Theta}(\mathbf{L})$ is a quadratic estimator in the temperature map, Eq. (2.32) includes the following integral over the eight-point correlation function in Fourier space

$$\begin{aligned}\left\langle[\mathbf{d}_{\Theta\Theta}(\mathbf{l}_1{}')\cdot\mathbf{d}_{\Theta\Theta}(-\mathbf{l}_1{}')]\,[\mathbf{d}_{\Theta\Theta}(\mathbf{l}_2{}')\cdot\mathbf{d}_{\Theta\Theta}(-\mathbf{l}_2{}')]\right\rangle = {} & \frac{A^2_{\Theta\Theta}(l_1')A^2_{\Theta\Theta}(l_2')}{(l_1')^2(l_2')^2} \\ \times\int\frac{d^2\mathbf{k}_1}{(2\pi)^2}\int\frac{d^2\mathbf{k}_3}{(2\pi)^2}\int\frac{d^2\mathbf{k}_5}{(2\pi)^2}\int\frac{d^2\mathbf{k}_7}{(2\pi)^2}\Big\{ & F_{\Theta\Theta}(\mathbf{k}_1,\mathbf{k}_2)F_{\Theta\Theta}(\mathbf{k}_3,\mathbf{k}_4)F_{\Theta\Theta}(\mathbf{k}_5,\mathbf{k}_6) \\ \times F_{\Theta\Theta}(\mathbf{k}_7,\mathbf{k}_8)\langle\Theta^{\rm t}(\mathbf{k}_1)\Theta^{\rm t}(\mathbf{k}_2)\Theta^{\rm t}(\mathbf{k}_3)\Theta^{\rm t}(\mathbf{k}_4)\Theta^{\rm t}(\mathbf{k}_5)\Theta^{\rm t}(\mathbf{k}_6)\Theta^{\rm t}(\mathbf{k}_7)\Theta^{\rm t}(\mathbf{k}_8)\rangle\Big\} & , \end{aligned} \tag{2.33}$$

where $\mathbf{k}_2 = \mathbf{l}_1{}' - \mathbf{k}_1$, $\mathbf{k}_4 = -\mathbf{l}_1{}' - \mathbf{k}_3$, $\mathbf{k}_6 = \mathbf{l}_2{}' - \mathbf{k}_5$, and $\mathbf{k}_8 = -\mathbf{l}_2{}' - \mathbf{k}_7$. A fully general eight-point correlation function consists of a connected part, as well as terms proportional to the product of lower-order correlation functions. Under the assumption that both the primordial CMB and the lensing potential are governed by Gaussian statistics, all correlation functions higher than the four-point have connected parts that are suppressed to high order in the lensing potential (Kesden et al., 2002a). The temperature eight-point correlation function will therefore be composed of three groups of terms, with membership in a group being determined by whether the term contains zero, one, or two factors of the trispectrum. Since the trispectrum given in Eq. (2.10) is first order in the lensing-potential power spectrum $C_L^{\phi\phi}$, terms of these three groups are zeroth, first, and second order, respectively, in $C_L^{\phi\phi}$. Combinatorics determines the number of terms in each group. There are: $\frac{1}{4!}\binom{8}{2}\binom{6}{2}\binom{4}{2}\binom{2}{2} = 105$ different ways of dividing $(\mathbf{k}_1, \ldots, \mathbf{k}_8)$ into four pairs, and hence, there will be 105 terms in the group containing no trispectra. Similar calculations reveal that there are $\frac{1}{2!}\binom{8}{4}\binom{4}{2}\binom{2}{2} = 210$ terms in the second group and $\frac{1}{2!}\binom{8}{4}\binom{4}{4} = 35$ terms in the third group. Many terms in all three groups will vanish for the same reason that the first term vanished in the four-point correlation function of Eq. (2.21); these terms are proportional to a Dirac delta function evaluated at nonzero argument. Consider now the first group of terms, those that are zeroth order in $C_L^{\phi\phi}$. The 60 nonvanishing terms in this group each contain four Dirac delta functions; they can be further segregated into the 12 terms that allow two of the integrals over $\mathbf{k}_i$ appearing in Eq. (2.33) to be immediately evaluated via Dirac delta functions, and the 48 terms that allow evaluation of three $\mathbf{k}_i$ integrals. The first 12 terms, inserted into Eq. (2.33) and appropriately evaluated using the normalization of Eq. (2.15), yield

$$\langle (D_L)^2 \rangle = \left(\frac{N^{(0)}_{\Theta\Theta,\Theta\Theta}(L)}{L^2} \right)^2 \left[1 + 2\frac{(2\pi)^2}{A} \frac{1}{2\pi L \Delta L} \right] + \ldots , \tag{2.34}$$

while the remaining 48 terms give the final result to zeroth order in $C_L^{\phi\phi}$,

$$\langle (D_L)^2 \rangle = \left(\frac{N^{(0)}_{\Theta\Theta,\Theta\Theta}(L)}{L^2} \right)^2 \left[1 + 2\frac{(2\pi)^2}{A} \frac{1}{2\pi L \Delta L} \right] + \frac{(2\pi)^2}{AL^4} \frac{2}{(2\pi L \Delta L)^2} \int_{a_L} \frac{d^2\mathbf{l}_1{}'}{(2\pi)^2} \int_{a_L} \frac{d^2\mathbf{l}_2{}'}{(2\pi)^2} \frac{A^2_{\Theta\Theta}(l_1')A^2_{\Theta\Theta}(l_2')}{(l_1')^2(l_2')^2} \int \frac{d^2\mathbf{k}_1}{(2\pi)^2} f_{\Theta\Theta}(\mathbf{k}_1, \mathbf{k}_2) P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{l}_1{}', \mathbf{l}_2{}') , \tag{2.35}$$

where

$$\begin{aligned}
&P(\mathbf{k}_1,\mathbf{k}_2,\mathbf{l}_1{}',\mathbf{l}_2{}') = \\
&\qquad f_{\Theta\Theta}(-\mathbf{l}_2{}'-\mathbf{k}_1,\mathbf{l}_2{}'-\mathbf{k}_2)F_{\Theta\Theta}(-\mathbf{k}_1,\mathbf{l}_2{}'+\mathbf{k}_1)F_{\Theta\Theta}(-\mathbf{k}_2,-\mathbf{l}_2{}'+\mathbf{k}_2) \\
&\qquad +f_{\Theta\Theta}(\mathbf{l}_2{}'-\mathbf{k}_1,-\mathbf{l}_2{}'-\mathbf{k}_2)F_{\Theta\Theta}(-\mathbf{k}_1,-\mathbf{l}_2{}'+\mathbf{k}_1)F_{\Theta\Theta}(-\mathbf{k}_2,\mathbf{l}_2{}'+\mathbf{k}_2) \\
&\qquad +f_{\Theta\Theta}(-\mathbf{l}_2{}'-\mathbf{k}_1,\mathbf{l}_1{}'-\mathbf{k}_2)F_{\Theta\Theta}(-\mathbf{k}_1,\mathbf{l}_2{}'+\mathbf{k}_1)F_{\Theta\Theta}(-\mathbf{k}_2,-\mathbf{l}_1{}'+\mathbf{k}_2) \\
&\qquad +f_{\Theta\Theta}(\mathbf{l}_1{}'-\mathbf{k}_1,-\mathbf{l}_2{}'-\mathbf{k}_2)F_{\Theta\Theta}(-\mathbf{k}_1,-\mathbf{l}_1{}'+\mathbf{k}_1)F_{\Theta\Theta}(-\mathbf{k}_2,\mathbf{l}_2{}'+\mathbf{k}_2) \\
&\qquad +f_{\Theta\Theta}(\mathbf{l}_1{}'-\mathbf{k}_1,\mathbf{l}_2{}'-\mathbf{k}_2)F_{\Theta\Theta}(-\mathbf{k}_1,-\mathbf{l}_1{}'+\mathbf{k}_1)F_{\Theta\Theta}(-\mathbf{k}_2,-\mathbf{l}_2{}'+\mathbf{k}_2) \\
&\qquad +f_{\Theta\Theta}(\mathbf{l}_2{}'-\mathbf{k}_1,\mathbf{l}_1{}'-\mathbf{k}_2)F_{\Theta\Theta}(-\mathbf{k}_1,-\mathbf{l}_2{}'+\mathbf{k}_1)F_{\Theta\Theta}(-\mathbf{k}_2,-\mathbf{l}_1{}'+\mathbf{k}_2)\,.
\end{aligned} \tag{2.36}$$

When $\langle D_L\rangle^2$ is subtracted from $\langle (D_L)^2\rangle$ in Eq. (2.31), the first term in the square brackets of Eq. (2.35) will be eliminated. Minimizing the variance associated with this estimator then consists of making an optimal choice of $\Delta L(L)$. The first noise term is proportional to $\frac{(2\pi)^2/A}{2\pi L\Delta L}$. For a survey of area A, $(2\pi)^2/A$ is the specific area of an individual mode in $\mathbf{L}$ space and $2\pi L\Delta L$ is the area in $\mathbf{L}$ space over which the power-spectrum estimator takes an average. This ratio is therefore the inverse of the number of individual $\mathbf{d}_{\Theta\Theta}(\mathbf{L})$ modes whose inverse variances are added to determine the inverse variance of $\widehat{C}_L^{\phi\phi}$. It is obviously minimized by choosing $(2\pi)^2/A \ll 2\pi L\Delta L(L)$. The second term – that involving $P(\mathbf{k}_1,\mathbf{k}_2,\mathbf{l}_1{}',\mathbf{l}_2{}')$ – differs from the first noise term in that a Dirac delta function has been used to evaluate an additional $\mathbf{k}_i$ integral rather than an annulus integral. Since the integrands are of the same order, we expect the second noise term to be suppressed relative to the first by a factor of $2\pi L\Delta L(L)/\pi l_{\max}^2$ where $l_{\max} \simeq \pi/\theta$ is set by the resolution θ of the survey. Under the conservative assumption $L\Delta L(L) \ll 1/\theta^2$, namely, that we are probing scales well above our resolution, this term is assured to be small. We neglect such terms for the remainder of this paper. If we insert the portions of the eight-point correlation function that are first and second order in $C_L^{\phi\phi}$ into Eq. (2.33) and evaluate using

Eq. (2.25), we find

$$\langle (D_L)^2 \rangle = L^{-4} \left(C_L^{dd} + N^{(0)}_{\Theta\Theta,\Theta\Theta}(L) + N^{(1)}_{\Theta\Theta,\Theta\Theta}(L) \right)^2 \left[1 + 2\frac{(2\pi)^2}{A}\frac{1}{2\pi L \Delta L} \right] , \tag{2.37}$$

and

$$\sigma^2_{\widehat{C}_L^{\phi\phi}} = 2\frac{(2\pi)^2}{A}\frac{1}{2\pi L \Delta L} L^{-4} \left(C_L^{dd} + N^{(0)}_{\Theta\Theta,\Theta\Theta}(L) + N^{(1)}_{\Theta\Theta,\Theta\Theta}(L) \right)^2 . \tag{2.38}$$

This result agrees with that given in Ref. Hu & Okamoto (2002) after subtracting our newly derived term, $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$.

The term $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$ and corresponding terms for polarization-based estimators have two principal effects on power spectrum estimation. As shown in Eq. (2.38), they provide a fractional contribution to the variance of roughly $2N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)/C_L^{dd}$ when C_L^{dd} dominates the variance as in the reference experiment in the right-hand panel of Fig. 2.2. For the ΛCDM cosmological model considered here, this represents an increase of $5-15\%$ in the variance of the EB estimator for $L \lesssim 1000$. More importantly, $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$ acts as a bias for the naive estimator D_L as shown by Eq. (2.29). If this bias is not calculated and subtracted iteratively to form the unbiased estimator $\widehat{C}_L^{\phi\phi}$ as in Eq. (2.30), C_L^{dd} will be *systematically* overestimated by $5-10\%$ at low L and by increasingly larger amounts at $L \gtrsim 100$ as the signal $L(L+1)C_L^{dd}/2\pi$ begins to plummet while $N^{(1)}_{\Theta\Theta,\Theta\Theta}(L)$ remains comparatively flat.

Having evaluated the variance of our estimator $\widehat{C}_L^{\phi\phi}$, we consider whether this estimator has a substantial covariance

$$\sigma_{\widehat{C}_L^{\phi\phi}\widehat{C}_{L'}^{\phi\phi}} \equiv \langle (\widehat{C}_L^{\phi\phi} - C_L^{\phi\phi})(\widehat{C}_{L'}^{\phi\phi} - C_{L'}^{\phi\phi}) \rangle = \langle D_L D_{L'} \rangle - \langle D_L \rangle \langle D_{L'} \rangle . \tag{2.39}$$

The estimator D_L as defined in Eq. (2.26) implies that:

$$\langle D_L D_{L'} \rangle = \left(\frac{2\pi}{ALL'} \right)^2 \frac{1}{L\Delta L L'\Delta L'} \times \int_{a_L} \frac{d^2\mathbf{l}_1{}'}{(2\pi)^2} \int_{a_{L'}} \frac{d^2\mathbf{l}_2{}'}{(2\pi)^2} \langle [\mathbf{d}_{\Theta\Theta}(\mathbf{l}_1{}') \cdot \mathbf{d}_{\Theta\Theta}(-\mathbf{l}_1{}')] \, [\mathbf{d}_{\Theta\Theta}(\mathbf{l}_2{}') \cdot \mathbf{d}_{\Theta\Theta}(-\mathbf{l}_2{}')] \rangle , \tag{2.40}$$

which can be evaluated using the same integral over the eight-point correlation function described in Eq. (2.33). Whereas 60 of the 105 zeroth-order terms in $C_L^{\phi\phi}$ coming from this equation were nonvanishing for the variance, for the covariance, only 52 terms are nonzero provided that the widths ΔL and $\Delta L'$ are chosen so that the annuli a_L and $a_{L'}$ do not overlap. This leads to a result analogous to Eq. (2.35),

$$\begin{aligned}\langle D_L D_{L'} \rangle = & \frac{N^{(0)}_{\Theta\Theta,\Theta\Theta}(L)}{L^2} \frac{N^{(0)}_{\Theta\Theta,\Theta\Theta}(L')}{L'^2} + \frac{2}{AL^3 \Delta L L'^3 \Delta L'} \\ & \times \int_{a_L} \frac{d^2\mathbf{l}_1{}'}{(2\pi)^2} \int_{a_{L'}} \frac{d^2\mathbf{l}_2{}'}{(2\pi)^2} \frac{A^2_{\Theta\Theta}(l'_1) A^2_{\Theta\Theta}(l'_2)}{(l'_1)^2 (l'_2)^2} \int \frac{d^2\mathbf{k}_1}{(2\pi)^2} f_{\Theta\Theta}(\mathbf{k}_1, \mathbf{k}_2) P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{l}_1{}', \mathbf{l}_2{}') \,.\end{aligned} \tag{2.41}$$

Note that the eight terms missing from the covariance when compared to the variance have altered the first term of Eq. (2.41), and that it was precisely these terms that provided the dominant contribution to the variance of Eq. (2.38) when $\Delta L(L)$ was chosen appropriately. We therefore find that to zeroth order in $C_L^{\phi\phi}$, the covariance is given by

$$\begin{aligned}\sigma_{\widehat{C}_L^{\phi\phi} \widehat{C}_{L'}^{\phi\phi}} = & \frac{2}{AL\Delta L L' \Delta L'} \left(\frac{N^{(0)}_{\Theta\Theta,\Theta\Theta}(L)}{L^2} \frac{N^{(0)}_{\Theta\Theta,\Theta\Theta}(L')}{L'^2} \right)^2 \\ & \times \int_{a_L} \frac{d^2\mathbf{l}_1{}'}{(2\pi)^2} \int_{a_{L'}} \frac{d^2\mathbf{l}_2{}'}{(2\pi)^2} \int \frac{d^2\mathbf{k}_1}{(2\pi)^2} f_{\Theta\Theta}(\mathbf{k}_1, \mathbf{k}_2) P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{l}_1{}', \mathbf{l}_2{}') \,,\end{aligned} \tag{2.42}$$

where we have extracted the $A_{\Theta\Theta}(l'_i)$ from the annular integrals since they are slowly varying over the widths ΔL and $\Delta L'$. For the same reasons that terms of this form were a subdominant contribution to the variance as discussed previously, we expect the covariance to be suppressed as well. If we define the ratio

$$R_{LL'} \equiv \frac{\sigma_{\widehat{C}_L^{\phi\phi} \widehat{C}_{L'}^{\phi\phi}}}{\sqrt{\sigma^2_{\widehat{C}_L^{\phi\phi}} \sigma^2_{\widehat{C}_{L'}^{\phi\phi}}}} \,, \tag{2.43}$$

we can quantify this suppression. The triple integral of Eq. (2.42), appearing in the numerator of $R_{LL'}$, involves integration over annuli with radii L and L' and one

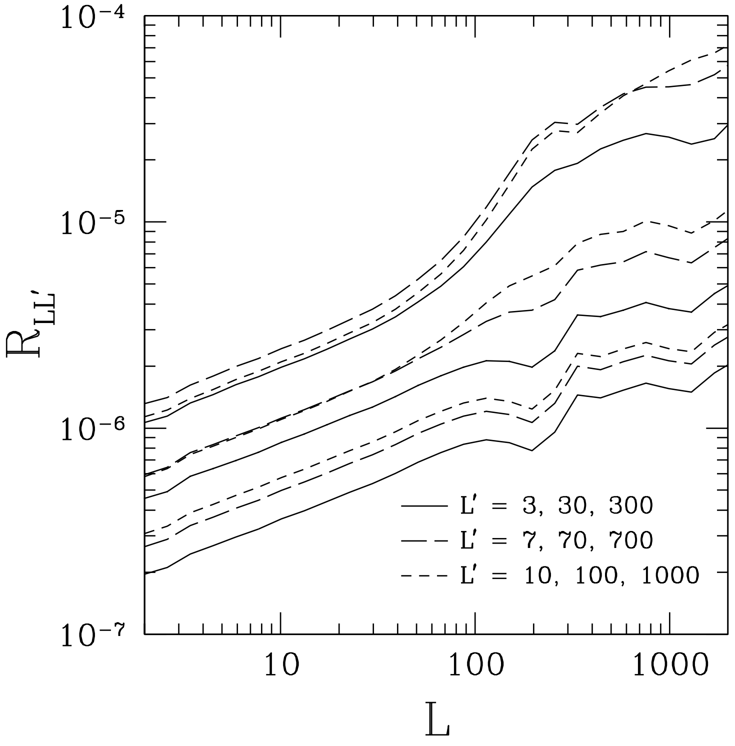

Figure 2.3: The ratio $R_{LL'}$ for Planck as a function of L for fixed values of L'. The solid curves correspond to $L' = 3, 30, 300$ ascending from bottom to top, while the long-dashed and short-dashed curves correspond to $L' = 7, 70, 700$ and $L' = 10, 100, 1000$ respectively, again with curves in each sequence appearing from bottom to top in the figure.

integration over all Fourier space. The triple integrals in the variances $\sigma^2_{\widehat{C}^{\phi\phi}_L}$ and $\sigma^2_{\widehat{C}^{\phi\phi}_{L'}}$ appearing in the denominator of $R_{LL'}$ each consist of a single integration over an annulus of radius L and L', respectively, and two integrations over all Fourier space. If we make the crude assumption that the integrand is constant, the ratio $R_{LL'}$ will simply be the ratio of these areas,

$$R_{LL'} \simeq \sqrt{(2\pi L \Delta L)(2\pi L' \Delta L')}/\pi l^2_{\max}. \tag{2.44}$$

The ratio $R_{LL'}$ is evaluated numerically for Planck in Fig. 2.3 as a function of L for various fixed values of L'. The estimators $\widehat{C}^{\phi\phi}_L$ and $\widehat{C}^{\phi\phi}_{L'}$ were chosen such that $\Delta L = \Delta L' = 1$, while integrals over Fourier space were cut off at $l_{\max} = 5000$.

Substituting these values into Eq. (2.44), we expect $R_{LL'} \simeq 8.0 \times 10^{-8}\sqrt{LL'}$. This crude estimate is surprisingly close to the numerically obtained results of Fig. 2.3; in particular, the slope of the curves is approximately 1/2 on this log-log plot. Even at $L \simeq 1000$, $R_{LL'} \lesssim 4.0 \times 10^{-4}$ suggesting that covariance in power-spectrum estimation can safely be neglected for Planck. The estimate of Eq. (2.44) implicitly depends on the experimental resolution θ because the integrand appearing in expressions for the variance and covariance decreases rapidly for $L \gtrsim l_{\max} \simeq \pi/\theta$. For future experiments with better resolution than Planck, $l_{\max}$ will be higher, implying by Eq. (2.44) that covariance will be even more negligible.

2.5 Discussion

Weak gravitational lensing induces non-Gaussian correlations between modes of the observed CMB temperature map as shown in Eq. (2.11). These correlations, and assumptions about the Gaussian nature of the primordial CMB, can be used to construct several temperature and polarization-based estimators of the Fourier modes $\mathbf{d}(\mathbf{L})$ of the deflection field. This procedure was outlined in Hu & Okamoto (2002), however in calculating the noise associated with this reconstruction, an assumption was made that the observed temperature map was Gaussian. In the presence of lensing, this assumption is invalid; when calculating the variance of quadratic estimators all permutations of the observed trispectrum must be taken into account. One such permutation reflects the desired correlation making our estimator sensitive to $\mathbf{d}(\mathbf{L})$, but the remaining two permutations induce additional variance proportional to the lensing-potential power spectrum $C_L^{\phi\phi}$. While subdominant, this variance will become increasingly significant for future experiments as shown in the right-hand panel of Fig. 2.2. Since the power spectrum $C_L^{\phi\phi}$ is itself a measure of uncertainty in the deflection field, this additional variance in lensing reconstruction acts as a bias during power-spectrum estimation because there is no *a priori* way to distinguish it from the intrinsic variance of the underlying distribution. Our calculation of the dependence of this variance on $C_L^{\phi\phi}$ allows it, in principle, to be subtracted iteratively, which will prevent a systematic $5 - 10\%$ overestimate of $C_L^{\phi\phi}$ at low L.

We close by considering several possible observational obstacles to the scheme for lensing reconstruction and power-spectrum estimation presented above. One hindrance is secondary contributions to the CMB such as the SZ and ISW effects. These effects increase the total temperature power spectrum appearing in the denominator of the optimum filter $F_{\Theta\Theta}(\mathbf{l}_1, \mathbf{l}_2)$ of Eq. (2.14) as would additional instrumental noise. They also correlate with the large-scale structure at low redshifts inducing further non-Gaussian couplings and additional variance to lensing reconstruction. Fortunately for our purposes, the frequency dependence of the thermal SZ effect differs from that of a blackbody. It can therefore be separated in principle from the lensed primordial CMB by an experiment with several frequency channels (Cooray et al., 2000). The ISW effect cannot be removed in this manner, but is too small to significantly inhibit lensing reconstruction. Polarization-dependent secondary effects are expected to appear at higher orders in the density contrast (Hu, 2000b), and we therefore anticipate that they will not make a contribution at the levels considered here. A potentially more serious problem is that of galactic foregrounds which, though uncorrelated with the lensing signal, may be substantial at certain frequencies. Significant polarization has also been observed in some of these sources (Giardino et al., 2002; Bouchet et al., 1999; Prunet et al., 2000). We hope to understand and minimize the effects of galactic foregrounds in future work, and to pursue further refinements of lensing reconstruction.

We thank Wayne Hu for useful discussions. This work was supported in part by NASA NAG5-11985 and DoE DE-FG03-92-ER40701. Kesden acknowledges the support of an NSF Graduate Fellowship and AC acknowledges support from the Sherman Fairchild Foundation.

2.6 Polarization-based Estimators

Here we provide the appropriate formulas for deriving the variance associated with polarization-based estimators of the deflection field $\mathbf{d}(\mathbf{L})$. The CMB polarization can be decomposed into E and B-modes (Kamionkowski et al., 1997a; Seljak & Zaldarriaga, 1997). These modes are mixed by weak lensing such that to linear order in $\phi(\mathbf{L})$,

$$\begin{aligned}
\tilde{E}(\mathbf{l}) &= E(\mathbf{l}) \\
&- \int \frac{d^2\mathbf{l}_1}{(2\pi)^2} \left[E(\mathbf{l}_1)\cos 2(\varphi_{\mathbf{l}_1} - \varphi_{\mathbf{l}}) - B(\mathbf{l}_1)\sin 2(\varphi_{\mathbf{l}_1} - \varphi_{\mathbf{l}})\right] \phi(\mathbf{l}-\mathbf{l}_1)\left[(\mathbf{l}-\mathbf{l}_1)\cdot\mathbf{l}_1\right] , \\
\tilde{B}(\mathbf{l}) &= B(\mathbf{l}) \\
&- \int \frac{d^2\mathbf{l}_1}{(2\pi)^2} \left[E(\mathbf{l}_1)\sin 2(\varphi_{\mathbf{l}_1} - \varphi_{\mathbf{l}}) + B(\mathbf{l}_1)\cos 2(\varphi_{\mathbf{l}_1} - \varphi_{\mathbf{l}})\right] \\
&\quad \times \phi(\mathbf{l}-\mathbf{l}_1)\left[(\mathbf{l}-\mathbf{l}_1)\cdot\mathbf{l}_1\right] . \qquad (2.45)
\end{aligned}$$

We can exploit the sensitivity of the polarization modes to the lensing potential to construct lensing estimators from quadratic combinations of polarization modes. We generalize Eq. (2.11) to arbitrary combinations $\{X, X'\}$ of Θ, E, and B-modes as first derived in Hu & Okamoto (2002)

$$\langle X^{\rm t}(\mathbf{l}) X'^{\rm t}(\mathbf{l}')\rangle_{\rm CMB} = f_{XX'}(\mathbf{l},\mathbf{l}')\phi(\mathbf{L})\,, \qquad (2.46)$$

where

$$\begin{aligned}
f_{\Theta E}(\mathbf{l},\mathbf{l}') &= C_l^{\Theta E}\cos 2(\varphi_{\mathbf{l}} - \varphi_{\mathbf{l}'})(\mathbf{L}\cdot\mathbf{l}) + C_{l'}^{\Theta E}(\mathbf{L}\cdot\mathbf{l}') = f_{E\Theta}(\mathbf{l}',\mathbf{l})\,, \\
f_{\Theta B}(\mathbf{l},\mathbf{l}') &= C_l^{\Theta E}\sin 2(\varphi_{\mathbf{l}} - \varphi_{\mathbf{l}'})(\mathbf{L}\cdot\mathbf{l}) = f_{B\Theta}(\mathbf{l}',\mathbf{l})\,, \\
f_{EE}(\mathbf{l},\mathbf{l}') &= \left[C_l^{EE}(\mathbf{L}\cdot\mathbf{l}) + C_{l'}^{EE}(\mathbf{L}\cdot\mathbf{l}')\right]\cos 2(\varphi_{\mathbf{l}} - \varphi_{\mathbf{l}'})\,, \\
f_{EB}(\mathbf{l},\mathbf{l}') &= \left[C_l^{EE}(\mathbf{L}\cdot\mathbf{l}) + C_{l'}^{BB}(\mathbf{L}\cdot\mathbf{l}')\right]\sin 2(\varphi_{\mathbf{l}} - \varphi_{\mathbf{l}'}) = f_{BE}(\mathbf{l}',\mathbf{l})\,, \\
f_{BB}(\mathbf{l},\mathbf{l}') &= \left[C_l^{BB}(\mathbf{L}\cdot\mathbf{l}) + C_{l'}^{BB}(\mathbf{L}\cdot\mathbf{l}')\right]\cos 2(\varphi_{\mathbf{l}} - \varphi_{\mathbf{l}'})\,.
\end{aligned}$$

$$(2.47)$$

In deriving these results, we used parity considerations to demand $C_l^{\Theta B} = C_l^{EB} = 0$. Using these relations, we follow the approach of Ref. Hu & Okamoto (2002) to derive symmetric lensing estimators

$$\mathbf{d}_{XX}(\mathbf{L}) \equiv \frac{i\mathbf{L}A_{XX}(L)}{L^2}\int \frac{d^2\mathbf{l}_1}{(2\pi)^2} X^{\rm t}(\mathbf{l}_1)X^{\rm t}(\mathbf{l}_2)F_{XX}(\mathbf{l}_1,\mathbf{l}_2)\,, \tag{2.48}$$

and

$$\mathbf{d}_{XX'}(\mathbf{L}) \equiv \frac{i\mathbf{L}A_{XX'}(L)}{L^2}\int \frac{d^2\mathbf{l}_1}{(2\pi)^2}\frac{1}{2}\left[X^{\rm t}(\mathbf{l}_1)X'^{\rm t}(\mathbf{l}_2) + X'^{\rm t}(\mathbf{l}_1)X^{\rm t}(\mathbf{l}_2)\right]F_{XX'}(\mathbf{l}_1,\mathbf{l}_2)\,. \tag{2.49}$$

We have explicitly symmetrized our estimators for $X \neq X'$ to simplify the form of the optimal filter. The normalization bias of the estimators is removed by choosing

$$A_{XX}(L) \equiv L^2\left[\int \frac{d^2\mathbf{l}_1}{(2\pi)^2} f_{XX}(\mathbf{l}_1,\mathbf{l}_2)F_{XX}(\mathbf{l}_1,\mathbf{l}_2)\right]^{-1}\,, \tag{2.50}$$

and

$$A_{XX'}(L) \equiv L^2\left[\int \frac{d^2\mathbf{l}_1}{(2\pi)^2}\frac{1}{2}\left[f_{XX'}(\mathbf{l}_1,\mathbf{l}_2) + f_{XX'}(\mathbf{l}_2,\mathbf{l}_1)\right]F_{XX'}(\mathbf{l}_1,\mathbf{l}_2)\right]^{-1}\,. \tag{2.51}$$

The minimum-variance filters $F_{XX'}(\mathbf{l}_1,\mathbf{l}_2)$ for the various cases $\{X, X'\}$ are given by

$$\begin{aligned}
F_{\Theta E}(\mathbf{l}_1,\mathbf{l}_2) &= \frac{f_{\Theta E}(\mathbf{l}_1,\mathbf{l}_2) + f_{\Theta E}(\mathbf{l}_2,\mathbf{l}_1)}{C_{l_1}^{\Theta\Theta{\rm t}}C_{l_2}^{EE{\rm t}} + 2C_{l_1}^{\Theta E{\rm t}}C_{l_2}^{\Theta E{\rm t}} + C_{l_1}^{EE{\rm t}}C_{l_2}^{\Theta\Theta{\rm t}}}\,,\\
F_{\Theta B}(\mathbf{l}_1,\mathbf{l}_2) &= \frac{f_{\Theta B}(\mathbf{l}_1,\mathbf{l}_2) + f_{\Theta B}(\mathbf{l}_2,\mathbf{l}_1)}{C_{l_1}^{\Theta\Theta{\rm t}}C_{l_2}^{EE{\rm t}} + C_{l_1}^{EE{\rm t}}C_{l_2}^{\Theta\Theta{\rm t}}}\,,\\
F_{EE}(\mathbf{l}_1,\mathbf{l}_2) &= \frac{f_{EE}(\mathbf{l}_1,\mathbf{l}_2)}{2C_{l_1}^{EE{\rm t}}C_{l_2}^{EE{\rm t}}}\,,\\
F_{EB}(\mathbf{l}_1,\mathbf{l}_2) &= \frac{f_{EB}(\mathbf{l}_1,\mathbf{l}_2) + f_{EB}(\mathbf{l}_2,\mathbf{l}_1)}{C_{l_1}^{EE{\rm t}}C_{l_2}^{BB{\rm t}} + C_{l_1}^{BB{\rm t}}C_{l_2}^{EE{\rm t}}}\,,\\
F_{BB}(\mathbf{l}_1,\mathbf{l}_2) &= \frac{f_{BB}(\mathbf{l}_1,\mathbf{l}_2)}{2C_{l_1}^{BB{\rm t}}C_{l_2}^{BB{\rm t}}}\,.
\end{aligned} \tag{2.52}$$

Using these optimal filters for the estimators defined in Eqs. (2.48) and (2.49), we can calculate the variances for these estimators in a fashion entirely analogous to

Eq. (2.19),

$$\begin{aligned}
&\Big\langle \langle \mathbf{d}^*_{XX}(\mathbf{L}) \cdot \mathbf{d}_{XX}(\mathbf{L}')\rangle_{\rm CMB} - \langle \mathbf{d}^*_{XX}(\mathbf{L})\rangle_{\rm CMB}\cdot\langle \mathbf{d}_{XX}(\mathbf{L}')\rangle_{\rm CMB}\Big\rangle_{\rm LSS} = \\
&\quad \frac{A_{XX}(L)}{L}\frac{A_{XX}(L')}{L'}\int\frac{d^2\mathbf{l}_1}{(2\pi)^2}\int\frac{d^2\mathbf{l}'_1}{(2\pi)^2}\langle X^{\rm t}(-\mathbf{l}_1)X^{\rm t}(-\mathbf{l}_2)X^{\rm t}(\mathbf{l}_1{}')X^{\rm t}(\mathbf{l}_2{}')\rangle \\
&\quad \times F_{XX}(\mathbf{l}_1,\mathbf{l}_2)F_{XX}(\mathbf{l}_1{}',\mathbf{l}_2{}') - (2\pi)^2\delta_{\rm D}(\mathbf{L}-\mathbf{L}')C_L^{dd}\,,
\end{aligned} \tag{2.53}$$

$$\begin{aligned}
&\Big\langle \langle \mathbf{d}^*_{XX'}(\mathbf{L}) \cdot \mathbf{d}_{XX'}(\mathbf{L}')\rangle_{\rm CMB} - \langle \mathbf{d}^*_{XX'}(\mathbf{L})\rangle_{\rm CMB}\cdot\langle \mathbf{d}_{XX'}(\mathbf{L}')\rangle_{\rm CMB}\Big\rangle_{\rm LSS} = \\
&\quad \frac{A_{XX'}(L)}{L}\frac{A_{XX'}(L')}{L'}\int\frac{d^2\mathbf{l}_1}{(2\pi)^2}\int\frac{d^2\mathbf{l}'_1}{(2\pi)^2}\frac{1}{4}\langle\big[X^{\rm t}(-\mathbf{l}_1)X'^{\rm t}(-\mathbf{l}_2)+X'^{\rm t}(-\mathbf{l}_1)X^{\rm t}(-\mathbf{l}_2)\big] \\
&\quad \times\big[X^{\rm t}(\mathbf{l}_1{}')X'^{\rm t}(\mathbf{l}_2{}')+X'^{\rm t}(\mathbf{l}_1{}')X^{\rm t}(\mathbf{l}_2{}')\big]\rangle F_{XX'}(\mathbf{l}_1,\mathbf{l}_2)F_{XX'}(\mathbf{l}_1{}',\mathbf{l}_2{}') \\
&\qquad -(2\pi)^2\delta_{\rm D}(\mathbf{L}-\mathbf{L}')C_L^{dd}\,.
\end{aligned} \tag{2.54}$$

As for that of the temperature estimator, these variances will consist of zeroth-order terms in $C_L^{\phi\phi}$, $N^{(0)}_{XX,XX}(L) = A_{XX}(L)$ and $N^{(0)}_{XX',XX'}(L) = A_{XX'}(L)$, and first-order terms,

$$\begin{aligned}
N^{(1)}_{XX,XX}(L) &= \frac{A^2_{XX}(L)}{L^2}\int\frac{d^2\mathbf{l}_1}{(2\pi)^2}\int\frac{d^2\mathbf{l}'_1}{(2\pi)^2}F_{XX}(\mathbf{l}_1,\mathbf{l}_2)F_{XX}(\mathbf{l}_1{}',\mathbf{l}_2{}') \\
&\quad\times\Big\{C^{\phi\phi}_{|\mathbf{l}_1-\mathbf{l}_1{}'|}f_{XX}(-\mathbf{l}_1,\mathbf{l}_1{}')f_{XX}(-\mathbf{l}_2,\mathbf{l}_2{}')+C^{\phi\phi}_{|\mathbf{l}_1-\mathbf{l}_2{}'|}f_{XX}(-\mathbf{l}_1,\mathbf{l}_2{}')f_{XX}(-\mathbf{l}_2,\mathbf{l}_1{}')\Big\}\,,
\end{aligned} \tag{2.55}$$

$$
\begin{aligned}
N^{(1)}_{XX',XX'}(L) = \frac{A^2_{XX'}(L)}{L^2} \int \frac{d^2\mathbf{l}_1}{(2\pi)^2} \int \frac{d^2\mathbf{l}_1'}{(2\pi)^2} & F_{XX'}(\mathbf{l}_1,\mathbf{l}_2) F_{XX'}(\mathbf{l}_1{}',\mathbf{l}_2{}') \\
\times \frac{1}{4}\Big\{ C^{\phi\phi}_{|\mathbf{l}_1-\mathbf{l}_1{}'|} \big[& f_{XX}(-\mathbf{l}_1,\mathbf{l}_1{}') f_{X'X'}(-\mathbf{l}_2,\mathbf{l}_2{}') + f_{XX'}(-\mathbf{l}_1,\mathbf{l}_1{}') f_{X'X}(-\mathbf{l}_2,\mathbf{l}_2{}') \\
& + f_{X'X}(-\mathbf{l}_1,\mathbf{l}_1{}') f_{XX'}(-\mathbf{l}_2,\mathbf{l}_2{}') + f_{X'X'}(-\mathbf{l}_1,\mathbf{l}_1{}') f_{XX}(-\mathbf{l}_2,\mathbf{l}_2{}') \big] \\
+ C^{\phi\phi}_{|\mathbf{l}_1-\mathbf{l}_2{}'|} \big[& f_{XX'}(-\mathbf{l}_1,\mathbf{l}_2{}') f_{X'X}(-\mathbf{l}_2,\mathbf{l}_1{}') + f_{XX}(-\mathbf{l}_1,\mathbf{l}_2{}') f_{X'X'}(-\mathbf{l}_2,\mathbf{l}_1{}') \\
& + f_{X'X'}(-\mathbf{l}_1,\mathbf{l}_2{}') f_{XX}(-\mathbf{l}_2,\mathbf{l}_1{}') + f_{X'X}(-\mathbf{l}_1,\mathbf{l}_2{}') f_{XX'}(-\mathbf{l}_2,\mathbf{l}_1{}') \big] \Big\} ,
\end{aligned}
\tag{2.56}
$$

The filters given in Eq. (2.52) are no longer optimal in the presence of this additional noise, but the difference between these filters and the optimal filters should be negligible, provided that $N^{(1)}_{XX,XX}(L) \ll N^{(0)}_{XX,XX}(L)$, $N^{(1)}_{XX',XX'}(L) \ll N^{(0)}_{XX',XX'}(L)$. For the purposes of power-spectrum estimation, the terms $N^{(1)}_{XX,XX}(L)$ and $N^{(1)}_{XX',XX'}(L)$ are not only an additional contribution to the variance, but are also a systematic bias if not subtracted iteratively, following Eq. (2.29).

A final point to consider is that the six different estimators $\mathbf{d}_{XX'}(\mathbf{L})$ defined in this paper are not independent, as they are constructed from only three distinct maps. The covariance matrix for the six estimators will therefore not be diagonal, and this needs to be taken into account if the estimators are to be linearly combined to produce a single minimum-variance estimator. The off-diagonal elements of the covariance matrix can be evaluated in a straightforward manner involving pairs of double integrals similar to those of Eqs. (2.53) and (2.54).

Chapter 3

Cumulants of the Probability Distribution Function

Reprinted with permission from M. Kesden, A. Cooray, and M. Kamionkowski Phys. Rev. D **66**, 083007 (2002).

We discuss the real-space moments of temperature anisotropies in the cosmic microwave background (CMB) due to weak gravitational lensing by intervening large-scale structure. We show that if the probability distribution function of primordial temperature anisotropies is Gaussian, then it remains unchanged after gravitational lensing. With finite resolution, however, non-zero higher-order cumulants are generated both by lensing auto-correlations and by cross-correlations between the lensing potential and secondary anisotropies in the CMB, such as the Sunayev-Zel'dovich (SZ) effect. Skewness is produced by these lensing-SZ correlations, while kurtosis receives contributions from both lensing alone and lensing-SZ correlations. We show that if the projected lensing potential is Gaussian, all cumulants of higher-order than the kurtosis vanish. While recent results raise the possibility of detection of the skewness in upcoming data, the kurtosis will likely remain undetected.

3.1 Introduction

Weak gravitational lensing deflects the paths of cosmic microwave background (CMB) photons propagating from the surface of last scattering. One result of this lensing is the transfer of power from large angular scales associated with acoustic-peak structures to small angular scales in the damping tail of the anisotropy power spectrum (Blanchard & Schneider, 1987; Kashlinsky, 1988; Linder, 1988; Cayon et al., 1993; Seljak, 1996; Hu, 2000a). This transfer only results in a few-percent modification of the power associated with the acoustic-peak structure, and the increase in power along the damping tail is significantly smaller than that generated by secondary anisotropies due to reionization (Cooray, 2001b). To indentify the effect of gravitational lensing on CMB data, it is necessary to consider signatures beyond that in the angular power spectrum of temperature fluctuations. The existence of non-vanishing higher order cumulants is one such non-Gaussian signature lensing can generate.

Since gravitational lensing conserves surface brightness, CMB fluctuations from lensing are at the second order in temperature fluctuations and result in non-Gaussian behavior through non-linear mode coupling. Though lensing alone does not lead to a three-point correlation function, the correlation between lensing and other secondary anisotropies can lead to such a contribution. This three-point correlation has been widely discussed in the literature in terms of its Fourier-space analogue, the bispectrum (Goldberg & Spergel, 1999; Zaldarriaga & Seljak, 1999; Peiris & Spergel, 2000; Cooray & Hu, 2000). Weak lensing of the primary anisotropies can produce a four-point correlation due to its non-linear mode-coupling nature (Bernardeau, 1997; Zaldarriaga, 2000; Hu, 2001b; Cooray, 2002), as can correlations between lensing and secondary effects (Cooray, 2002). When probed appropriately through quadratic statistics such as the power spectrum of the squared-temperature map, the trispectrum due to lensing alone can be used for a model-independent recovery of the projected mass distribution out to the last scattering surface (Hu, 2001a; Hu & Okamoto, 2002; Benabed et al., 2001; Cooray & Kesden, 2003). Though these statistics have been shown to be interesting and potentially detectable, measurement of these Fourier-based statistics is challenging and techniques are still underdeveloped for this purpose.

Here, we discuss real-space moments of the lensed CMB temperature anisotropies. Real-space statistics are easily measurable from data. The only drawbacks are that they are unlikely to be optimal and only provide limited knowledge of the full non-

Gaussian aspect of the temperature distribution. The first attempts to measure non-Gaussianity in the COBE data relied on real-space cumulants (Hinshaw et al., 1995; Seljak & Zaldarriaga, 1999a), as will attempts using data from its successor experiments such as MAP and Planck. This motivates our emphasis here on the real-space cumulants such as the skewness and kurtosis; we make several remarks on higher-order cumulants as well.

As part of this calculation, we extend a previous discussion of the kurtosis due to lensing in Bernardeau (1997), and also consider effects related to correlations between lensing and secondary effects such as the Sunyaev-Zel'dovich effect (Sunyaev & Zel'dovich, 1980).

Real-space moments can be derived from the one-point probability distribution function (PDF) of temperature fluctuations, and can conversely be used to constrain the form of this function. In the case of infinite angular resolution, we conclude that lensing does not modify the PDF of temperature anisotropies produced at the last scattering surface, which is a reflection on the fact that lensing does not create new power, but rather, transfers power from large to small angular scales. The higher-order moments are only generated in a temperature map by finite-resolution effects such as beam smoothing introduced either experimentally or artificially by explicit filtering.

The paper is organized as follows. In § 3.2, we introduce formalism concerning the weak-lensing approximation and define the bispectrum, trispectrum, and corresponding higher-order quantities. The bispectrum and trispetrum induced in the CMB by lensing and secondary anisotropies are derived in § 3.3, and some remarks are made concerning higher-order cumulants as well. The nonzero bispectrum and trispectrum yield a skewness and kurtosis, respectively, in the one-point distribution function of the CMB as shown in § 3.4. We refer the reader to Ref. Cooray & Kesden (2003) for additional details related to the effect of lensing on CMB anisotropies. Though we present a general discussion, we illustrate our results in § 3.5 using the currently favored ΛCDM cosmological model with $\Omega_b = 0.05$, $\Omega_m = 0.35$, $\Omega_\Lambda = 0.65$, h=0.65, and $\sigma_8 = 0.9$. Results for a model with $\sigma_8 = 1.2$ as suggested by CBI are also considered.

3.2 Lensing Contribution to CMB Fluctuations

In order to derive the effects of weak lensing on the CMB, we follow Hu (2000a) and Zaldarriaga (2000) and adopt a flat-sky approximation. As discussed in prior papers (Hu, 2000a; Goldberg & Spergel, 1999), weak lensing remaps temperature through angular deflections along the photon path:

$$\begin{aligned}\tilde{\Theta}(\hat{\mathbf{n}}) &= \Theta(\hat{\mathbf{n}} + \nabla\phi) \\ &= \Theta(\hat{\mathbf{n}}) + \nabla_i\phi(\hat{\mathbf{n}})\nabla^i\Theta(\hat{\mathbf{n}}) + \frac{1}{2}\nabla_i\phi(\hat{\mathbf{n}})\nabla_j\phi(\hat{\mathbf{n}})\nabla^i\nabla^j\Theta(\hat{\mathbf{n}}) + \ldots . \qquad (3.1)\end{aligned}$$

Here, $\Theta(\hat{\mathbf{n}})$ is the unlensed primary component of the CMB in direction $\hat{\mathbf{n}}$ at the last-scattering surface, $\tilde{\Theta}(\hat{\mathbf{n}})$ is the lensed map, $\phi(\hat{\mathbf{n}})$ is the projected gravitational potential, and $\nabla\phi$ is the lensing deflection angle. It should be understood that in the presence of low-redshift contributions to CMB fluctuations resulting from large-scale structure, the total map includes secondary contributions which we denote by $\Theta^{\mathrm{s}}(\hat{\mathbf{n}})$. Since the weak-lensing deflection angles $\nabla\phi$ also trace the large-scale structure at low redshifts, secondary effects which are first order in density fluctuations correlate with the lensing deflection angles. These secondary effects include the integrated Sachs-Wolfe (ISW; (Sachs & Wolfe, 1967)) and the SZ (Sunyaev & Zel'dovich, 1980) effects (Goldberg & Spergel, 1999). In all real cases, a noise component denoted by $\Theta^{\mathrm{n}}(\hat{\mathbf{n}})$ due to finite experimental sensitivity must be included as well. Thus, the total observed CMB anisotropy will be $\Theta^{\mathrm{t}}(\hat{\mathbf{n}}) = \tilde{\Theta}(\hat{\mathbf{n}}) + \Theta^{\mathrm{s}}(\hat{\mathbf{n}}) + \Theta^{\mathrm{n}}(\hat{\mathbf{n}})$. In the following discussion, secondary anisotropies $\Theta^{\mathrm{s}}(\hat{\mathbf{n}})$ will be neglected until subsection 3.3.2, while the effects of instrumental noise $\Theta^{\mathrm{n}}(\hat{\mathbf{n}})$ on the PDF are discussed in § 3.4.

Taking the Fourier transform, as appropriate for a flat sky, we write

$$\begin{aligned}\tilde{\Theta}(\mathbf{l}) &= \int d\hat{\mathbf{n}}\,\tilde{\Theta}(\hat{\mathbf{n}})e^{-i\mathbf{l}\cdot\hat{\mathbf{n}}} \\ &= \Theta(\mathbf{l}) - \int \frac{d^2\mathbf{l}'}{(2\pi)^2}\Theta(\mathbf{l}')L(\mathbf{l},\mathbf{l}'), \qquad (3.2)\end{aligned}$$

where

$$\begin{aligned} L(\mathbf{l},\mathbf{l}') &= \phi(\mathbf{l}-\mathbf{l}')\ [(\mathbf{l}-\mathbf{l}')\cdot\mathbf{l}'] + \frac{1}{2}\int \frac{d^2\mathbf{l}''}{(2\pi)^2}\phi(\mathbf{l}'') \\ &\quad \times \phi^*(\mathbf{l}''+\mathbf{l}'-\mathbf{l})\,(\mathbf{l}''\cdot\mathbf{l}')\,[(\mathbf{l}''+\mathbf{l}'-\mathbf{l})\cdot\mathbf{l}'] + \ldots\,. \end{aligned} \tag{3.3}$$

We define the power spectrum, bispectrum, trispectrum and the n-point correlator in Fourier space in the usual way:

$$\begin{aligned} \left\langle \tilde\Theta(\mathbf{l}_1)\tilde\Theta(\mathbf{l}_2)\right\rangle &= (2\pi)^2\delta_{\rm D}(\mathbf{l}_{12})\tilde C^{\Theta}_{l_1}\,, \\ \left\langle \tilde\Theta(\mathbf{l}_1)\ldots\tilde\Theta(\mathbf{l}_3)\right\rangle_c &= (2\pi)^2\delta_{\rm D}(\mathbf{l}_{123})\tilde B^{\Theta}(\mathbf{l}_1,\mathbf{l}_2,\mathbf{l}_3)\,, \\ \left\langle \tilde\Theta(\mathbf{l}_1)\ldots\tilde\Theta(\mathbf{l}_4)\right\rangle_c &= (2\pi)^2\delta_{\rm D}(\mathbf{l}_{1234})\tilde T^{\Theta}(\mathbf{l}_1,\mathbf{l}_2,\mathbf{l}_3,\mathbf{l}_4)\,, \\ \left\langle \tilde\Theta(\mathbf{l}_1)\ldots\tilde\Theta(\mathbf{l}_n)\right\rangle_c &= (2\pi)^2\delta_{\rm D}(\mathbf{l}_{1\ldots n})\tilde T^{\Theta}_n(\mathbf{l}_1,\ldots,\mathbf{l}_n)\,. \end{aligned} \tag{3.4}$$

where $\mathbf{l}_{1\ldots n} \equiv \mathbf{l}_1 + \ldots + \mathbf{l}_n$, and the subscript c denotes the connected portion of the correlation function. We make the assumption that primary anisotropies at the last scattering surface are Gaussian, implying that all cumulants higher than the power spectrum vanish: $\langle\Theta(\mathbf{l}_1)\ldots\Theta(\mathbf{l}_n)\rangle_c = 0$, when $n > 2$.

The power spectra for lensing correlations and lensing-secondary cross-correlations are defined analogously:

$$\begin{aligned} \langle\phi(\mathbf{l}_1)\phi(\mathbf{l}_2)\rangle &= (2\pi)^2\delta_{\rm D}(\mathbf{l}_{12})\,C^{\phi\phi}_{l_1}\,, \\ \langle\phi(\mathbf{l}_1)\Theta^{\rm s}(\mathbf{l}_2)\rangle &= (2\pi)^2\delta_{\rm D}(\mathbf{l}_{12})\,C^{\phi{\rm s}}_{l_1}\,. \end{aligned} \tag{3.5}$$

Primary CMB anisotropies $\Theta(\hat{\mathbf{n}})$ are generated at the surface of last scatter at $z \simeq 1,100$, while the lensing potential $\phi(\hat{\mathbf{n}})$ and secondary contributions $\Theta^{\rm s}(\hat{\mathbf{n}})$ arise from large-scale structure at much lower redshifts ($z \simeq 3$). As such, correlations between these quantities vanish: $\langle\Theta(\mathbf{l}_1)\phi(\mathbf{l}_2)\rangle = \langle\Theta(\mathbf{l}_1)\Theta^{\rm s}(\mathbf{l}_2)\rangle = 0$.

The nth cumulant of the temperature anisotropies is defined in the usual manner,

$$C^n(\theta) = \int \frac{d^2\mathbf{l}_1}{(2\pi)^2} \dots \frac{d^2\mathbf{l}_n}{(2\pi)^2} \left\langle \Theta^{\rm t}(\mathbf{l}_1) \dots \Theta^{\rm t}(\mathbf{l}_n) \right\rangle_c W(l_1\theta) \dots W(l_n\theta), \tag{3.6}$$

where θ is the smoothing scale of the map from which the cumulants are determined, and $W(l\theta)$ is the smoothing window function. We will use Gaussian window functions throughout this paper. In general, the finite resolution of real CMB anisotropy experiments induces Gaussian smoothing at the angular scale of the experimental beam size. For infinite resolution, we take $\theta \to 0$ such that $W(l\theta) \to 1$. The variance, skewness, and kurtosis defined later in this paper can all be expressed in terms of cumulants:

$$\sigma^2(\theta) = C^2(\theta)\,, \quad S(\theta) = \frac{C^3(\theta)}{[C^2(\theta)]^{3/2}}\,, \quad K(\theta) = \frac{C^4(\theta)}{[C^2(\theta)]^2}\,. \tag{3.7}$$

3.3 Power spectrum, Bispectrum and Trispectrum

Using the formalism introduced in the previous Section, we can calculate the moments of the CMB fluctuations generated by lensing assuming Gaussian fluctuations at the surface of last scatter. Because ϕ is a small parameter, terms beyond linear order in $C_l^{\phi\phi}$ are neglected in these calculations. The power spectrum for the lensed map is (Blanchard & Schneider, 1987; Kashlinsky, 1988; Linder, 1988; Cayon et al., 1993; Seljak, 1996; Hu, 2000a)

$$\tilde{C}_l^{\Theta} = \left[1 - \int \frac{d^2\mathbf{l}_1}{(2\pi)^2} C_{l_1}^{\phi\phi} \left(\mathbf{l}_1 \cdot \mathbf{l}\right)^2\right] C_l^{\Theta} + \int \frac{d^2\mathbf{l}_1}{(2\pi)^2} C_{|\mathbf{l}-\mathbf{l}_1|}^{\Theta} C_{l_1}^{\phi\phi} [(\mathbf{l}-\mathbf{l}_1)\cdot\mathbf{l}_1]^2\,. \tag{3.8}$$

The variance, or the second moment of the temperature, can be obtained following Eq. (3.6)

$$\sigma^2(\theta) = \int \frac{d^2\mathbf{l}}{(2\pi)^2} \tilde{C}_l^{\Theta} W^2(l\theta)\,. \tag{3.9}$$

Substituting Eq. (3.8) in here, we find that in the case of infinite resolution ($W(l\theta) = 1$), the variance of the lensed temperature map coincides with that of the unlensed map. Thus, as expected, lensing conserves the total power associated with the temperature fluctuations. This is consistent with our basic expectation that lensing only results in a transfer of power from large angular scales to small angular scales. With finite resolution at levels considered here, the variance of the lensed temperature field differs from that of the unlensed field by a few percent at most.

We will now discuss higher-order correlations of temperature due to gravitational lensing. We consider first contributions due to lensing alone, and then discuss additional contributions created by lensing-secondary correlations.

3.3.1 Lensing Correlations

We will first discuss the temperature bispectrum and show that it is zero in the absence of secondary anisotropies. To understand why there is no contribution to the bispectrum, consider the moments involving three temperature terms in Fourier space:

$$\begin{aligned}
\left\langle \tilde\Theta(\mathbf{l}_1)\tilde\Theta(\mathbf{l}_2)\tilde\Theta(\mathbf{l}_3)\right\rangle_c &= \\
&= \Big\langle \left(\Theta(\mathbf{l}_1) - \int \frac{d^2\mathbf{l}_1'}{(2\pi)^2}\Theta(\mathbf{l}_1{}')L(\mathbf{l}_1,\mathbf{l}_1{}')\right) \\
&\quad \times \left(\Theta(\mathbf{l}_2) - \int \frac{d^2\mathbf{l}_2'}{(2\pi)^2}\Theta(\mathbf{l}_2{}')L(\mathbf{l}_2,\mathbf{l}_2{}')\right) \\
&\quad \times \left(\Theta(\mathbf{l}_3) - \int \frac{d^2\mathbf{l}_3'}{(2\pi)^2}\Theta(\mathbf{l}_3{}')L(\mathbf{l}_3,\mathbf{l}_3{}')\right)\Big\rangle \\
&= \Big\langle \Theta(\mathbf{l}_1)\Theta(\mathbf{l}_2)\Theta(\mathbf{l}_3)\Big\rangle \\
&\quad - \Big\langle \Theta(\mathbf{l}_1)\Theta(\mathbf{l}_2)\left(\int \frac{d^2\mathbf{l}_3'}{(2\pi)^2}\Theta(\mathbf{l}_3{}')L(\mathbf{l}_3,\mathbf{l}_3{}')\right)\Big\rangle + \text{Perm.} \\
&\quad + \Big\langle \Theta(\mathbf{l}_1)\left(\int \frac{d^2\mathbf{l}_2'}{(2\pi)^2}\Theta(\mathbf{l}_2{}')L(\mathbf{l}_2,\mathbf{l}_2{}')\right) \\
&\quad \times \quad \left(\int \frac{d^2\mathbf{l}_3'}{(2\pi)^2}\Theta(\mathbf{l}_3{}')L(\mathbf{l}_3,\mathbf{l}_3{}')\right)\Big\rangle + \text{Perm.} \\
&\quad - \Big\langle \left(\int \frac{d^2\mathbf{l}_1'}{(2\pi)^2}\Theta(\mathbf{l}_1{}')L(\mathbf{l}_1,\mathbf{l}_1{}')\right)\left(\int \frac{d^2\mathbf{l}_2'}{(2\pi)^2}\Theta(\mathbf{l}_2{}')L(\mathbf{l}_2,\mathbf{l}_2{}')\right) \\
&\quad \times \left(\int \frac{d^2\mathbf{l}_3'}{(2\pi)^2}\Theta(\mathbf{l}_3{}')L(\mathbf{l}_3,\mathbf{l}_3{}')\right)\Big\rangle . \qquad (3.10)
\end{aligned}$$

All these terms, and the necessary permutations, involve an expectation value of three primary temperature anisotropies. Under our assumption of Gaussian primary temperature fluctuations, such expectation values vanish and thus, there is no contribution to the bispectrum or the skewness.

The trispectrum due to lensing alone can be calculated in a similar fashion. Introducing the power spectrum of lensing potentials, following Zaldarriaga (2000) and Cooray & Kesden (2003), we obtain the CMB trispectrum due to gravitational lensing as:

$$\begin{aligned}\tilde{T}^{\Theta}(\mathbf{l}_1,\mathbf{l}_2,\mathbf{l}_3,\mathbf{l}_4) &= -C^{\Theta}_{l_3}C^{\Theta}_{l_4}\Big\{C^{\phi\phi}_{|\mathbf{l}_1+\mathbf{l}_3|}\left[(\mathbf{l}_1+\mathbf{l}_3)\cdot\mathbf{l}_3\right]\left[(\mathbf{l}_1+\mathbf{l}_3)\cdot\mathbf{l}_4\right] \\ &\quad + C^{\phi\phi}_{|\mathbf{l}_2+\mathbf{l}_3|}\left[(\mathbf{l}_2+\mathbf{l}_3)\cdot\mathbf{l}_3\right]\left[(\mathbf{l}_2+\mathbf{l}_3)\cdot\mathbf{l}_4\right]\Big\} + \text{Perm.}\,, \qquad (3.11)\end{aligned}$$

where the permutations now contain five additional terms with the replacement of (l_3, l_4) by any other pair.

We can generalize our discussion of the power spectrum, bispectrum, and trispectrum to that of the n-point correlation function in Fourier space. In the absence of secondary anisotropies that correlate directly with the lensing potential, the n-point correlation function will vanish for odd n for the same reason that lensing alone did not generate a bispectrum. All such terms would involve the expectation value of an odd number of temperature fluctuations, and under the assumption of Gaussian primary anisotropies, such expectation values must vanish. This statement applies in particular to the case when measurements of non-Gaussianity are made using CMB maps which have been cleaned *a priori* of secondary fluctuations using information such as the nonthermal frequency dependence of these fluctuations. We will discuss the case of secondary anisotropies in the next subsection.

The lowest even nth correlator after the trispectrum is the six-point correlation function in Fourier space. We can write the portion of the connected part of this correlation function containing the lowest-order contribution in ϕ as

$$\begin{aligned}\left\langle\tilde{\Theta}(\mathbf{l}_1)\ldots\tilde{\Theta}(\mathbf{l}_6)\right\rangle_c &= \Big\langle\left(\Theta(\mathbf{l}_1) - \int\frac{d^2\mathbf{l}_1'}{(2\pi)^2}\Theta(\mathbf{l}_1{}')L(\mathbf{l}_1,\mathbf{l}_1{}')\right)\ldots\\ &\quad\times\left(\Theta(\mathbf{l}_3) - \int\frac{d^2\mathbf{l}_3'}{(2\pi)^2}\Theta(\mathbf{l}_3{}')L(\mathbf{l}_3,\mathbf{l}_3{}')\right)\\ &\quad\times\Theta(\mathbf{l}_4)\Theta(\mathbf{l}_5)\Theta(\mathbf{l}_6)\Big\rangle + \text{Perm.}\\ &= -\Big\langle\int\frac{d^2\mathbf{l}_1'}{(2\pi)^2}\Theta(\mathbf{l}_1{}')L(\mathbf{l}_1,\mathbf{l}_1{}')\ldots\\ &\quad\times\int\frac{d^2\mathbf{l}_3'}{(2\pi)^2}\Theta(\mathbf{l}_3{}')L(\mathbf{l}_3,\mathbf{l}_3{}')\Theta(\mathbf{l}_4)\Theta(\mathbf{l}_5)\Theta(\mathbf{l}_6)\Big\rangle\\ &\quad+\text{Perm.}\end{aligned} \tag{3.12}$$

Simplifying further, we see that the lowest order contribution in ϕ thus involves

$$\begin{aligned}\left\langle\tilde{\Theta}(\mathbf{l}_1)\ldots\tilde{\Theta}(\mathbf{l}_6)\right\rangle_c &= C^{\Theta}_{l_4}C^{\Theta}_{l_5}C^{\Theta}_{l_6}\Big\langle\phi(\mathbf{l}_1+\mathbf{l}_4)\phi(\mathbf{l}_2+\mathbf{l}_5)\phi(\mathbf{l}_3+\mathbf{l}_6)\Big\rangle\\ &\quad\times\left[(\mathbf{l}_1+\mathbf{l}_4)\cdot\mathbf{l}_4\right]\left[(\mathbf{l}_2+\mathbf{l}_5)\cdot\mathbf{l}_5\right]\left[(\mathbf{l}_3+\mathbf{l}_6)\cdot\mathbf{l}_6\right]\\ &\quad+\text{Perm.}\end{aligned} \tag{3.13}$$

The connected part of the six-point correlation function in Fourier space is thus proportional to the bispectrum of lensing potentials. We can write

$$\begin{aligned}\tilde{T}^{\Theta}_6(\mathbf{l}_1,\mathbf{l}_2,\mathbf{l}_3,\mathbf{l}_4,\mathbf{l}_5,\mathbf{l}_6) &= C^{\Theta}_{l_4}C^{\Theta}_{l_5}C^{\Theta}_{l_6}\Big[B^{\phi}(\mathbf{l}_1+\mathbf{l}_4,\mathbf{l}_2+\mathbf{l}_5,\mathbf{l}_3+\mathbf{l}_6)\\ &\quad\times\left[(\mathbf{l}_1+\mathbf{l}_4)\cdot\mathbf{l}_4\right]\left[(\mathbf{l}_2+\mathbf{l}_5)\cdot\mathbf{l}_5\right]\left[(\mathbf{l}_3+\mathbf{l}_6)\cdot\mathbf{l}_6\right]\Big]\\ &\quad+\text{Perm.}\end{aligned} \tag{3.14}$$

There are in total 120 such terms appearing in the six-point correlator when we include all permutations, coming from the 20 different triplets (l_i, l_j, l_k) and the six permutations of each triplet.

We can generalize these derivations to the n-point temperature correlation in Fourier space under gravitational lensing. In the following, note that contributions to n-point temperature correlations in Fourier space come from $(n/2)$-point correlations in the lensing potential. We can thus write the connected part of the n-point temperature correlator, when $n > 2$, as

$$\begin{aligned}
\tilde{T}_n^{\Theta}(\mathbf{l}_1, \ldots, \mathbf{l}_n) &= C^{\Theta}_{l_{n/2+1}} \ldots C^{\Theta}_{l_n} \Big[T^{\phi}_{n/2}(\mathbf{l}_1 + \mathbf{l}_{(n/2)+1}, \ldots, \mathbf{l}_{n/2} + \mathbf{l}_n) \\
&\quad \times (\mathbf{l}_1 + \mathbf{l}_{(n/2)+1}) \cdot \mathbf{l}_{(n/2)+1} \ldots (\mathbf{l}_{n/2} + \mathbf{l}_n) \cdot \mathbf{l}_n \Big] \\
&\quad + \text{Perm.}\,, \qquad (3.15)
\end{aligned}$$

where $T_n^{\phi}(\mathbf{l}_1, \ldots, \mathbf{l}_n)$ is the n-point correlator of the lensing potential in Fourier space. The permutations here now involve $n!/(n/2)!$ terms corresponding to the replacement of $(l_{(n/2)+1}, ..., l_n)$ with one of the other $n!/[(n/2)!(n/2)!]$ combinations and the $(n/2)!$ permutations of each combination. As we have discussed, note that $\tilde{T}_n^{\Theta}(\mathbf{l}_1, \ldots, \mathbf{l}_n) = 0$ when n is odd.

In the limit that the lensing potentials are Gaussian distributed, $T_n^{\phi}(\mathbf{l}_1, \ldots, \mathbf{l}_n) = 0$ when $n > 2$. Thus, lensing of CMB anisotropies can only generate a trispectrum and, with smoothing, a kurtosis. The non-Gaussianity associated with the large-scale structure, however, will induce non-Gaussian contributions to the distribution of projected potentials such that $T_n^{\phi}(\mathbf{l}_1, \ldots, \mathbf{l}_n) \neq 0$ for some n. Since large-scale structure most efficiently lenses the CMB at redshifts close to 3, where the non-Gaussianity is mild, we ignore the higher-order correlations of lensing potentials and only consider the dominant power spectrum, $C_l^{\phi\phi}$, which contributes to the trispectrum only.

Although theoretical predictions are made in terms of ensemble-averaged correlation functions, observationally we have access to only one realization of the CMB and one realization of the large-scale structure. The arbitrariness of the observed realization of the large-scale structure induces additional cosmic variance beyond that normally associated with the surface of last-scatter. One consequence is that when measured on a small patch of the sky, the observed two-point correlation function of the lensed map is more anisotropic than that of the unlensed map, though isotropy holds when a sufficiently large region of the sky is considered. The excess anisotropy is induced by cosmic shear, and allows us to reconstruct the lensing deflection angle from quadratic maps involving the CMB temperature and polarization (Cooray & Kesden, 2003). While we emphasize one-point statistics in this paper, a more detailed account of how higher-order statistics probe the local anisotropy induced by lensing may prove fruitful in the future.

3.3.2 Lensing-Secondary Correlations

The above discussion applies to the case where other secondary fluctuations do not contribute to temperature anisotropies. In practice, such a situation can be achieved when thermal CMB fluctuations are separated from dominant secondary effects like the SZ contribution. In experiments where this is not possible, say due to a lack of multifrequency data, additional non-Gaussianities will be present in the CMB map due to correlations between lensing potentials and the secondary anisotropies. The most significant of these contributions is to the three-point correlation function. We can calculate this by replacing the $\tilde{\Theta}(\mathbf{l})$ terms in Eq. (3.10) with $\Theta^t(\mathbf{l})$. By assumption, Gaussian instrumental noise cannot generate a bispectrum, and as shown above, neither does lensing alone. The total observed bispectrum is therefore that due to lensing-secondary correlations (Goldberg & Spergel, 1999; Zaldarriaga, 2000)

$$B^{\Theta\mathrm{t}}(\mathbf{l}_1,\mathbf{l}_2,\mathbf{l}_3) = -C_{l_1}^{\phi\mathrm{s}}\left[C_{l_2}^{\Theta}(\mathbf{l}_2\cdot\mathbf{l}_1) + C_{l_3}^{\Theta}(\mathbf{l}_3\cdot\mathbf{l}_1)\right] + \text{ Perm.}\,, \tag{3.16}$$

where permutations involve two additional terms with the replacement of l_1 with l_2 and l_3. Here, $C_{l_1}^{\phi\mathrm{s}}$ is the power spectrum describing correlations between secondary anisotropies and the lensing potential generated by large-scale structure. These correlations were discussed in detail in Goldberg & Spergel (1999) where it was found that the most significant correlation is the one between lensing potentials and the SZ effect. We will use this correlation in illustrating our results.

The presence of secondary effects also modifies the trispectrum and generates an additional contribution beyond the one discussed in Eq. (3.11). Following Cooray &

Kesden (2003), we can write this contribution as

$$\begin{aligned} T^{\Theta\mathrm{s}}(\mathbf{l}_1, \mathbf{l}_2, \mathbf{l}_3, \mathbf{l}_4) \;\; &= \;\; C^{\phi\mathrm{s}}_{l_3} C^{\phi\mathrm{s}}_{l_4} \Big\{ C^{\Theta}_{l_1} (\mathbf{l}_3 \cdot \mathbf{l}_1)(\mathbf{l}_4 \cdot \mathbf{l}_1) + C^{\Theta}_{l_2} (\mathbf{l}_3 \cdot \mathbf{l}_2)(\mathbf{l}_4 \cdot \mathbf{l}_2) \\ &\quad + [\mathbf{l}_3 \cdot (\mathbf{l}_1 + \mathbf{l}_3)]\, [\mathbf{l}_4 \cdot (\mathbf{l}_2 + \mathbf{l}_4)]\, C^{\Theta}_{|\mathbf{l}_1 + \mathbf{l}_3|} \\ &\quad + [\mathbf{l}_4 \cdot (\mathbf{l}_1 + \mathbf{l}_4)]\, [\mathbf{l}_3 \cdot (\mathbf{l}_2 + \mathbf{l}_3)]\, C^{\Theta}_{|\mathbf{l}_1 + \mathbf{l}_4|} \Big\} + \text{Perm.} \end{aligned} \tag{3.17}$$

where permutations involve five additional terms involving the pairings of (l_3, l_4).

Due to an increase in terms as one goes to higher order, we failed to obtain a general expression for the n-point correlator of temperature fluctuations in Fourier space due to lensing-secondary correlations. As we will soon discuss, cumulants beyond the skewness are unlikely to be important as we find kurtosis to be undetectable even for a perfect experiment with no noise and all-sky observations. We expect this to hold true even when considering higher-order moments beyond the kurtosis.

3.4 Skewness and Kurtosis

A simple way to identify the non-Gaussianity induced in the CMB by gravitational lensing is to measure the higher-order cumulants of its one-point probability distribution function $P_{\mathrm{obs}}(\Theta^{\mathrm{t}}; \theta)$ smoothed with beamwidth θ. This observed one-point probability distribution function (PDF) is actually a convolution of the signal PDF $P_{\mathrm{sig}}(\Theta^{\mathrm{sig}}; \theta)$ with the noise PDF $P_{\mathrm{noise}}(\Theta^{\mathrm{n}}; \theta)$ as described below, where $\Theta^{\mathrm{sig}} = \tilde{\Theta} + \Theta^{\mathrm{s}}$ is the total of both lensed primary and secondary contributions to the signal. The signal PDF can be expressed in terms of its cumulants, which we now proceed to calculate. The third and fourth cumulants are proportional to the dimensionless quantities known as the skewness, S, and the kurtosis, K, respectively:

$$\begin{aligned} S(\theta) \;\; &\equiv \;\; [\sigma(\theta)]^{-3} \int (\Theta^{\mathrm{sig}})^3 P_{\mathrm{sig}}(\Theta^{\mathrm{sig}}; \theta)\, d\Theta^{\mathrm{sig}} \,, \\ K(\theta) \;\; &\equiv \;\; [\sigma(\theta)]^{-4} \int (\Theta^{\mathrm{sig}})^4 P_{\mathrm{sig}}(\Theta^{\mathrm{sig}}; \theta)\, d\Theta^{\mathrm{sig}} - 3 \,. \end{aligned} \tag{3.18}$$

They can be expressed as integrals over the bispectrum and trispectrum derived in the preceding Section according to Eq. (3.6):

$$\begin{aligned} S(\theta) &= [\sigma(\theta)]^{-3} \int \frac{d^2\mathbf{l}_1}{(2\pi)^2}\frac{d^2\mathbf{l}_2}{(2\pi)^2}\frac{d^2\mathbf{l}_3}{(2\pi)^2}(2\pi)^2\delta_{\rm D}(\mathbf{l}_{123})B^{\rm t}(\mathbf{l}_1,\mathbf{l}_2,\mathbf{l}_3) \\ &\quad \times W(l_1\theta)W(l_2\theta)W(l_3\theta)\,, \\ K(\theta) &= [\sigma(\theta)]^{-4} \int \frac{d^2\mathbf{l}_1}{(2\pi)^2}\frac{d^2\mathbf{l}_2}{(2\pi)^2}\frac{d^2\mathbf{l}_3}{(2\pi)^2}\frac{d^2\mathbf{l}_4}{(2\pi)^2}(2\pi)^2\delta_{\rm D}(\mathbf{l}_{1234})T^{\rm t}(\mathbf{l}_1,\mathbf{l}_2,\mathbf{l}_3,\mathbf{l}_4) \\ &\quad \times W(l_1\theta)W(l_2\theta)W(l_3\theta)W(l_4\theta)\,. \end{aligned} \tag{3.19}$$

Inserting Eqs. (3.16), (3.11), and (3.17) into the above expressions, and adopting a Gaussian window function $W(l\theta) = e^{-(l\sigma_b)^2/2}$ with $\sigma_b = \theta/\sqrt{8\ln 2}$, we obtain:

$$\begin{aligned} S(\theta) &= \frac{6}{(2\pi)^2[\sigma(\theta)]^3}\int l_1^2\,dl_1\,l_2^2\,dl_2\,C^{\Theta}_{l_1}C^{\phi{\rm s}}_{l_2}I_1(\sigma_b^2 l_1 l_2)\,e^{-\sigma_b^2(l_1^2+l_2^2)}\,, \\ K^{\phi\phi}(\theta) &= \frac{12}{(2\pi)^3[\sigma(\theta)]^4}\int dl_1\,l_1^3\,C^{\phi\phi}_{l_1}e^{-\sigma_b^2 l_1^2}\left[\int dl_2\,l_2^2\,C^{\Theta}_{l_2}I_1(\sigma_b^2 l_1 l_2)\,e^{-\sigma_b^2 l_2^2}\right]^2\,, \\ K^{\phi{\rm s}}(\theta) &= \frac{12}{(2\pi)^3[\sigma(\theta)]^4}\int dl_1\,l_1^3\,C^{\Theta}_{l_1}e^{-\sigma_b^2 l_1^2}\Bigg\{\left[\int dl_2\,l_2^2\,C^{\phi{\rm s}}_{l_2}I_1(\sigma_b^2 l_1 l_2)\,e^{-\sigma_b^2 l_2^2}\right]^2 \\ &\quad -\frac{1}{2\pi}\int l_2^2\,dl_2\,l_3^2\,dl_3\,d\varphi\,C^{\phi{\rm s}}_{l_2}C^{\phi{\rm s}}_{l_3}I_1\left(\sigma_b^2 l_2\sqrt{l_1^2+l_3^2+2l_1l_3\cos\varphi}\right) \\ &\quad \times e^{-\sigma_b^2(l_2^2+l_3^2+l_1l_3\cos\varphi)}\frac{l_1\cos\varphi+l_3\cos^2\varphi}{\sqrt{l_1^2+l_3^2+2l_1l_3\cos\varphi}}\Bigg\}\,. \end{aligned} \tag{3.20}$$

Here $I_1(x)$ is a modified Bessel function of the first kind.

In the presence of instrumental noise, the observed one-point probability distribution function (PDF) will be a convolution of the signal PDF characterized by the skewness and kurtosis given above, and a Gaussian noise PDF: $P_{\rm obs}(\Theta^{\rm t}) = \int d\tau\,P_{\rm sig}(\tau)P_{\rm noise}(\Theta^{\rm t}-\tau)$. In order to perform this convolution, we must first determine the explicit form of the signal PDF that will have nonzero skewness or kurtosis, but vanishing higher cumulants. To do this, we follow the formalism discussed in Juszkiewicz et al. (1995) and references therein. The PDF of a random variable δ with zero mean and variance σ^2 can be expressed as a Gram-Charlier series in the normalized variable $\nu \equiv \delta/\sigma$:

$$p(\nu) = c_0\phi(\nu) + \frac{c_1}{1!}\phi^{(1)}(\nu) + \frac{c_2}{2!}\phi^{(2)}(\nu) + \ldots, \tag{3.21}$$

where $\phi(\nu) \equiv (2\pi)^{-1/2}e^{-\nu^2/2}$ is a Gaussian distribution. The $\phi^{(l)}(\nu)$ are derivatives of the Gaussian distribution with respect to ν:

$$\phi^{(l)}(\nu) \equiv \frac{d^l\phi}{d\nu^l} = (-1)^l H_l(\nu)\phi(\nu), \tag{3.22}$$

and the $H_l(\nu)$ are Hermite polynomials with the unconventional normalization,

$$\int_{-\infty}^{\infty} H_l(\nu)H_m(\nu)\phi(\nu)d\nu = l!\,\delta_{lm}\,. \tag{3.23}$$

The central moments of the PDF are defined as

$$\mu_l \equiv \sigma^l \int_{-\infty}^{\infty} p(\nu)\nu^l\,d\nu\,, \tag{3.24}$$

while the cumulants or "connected" portions of these moments can be derived from the relation

$$M_l \equiv \frac{d^l \ln\langle e^{t\delta}\rangle}{dt^l}\,. \tag{3.25}$$

Using the expansion (3.21) and the orthogonality relation (3.23), the coefficients of the Gram-Charlier series can be expressed in terms of the central moments. By inverting Eq. (3.25), the central moments can then be re-expressed in terms of cumulants. As discussed in the previous Section, the assumption that the lensing potential is Gaussian implies that all cumulants of higher order than the kurtosis must vanish. Using this result, we can rewrite the Gram-Charlier expansion as a power series in the skewness S or kurtosis K, which in the case of lensing will be small quantities.

$$
\begin{aligned}
p(\nu) &= \left[1 + \frac{1}{3!} S H_3(\nu) + \frac{10}{6!} S^2 H_6(\nu) + \ldots\right] \phi(\nu)\,, \\
p(\nu) &= \left[1 + \frac{1}{4!} K H_4(\nu) + \frac{35}{8!} K^2 H_8(\nu) + \ldots\right] \phi(\nu)\,.
\end{aligned}
\tag{3.26}
$$

These power series can be convolved with a Gaussian noise PDF of variance $\sigma^2_{\rm noise}(\theta)$ to obtain the observed PDF $P_{\rm obs}(\Theta^{\rm t})$. To linear order in the true skewness and kurtosis, we find:

$$
\begin{aligned}
S_{\rm obs}(\theta) &= S(\theta) \left\{ \frac{\sigma^2(\theta)}{\sigma^2(\theta) + \sigma^2_{\rm noise}(\theta)} \right\}^{3/2}\,, \\
K_{\rm obs}(\theta) &= \frac{K(\theta)}{8} \left\{ \frac{5\sigma^4(\theta) - 3\sigma^4_{\rm noise}(\theta) - 6\sigma^2(\theta)\sigma^2_{\rm noise}(\theta)}{\left[\sigma^2(\theta) + \sigma^2_{\rm noise}(\theta)\right]^2} + 3 \right\}\,.
\end{aligned}
\tag{3.27}
$$

As expected, the observed skewness and kurtosis converge to the signal values in the absence of noise and to zero in the case when the Gaussian noise is dominant. To actually observe skewness or kurtosis in an experimental sky map, we must construct estimators for these quantities using our data points, the $N = 4\pi f_{\rm sky}/\pi(\theta/2)^2$ pixels in the map. We can write estimators for the skewness and kurtosis as

$$
\begin{aligned}
\widehat{S\sigma^3} &\equiv \frac{1}{N} \sum_{i=1}^{N} (x_i - \bar{x})^3\,, \\
\widehat{K\sigma^4} &\equiv \frac{1}{N} \sum_{i=1}^{N} (x_i - \bar{x})^4 - 3 \left[\frac{1}{N} \sum_{i=1}^{N} (x_i - \bar{x})^2 \right]^2\,,
\end{aligned}
\tag{3.28}
$$

where $\bar{x} = \frac{1}{N}\sum_{i=1}^{N} x_i$ is the traditional estimator for the mean of a distribution. For a distribution like that of the CMB anisotropies which is *a priori* defined to have a zero mean, we find:

$$\begin{aligned}\langle\widehat{S\sigma^3}\rangle &= (1-\frac{3}{N}+\frac{2}{N^2})S\sigma^3\,,\\ \langle\widehat{K\sigma^4}\rangle &= (1-\frac{7}{N}+\frac{12}{N^2}-\frac{6}{N^3})K\sigma^4-\frac{6}{N}(1-\frac{1}{N})\sigma^4\,.\end{aligned} \tag{3.29}$$

These are biased estimators, as has been noted elsewhere under a different context (Hui & Gaztanaga, 1999), but in the large-N limit, they converge to the desired quantities. Assuming that the underlying PDF is Gaussian, the variance of these estimators to lowest order in $1/N$ is given by:

$$\begin{aligned}\sigma^2_{\widehat{S\sigma^3}} &\equiv \langle(\widehat{S\sigma^3})^2\rangle-\langle\widehat{S\sigma^3}\rangle^2=\frac{3!}{N}\sigma^3\,,\\ \sigma^2_{\widehat{K\sigma^4}} &\equiv \langle(\widehat{K\sigma^4})^2\rangle-\langle\widehat{K\sigma^4}\rangle^2=\frac{4!}{N}\sigma^4\,.\end{aligned} \tag{3.30}$$

An alternate derivation of these variances can be obtained from the explicit form of the PDFs following Eq. (3.26). If N pixels or data points are collected and binned such that p_i is the probability that a data point will fall within bin i and σ_i is the standard deviation of that probability, then the best variance of a parameter ϵ characterizing the PDF is given by the Cramér-Rao bound (Kendall & Stuart, 1969),

$$\frac{1}{\sigma^2_\epsilon}=\sum_i\left(\frac{\partial p_i}{\partial\epsilon}\right)^2\frac{1}{\sigma_i^2}\,. \tag{3.31}$$

If the error on each bin is assumed to be Poisson, then $\sigma_i^2=p_i/N$. In the limit of a continuous PDF, $p_i\rightarrow p(\nu)\,d\nu$ and the discrete sum (3.31) becomes an integral:

$$\frac{1}{\sigma^2_\epsilon}=N\int\left(\frac{\partial p}{\partial\epsilon}\right)^2p^{-1}\,d\nu\,. \tag{3.32}$$

Inserting Eq. (3.26) into Eq. (3.32) under the Gaussian null hypothesis $S=K=0$, we find lowest attainable errors as $\sigma^2_\epsilon=3!/N,4!/N$ for $\epsilon=S,K$ in agreement with

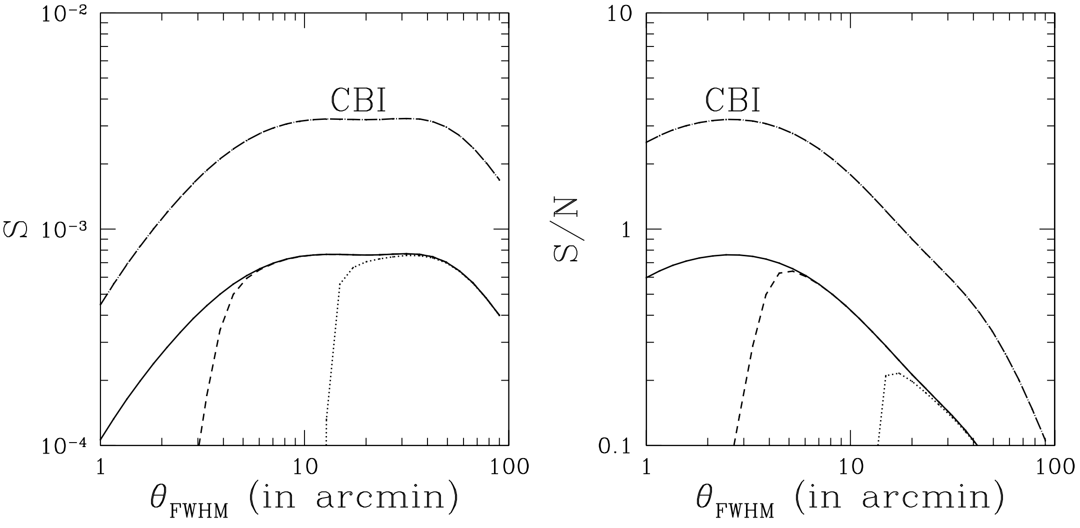

Figure 3.1: *Left:* The skewness due to lensing-SZ correlations for a perfect (no-noise) experiment (solid line), Planck (dashed line), and MAP (dotted line) for $\sigma_8 = 0.9$. The CBI 1σ upper bound of $\sigma_8 \leqslant 1.2$ leads to a higher value for the skewness as indicated by the dot-dashed line. *Right:* The signal-to-noise ratio for the detection of skewness in CMB data with curves labeled as in the left figure. We assume full sky-coverage; for partial sky coverage, the signal-to-noise ratio scales as $\sqrt{f_{\rm sky}}$, where $f_{\rm sky}$ is the fraction of sky covered.

the explicit calculation of the variance of our estimators noted in Eq. (3.30). Further discussion of the variance associated with different estimators for the skewness and kurtosis is included in the Appendix.

3.5 Results & Discussion

3.5.1 Skewness

We illustrate in Fig. (3.1) our results for skewness due to the correlation between lensing and the SZ effect. We calculate this correlation following Cooray (2001a) using the halo approach to large-scale structure (Cooray & Sheth, 2002). The skewness approaches zero at small values of the smoothing scale, consistent with our conclusion that no non-Gaussian signatures exist in the PDF in the limit of infinite resolution. As shown, skewness due to the lensing-SZ correlation peaks at an angular scale of tens of arcminutes, which is in the range of interest to upcoming experiments such as MAP and Planck. When calculating expected signal-to-noise ratios for these experiments, we use detector sensitivities and resolutions tabulated in Cooray et al. (2000). For simplicity, we combine information from individual frequency channels to form one estimate of temperature with an overall noise given by inversely weighting individual noise contributions.

The skewness as shown has signal-to-noise ratios slightly less than unity suggesting that its detection may be hard and potentially affected by noise. However, recent small-scale excess-power detections by experiments such as CBI (Mason et al., 2003) raise the possibility that we may have underestimated the lensing-SZ correlation and thus, the skewness. The lensing-SZ power spectrum $C_l^{\phi \mathrm{s}}$ is roughly proportional to the fifth power of σ_8, the standard deviation of linear mass fluctuations within an $8h^{-1}$ Mpc sphere. If we adopt the CBI 1σ upper bound of $\sigma_8 \leqslant 1.2$ (Mason et al., 2003) as opposed to the value $\sigma_8 = 0.9$ suggested by previous studies, our signal increases by a factor of 4.21. In this case, Planck could conceivably detect skewness with a signal-to-noise of 2.5. The potential for detection of the temperature skewness is consistent with previous expectations that the temperature anisotropy bispectrum due to lensing-SZ correlation can be detected in future data (Goldberg & Spergel, 1999). The cumulative signal-to-noise for skewness, however, is significantly smaller than that for the full bispectrum because the skewness is a single number while the bispectrum contains all information related to non-Gaussianities at the three-point level. As described below, we find a similar reduction in signal-to noise for kurtosis when compared to the full trispectrum.

The frequency dependence of the SZ effect allows us to construct an SZ map of the sky as well as a temperature map with the SZ effect removed. This provides us a unique opportunity to test our understanding of non-Gaussianity at the three-point level. If skewness is purely a consequence of lensing-SZ correlations as posited in this paper, then the skewness obtained by combining one measurement of the SZ map with two measurements of the SZ-cleaned temperature map at the same location using the estimator in Eq. (3.28) should be precisely one third that produced

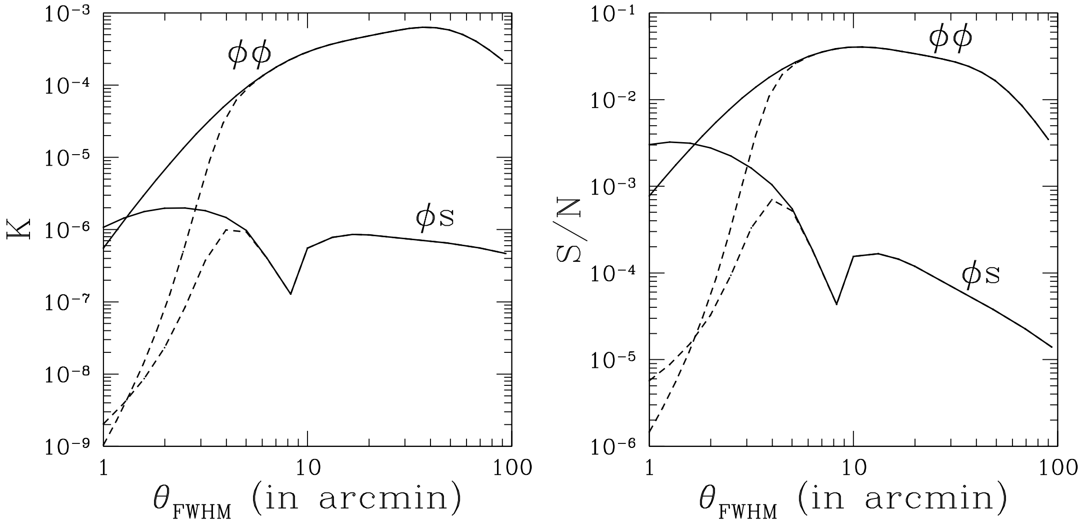

Figure 3.2: *Left:* The kurtosis $K^{\phi\phi}$ due to lensing auto-correlations and $K^{\phi\mathrm{s}}$ due to lensing-SZ cross-correlations for a perfect (no-noise) experiment (solid line) and Planck (dashed line). The kurtosis due to lensing-SZ correlations is negative at smoothing scales below the kink at ~ 8 arcminutes and positive thereafter; its absolute value is shown here. *Right:* The signal-to-noise ratio for the detection of kurtosis in CMB data with curves labeled as in the left figure. We assume full sky-coverage; for partial sky coverage, the signal-to-noise ratio scales as $\sqrt{f_{\rm sky}}$, where $f_{\rm sky}$ is the fraction of sky covered.

by three measurements of the total anisotropy map. This corresponds to the fact that the composite map will sample only one of the three permutations appearing in Eq. (3.16).

3.5.2 Kurtosis

Both lensing kurtosis $K^{\phi\phi}$ and the kurtosis $K^{\phi\mathrm{s}}$ due to lensing-SZ correlations are undetectable even for a perfect no-noise experiment as illustrated in Fig. (3.2). Since the cumulative signal-to-noise ratio for $K^{\phi\mathrm{s}}$ is well below one, we expect it to remain undetectable despite any uncertainty in our calculation of the SZ effect. Note our prediction of the lensing kurtosis $K^{\phi\phi}$ is likely to be more certain since it only depends on the matter power spectrum, with contributions coming mainly from the linear regime. Thus, uncertainties in non-linear aspects of clustering are unlikely to affect our conclusion.

The signal-to-noise value for $K^{\phi\phi}$ can be compared to the cumulative signal-to-noise ratio for the direct detection of the full trispectrum due to lensing which, in the case of Planck, can be as high as ~ 55 (Zaldarriaga, 2000). Consequently, although the lensing kurtosis cannot be detected directly from the data, lensing effects associated with this kurtosis can be used to reconstruct the lensing deflection angle as described in Hu (2001a) and Cooray & Kesden (2003), again, with cumulative signal-to-noise ratios significantly greater than that for the kurtosis itself. The higher signal-to-noise ratio in lensing reconstruction is possible for two reasons. Unlike the kurtosis, which averages indiscriminately over all configurations of the trispectrum as shown in Eq. (3.19), lensing reconstruction is sensitive to certain configurations of the trispectrum – mainly those that contribute to the power spectrum of squared temperature. This avoids severe positive-negative cancellations that significantly reduce the signature of non-Gaussianity. Secondly, the noise contribution associated with lensing reconstruction is also reduced *a priori* through a filter which is designed to optimally extract information on the lensing potentials.

The low signal-to-noise associated with the kurtosis is also consistent with the fact that real-space moments, in general, suffer from excess noise. Though such statistics are easily measurable in data, they do not provide the most optimal methods to search for the existence of non-Gaussian signatures. While we recommend construction of cumulants such as skewness and kurtosis as a first step in understanding non-Gaussianity from effects such as lensing, we suggest that full measures of quantities such as bispectrum and trispectrum will be necessary to fully understand the non-Gaussian behavior of lensing. If measurement of such statistics are still cumbersome, we suggest the use of quadratic statistics in real space, such as the squared-temperature–temperature (Cooray, 2002) and the squared-temperature–squared-temperature (Cooray & Kesden, 2003) power spectra which probe certain configurations of the bispectrum and trispectrum.

This work was supported in part by NSF AST-0096023, NASA NAG5-8506, and DoE DE-FG03-92-ER40701. Kesden acknowledges the support of an NSF Graduate Fellowship, and AC acknowledges support from the Sherman Fairchild Foundation.

3.6 Variance of Skewness and Kurtosis Estimators

A question arose during the composition of this paper as to the appropriate variance for estimates of the skewness and kurtosis of a Gaussian distribution. The true skewness and kurtosis of a Gaussian distribution are necessarily zero, but given N data points x_i drawn from this distribution, even unbiased estimators will yield results distributed about zero with some variance. Some sources (e.g., Press et al. (1992)) indicate variances of $15/N$ and $96/N$, respectively, for the skewness and kurtosis estimators defined in Eq. (3.28) as opposed to our values of $6/N$ and $24/N$. This discrepancy prompted us to investigate further. The estimators of Eq. (3.28) differ from those given in Press et al. (1992) in that they are estimators for the third and fourth cumulants rather than the dimensionless skewness and kurtosis to which they are proportional. Assuming an underlying Gaussian distribution with a variance of unity, standard propagation of errors reveals that the two pairs of estimators have the same variances to lowest order in $1/N$. However, the naïve estimators

$$\begin{aligned} \widehat{S\sigma^3}' &\equiv \frac{1}{N}\sum_{i=1}^{N} x_i^3 \text{ and} \\ \widehat{K\sigma^4}' &\equiv \frac{1}{N}\sum_{i=1}^{N} x_i^4 - 3\left[\frac{1}{N}\sum_{i=1}^{N} x_i^2\right]^2 \end{aligned} \tag{3.33}$$

do indeed have variances of $15/N$ and $96/N$ for skewness and kurtosis respectively. We show this explicitly for the naïve skewness estimator $\widehat{S\sigma^3}'$. The ensemble average of this estimator is simply $S\sigma^3$, so it is truly an unbiased estimator for the skewness. However, taking the ensemble average of $(\widehat{S\sigma^3}')^2$ we find

$$\langle(\widehat{S\sigma^3}')^2\rangle = \frac{1}{N}\left[\mu_6 + (N-1)S^2\sigma^6\right] \,, \tag{3.34}$$

leading to a variance

$$\sigma^2_{\widehat{S\sigma^3}'} \equiv \langle (\widehat{S\sigma^3}')^2 \rangle - \langle \widehat{S\sigma^3}' \rangle^2 = \frac{1}{N}(\mu_6 - S^2\sigma^6)\,. \tag{3.35}$$

For a Gaussian distribution, $\mu_6 = 15\sigma^6$ and $S = 0$, implying that this estimator measures skewness with a variance of $15/N$ and is therefore less sensitive than $\widehat{S\sigma^3}$ defined in Eq. (3.28) which was shown to have a variance of $6/N$. An entirely analogous calculation shows that the naïve kurtosis estimator in Eq. (3.33) has a variance of $96/N$, not $24/N$.

Why do the estimators of Eq. (3.28) outperform those of Eq. (3.33)? Although the true mean of the underlying Gaussian distribution has been chosen to be zero, the estimated mean $\bar{x} = \frac{1}{N}\sum_{i=1}^{N} x_i$ of N data points will not necessarily vanish. The more sophisticated estimators of Eq. (3.28) take this into account by subtracting the estimated mean from each data point, and are therefore able to provide lower-variance estimates of the skewness and kurtosis. These lower values for the variances are adopted for all results concerning signal-to-noise mentioned in this paper.

Chapter 4

Separation of Gravitational-Wave and Cosmic-Shear Contributions to CMB Polarization

Reprinted with permission from M. Kesden, A. Cooray, and M. Kamionkowski Phys. Rev. Lett. **89**, 011304 (2002).

Inflationary gravitational waves generate a distinct signature in the cosmic microwave background (CMB) through a contribution to the curl, or magnetic-like, component in the polarization pattern. Though there is no direct contribution to these curl modes from scalar (density) perturbations, cosmic shear – gravitational lensing of the CMB by large-scale structure along the line of sight – converts a fraction of the polarization in the dominant gradient, or electric-like, component to the curl component. Measurements of higher-order correlations in the CMB temperature and polarization can be used to map the cosmic shear as a function of position on the sky, and with this cosmic-shear map, the lensing contribution to the curl can, in principle, be subtracted. We study how well this cosmic-shear subtraction can be accomplished. We find that CMB information down to arcminute resolution will be required to pursue gravitational-wave amplitudes significantly smaller than those accessible with the Planck satellite. The existence of a finite cutoff in the CMB power spectrum at small scales leads to a minimum detectable gravitational-wave-background amplitude that corresponds to an inflationary energy scale near 10^{15} GeV.

Observation of acoustic oscillations in the temperature anisotropies of the cosmic microwave background (Miller et al., 1999; de Bernardis et al., 2000; Hanany et al., 2000; Halverson et al., 2002) strongly suggests an inflationary origin for primordial perturbations (Guth, 1981; Linde, 1982a; Albrecht & Steinhardt, 1982). It has been argued that a new smoking-gun signature for inflation would be detection of stochastic background of inflationary gravitational waves (IGWs) (Kamionkowski & Kosowsky, 1999). These IGWs produce a distinct signature in the CMB in the form of a contribution to the curl, or magnetic-like, component of the polarization (Kamionkowski et al., 1997a; Seljak & Zaldarriaga, 1997). Since there is no scalar, or density-perturbation, contribution to these curl modes, curl polarization was considered to be a direct probe of IGWs.

There is, however, another source of a curl component. Cosmic shear (CS) – weak gravitational lensing of the CMB due to large-scale structure along the line of sight – results in a fractional conversion of the gradient mode from density perturbations to the curl component (Zaldarriaga & Seljak, 1998). The amplitude of the IGW background varies quadratically with the energy scale $E_{\rm infl}$ of inflation, and so the prospects for detection also depend on this energy scale. In the absence of CS, the smallest detectable IGW background scales simply with the sensitivity of the CMB experiment; as the instrumental sensitivity is improved, smaller values of $E_{\rm infl}$ become accessible (Kamionkowski & Kosowsky, 1999; Jaffe et al., 2000). More realistically, however, the CS-induced curl introduces a noise from which IGWs must be distinguished. If the IGW amplitude (or $E_{\rm infl}$) is sufficiently large, the CS-induced curl will be no problem. However, as $E_{\rm infl}$ is reduced, the IGW signal becomes smaller and will get lost in the CS-induced noise. This confusion leads to a minimum detectable IGW amplitude (Lewis et al., 2002).

In addition to producing a curl component, CS also introduces distinct higher-order correlations in the CMB temperature pattern. Roughly speaking, lensing can stretch the image of the CMB on a small patch of sky and thus, lead to something akin to anisotropic correlations on that patch of sky, even though the CMB pattern at the surface of last scatter had isotropic correlations. By mapping these effects, the CS can be mapped as a function of position on the sky (Seljak & Zaldarriaga, 1999a).

The observed CMB polarization can then be corrected for these lensing deflections to reconstruct the intrinsic CMB polarization at the surface of last scatter (in which the only curl component would be that due to IGWs). In this *Letter*, we evaluate how well this subtraction can be accomplished and study the impact of CS on experimental strategies for detection of IGWs.

To begin, we review the determination of the smallest detectable IGW amplitude in the absence of CS. Following Jaffe et al. (2000), we consider a CMB-polarization experiment of some given instrumental sensitivity quantified by the noise-equivalent temperature (NET) s, angular resolution $\theta_{\rm FWHM}$, duration $t_{\rm yr}$ in years, and fraction of sky covered $f_{\rm sky}$. We then make the null hypothesis of no IGWs and determine the largest IGW amplitude $\mathcal{T}$, defined as $\mathcal{T} = 9.2V/m_{\rm Pl}^4$, where $V = E_{\rm infl}^4$ is the inflaton-potential height, that would be consistent at the 1σ level with the null detection. We then obtain the smallest detectable IGW amplitude $\sigma_{\mathcal{T}}$ from

$$\sigma_{\mathcal{T}}^{-2} = \sum_{l>180/\theta} \left(\partial C_l^{BB,GW}/\partial \mathcal{T}\right)^2 (\sigma_l^{BB})^{-2}, \tag{4.1}$$

where $C_l^{BB,GW}$ is the IGW contribution to the curl power spectrum, and

$$\sigma_l^{BB} = \sqrt{\frac{2}{f_{sky}(2l+1)}} \left(C_l^{BB} + f_{sky} w^{-1} e^{l^2\sigma_b^2}\right) , \tag{4.2}$$

is the standard error with which each multipole moment C_l^{BB} can be determined. Here, $w^{-1} = 4\pi(s/T_{\rm CMB})^2/(t_{\rm pix}N_{\rm pix})$ (Knox, 1995) is the variance per unit area on the sky for polarization observations when $t_{\rm pix}$ is the time spent on each of $N_{\rm pix}$ pixels with detectors of NET s, and $\theta \simeq 203\, f_{\rm sky}^{1/2}$ is roughly the width (in degrees) of the survey. In restricting the sum to $l > 180/\theta$, we have assumed that no information from modes with wavelengths larger than the survey size can be obtained; in fact, some information can be obtained, and our results should thus be viewed as conservative (Lewis et al., 2002).

The second term in Eq. (4.2) is due to instrumental noise, and the first is due to cosmic variance. In the absence of CS, and for the null hypothesis of no IGWs, we set $C_l^{BB}=0$, and the results for the smallest detectable IGW amplitude are shown as the solid curves in Fig. 4.1 for an experiment with detectors of comparable sensitivity to Planck's (left) and a hypothetical experiment (right) with better sensitivity. The smallest detectable IGW amplitude $\mathcal{T}$ scales as $s^2 t_{\rm yr}^{-1}$. For large survey widths, it scales as θ, but at survey widths smaller than $\sim 5^\circ$, it increases because information from the larger-angle modes in the IGW-induced curl power spectrum is lost (cf. the IGW power spectrum in Fig. 4.2).

It is now easy to see how inclusion of CS affects these results. As discussed above, lensing of the gradient polarization at the surface of last scatter due to density perturbations leads to a CMB curl component with a power spectrum

$$\tilde{C}_l^{BB} = \frac{1}{2}\int \frac{d^2\mathbf{l}_1}{(2\pi)^2}\left[\mathbf{l}_2\cdot\mathbf{l}_1\right]^2 (1-\cos 4\phi_{l_1}) C_{l_2}^{\phi\phi} C_{l_1}^{EE}\,, \tag{4.3}$$

where $\mathbf{l}_2 = \mathbf{l}-\mathbf{l}_1$. Here and throughout, C_l^{EE} is the power spectrum of the gradient component of polarization and $C_l^{\phi\phi}$ is the power spectrum of the projected lensing potential (Hu, 2000a). The latter is defined in terms of the potential fluctuations Φ along the line of sight such that

$$\phi(\hat{\mathbf{n}}) \;=\; -2\int_0^{r_0} dr \frac{d_A(r_0-r)}{d_A(r)d_A(r_0)}\Phi(r,\hat{\mathbf{n}}r)\,, \tag{4.4}$$

where r is the comoving radial distance, or conformal look-back time, with r_0 at the last scattering surface, and $d_A(r)$ is the comoving angular diameter distance. The CS-induced curl power spectrum is shown as the dashed curve in Fig. 4.2.

By the time these measurements are made, the cosmological parameters that determine this lensed curl power spectrum should be sufficiently well determined such that this power spectrum can be predicted with some confidence. In that case, the CS-induced curl component can be treated simply as a well-understood noise for the IGW background. The smallest detectable IGW amplitude can then be calculated as above, by inserting the lensed power spectrum, Eq. (4.3), in Eq. (4.2) for C_l^{BB}. The results are shown as the short-dash curves in Fig. 4.1. When lensing is included, the results no longer scale simply with $f_{\rm sky}$, s, or $t_{\rm yr}$, as there is now a trade-off between the instrumental-noise and CS-noise terms in Eq. (4.2). The left panel shows that the IGW sensitivity for an experiment with NET similar to Planck's should not be affected by CS. This is because the IGW amplitudes that could be detectable by such experiments are still relatively large compared with the expected CS signal, especially at the larger angles that will be best accessed by Planck. However, CS will affect the ability of experiments more sensitive than Planck to detect unambiguously IGWs, as shown in the right panel. CS also shifts the preferred survey region to larger areas, as the IGW power spectrum peaks at larger angles than the CS power spectrum (cf. Fig. 4.2). Finally, note that if the CS curl is treated as an unsubtracted noise, it leads, assuming a no-noise polarization map, to a smallest detectable IGW amplitude, corresponding to an inflaton-potential height, $V^{1/4} \sim 4 \times 10^{15}$ GeV.

Now we arrive at the main point of this *Letter*; i.e., how well can the CS-induced curl be subtracted by mapping the CS as a function of position on the sky? One possibility is that the primordial polarization pattern might be reconstructed from that observed by using CS maps obtained with correlations of galaxy ellipticities (Miralda-Escudé, 1991; Blandford et al., 1991; Kaiser, 1992; Bartelmann & Schneider, 1992). However, the source galaxies for these CS surveys are at redshifts $z \sim 1$, while only a small fraction of the CS-induced curl comes from these redshifts, as indicated in Fig. 4.2. An alternative possibility is to use higher-order correlations in the CMB (Seljak & Zaldarriaga, 1999a) to map the CS-induced curl all the way back to the surface of last scatter.

CS modifies the temperature and polarization pattern, giving rise to anisotropic correlations on small scales where the image of the CMB surface of last scatter is sheared by weak lensing. According to Hu (2001b) and Hu (2001a), the quantity, $\nabla \cdot [T(\hat{\mathbf{n}})\nabla T(\hat{\mathbf{n}})]$, provides the best indicator, given a temperature map, of the deflection angle at position ∇ on the sky. In Fourier space, we can write this quadratic estimator for the deflection angle as

$$\hat{\alpha}(\mathbf{l}) = \frac{N_l}{l} \int \frac{d^2\mathbf{l}_1}{(2\pi)^2} \left(\mathbf{l}\cdot\mathbf{l}_1 C_{l_1}^{\rm CMB} + \mathbf{l}\cdot\mathbf{l}_2 C_{l_2}^{\rm CMB}\right) \frac{T(l_1)T(l_2)}{2C_{l_1}^{\rm tot} C_{l_2}^{\rm tot}}\,, \tag{4.5}$$

where $C_l^{\rm CMB}$ is the unlensed CMB power spectrum and

$$C_l^{\rm tot} = C_l^{\rm lensed-CMB} + f_{\rm sky} w^{-1} e^{l^2\sigma_b^2} \tag{4.6}$$

includes all contributions to the CMB temperature power spectrum. The ensemble average over CMB realizations, $\langle\hat{\alpha}(\mathbf{l})_{\rm CMB}\rangle$, is equal to the deflection angle, $l\phi(\mathbf{l})$, when

$$N_l^{-1} = \frac{1}{l^2}\int \frac{d^2\mathbf{l}_1}{(2\pi)^2} \frac{\left(\mathbf{l}\cdot\mathbf{l}_1 C_{l_1}^{\rm CMB} + \mathbf{l}\cdot\mathbf{l}_2 C_{l_2}^{\rm CMB}\right)^2}{2C_{l_1}^{\rm tot}C_{l_2}^{\rm tot}} \,. \tag{4.7}$$

It can also be shown (Hu, 2001b,a) that N_l is the noise power spectrum associated with the reconstructed deflection angle power spectrum

$$\langle\hat{\alpha}(\mathbf{l})\hat{\alpha}(\mathbf{l}')\rangle = (2\pi)^2\delta_D(\mathbf{l}+\mathbf{l}')\left(l^2 C_l^{\phi\phi} + N_l\right) \,. \tag{4.8}$$

Here, the ensemble average is taken independently over realizations of both the CMB and the intervening large-scale structure. In addition to these temperature estimators for the deflection angle, we also use analogous ones constructed from the polarization, as discussed in Hu & Okamoto (2002), although we do not reproduce those formulas here. The total noise in the estimator for the deflection angle can then be constructed by summing the inverses of the individual noise contributions. We thus determine the variance with which each Fourier mode of ϕ can be reconstructed.

With the deflection angle obtained this way as a function of position on the sky, the polarization at the CMB surface of last scatter can be reconstructed (details to be presented elsewhere (Kesden et al., 2003)). In the ideal case, there would be no error in the CS reconstruction leading to no residual lensing-induced curl component. Realistically, however, there will be some error in the CS reconstruction, from measurement error and also from cosmic variance. The exponential reduction of CMB power on small scales from Silk damping means that there will be a finite number of small-scale coherence patches in CMB with which to reconstruct the CS. There will therefore be some residual curl component in the CMB with a noise power spectrum, $C_l^{\phi\phi,{\rm noise}} = N_l/l^2$. The lensing reconstruction from CMB data only allows the extraction of $C_l^{\phi\phi}$ to a multipole of $\lesssim 1000$ (Hu & Okamoto, 2002), but there is substantial contribution to the CS-induced curl component from lensing at smaller angular scales. We thus replace N_l/l^2 by $C_l^{\phi\phi}$ when the former exceeds the latter at large l. This provides an estimate to the noise expected in the reconstructed curl component that follows from implementing a filtering scheme where high-frequency noise in the CS reconstruction is removed to the level of the expected CS signal. The dot-dash curve in Fig. 4.2 shows the residual CS-induced curl component that remains after subtraction.

We can now anticipate the smallest IGW amplitude detectable by a CS-corrected polarization map by simply using this residual noise power spectrum in Eq. (4.3). Results are shown as dotted curves in Fig. 4.1. The left panel shows results as a function of survey size for an experiment with NET similar to Planck, while the right-hand panel shows results for experiments with better sensitivity and resolution. Since the dotted curves are just below the dashed curve in the left-hand panel of Fig. 4.1, we learn that Planck's sensitivity will not be sufficient to warrant an effort to reconstruct the primordial curl and we would do just as well to simply treat the CS-induced curl as a noise component of known amplitude. We can expect to improve the discovery reach for IGWs by increasing the sensitivity and resolution. The right-hand panel of Fig. 4.1 shows results for a hypothetical experiment with $s = 1\,\mu\mathrm{K}\,\sqrt{\mathrm{sec}}$ and angular resolutions of $5'$, $2'$ and $1'$. We now see there is a significant difference between the dashed curve and the dotted lines suggesting the increasing improvement with increasing resolution.

Finally, to indicate the ultimate limits of this class of experiments, the long-dash curve in Fig. 4.1 shows the results assuming perfect detectors (i.e., $s = 0$). If there were no CS-induced curl, then we would have sensitivity to an arbitrarily small IGW amplitude, but the existence of a CS-induced curl provides an ultimate limit of $V^{1/4} \simeq 4 \times 10^{15}$ GeV, as discussed above. Correction for the effects of CS with the CS map inferred from higher-order correlations would allow us to access lower IGW amplitudes, but eventually, such a correction is ultimately limited by the existence of only a finite number of small-scale coherence patches with which to reconstruct the CS. If there is no instrumental-noise limitation, the sensitivity to an IGW signal is maximized by covering as much sky as possible, and the lowest accessible inflaton potential, $\sim 10^{15}$ GeV, is obtained with a nearly all-sky experiment.

To conclude, we have studied the IGW amplitudes accessible by mapping the curl component of the CMB polarization, taking into account the effects of a CS-induced curl that is either modeled as an unsubtracted noise or subtracted with a CS map obtained with higher-order correlations. We find that the CS reconstruction is unlikely to improve the IGW discovery reach of Planck. To go beyond Planck, however, a CS map will need to be constructed with temperature and polarization maps of higher sensitivity and resolution than Planck. An ultimate limit of roughly $V^{1/4} \sim 10^{15}$ GeV to the detectable IGW amplitude using the techniques considered here comes from the existence of finite CMB power on small scales. There are several possible ways this lower limit may be improved upon. First of all, we have used only the lowest-order temperature-polarization correlations to reconstruct the CS. The inclusion of the complete temperature-polarization four-point correlation functions and higher-order correlations may possibly improve the CS reconstruction. We have neglected reionization in our analysis. Reionization will boost the large-angle CMB polarization (Zaldarriaga, 1997), thus improving the detectability of IGWs. Another improvement in this limit may be achieved by relaxing our assumption that power in the CMB drops exponentially at small scales. If the excess small-scale power recently detected by CBI (Sievers et al., 2001) comes from high redshifts, then there will be more small-scale coherence patches with which to reconstruct the CS. In this case, it is imaginable that a far more precise CS map can be reconstructed, but this might require even better angular resolution and sensitivity.

During the preparation of the paper, we learned of other very recently completed work by Knox and Song (Knox & Song, 2002) that performs a very similar calculation and reaches similar conclusions. This work was supported in part by NSF AST-0096023, NASA NAG5-8506, and DoE DE-FG03-92-ER40701. Kesden acknowledges the support of an NSF Graduate Fellowship and AC acknowledges support from the Sherman Fairchild Foundation.

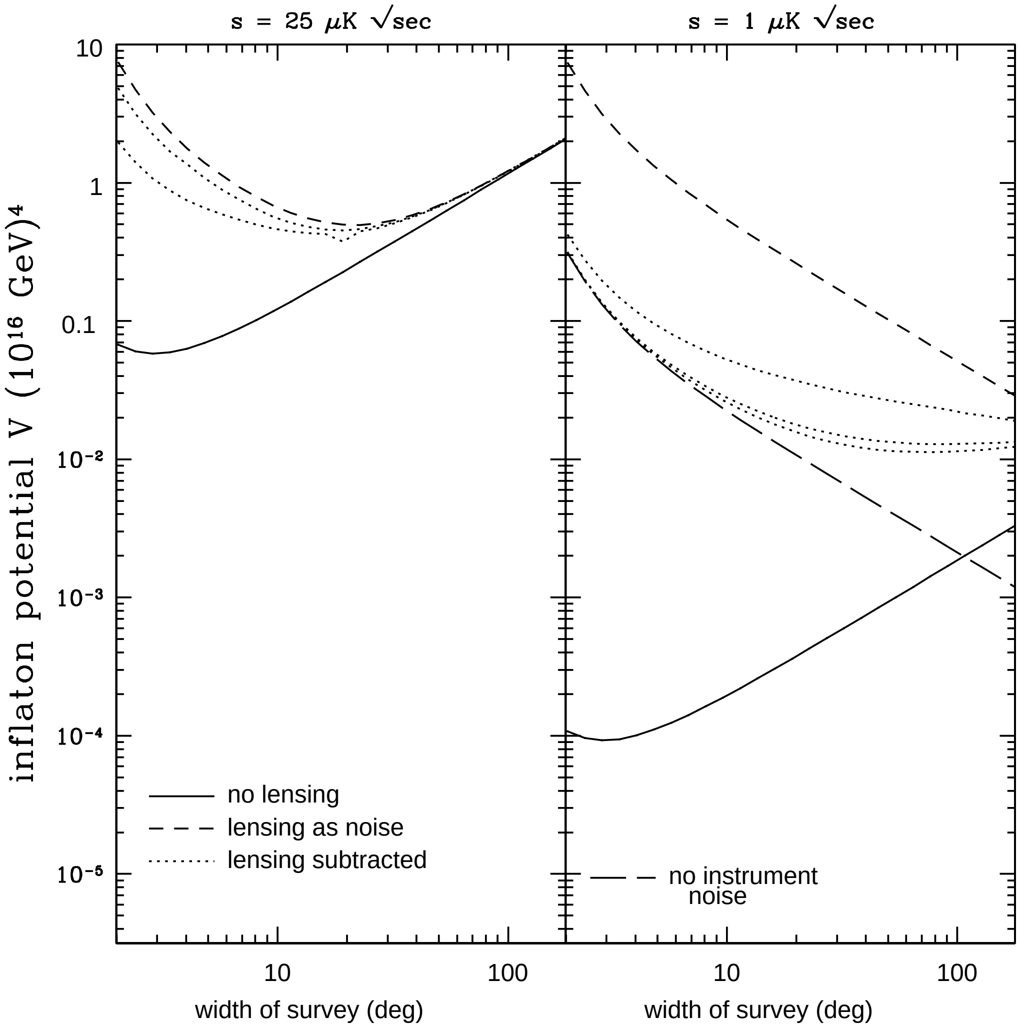

Figure 4.1: Minimum inflation potential observable at 1σ as a function of survey width for a one-year experiment. The left panel shows an experiment with NET $s = 25\,\mu\mathrm{K}\ \sqrt{\mathrm{sec}}$. The solid curve shows results assuming no CS, while the dashed curve shows results including the effects of an unsubtracted CS; we take $\theta_{\mathrm{FWHM}} = 5'$ in these two cases. The dotted curves assume the CS is subtracted with $\theta_{\mathrm{FWHM}} = 10'$ (upper curve) and $5'$ (lower curve). Since the dotted curves are close to the dashed curve, it shows that these higher-order correlations will not be significantly useful in reconstructing the primordial curl for an experiment similar to Planck's sensitivity and resolution. The right panel shows results for hypothetical improved experiments. The dotted curves show results with CS subtracted and assuming $s = 1\,\mu\mathrm{K}\ \sqrt{\mathrm{sec}}$, $\theta_{\mathrm{FWHM}} = 5'$, $2'$, and $1'$ (from top to bottom). The solid curve assumes $\theta_{\mathrm{FWHM}} = 1'$ and $s = 1\,\mu\mathrm{K}\ \sqrt{\mathrm{sec}}$, and no CS, while the dashed curve treats CS as an additional noise. The long-dash curve assumes CS subtraction with no instrumental noise ($s = 0$).

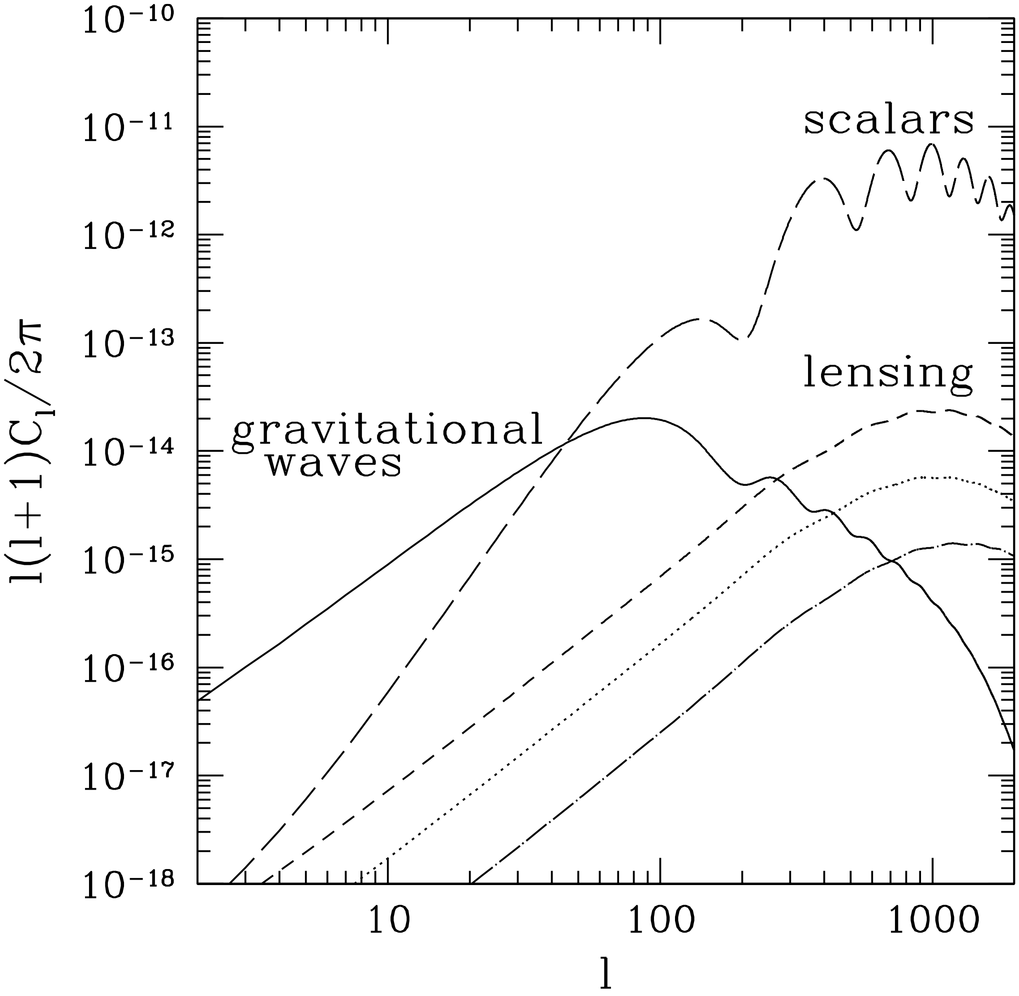

Figure 4.2: Contributions to the CMB polarization power spectra. The long-dashed curve shows the dominant polarization signal in the gradient component due to scalar perturbations. The solid line shows the maximum allowed curl polarization signal from the gravitational-wave background, which will be smaller if the inflationary energy scale is smaller than the maximum value allowed by COBE of 3.47×10^{16} GeV. The dashed curve shows the power spectrum of the curl component of the polarization due to CS. The dotted curve is the CS contribution to the curl component that comes from structures out to a redshift of 1; this is the level at which low-redshift lensing surveys can be used to separate the CS-induced polarization from the IGW signal. The dot-dashed line is the residual when lensing contribution is separated with a no-noise experiment and 80% sky coverage.

Chapter 5

Can Cosmic Shear Shed Light on Low Cosmic Microwave Background Multipoles?

Reprinted with permission from M. Kesden, M. Kamionkowski, and A. Cooray, Phys. Rev. Lett. **91**, 221302 (2003).

The lowest multipole moments of the cosmic microwave background (CMB) are smaller than expected for a scale-invariant power spectrum. One possible explanation is a cutoff in the primordial power spectrum below a comoving scale of $k_c \simeq 5.0 \times 10^{-4}$ Mpc^{-1}. This would affect not only the CMB, but also the cosmic-shear (CS) distortion of the CMB. Such a cutoff increases significantly the cross-correlation between the large-angle CMB and cosmic-shear patterns. The cross-correlation may be detectable at $> 2\sigma$ which, when combined with the low CMB moments, may tilt the balance between a 2σ result and a firm detection of a large-scale power-spectrum cutoff. As an aside, we also note that the cutoff increases the large-angle cross-correlation between the CMB and low-redshift tracers of the mass distribution.

One of the more intriguing results of Wilkinson Microwave Anisotropy Probe (WMAP) (Bennett et al., 2003; Spergel et al., 2003) is confirmation of the absence of large-scale temperature correlations in the cosmic microwave background (CMB), or equivalently, a suppression of power in the quadrupole and octupole moments, found earlier by the Cosmic Background Explorer (Bennett et al., 1996). A variety of measures of the power spectrum – from the $l = 4$ moment of the CMB power spectrum, which probes wavelengths $\sim 10^4$ Mpc, to galaxy surveys and the Lyman-alpha forest, which probe down to 1–10 Mpc – show consistency with a scale-invariant spectrum of primordial perturbations. Thus, the suppression of the $l = 2$ and $l = 3$ moments of the CMB power spectrum come as a bit of a surprise.

Is this simply a statistical fluke? Or is something novel occurring just beyond our observable cosmological horizon? Possibilities include remnants of a pre-inflationary Universe, a curvature scale just larger than the horizon, and/or exotic inflation (Efstathiou, 2003b; Blanchard et al., 2003; Feng & Zhang, 2003; Kawasaki & Takahashi, 2003; DeDeo et al., 2003; Contaldi et al., 2003; Gaztanaga et al., 2003; Efstathiou, 2003a; Cline et al., 2003). If there is indeed a suppression of large-scale power, it occurs at distance scales $\sim 10^4$ Mpc (Contaldi et al., 2003; Gaztanaga et al., 2003; Efstathiou, 2003a; Cline et al., 2003), larger than those typically probed by galaxy surveys. Future experiments to determine the lowest moments of the CMB power spectrum are also of limited value because current measurements are already dominated by cosmic variance rather than instrumental noise. Thus, although the current evidence for new super-horizon physics is tantalizing, the prospects for further testing it are limited.

In this paper, we point out that there exists another probe of the mass distribution on these largest distance scales. Cosmic shear (CS) – weak gravitational lensing by density perturbations along the line of sight – will produce identifiable distortions in the temperature-polarization pattern of the CMB. When observed, these distortions map the gravitational potential projected along a given line of sight. Here we show that a power-spectrum cutoff enhances significantly (roughly a factor of four) the cross-correlation between the CMB and CS distortion of the CMB on the largest scales. This cross-correlation may be detectable at the $> 2\sigma$ level and may thus provide a valuable cross-check to the current $\sim 2\sigma$ evidence for a dearth of large-scale CMB power. As an aside, we also show that the large-angle cross-correlation between the CMB and low-redshift tracers of large-scale structure (Crittenden & Turok, 1996; Kamionkowski, 1996; Boughn et al., 1998; Kinkhabwala & Kamionkowski, 1999; Seljak & Zaldarriaga, 1999b) is roughly doubled if the large-scale cutoff is real. Although recent detections (Boughn & Crittenden, 2003; Nolta et al., 2004; Multamaki et al., 2004; Fosalba & Gaztanaga, 2004; Scranton et al., 2003; Afshordi et al., 2004) of this effect are at smaller scales than would be affected by a large-scale cutoff, correlations on larger scales might be probed by future experiments.

Below, we first discuss the large-scale CS power spectra, as well as the cross-correlation of the CS pattern with the CMB temperature pattern. We then construct an estimator for the cross-correlation, and show that it can distinguish the cross-correlation with and without a cutoff at roughly the 2σ level. When combined with the already suspiciously low $l = 2$ and $l = 3$ moments of the CMB power spectrum, this finding may tilt the balance between a 2σ result and a 3σ discrepancy with scale invariance.

Perturbations in the matter density induce perturbations to the gravitational potential $\Phi(\mathbf{r}, z)$ which then induce temperature perturbations in the CMB through the Sachs-Wolfe effect

$$\Theta(\hat{\mathbf{n}}) = \frac{1}{3}\Phi(\mathbf{r}_0, z_0) - 2\int_0^{r_0} \frac{d\Phi}{dr}(\mathbf{r}, z(r))\, dr\,, \tag{5.1}$$

where $\mathbf{r}$ and z are the physical comoving distance and redshift, respectively, and the subscript 0 denotes these quantities at the last-scattering surface. The position vector $\mathbf{r}$ points in the direction $\hat{\mathbf{n}}$ on the sky. The potential at redshift z can be related to its present-day value with the linear-theory growth factor $G(z)$ (normalized to unity today) through $\Phi(\mathbf{r}, z) = (1 + z)G(z)\Phi(\mathbf{r}, 0)$. The first term in Eq. (5.1) comes from density perturbations at the surface of last scatter, while the second term (the integrated Sachs-Wolfe effect; ISW) comes from density perturbations along the line of sight.

Relating the potential to the matter perturbation through the Poisson equation, if the three-dimensional matter power spectrum is $P(k)$ as a function of wavenumber k, then the angular power spectrum for temperature fluctuations is

$$C_l^{\Theta\Theta} \propto \int dk\, k^{-2} P(k) [\widetilde{\Theta_l}(k)]^2, \tag{5.2}$$

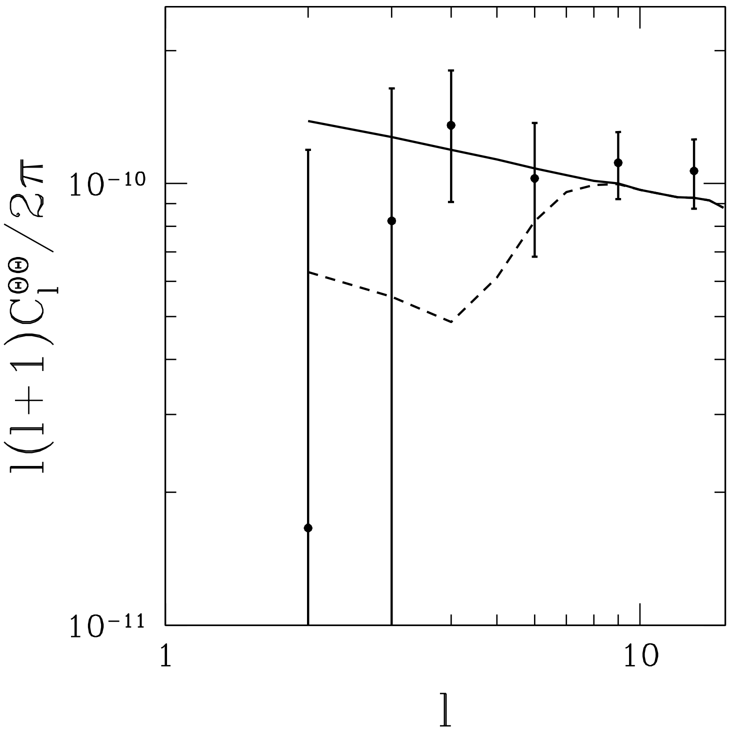

Figure 5.1: The CMB temperature power spectrum. The solid curve is the power spectrum of Eq. (5.2) for the ISW effect without a cutoff. The dashed curve has a cutoff in $P(k)$ below $k_c = 5.0 \times 10^{-4}$ Mpc^{-1}. The binned error bars represent actual WMAP data.

as a function of multipole moment l, where

$$\begin{aligned} \widetilde{\Theta}_l(k) &= \frac{1}{3}(1+z_0)G(z_0)j_l(kr_0) \\ &- 2\int_0^{z_0} dz\,[(1+z)G(z)]'\,j_l\,(kr(z))\ , \end{aligned} \tag{5.3}$$

and the prime denotes derivative with respect to redshift z. The main contributions to the integral in Eq. (5.2) come from wavenumbers k near 10^{-4} Mpc^{-1} (see, e.g., Fig. 7 in Kamionkowski & Spergel (1994)).

If scale invariance holds out to super-horizon scales (as predicted by the generic inflationary model), then the power spectrum $P(k)$ at distance scales relevant for $l \lesssim 10$ is simply $P(k) \propto k^n$, with n near unity. This is certainly what the CMB data show at multipole moments $l \geqslant 4$ and it is consistent with determinations of the power spectrum from the CMB and large-scale structure out to scales as small as ∼few Mpc. Thus, the observed suppression of $C_2^{\Theta\Theta}$ and $C_3^{\Theta\Theta}$ shown in Fig. 5.1 is a bit of a surprise.

The same potential perturbations $\Phi(\mathbf{r}, z)$ that contribute to the Sachs-Wolfe effect also give rise to weak gravitational lensing described by the projected potential,

$$\phi(\hat{\mathbf{n}}) = -2 \int_0^{r_0} dr \frac{r_0 - r}{r_0 r} \Phi(\mathbf{r}, z(r)) \,. \tag{5.4}$$

The angular power spectrum of the lensing potential is then

$$C_l^{\phi\phi} \propto \int dk \, k^{-2} P(k) [\widetilde{\phi_l}(k)]^2 \,. \tag{5.5}$$

In fact, the only difference between this expression and its SW counterpart is the replacement of $\widetilde{\Theta_l}(k)$ by

$$\widetilde{\phi_l}(k) = -2 \int_0^{z_0} \frac{c\,dz}{H(z)} \frac{r_0 - r(z)}{r_0 r(z)} (1+z) G(z) j_l(kr(z)). \tag{5.6}$$

The projected potential receives contributions from a wide variety of distances, peaked at roughly half the comoving distance to the surface of last scatter. The lowest multipole moments of the CS power spectrum come from wavenumbers k near 10^{-4} Mpc^{-1}. Since the small-k Fourier modes of the potential that give rise to low-l CS moments are the same as those that give rise to the low-l CMB moments, we anticipate that the CS power spectrum should also reflect the suppression of large-scale power. Evaluating these expressions numerically, however, we find that the cutoff suppresses $C_2^{\phi\phi}$ by no more than $\sim 10\%$, too small to be detected.

However, the CMB and CS multipole moments are generated by the same underlying potential fluctuations, and so there should be some cross-correlation between the two. And as we show, this cross-correlation turns out to be increased significantly if there is a cutoff. The cross-correlation power spectrum $C_l^{\Theta\phi}$ is

$$C_l^{\Theta\phi} \propto \int dk \, k^{-2} P(k) \widetilde{\Theta_l}(k) \widetilde{\phi_l}(k) \,. \tag{5.7}$$

We can define a dimensionless cross-correlation coefficient, $r_l \equiv (C_l^{\Theta\phi})^2/C_l^{\Theta\Theta}C_l^{\phi\phi}$. If $\widetilde{\Theta}_l(k)$ and $\widetilde{\phi}_l(k)$ had precisely the same k dependence, then the CMB maps would be maximally correlated ($r_l = 1$). In this case, we would be able to predict precisely that the CS spherical-harmonic coefficients should be $\phi_{lm} = (C_l^{\Theta\phi}/C_l^{\Theta\Theta})\Theta_{lm}$ in terms of the measured temperature coefficients Θ_{lm}. Moreover, if r_l were equal to unity, then a CS map might be used to confirm the CMB measurements, but it would add no additional statistically-independent information on the large-scale power spectrum.

If, on the other hand, there was no overlap between $\widetilde{\Theta}_l(k)$ and $\widetilde{\phi}_l(k)$ whatsoever, then there would be no cross-correlation, $r_l = 0$. In this case, the CS pattern could not confirm the CMB measurement, but it would provide a statistically independent probe of the large-scale power spectrum.

Most generally, $0 < r_l < 1$, and the lensing spherical-harmonic coefficients will be

$$\phi_{lm} = (C_l^{\Theta\phi}/C_l^{\Theta\Theta})\Theta_{lm} + \left[C_l^{\phi\phi} - (C_l^{\Theta\phi})^2/C_l^{\Theta\Theta}\right]^{1/2}\zeta\,, \tag{5.8}$$

where ζ is a Gaussian random variable with zero mean and unit variance; i.e., there is a correlated part determined by the CMB pattern and an uncorrelated part.

Fig. 5.2 shows our central result: the cross-correlation coefficient for a scale-invariant spectrum and one in which $P(k) = 0$ for $k < k_c = 5 \times 10^{-4}$ Mpc^{-1}. The dramatic increase in the cross-correlation for the lowest l in the presence of a cutoff can be understood by examining the two terms of Eq. (5.1). The first of these terms, generated at the last-scattering surface, is uncorrelated with the lensing potential; the second term is a line-of-sight integral like the projected potential of Eq. (5.4). Since the contribution of the first term comes from a larger distance from the observer than that of the second term, correspondingly larger structures with lower wavenumber k will be projected onto the angular scale set by the multipole moment l. The lowest multipole moments will correspond to structures at the last-scattering surface with $k < k_c$, implying that in the presence of a cutoff, only the second term of Eq. (5.1) will be nonvanishing for the lowest multipole moments. Since it is only this term that is correlated to the lensing potential, the dimensionless cross-correlation

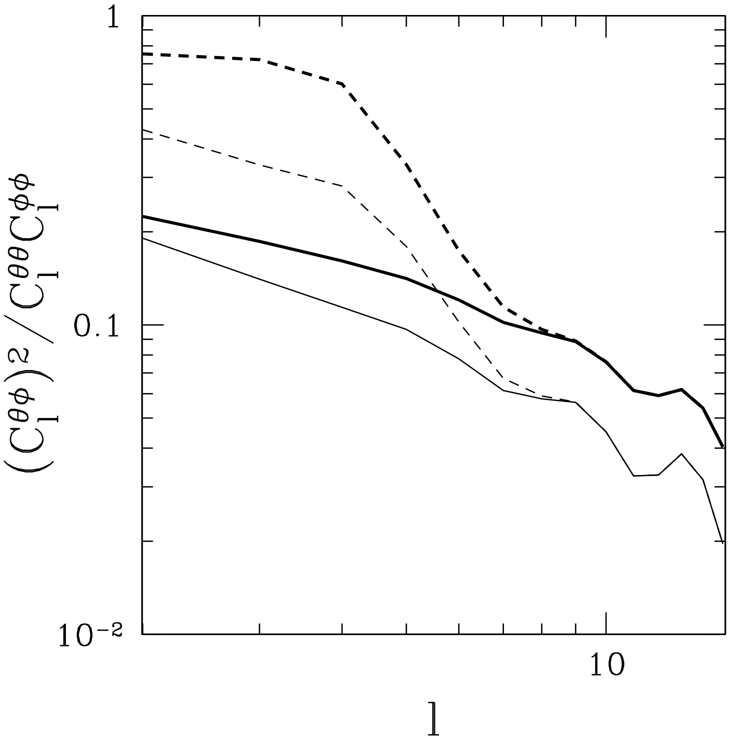

Figure 5.2: The dimensionless correlation $(C_l^{\Theta\phi})^2/C_l^{\Theta\Theta}C_l^{\phi\phi}$ between maps of the CMB temperature and lensing potential. The solid curve shows this correlation in the absence of a cutoff, while the dashed curve is for a cutoff $k_c = 5.0 \times 10^{-4}$ Mpc^{-1}. The upper, darker curves correspond to the lensing potential seen by sources at the CMB last-scattering surface, while the lower, lighter curves correspond to lensing sources at redshift $z = 1.0$.

will be significantly higher in the presence of a cutoff. It is important to note that this increase in the dimensionless cross-correlation in the presence of a cutoff is an independent prediction and not merely a consequence of the observed suppression of $C_l^{\Theta\Theta}$ for low l. If the CMB and CS multipole moments Θ_{lm} and ϕ_{lm} were multiplied by an l-dependent normalization to suppress power on large scales, r_l itself would remain unaffected because it is dimensionless. The independence of this prediction allows estimates of r_l from CMB and CS maps to constrain k_c with greater statistical significance than measurements of $C_l^{\Theta\Theta}$ alone.

We now determine how well measurements of r_l can discriminate between a model with a scale-invariant power spectrum and one with a cutoff (i.e., $P(k) = 0$ for

$k < k_c = 5.0 \times 10^{-4}$ Mpc^{-1}). Higher-order correlations in a high-resolution low-noise CMB temperature-polarization map, can be used to construct estimators for the projected potential (Seljak & Zaldarriaga, 1999a; Hu, 2001b,a; Hu & Okamoto, 2002; Kesden et al., 2002b; Knox & Song, 2002; Kesden et al., 2003; Hirata & Seljak, 2003b) and thus, estimators $\widehat{C}_l^{\Theta\phi}$ and $\widehat{C}_l^{\phi\phi}$, in addition to those $\widehat{C}_l^{\Theta\Theta}$ for the temperature already obtained by WMAP. An estimator $\widehat{r}_l \equiv (\widehat{C}_l^{\Theta\phi})^2/\widehat{C}_l^{\Theta\Theta}\widehat{C}_l^{\phi\phi}$ for r_l can then be formed. Although it is not an unbiased estimator, it is sensitive to r_l and converges to r_l for $l \gg 1$. For a given realization of the CMB and CS patterns, r_l can be estimated independently for each value of l. We have calculated the probability distributions for $\widehat{r}_l$ for each l from many different Monte Carlo realizations of the two models for r_l described above. The CMB coefficients Θ_{lm} are set to the values consistent with the WMAP power spectrum shown in Fig. 5.1, while the uncorrelated part of ϕ_{lm} is determined for each realization of the two models in accordance with Eq. (5.8). We assume the CS projected potential is reconstructed from a full-sky CMB temperature-polarization map with $7'$ angular resolution and noise-equivalent temperature of 0.46 $\mu\mathrm{K}\sqrt{\mathrm{sec}}$. The different predictions for $\widehat{r}_3$ for the two models are shown in Fig. 5.3; the predictions for $\widehat{r}_2$ and $\widehat{r}_4$ are qualitatively similar, while for $\widehat{r}_5$, the two probability distributions begin to merge, and for $\widehat{r}_6$, they are almost indistinguishable.

Assuming that the first model (no cutoff) is correct, we calculated the fraction of realizations in which the measured values of $\widehat{r}_l$ would lead us to conclude that they were more likely drawn from the probability distributions of the second model. This occurs only 0.7% of the time, implying that only 0.7% of the time would cosmic variance mislead us into thinking that a cutoff as large as $k_c = 5.0 \times 10^{-4}$ Mpc^{-1} was favored over the no-cutoff model. Since the CMB coefficients Θ_{lm} are constrained to a single realization consistent with WMAP in both models, this measurement would be statistically independent of the low observed CMB multipole moments themselves for the purpose of distinguishing the two models. It could thus increase the $\sim 2\sigma$ discrepancy of that measurement into a $> 3\sigma$ detection of a large-scale cutoff in the power spectrum.

Finally, we consider the cross-correlation between the CMB and low-redshift

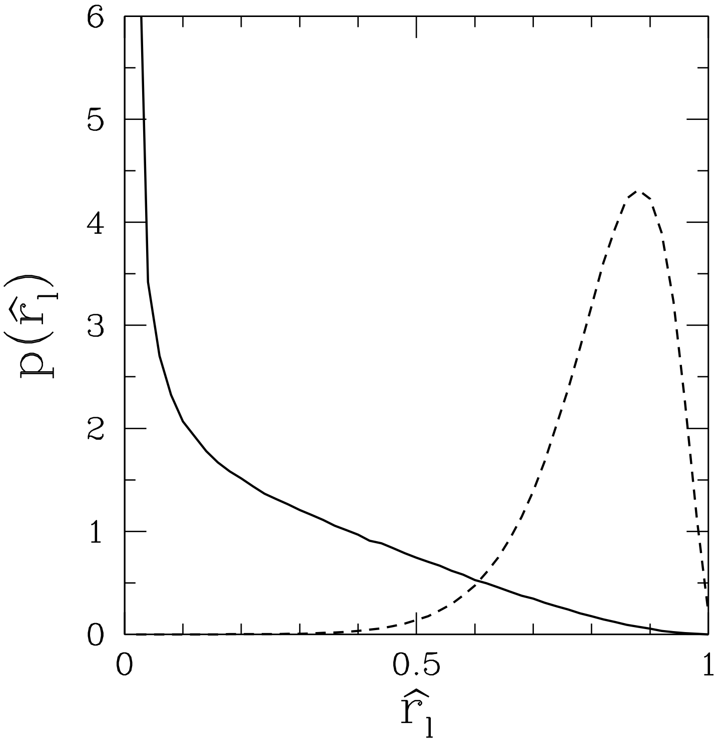

Figure 5.3: Probability distributions for $\widehat{r}_l$ for $l = 3$ for the two models described in the text; the solid curve corresponds to the model without a cutoff in $P(k)$ while the dashed curve has a cutoff $k_c = 5.0 \times 10^{-4}$ Mpc^{-1}.

($z \lesssim 1$) tracers of large-scale structure (Crittenden & Turok, 1996; Kamionkowski, 1996; Boughn et al., 1998; Kinkhabwala & Kamionkowski, 1999; Seljak & Zaldarriaga, 1999b) that several groups have already claimed to detect (Boughn & Crittenden, 2003; Nolta et al., 2004; Multamaki et al., 2004; Fosalba & Gaztanaga, 2004; Scranton et al., 2003; Afshordi et al., 2004). To estimate the effect of a cutoff, we have calculated the cross-correlation between the CMB and CS of a hypothetical population of sources at redshift $z = 1$ for models with and without a cutoff. As indicated by the lighter curves in Fig. 5.2, a cutoff boosts the cross-correlation coefficient for these low-redshift sources by roughly a factor of two. Similar results are obtained for cross-correlation with the galaxy distribution. Although recent detections of the cross-correlation occur at smaller scales ($l \simeq 30$) than those affected by the cutoff, future large-scale surveys could be sensitive to a cutoff-induced enhancement.

The WMAP observations clearly support the ΛCDM concordance model, but they do present a few tantalizing discrepancies. Perhaps the most intriguing is the sharp decrease in observed power at the lowest multipole moments shown in Fig. 5.1. A variety of fundamental causes for this large-scale suppression can be modeled empirically by an effective cutoff k_c in the primordial power spectrum $P(k)$. Though the WMAP team found that only 0.15% of simulated CMB maps had less power on large scales (Bennett et al., 2003; Spergel et al., 2003), it would be highly desirable to find corroborating evidence to confirm that this observation is not merely a statistical anomaly. One possibility is measurements of the polarization signal from nearby galaxy clusters, which are proportional to their local CMB quadrupole moment (Cooray & Baumann, 2003; Kamionkowski & Loeb, 1997). We have proposed here that the cross-correlation $C_l^{\Theta\phi}$ could provide additional evidence. Although the experimental requirements for measurements of the CS distortions to the CMB are ambitious, they are closely aligned with those for the CMBPOL experiment that appears in NASA's roadmap. Thus, this measurement, like the effect it seeks to study, is on the horizon.

This work was supported in part by NASA NAG5-11985 and DoE DE-FG03-92-ER40701. Kesden acknowledges the support of an NSF Graduate Fellowship.

Chapter 6

Gravitational-Wave Signature of an Inspiral into a Supermassive Horizonless Object

Reprinted with permission from M. Kesden, J. Gair, and M. Kamionkowski, Phys. Rev. D **71**, 044015 (2005).

Event horizons are among the most intriguing of general relativity's predictions. Although on firm theoretical footing, direct indications of their existence have yet to be observed. With this motivation in mind, we explore here the possibility of finding a signature for event horizons in the gravitational waves (GWs) produced during the inspiral of stellar-mass compact objects (COs) into the supermassive ($\sim 10^6 M_\odot$) objects that lie at the center of most galaxies. Such inspirals will be a major source for LISA, the future space-based GW observatory. We contrast supermassive black holes with models in which the central object is a supermassive boson star (SMBS). Provided the COs interact only gravitationally with the SMBS, stable orbits exist not just outside the Schwarzschild radius, but also inside the surface of the SMBS as well. The absence of an event horizon allows GWs from these orbits to be observed. Here, we solve for the metric in the interior of a fairly generic class of SMBS and evolve the trajectory of an inspiraling CO from the Schwarzschild exterior through the plunge into the exotic SMBS interior. We calculate the approximate waveforms for GWs emitted during this inspiral. Geodesics within the SMBS surface will exhibit extreme pericenter precession and other features making the emitted GWs readily distinguishable from those emitted during an inspiral into a black hole.

6.1 Introduction

The black-hole event horizon – a surface from beyond which no information can be received – is one of the most intriguing predictions of general relativity. From the theoretical point of view, the prediction is on a fairly firm footing. However, event horizons have become so central to physics and astrophysics that efforts to obtain direct empirical evidence for their existence are certainly warranted. In other words, is there any observation or measurement that we can make that would allow us to "see" the event horizon?

This question has motivated a small body of work in which the phenomenology of accretion onto alternatives to stellar-mass or supermassive black holes (SMBHs) has been worked out (Torres et al., 2000; Yuan et al., 2004). X-ray spectroscopy or μ-arcsec imaging might distinguish the accretion disks of black holes from those surrounding more exotic alternatives (Schunck & Liddle, 1997). A second approach involves the "shadow" cast on background sources by the black hole when acting as a strong gravitational lens (Falcke et al., 2000). Current observations are within a factor of two of being able to resolve the shadow of Sgr A^* using very long baseline interferometry (VLBI) in the submillimeter. Although these calculations may not be applicable if the predictions of general relativity hold true on astrophysical scales, such work can be interesting in its own right and sometimes illuminates certain aspects of the phenomenology of the standard black-hole spacetime.

With these motivations in mind, we study here a third possibility, namely whether gravitational-wave (GW) signals from the inspiral of stellar-mass compact objects (COs) into the supermassive objects at galactic centers may ultimately be used to ascertain the existence of event horizons. There are in fact good prospects of detecting GWs from such extreme-mass-ratio inspirals (EMRIs). Measurements of stellar velocities within galactic cusps imply that the dynamics within a few parsecs of the center are dominated by a supermassive object (Richstone et al., 1998; Kormendy &

Gebhardt, 2002). Stellar clusters sufficiently dense to enclose enough mass within the observationally determined volume would have lifetimes much less than the age of the Galaxy (Maoz, 1995), leading most to believe that the object must be compact (i.e., a SMBH). Still, stellar-dynamical measurements probe distances much larger than the Schwarzschild radius, and thus, cannot distinguish SMBHs from sufficiently compact alternatives. Although future stellar-dynamical observations will probe smaller radii, they will never probe the spacetime structure anywhere even close to the horizon (Weinberg et al., 2004).

These supermassive compact objects are expected to capture stars from the surrounding galactic cusps (Sigurdsson & Rees, 1997; Sigurdsson, 1997). Main-sequence stars will generally be disrupted at a tidal radius $r_t \approx 2(M/m_*)^{1/3} r_*$, where m_* and r_* are the stellar mass and radius and M the mass of the supermassive central object. For stars with a solar mass and radius, this implies $r_t \approx 50 R_s M_6^{-2/3}$ with M_6 the mass of the central object in $10^6 M_\odot$. Thus, main-sequence stars are tidally disrupted long before they reach the Schwarzschild radius, $R_s = 2GM/c^2$. However an evolved stellar population will also contain a fraction of white dwarfs, neutron stars, and stellar-mass black holes that, owing to their smaller radii, can maintain their integrity down to the innermost stable orbit (ISO) and beyond. Under the extreme–mass-ratio approximation – well justified for the case $\eta \equiv m_*/M \lesssim 10^{-6}$ – these COs will travel along geodesics of the central object. However, the galactic cusp is a very crowded environment, and two-body scattering will change the orbital parameters of the COs over a relaxation time t_r. Dynamical friction will cause heavy objects such as neutron stars and stellar-mass black holes to sink to the center of the cusp. Eventually, the COs will enter the critical region of phase space known as the loss cone in which the time-scale for the loss of energy due to gravitational radiation is less than the time-scale for two-body scattering out of the loss cone. After this point, the evolution of the orbit should be entirely determined by the loss of energy and angular momentum in GWs, ultimately leading to capture by the central object. The expected rate of such captures is highly model-dependent, subject to uncertainties in the galaxy luminosity function, the mass function of central objects, and the initial stellar-mass function

(IMF) among other variables. Conservatively, event rates of order 10^{-8} per galaxy per year are anticipated, implying 0.1 captures per year out to 1 Gpc (Sigurdsson & Rees, 1997). This result could be enhanced by an order of magnitude by a top-heavy IMF, either due to low metallicity in the early universe or starbursts in the high-density environment of the galactic cusp itself (Sigurdsson, 1997). EMRIs about a $10^6 M_\odot$ central object would produce GWs in the frequency band 10^{-4} to 10^{-2} Hz probed by LISA, making them an interesting subject for theoretical investigation.

To perform this investigation, we need to contrast a Schwarzschild black hole with a specific candidate whose spacetime is identical to Schwarzschild at large distances, but merges smoothly onto a horizonless solution in an interior region. For this purpose, we adopt as a "straw man" the spacetime of a supermassive boson star (SMBS) whose radius is only a few times the Schwarzschild radius. Such a star consists of a coherent scalar-field configuration supported against gravity by its own self-interaction. Although no fundamental scalar fields have yet been discovered, the fertile imaginations of particle theorists have provided no shortage of candidates (e.g., the standard-model Higgs field, the squark, slepton, and sneutrino fields in supersymmetric models, the axion field, and the dilaton in supergravity models). We restrict our attention to nontopological soliton stars, which are characterized by interaction potentials for which bound, stable solutions exist even in the absence of gravity (Lee & Pang, 1992). This model allows us to choose parameters for the scalar-field Lagrangian so that a $\sim 10^6 M_\odot$ SMBS emerges for massive parameters of order 100 GeV. This is not to suggest, however, that such a SMBS is necessarily likely to arise within the standard model or its most natural extensions. SMBSs with similar structure can just as easily arise with vastly different mass scales in other models, including "mini-boson stars" (Friedberg et al., 1987a), nonspherical but axially symmetric scalar-field configurations partially supported against gravitational collapse by angular momentum (Ryan, 1997a), and non-solitonic boson stars in which a massive scalar field is held up against gravitational collapse by a quartic self-interaction (Colpi et al., 1986).

How might the EMRI into such a SMBS differ from that into a SMBH? The famous "no hair" theorem states that the properties of an uncharged black hole are uniquely determined by the hole's mass M and spin a. Any stationary, axisymmetric metric can be expanded in terms of mass and current multipole moments M_l and S_l (Hansen, 1974). For the Kerr metric, the "no hair" theorem implies that all multipole moments can be expressed in terms of the mass and spin, $M_l + iS_l = M(ia)^l$. By contrast, no such strict relation need exist between the multipole moments of a generic boson-star metric. For example, Ryan (Ryan, 1995) showed that all the multipole moments of the central object can be extracted from the gravitational waveform produced during an EMRI, even in the restricted case of circular orbits in the equatorial plane. As the orbital frequency Ω increases during inspiral, the number of radians of orbital motion per logarithmic frequency interval $d\Phi/d(\ln\Omega)$ can be calculated in a power series in Ω. The coefficients of this power series are simple polynomial functions of the multipole moments. Ryan applied this formalism to spinning boson stars whose mass quadrupole moment $|M_2|$ greatly exceeded Ma^2, the value expected for a black hole of comparable mass and spin (Ryan, 1997a). The spherical boson stars we consider here have perfectly Schwarzschild spacetimes outside of their surfaces, making them indistinguishable from black holes during the early stages of an EMRI. We rely instead on the GWs produced following the final plunge from the ISO into the central object itself. For a black hole, the presence of an event horizon precludes all observations of the inspiraling CO subsequent to the final plunge. After a brief "ringdown" period, the black hole ceases to be a significant source of gravitational radiation. For boson-star inspirals however, many orbits within the boson-star interior are expected, provided that the CO interacts only gravitationally with the scalar field. For compact boson stars with surfaces interior to the ISO, circular orbits can develop an extremely large post-plunge eccentricity, leading to a sudden excitation of higher-order harmonics of the fundamental frequency $f \equiv \Omega/\pi$. As the CO spirals deeper into the boson-star potential well, we expect the fundamental frequency to decrease as less mass is enclosed by smaller orbits. By explicitly comparing the waveforms produced during black-hole and boson-star inspirals, we hope to determine how effectively they can be differentiated.

A final caveat to consider is whether the accumulation of a large mass at the center of the boson star either through inspirals or accretion will cause it to collapse into a black hole. While we expect such a collapse to occur beyond a certain mass limit, the calculation of this collapse is highly model-dependent and beyond the scope of this paper. Such collapses may affect the event rates of boson-star inspirals, but will not alter either the dymanics or waveforms which are the subject of this paper.

The paper is organized as follows. In §6.2, we describe the particular boson-star model examined in this paper as originally formulated by Friedberg et al. (1987b). Spherically symmetric solutions for the metric and scalar field are obtained from the Euler-Lagrange and Einstein equations. In §6.3, we determine the geodesics of the boson-star metric, and consider qualitatively the possible allowed trajectories for the CO during inspiral. We then quantify this approach in §6.4 by presenting a model for the loss of energy and angular momentum to gravitational radiation. The CO's orbital parameters are then evolved in light of these radiative losses in §6.5. Calculated trajectories and waveforms for several initial conditions are displayed in §6.6, and some final remarks on the limitations of our approach and open questions that need to be addressed are given in §6.7. A brief appendix examines the accuracy of the quadrupole approximation near the ISO.

6.2 Boson Star Model

Two requirements must be satisfied for a scalar field to have nontopological-soliton solutions (Lee & Pang, 1992). The first of these requirements is the existence of an additive conservation law, which by Noether's theorem, can be guaranteed by a symmetry of the Lagrangian. In the model of Friedberg, Lee, and Pang Friedberg et al. (1987b) adopted in this paper, the Lagangian density $\mathfrak{L}$ is invariant under a global phase transformation $\phi \rightarrow e^{i\theta}\phi$ of the complex scalar field ϕ,

$$\mathfrak{L} = -\phi^{\dagger\mu}\phi_{\mu} - U(\phi^{\dagger}\phi)\,, \tag{6.1}$$

where a dagger denotes Hermitian conjugation and $\phi_\mu \equiv \partial\phi/\partial x^\mu$. Such a Lagrangian will possess a conserved Noether current

$$j^\mu \equiv -i(\phi^\dagger\phi^\mu - \phi^{\dagger\mu}\phi)\,, \tag{6.2}$$

and a corresponding conserved global charge or particle number

$$N \equiv \int j^0|g|^{1/2}dx^1dx^2dx^3\,, \tag{6.3}$$

where $|g|$ is the absolute value of the determinant of the metric $g_{\mu\nu}$. A nonzero charge N implies a time-varying scalar field by the definition of the current in Eq. (6.2); this time dependence is given by

$$\phi(r,t) = \frac{1}{\sqrt{2}}\sigma(r)e^{-i\omega t}\,, \tag{6.4}$$

where the frequency ω is independent of position but varies with particle number N. The second requirement for the existence of soliton solutions is a constraint on the interaction potential $U(\sigma^2)$. For some forms of this potential, free particles or a black hole will be energetically favored over a boson star for all values of N. However, if $U(\sigma^2) - 1/2m^2\sigma^2$ is negative for some range of σ, then stable boson-star solutions are guaranteed to exist for some values of N. The simplest polynomial form of the potential that satisfies this criterion is

$$U = \frac{1}{2}m^2\sigma^2\left[1 - \left(\frac{\sigma}{\sigma_0}\right)^2\right]^2\,, \tag{6.5}$$

which allows for a false vacuum within the boson star for $\sigma = \sigma_0$.

Having specified the form of the Lagrangian density $\mathfrak{L}$, we can now determine the allowed boson-star solutions (Friedberg et al., 1987b). In the case of spherical symmetry, the spacetime metric of the boson star can be given in full generality by

$$ds^2 = -e^{2u(r)}dt^2 + e^{2\bar{v}(r)}dr^2 + r^2(d\theta^2 + \sin^2\theta d\phi^2)\,. \tag{6.6}$$

The two metric functions $u(r)$ and $\bar{v}(r)$, along with the scalar-field configuration $\sigma(r)$ fully specify a particular boson-star solution. They are chosen to satisfy three independent, ordinary differential equations: the Euler-Lagrange equation for the scalar field, and the tt and rr components of the Einstein equation. In general these equations are stiff; the scalar field drops from $\sigma \simeq \sigma_0$ to zero across a surface layer of thickness m^{-1}, while the metric functions vary over the much longer length scale $R \simeq GM$. If $\lambda \equiv \sigma_0/m_{\rm Pl} \ll 1$, the scalar field can be described by a step function

$$\begin{aligned} \sigma &= \sigma_0 \qquad r \leqslant R \\ &= 0 \qquad r > R\,. \end{aligned} \tag{6.7}$$

The system of equations then reduces to *two* coupled, first-order equations for the metric functions interior to the boson-star surface at $r = R$. With the definitions

$$\bar{r} \equiv \lambda^2 m r \quad \text{and} \quad e^{-\bar{u}} \equiv \frac{\omega}{\lambda m} e^{-u}\,, \tag{6.8}$$

these two equations are

$$\begin{aligned} 2\bar{r}\frac{d\bar{v}}{d\bar{r}} &= \left(\frac{1}{2}e^{-2\bar{u}}\bar{r}^2 - 1\right)e^{2\bar{v}} + 1\,, \\ 2\bar{r}\frac{d\bar{u}}{d\bar{r}} &= \left(\frac{1}{2}e^{-2\bar{u}}\bar{r}^2 + 1\right)e^{2\bar{v}} - 1\,. \end{aligned} \tag{6.9}$$

Physical boundary conditions imply that $\bar{v}(0) = 0$, while $\bar{u}(0)$ and R are chosen self-consistently to ensure that the metric function $u(r)$ smoothly matches onto the Schwarzschild solution $e^{2u} \to (1 - 2GM/r)$ for $r \geqslant R$. In this choice of coordinates, $\bar{v}(r)$ need not be continuous at $r = R$. The metric functions $e^{2u(r)}$ and $e^{2\bar{v}(r)}$ are depicted in Fig. 6.1.

The final result of this analysis is that for fixed values of m and σ_0, the boson-star solutions constitute a one-parameter family corresponding to different values of the particle number N. The boson-star mass M increases monotonically with N, and the boson star becomes increasingly compact. Eventually a critical limit of compactness

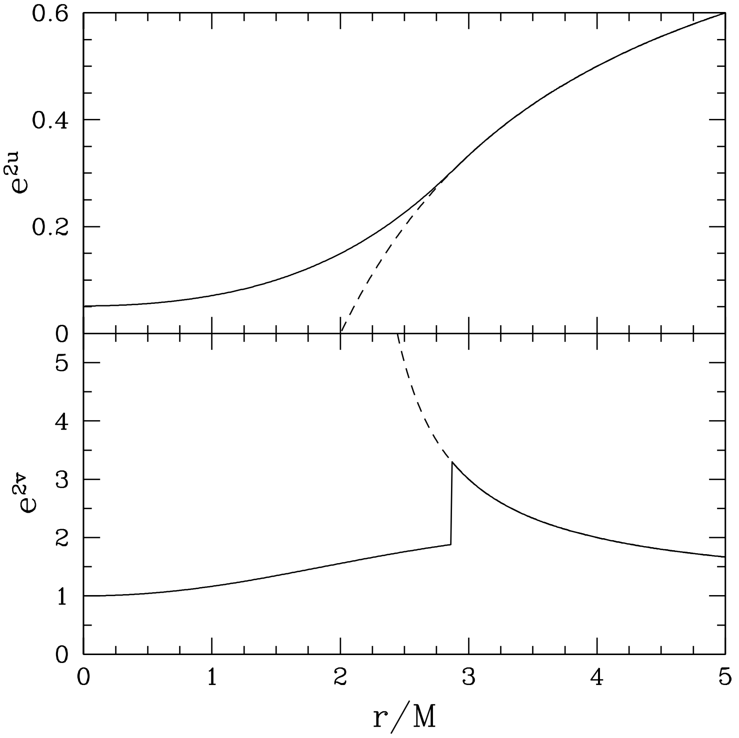

Figure 6.1: The metric functions $e^{2u(r)}$ and $e^{2\bar{v}(r)}$ as functions of radius. The solid curves correspond to a boson star with $R = 2.869M$, while the the dashed curves are for a black hole of the same mass. Outside the boson-star surface, the two curves are identical, while for $r < R$, the solid curves are a numerical solution to Eq. (6.9) and the dashed curves are the Schwarzschild metric functions $e^{2u(r)} = (1 - 2GM/r)$, $e^{2\bar{v}(r)} = (1 - 2GM/r)^{-1}$.

$R = 2.869GM$ is reached beyond which the boson star collapses into a black hole. We choose to investigate EMRIs into this critically compact boson star. We set $m = \sigma_0$ to the value that yields the desired SMBS mass M. For the nontopological soliton considered here, $M \sim m_{\rm Pl}^4/m^3$, implying that a boson of mass $m \sim 10$ TeV will lead to a SMBS of $M \sim 10^6 M_\odot$. The SMBS mass scales differently with m for other types of boson stars.

6.3 Geodesics of the Boson-Star Metric

In the extreme–mass-ratio limit $\eta \equiv \mu/M \ll 1$, where μ and M are the mass of the CO and boson star respectively, the CO will travel along geodesics of the boson-star metric. During the course of the EMRI, energy and angular momentum are radiated away on time-scales much longer than an orbital period. We can therefore make the adiabatic approximation that the CO will migrate smoothly between geodesics characterized by decreasing energy and angular momentum. A thorough description of the possible geodesics of the boson-star metric is thus essential to understanding the evolution of the CO's trajectory. This description can be simplified by noting that the metric of Eq. (6.6) is independent of t and ϕ, implying the existence of time-like and azimuthal Killing fields ξ^μ and ψ^μ. For a CO travelling along a geodesic with four-velocity $u^\nu = (dt/d\tau, dr/d\tau, d\theta/d\tau, d\phi/d\tau)$, the inner product of u^ν with a Killing field is conserved (Wald, 1984). This allows a formal definition of the conserved energy and angular momentum per unit mass,

$$\begin{aligned} E &\equiv -g_{\mu\nu}\xi^\mu u^\nu = e^{2u}\frac{dt}{d\tau}\,, \\ L &\equiv g_{\mu\nu}\psi^\mu u^\nu = r^2\frac{d\phi}{d\tau}\,. \end{aligned} \tag{6.10}$$

Note that we can restrict ourselves to orbits with $\theta = \pi/2$ without loss of generality because of spherical symmetry, in which case $d\theta/d\tau = 0$. The definitions of Eq. (6.10), coupled with the norm of the four-velocity

$$g_{\mu\nu}u^\mu u^\nu = -e^{2u}\left(\frac{dt}{d\tau}\right)^2 + e^{2\bar{v}}\left(\frac{dr}{d\tau}\right)^2 + r^2\left(\frac{d\phi}{d\tau}\right)^2 = -1\,, \tag{6.11}$$

provide three coupled, first-order differential equations that can be solved for the CO's orbit $\{t(\tau), r(\tau), \phi(\tau)\}$ as a function of proper time τ.

While this approach formally solves the problem of geodesic motion, further insight can be gained by recasting Eq. (6.11) to make an analogy with one-dimensional particle motion (Chandrasekhar, 1992),

$$\begin{aligned} E^2 &= e^{2u+2\bar{v}}\left(\frac{dr}{d\tau}\right)^2 + e^{2u}\left(\frac{L^2}{r^2}+1\right) \\ &= e^{2u+2\bar{v}}\left(\frac{dr}{d\tau}\right)^2 + V^2(r)\,. \end{aligned} \tag{6.12}$$

The first term on the right-hand side of Eq. (6.12) acts like a positive-definite kinetic energy, while the second term is the potential well in which the one-dimensional particle motion occurs. This analogy is not as complete as in the case of a Schwarzschild black hole, where the metric functions conspire to make the coefficient of the kinetic term independent of position. For boson stars, the term $e^{2u+2\bar{v}}$ acts like a position-dependent mass; while this will affect the particle motion quantitatively, it will not preclude the existence of bound orbits in the boson-star interior.

Plots of the effective potential $V(r)$ for different values of the angular momentum L are given in Fig. 6.2. The effective potential per unit mass and ratio r/M are dimensionless in units where $G = c = 1$. Outside of the boson-star surface at $R = 2.869M$, $V(r)$ is identical to that of Schwarzschild black holes. For $L^2 \geqslant 12M^2$, a local minimum exists at $r_1 \geqslant 6M$. Test particles with $E = V(r_1)$ can follow stable circular orbits at r_1, while those with slightly larger energies experience radial oscillations about this minimum corresponding to eccentric orbits. Particles with $E > 1$ are unbound, and will either escape to infinity or – if they have energies greater than the local maximum and $dr/dt < 0$ – penetrate to the boson-star interior. It is here that the story changes dramatically from that of the black hole, where $V(r)$ plunges to negative infinity at the singularity. Particles that enter the black hole's event horizon are directed inexorably to the center, never to return. By contrast, the potential well of the boson star has finite depth, implying that for $L > 0$, the effective potential is repulsive at short distances. This creates a second minimum in the effective potential at r_2, at which stable circular orbits and associated eccentric orbits can occur. Such orbits exist even for $E > 1$, though it is unclear in what astrophysical context a CO might find itself on such a geodesic. For a typical EMRI, a CO will be

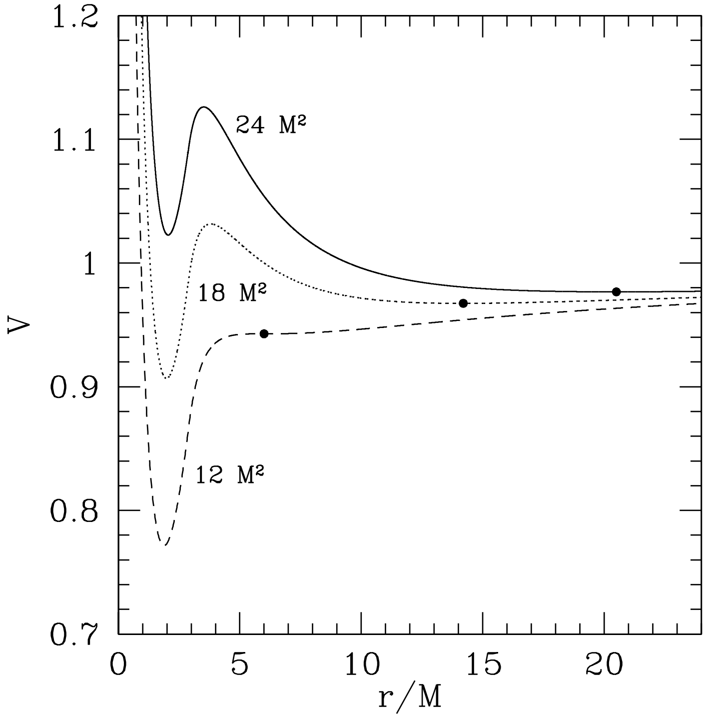

Figure 6.2: The effective potential $V(r)$ of the boson star for three different values of the angular momentum. The solid, short-dashed, and long-dashed curves have $L^2 = 24M^2$, $18M^2$, and $12M^2$, respectively. The large black dots on each curve show the location of the outer minimum. For $L^2 = 12M^2$ – the lowest angular momentum for which there are two minima – the black dot is located at the innermost stable circular orbit $r = 6M$.

scattered onto an eccentric orbit and slowly circularize and wander inward as energy and angular momentum are lost to gravitational radiation. Eventually, either before or after circularization depending on the initial eccentricity and semi-major axis, the CO will reach the ISO and plunge into the boson-star interior. At this point, the trajectory and corresponding gravitational waveform will diverge radically from the case of an EMRI into a black hole. Provided there is no non-gravitational interaction between the CO and the scalar field, the inspiral will continue on a highly eccentric geodesic about the inner minimum of the effective potential. Energy and angular momentum will evolve continuously through the plunge, while both the apocenter and pericenter will change on orbital time-scales as the center of radial oscillations changes from the outer minimum r_1 to the inner minimum r_2. High eccentricity at the ISO leads to even more complicated inspirals as the plunge takes place in a two-step process. The CO crosses over the local maximum of the effective potential several orbits before settling into the inner local minimum. The pericenter penetrates deep into the boson-star interior in the first stage of this two-step plunge, while in the second stage, the apocenter shifts from outside to inside the local maximum. This inspiral would undoubtably produce a very distinctive waveform, but only a quantitative analysis can determine whether such scenarios are possible. To make such a quantitative analysis, we must first develop a suitable approximation for the calculation of the gravitational-wave emission.

6.4 Gravitational Radiation

Gravitational-wave emission by arbitrary sources is a challenging, unsolved problem. In any asymptotically flat spacetime one can define the GW field,

$$\bar{h}^{\alpha\beta} \equiv -(-g)^{1/2} g^{\alpha\beta} + \eta^{\alpha\beta}\,, \tag{6.13}$$

where $g_{\alpha\beta}$ is the metric, g is its determinant, and $\eta^{\alpha\beta}$ is the flat-space Lorentz metric (Thorne, 1980). Far from the source, this quantity reduces to the trace-reversed metric perturbation $\bar{h}^{\alpha\beta} = h^{\alpha\beta} - 1/2\eta^{\alpha\beta}h$. The perturbation $\bar{h}^{\alpha\beta}$ satisfies the de Donder gauge condition

$$\bar{h}^{\alpha\beta}{}_{,\beta} = 0\,, \tag{6.14}$$

and the exact Einstein field equations

$$\Box \bar{h}^{\alpha\beta} = -16\pi\tau^{\alpha\beta}\,. \tag{6.15}$$

We see that the GW field is sourced by an "effective stress-energy tensor" $\tau^{\alpha\beta}$ of the form (Thorne, 1980)

$$\tau^{\alpha\beta} = (-g)(T^{\alpha\beta} + t^{\alpha\beta}_{\rm LL}) + (16\pi)^{-1}[\bar{h}^{\alpha\mu}{}_{,\nu}\,\bar{h}^{\beta\nu}{}_{,\mu} - \bar{h}^{\alpha\beta}{}_{,\mu\nu}\,\bar{h}^{\mu\nu}]\,, \tag{6.16}$$

where $T^{\alpha\beta}$ is the true stress-energy tensor and $t^{\alpha\beta}_{\rm LL}$, the Landau-Lifschitz pseudotensor given by Eq. (20.22) of Misner et al. (1973), is a highly nonlinear function of the full metric. We thus see that for highly relativistic, strong-field sources, the GW field $\bar{h}^{\alpha\beta}$ will itself provide a non-negligible contribution to its source $\tau^{\alpha\beta}$. In general Eq. (6.15) will lead to a system of tightly coupled, nonlinear differential equations that can only be integrated for brief periods numerically before becoming unstable. To make any analytical progress, we must rely on a series of simplifying approximations of varying degrees of validity.

The first approximation we make is that in the extreme-mass-ratio limit $\eta \equiv \mu/M \ll 1$, the boson star is at rest at the origin of our coordinate system and its metric is static. The inspiraling CO is thus the sole time-varying source that contributes to GW emission. In reality, the CO should raise tides on the surface of the boson star, which could then generate GWs themselves as well as back-react on the CO, altering its trajectory. In an inspiral into a black hole, the energy loss due to the tidal interaction is at greatest, only a few percent of the total flux, although this can still lead to several hundred cycles of phase difference in the gravitational waveform (Finn & Thorne, 2000). In the boson-star case, the nature of the tidal interaction will be different and model-dependent, and provides another way to identify boson-star inspirals. However, the tidal interaction should still be a lesser effect than the orbital dynamics. The back-reaction on the orbit is suppressed by an additional factor of η, implying that for EMRIs with $\eta \simeq 10^{-6}$, it will introduce a very subdominant source of error into our calculations. The second approximation we make is to ignore the curvature induced by the boson star as it affects the propagation of GWs to infinity. GWs propagating outwards from $r \lesssim (\lambda^2 M)^{1/3}$, where λ is the GW wavelength, will be distorted and backscattered by the background curvature of the boson star

(Thorne, 1980). The importance of this effect will depend on the inner boundary conditions, which should differ significantly from those of a black hole with an event horizon. We ignore this complication entirely by assuming that backscattering does not occur. A third approximation that we make is to compute waveforms using weak-field formulae. We wish to apply these to orbits that come very close to the boson star, a regime in which there is no natural flat-space coordinate system in which to evaluate the weak-field expressions. Our approach is to identify the Schwarzschild coordinates (t, r, θ, ϕ) of the inspiral trajectory with true flat-space spherical polar coordinates. This identification is exact for orbits out in the weak-field. There is no *a priori* reason why this identification is better than any other (e.g., identifying isotropic coordinates with the flat-space spherical polar coordinates). In practice, using different coordinate identifications gives different answers, but the differences are only a few percent. Taken together, these three approximations are equivalent to assuming the GWs from the EMRI will be the same as those that would be produced were the CO a "particle on a string", artificially constrained to move on an orbit $\{t(\tau), r(\tau), \phi(\tau)\}$ in flat space. In the case of black-hole inspirals, waveforms computed using this "hybrid" approach (Glampedakis et al., 2002; Babak et al., 2005; Creighton et al., 2005) have been found to compare quite well to waveforms computed using more accurate perturbative techniques.

In the flat-space, weak-field limit implied by the above approximations, the transverse traceless (TT) part of the GW field can be expanded in symmetric trace-free (STF) tensors (Thorne, 1980)

$$\begin{aligned} h^{\mathrm{TT}}_{jk} \;=\; & \Bigg[\sum_{l=2}^{\infty}\left(\frac{4}{l!}r^{-1}\right)\,{}^{(l)}\mathcal{I}_{jkA_{l-2}}(t-r)N_{A_{l-2}} \\ & + \sum_{l=2}^{\infty}\left(\frac{8l}{(l+1)!}\right)r^{-1}\epsilon_{pq(j}\;{}^{(l)}\mathcal{S}_{k)pA_{l-2}}(t-r)n_q N_{A_{l-2}}\Bigg]^{\mathrm{TT}}, \end{aligned} \tag{6.17}$$

where $A_l \equiv a_1...a_l$, $\mathcal{I}_{A_l}$, and $\mathcal{S}_{A_l}$ are the rank-l mass and current multipole-moment tensors, n_a is a unit radial vector, and $N_{A_l} \equiv n_{a_1}...n_{a_l}$. Parentheses on the indices denote taking the symmetric part, and a superscript (l) in front of a tensor indicates taking the l-th time derivative. The energy and angular momentum loss to gravitational radiation can also be expanded in multipole-moment tensors,

$$\begin{aligned} \frac{d(\mu E)}{dt} &= \sum_{l=2}^{\infty} \frac{(l+1)(l+2)}{(l-1)l} \frac{1}{l!(2l+1)!!} \langle {}^{(l+1)}\mathcal{I}_{A_l}\, {}^{(l+1)}\mathcal{I}_{A_l} \rangle \\ &\quad + \sum_{l=2}^{\infty} \frac{4l(l+2)}{(l-1)} \frac{1}{(l+1)!(2l+1)!!} \langle {}^{(l+1)}\mathcal{S}_{A_l}\, {}^{(l+1)}\mathcal{S}_{A_l} \rangle\,, \qquad (6.18) \\ \frac{d(\mu L_j)}{dt} &= \sum_{l=2}^{\infty} \frac{(l+1)(l+2)}{(l-1)l!(2l+1)!!} \langle \epsilon_{jpq}\, {}^{(l)}\mathcal{I}_{pA_{l-1}}\, {}^{(l+1)}\mathcal{I}_{qA_{l-1}} \rangle \\ &\quad + \sum_{l=2}^{\infty} \frac{4l^2(l+2)}{(l-1)(l+1)!(2l+1)!!} \langle \epsilon_{jpq}\, {}^{(l)}\mathcal{S}_{pA_{l-1}}\, {}^{(l+1)}\mathcal{S}_{qA_{l-1}} \rangle\,. \qquad (6.19) \end{aligned}$$

Here, $\langle \quad \rangle$ denotes averaging over an entire orbital period. A μ appears on the left-hand side of the equation in accordance with our definition of E and L as the energy and angular momentum per unit mass. This multipole-moment expansion is valid in principle for sources moving at arbitrarily relativistic velocities. We now further approximate that the sources are Newtonian; the maximum velocity is only mildly relativistic, and the internal stresses are small compared to the energy density (Thorne, 1980). Although this approximation is violated during the EMRI in the vicinity of the plunge, key features of the relativistic motion are preserved by constraining the CO to travel along geodesics that are *exact* even in the relativistic regime. Empirically, expansions in the post-Newtonian parameter $(L/\lambda)^2 \sim v^2$ are found to give physical results (no outspiral), even in regions where they do not converge (Barack & Cutler, 2004). Assuming Newtonian sources, the mass and current multipole tensors are given by (Thorne, 1980)

$$\begin{aligned} \mathcal{I}_{A_l} &= \left[\int \rho X_{A_l}\, d^3\vec{x} \right]^{\rm STF}\,, \qquad (6.20) \\ \mathcal{S}_{A_l} &= \left[\int (\epsilon_{a_l pq} x_p \rho v_q) X_{A_{l-1}}\, d^3\vec{x} \right]^{\rm STF}\,, \qquad (6.21) \end{aligned}$$

where $X_{A_l} \equiv x_{a_1}...x_{a_l}$, and ρ is the energy density. In keeping with our assumption that the static boson-star metric does not contribute to the production of GWs, the energy density is simply that of the CO

$$\rho(\vec{x}) = \mu\delta^3(\vec{x} - \vec{x}_{\rm CO}(t))\,, \tag{6.22}$$

where $\vec{x}_{\rm CO}(t)$ is the flat-space trajectory of the CO. The mass and current multipole moments $\mathcal{I}_{\mathcal{A}_\updownarrow}$ and $\mathcal{S}_{\mathcal{A}_\updownarrow}$ contribute to the GW field $h^{\rm TT}_{jk}$ at order $(M/r)(L/\lambda)^l$ and $(M/r)(L/\lambda)^{l+1}$, respectively, while the assumption of Newtonian sources induces errors of order $(L/\lambda)^2$ to each term (Thorne, 1980). It is therefore inconsistent to include terms higher than the mass quadrupole, current quadrupole, and mass octupole under this assumption. We consider only the mass quadrupole term, relegating an analytic calculation of the error associated with this approximation to the Appendix. In this case, Eqs. (6.17), (6.18), and (6.19) reduce to the familiar form of the "quadrupole approximation"

$$h^{\rm TT}_{jk} = \frac{2}{r}\frac{d^2\mathcal{I}_{jk}}{dt^2}\,, \tag{6.23}$$

$$\frac{d(\mu E)}{dt} = \frac{1}{5}\left\langle \frac{d^3\mathcal{I}_{jk}}{dt^3}\frac{d^3\mathcal{I}_{jk}}{dt^3}\right\rangle\,, \tag{6.24}$$

$$\frac{d(\mu L_j)}{dt} = \frac{2}{5}\left\langle \epsilon_{jpq}\frac{d^2\mathcal{I}_{pk}}{dt^2}\frac{d^3\mathcal{I}_{qk}}{dt^3}\right\rangle\,. \tag{6.25}$$

How can the humble quadrupole approximation be justified for the eccentric, highly relativistic final stages of an EMRI into a boson star? While the energy and angular momentum fluxes are indeed only approximate, the derivatives relating these fluxes to changes in the orbital parameters are *exact*, as are the equations of motion used in performing the orbital averages of Eqs. (6.24) and (6.25). Unlike direct post-Newtonian expansions for the time evolution of orbital elements, this "hybrid approximation" incorporates the exact orbital dynamics of geodesic motion (Glampedakis et al., 2002; Gair & Glampedakis, 2005). This approach is self-consistent in the sense that the energy and angular momentum carried away by gravitational waves (within our quadrupole approximation) is equal to the loss of angular momentum and energy of the orbit. This approach can reproduce features missing from direct post-

Newtonian expansions (which conserve energy and angular momentum only to $\mathcal{O}(v^2)$) such as the *increase* in orbital eccentricity just prior to plunge. Moreover, using exact geodesics in the flat-space quadrupole formula ensures that the fundamental frequencies of the orbit are reproduced in the gravitational waveforms. Although the distribution of power between harmonics is not correct, these approximate waveforms do encode the same orbital dynamics as the true waveform and therefore, the qualitative features of inspiral waveforms should be well represented. At substantial computational expense, gravitational perturbation theory can calculate waveforms at infinity using the Teukolsky-Sasaki-Nakamura (TSN) formalism (Teukolsky, 1973; Sasaki & Nakamura, 1982). Such calculations for EMRIs in the Schwarzschild spacetime show agreement with the hybrid approach to within 5 to 45% for the time derivatives of the orbital semi-latus rectum and eccentricity even for moderate eccentricity ($e \sim 0.4$) at $r \simeq 7M$ (Glampedakis et al., 2002; Babak et al., 2005). We hope that such accuracy will be retained in the case of boson stars, sparing us for now the additional computational expense of a numerical TSN approach. The details of adapting this hybrid approximation to boson-star EMRIs is the subject of the next Section.

6.5 Orbital Evolution

In a spherically symmetric spacetime, conservation of angular momentum implies that geodesic motion will be confined to an orbital plane (Murray & McDermott, 1999). Without loss of generality, this plane can be chosen to be the equatorial plane $\theta = \pi/2$. In this case, the position of the CO is fully specified by a radius r and a true longitude ϕ. A geodesic will be characterized by orbital elements like the semi-latus rectum p and eccentricity e, in terms of which the radius is given by

$$r = \frac{p}{1 + e\cos\psi}\,. \tag{6.26}$$

The true anomaly ψ can be specified independently of ϕ, since in general the pericenter will precess. In the hybrid approximation, EMRI orbits and waveforms are calculated in a two-step process that hinges on the assumption of adiabaticity (i.e., orbital elements like p and e vary on time-scales much longer than an orbital period over which ψ and ϕ change (Glampedakis et al., 2002; Gair & Glampedakis, 2005)). This allows the differential equations governing the evolution of p and e to be integrated with a much longer timestep than is required for evolving ψ and ϕ. We will first describe how the two stages of integration are accomplished with this assumption, and then we will discuss the corrections necessary to patch together a physically reasonable orbit through those regions where the assumption is violated.

In the first stage of our calculation, the EMRI's trajectory through a phase space of p and e is determined by relating these quantities to E and L, and then using this relation and Eqs. (6.24) and (6.25) to obtain differential equations for p and e. Particular values of p and e uniquely determine the pericenter $r_p = p/(1+e)$ and apocenter $r_a = p/(1-e)$ from Eq. (6.26). At both pericenter and apocenter, $dr/d\tau = 0$ implying $V(r_p) = V(r_a) = E$ by Eq. (6.12). This equation can be solved to obtain the energy E and L as functions of p and e,

$$\begin{aligned} L(p,e) &= \left[\left(e^{2u(r_p)} - e^{2u(r_a)}\right)\left(\frac{e^{2u(r_a)}}{r_a^2} - \frac{e^{2u(r_p)}}{r_p^2}\right)^{-1}\right]^{1/2} , \\ E(p,e) &= \left[e^{2u(r_a)}\left(\frac{L^2}{r_a^2} + 1\right)\right]^{1/2} . \end{aligned} \tag{6.27}$$

First-order differential equations for the time evolution of p and e can then be derived from the *exact* derivatives of $E(p,e)$ and $L(p,e)$ along with the quadrupole appoximations of Eqs. (6.24) and (6.25),

$$\begin{aligned} \frac{dp}{dt} &= \frac{\frac{dL}{dt} - \frac{\partial L/\partial e}{\partial E/\partial e}\frac{dE}{dt}}{\frac{\partial L}{\partial p} - \frac{\partial L/\partial e}{\partial E/\partial e}\frac{\partial E}{\partial p}} , \\ \frac{de}{dt} &= \frac{\frac{dL}{dt} - \frac{\partial L/\partial p}{\partial E/\partial p}\frac{dE}{dt}}{\frac{\partial L}{\partial e} - \frac{\partial L/\partial p}{\partial E/\partial p}\frac{\partial E}{\partial e}} . \end{aligned} \tag{6.28}$$

It is the combination of these exact derivatives with the quadrupole approximation for the fluxes that accounts for the superior performance of this "hybrid" approximation over the consistent (in that we include *all* terms up to a given order in v^2) post-Newtonian approach. The quadrupole fluxes themselves can be determined by inserting Eq. (6.22) for the CO's energy density into Eq. (6.20) for the Newtonian quadrupole moments. Using an overdot to denote a derivative with respect to coordinate time t, we find that the time derivatives of the quadrupole moment appearing in Eqs. (6.24) and (6.25) take the form

$$\begin{aligned}
\dddot{\mathcal{I}}_{jk}\dddot{\mathcal{I}}_{jk} &= 2\eta^2[(1/3)(\dddot{r} + 3\ddot{r}\dot{r})^2 + (6\ddot{r}\dot{\phi}r + 6\dot{r}\ddot{\phi}r + 6\dot{r}^2\dot{\phi} - 4\dot{\phi}^3r^2 + \dddot{\phi}r^2)^2 \\
&\quad + (\dddot{r}r + 3\ddot{r}\dot{r} - 12\dot{r}\dot{\phi}^2r - 6\ddot{\phi}\dot{\phi}r^2)^2]\,, \\
\epsilon_{jpq}\ddot{\mathcal{I}}_{pk}\dddot{\mathcal{I}}_{qk} &= 2\eta^2[(\ddot{r}r + \dot{r}^2 - 2r^2\dot{\phi}^2)(6\ddot{r}\dot{\phi}r + 6\dot{r}\ddot{\phi}r + 6\dot{r}^2\dot{\phi} - 4\dot{\phi}^3r^2 + \dddot{\phi}r^2) \\
&\quad -(4\dot{r}r\dot{\phi} + r^2\ddot{\phi})(\dddot{r}r + 3\ddot{r}\dot{r} - 12\dot{r}\dot{\phi}^2r - 6\ddot{\phi}\dot{\phi}r^2)]\,. \qquad (6.29)
\end{aligned}$$

The derivatives appearing in Eq. (6.29) can be evaluated analytically by solving Eqs. (6.10) and (6.11) for the equations of motion

$$\frac{dt}{d\tau} = Ee^{-2u}\,, \qquad (6.30)$$

$$\frac{d\phi}{d\tau} = \frac{L}{r^2}\,, \qquad (6.31)$$

$$\frac{dr}{d\tau} = \pm e^{-\bar{v}}\left[E^2e^{-2u} - \left(\frac{L^2}{r^2} + 1\right)\right]^{1/2}\,, \qquad (6.32)$$

dividing Eqs. (6.32) and (6.31) by Eq. (6.30), and then taking the appropriate derivatives. We then perform the orbital average appearing in Eqs. (6.24) and (6.25),

$$\langle f(r)\rangle = \frac{\int_{r_p}^{r_a} f(r)\,\frac{dr}{\dot{r}}}{\int_{r_p}^{r_a} \frac{dr}{\dot{r}}}\,. \qquad (6.33)$$

The products of these manipulations are first-order differential equations for $\dot{p}$ and $\dot{e}$ that can be integrated forward in time for arbitrary initial p_0 and e_0. This tabulated phase-space trajectory $\{p(t), e(t)\}$ will then be used in the second stage to produce the real-space trajectory $\{r(t), \phi(t)\}$.

The second stage of calculating the EMRI real-space trajectory is essentially no more complicated than integrating the equations of motion Eqs. (6.30), (6.31), and (6.32) with a timestep that is a sufficiently small fraction of the orbital period as to obtain the desired accuracy. The phase-trajectory determined in the first stage can be linearly interpolated to obtain $p(t)$ and $e(t)$, from which the energy and angular momentum follow from the relations $E(p, e)$ and $L(p, e)$. These are then inserted into the equations of motion along with r and ϕ at the beginning of the timestep. A minor technical difficulty arises for circular orbits ($e = 0$) and at the turning points of eccentric orbits ($\psi = n\pi$) where the relation

$$\frac{d\psi}{d\tau} = \frac{1 + e\cos\psi}{er\sin\psi}\frac{dr}{d\tau} \tag{6.34}$$

becomes undefined. This problem can be solved by a judicious Taylor expansion of Eq. (6.32) about the appropriate r values. Once the real-space trajectory $\{x(t)$, $y(t)$, $z(t)\}$ has been obtained, the quadrupole waveform of Eq. (6.23) can be calculated by taking the appropriate numerical derivatives. This portion of the calculation does not involve any special features of the boson-star spacetime, in keeping with the "particle on a string" approximation. The GW field $h_{ij}^{\rm TT}$ can then be decomposed into plus and cross polarizations by defining unit vectors in the plane of the sky. With $\hat{L}$ a unit vector parallel to the CO's angular momentum and $\hat{n}$ a unit vector pointed from the observer to the source, we define

$$\begin{aligned} \hat{p} &\equiv (\hat{n} \times \hat{L})/|\hat{n} \times \hat{L}| \,, \\ \hat{q} &\equiv \hat{n} \times \hat{p} \,. \end{aligned} \tag{6.35}$$

Note that the sign of $\hat{q}$ differs from that of Ref. Barack & Cutler (2004), but in other respects our conventions are equivalent. Our basis vectors $\hat{p}$ and $\hat{q}$ are constants because $\hat{L}$ does not precess in a spherically symmetric spacetime. The two polarization basis tensors can be defined as

$$\begin{aligned} H^{+}_{ij} &\equiv \hat{p}_i\hat{p}_j - \hat{p}_i\hat{p}_j \,, \\ H^{\times}_{ij} &\equiv \hat{p}_i\hat{q}_j + \hat{q}_i\hat{p}_j \,, \end{aligned} \tag{6.36}$$

allowing the GW field to be expressed as two amplitudes,

$$h^{\mathrm{TT}}_{ij}(t) = A^{+}(t)H^{+}_{ij} + A^{\times}(t)H^{\times}_{ij} \,. \tag{6.37}$$

In the next Section, we will examine the trajectories $\{x(t), y(t), z(t)\}$ and GW amplitudes $\{A^{+}(t), A^{\times}(t)\}$ for different EMRIs.

Before presenting these results, we must explain the kludges we have chosen where the hybrid approximation described above is no longer valid. Our philosophy is to adopt the simplest, least computationally intensive approach that yields physically reasonable waveforms. The numerous approximations involved in the "particle on a string" model introduce at least a 10% error into our trajectories and waveforms in the highly eccentric, relativistic limit (Glampedakis et al., 2002; Babak et al., 2005); to demand greater accuracy where the hybrid approximation is also violated would be inconsistent. Special provisions beyond the hybrid approximation must be made whenever p and e are changing on time-scales shorter than the orbital period. For extremely eccentric orbits, gravitational radiation will tend to be emitted in bursts near pericenter, while the CO will be found near apocenter for most of the orbital period. This is particularly true of "zoom-whirl" orbits like those found around rapidly spinning black holes, where the CO may whirl for many radians near pericenter before zooming back out to apocenter (Glampedakis & Kennefick, 2002). In such cases, the orbital averaging of Eq. (6.33) will lead to an artificially smooth phase-space trajectory $\{p(t), e(t)\}$. For boson-star EMRIs, this adiabatic approximation is most flagrantly violated during the plunge itself. Since the rapidly varying true anomaly $\psi(t)$ is only calculated in the second stage of our two-step approach, the CO could be anywhere along its orbit at the time the plunge is calculated to occur in the first stage.

This poses problems in both stages of the calculation. As described qualitatively in §6.3, a CO initially oscillating about the outer minimum of the effective potential $V(r)$ can plunge into the boson star with $L^2 > 12M^2$, in which case, two minima will still exist even after the plunge. For a time, the energy of the CO will exceed $V_{\rm max}(L)$, the local maximum of $V(r)$, so that *both* of the local minima will lie between the orbit's pericenter and apocenter. Such a highly eccentric orbit will radiate energy very efficiently, implying that the CO will quickly plunge into the inner minimum. However, if the initial conditions were such that the CO began on a geodesic with E was significantly above $V_{\rm max}$, there is no reason *a priori* why it could not first fall into the outer minimum before eventually inspiraling into the boson-star interior. Which minimum the CO ends up in depends on where the CO is in its orbit when E drops below $V_{\rm max}$. But the true anomaly $\psi(t)$ is not calculated in the first stage of the hybrid approximation, so we have no way of accurately knowing which minimum to choose. The semi-latus rectum p and eccentricity e change discontinuously when the CO falls into one of the two minima, a sure sign that the adiabatic condition no longer holds. We assume without rigorous proof that the CO *always* ends up in the inner minimum in such a situation. This assumption comes back to haunt us in the second stage, since the CO may well be near apocenter at the time determined for the plunge in the first stage. In this case, no possible instantaneous change in ψ would allow r to remain continuous. Instead, we fix p and e to their values at the time of the plunge, and allow the CO to travel along a geodesic until it reaches the local maximum at $r_{\rm max}$. This then becomes the apocenter of the CO's radial motion about the inner minimum, and we resume linearly interpolating p and e from the phase-space trajectory tabulated in the first stage. Having exhaustively described our technique, we will now examine the fruits of our labor.

6.6 Results

First we consider the phase-space trajectories $\{p(t), e(t)\}$ generated in the first part of our calculation. We choose for our fiducial model a $3M_\odot$ CO spiraling into a

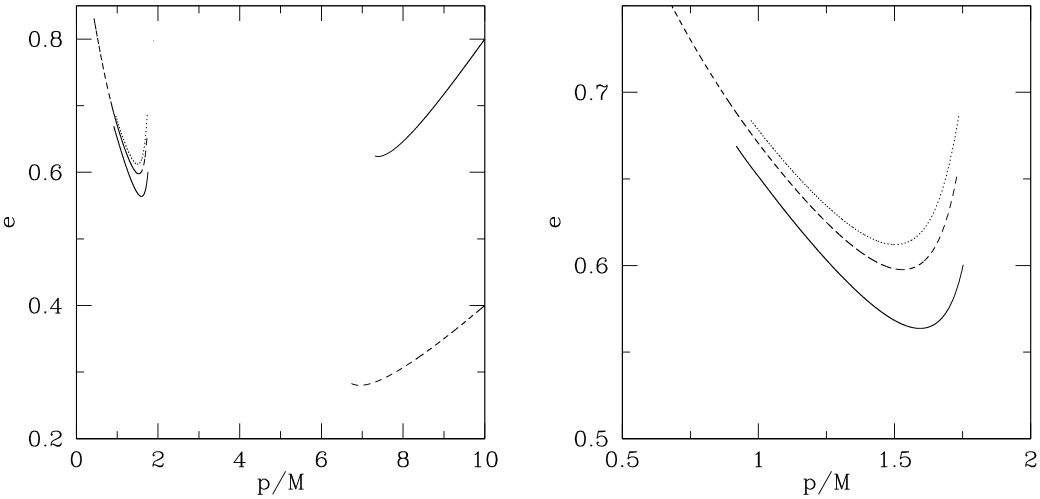

Figure 6.3: The phase-space trajectories $\{p(t), e(t)\}$ for three different EMRIs into a supermassive boson star. The solid, dashed, and dotted curves begin with eccentricities $e_0 = 0.8$, 0.4, and 0.0, respectively. The left panel shows the complete EMRI from $p_0/M = 10.0$, while the right panel is a close-up of the post-plunge phase of the EMRI that appears in the upper-left corner of the top panel. In both panels, time increases from right to left as p/M decreases monotonically.

$3 \times 10^6 M_\odot$ boson star. Three different trajectories are depicted in Fig. 6.3, each with $p_0 = 10M$, but with differing initial eccentricities ($e_0 = 0.0, 0.4, 0.8$). Several features of the phase-space evolution are particularly striking. Outside the boson-star surface at $R = 2.869GM$, the spacetime is identical to that of a Schwarzschild black hole and the behavior of the EMRI is well known. Orbits far from the boson star rapidly circularize, but those that retain an appreciable eccentricity exhibit an increase in the eccentricity in the last few orbits before the ISO. This behavior reflects the increasing shallowness of the outer minima as the ISO is approached, and is not captured by post-Newtonian expansions that ignore geodesic motion (Glampedakis et al., 2002; Glampedakis & Kennefick, 2002). The second important feature is the discontinuity in each of the three trajectories as the plunge is reached. As described in the previous Section, the actual position of the CO remains continuous; only the orbital elements p and e change at the instant of the plunge. This is accomplished by an appropriate instantaneous change in the true anomaly ψ. For the EMRIs with $e_0 = 0.8$ and

0.4, the orbit is actually divided into three distinct pieces, with the middle section belonging to geodesics where *both* minima lie between pericenter and apocenter. This portion cannot be seen in Fig. 6.3 because it only lasts for a few orbital periods, which is short compared to the time scales on which p and e are changing. It is interesting to note that the three curves have reversed their order in the plunge; the more eccentric the EMRI prior to plunge, the lower its eccentricity afterwards. This can be understood by realizing that orbits with higher residual eccentricity plunge into the boson star with higher angular momentum L, implying that the inner minimum will be much steeper as depicted in Fig. 6.2. The pericenter and apocenter will be closer together in this narrow potential well, leading to a correspondingly smaller post-plunge eccentricity. A final surprising feature of these EMRIs is the sharp increase in eccentricity for the final stages of the EMRI deep within the interior of the boson star. The energy and angular momentum are monotonically decreasing during this stage of the EMRI in keeping with GW-induced losses; only the relations $E(p, e)$ and $L(p, e)$ are unusual. As angular momentum is lost, the narrow inner minimum rapidly broadens, leading to a greater separation between pericenter and apocenter. The separation

$$r_a - r_p = \frac{2pe}{1 - e^2} \tag{6.38}$$

is a sharply increasing function of eccentricity, explaining why the broadening inner minimum leads to rising eccentricity. For completeness, a plot of these same trajectories in the phase space of angular momentum and energy $\{L(t)/M, E(t)\}$ is given in Fig. 6.4. Note that E and L are continuous at the plunge as is required for physical quantities.

We now consider the real-space trajectory $\{x(t), y(t), z(t)\}$ itself. Choosing the orbital and equatorial planes to coincide, $z = 0$. Several orbits of an initially circular EMRI on both sides of the plunge are shown in Fig. 6.5. The trajectory remains continuous during the plunge, and the pericenter precesses by an appreciable fraction of a radian on each post-plunge orbit. Of more interest is the EMRI of Fig. 6.6 with $e_0 = 0.4, p_0 = 10.0M$. We have omitted the EMRI with $e_0 = 0.8$ because the two

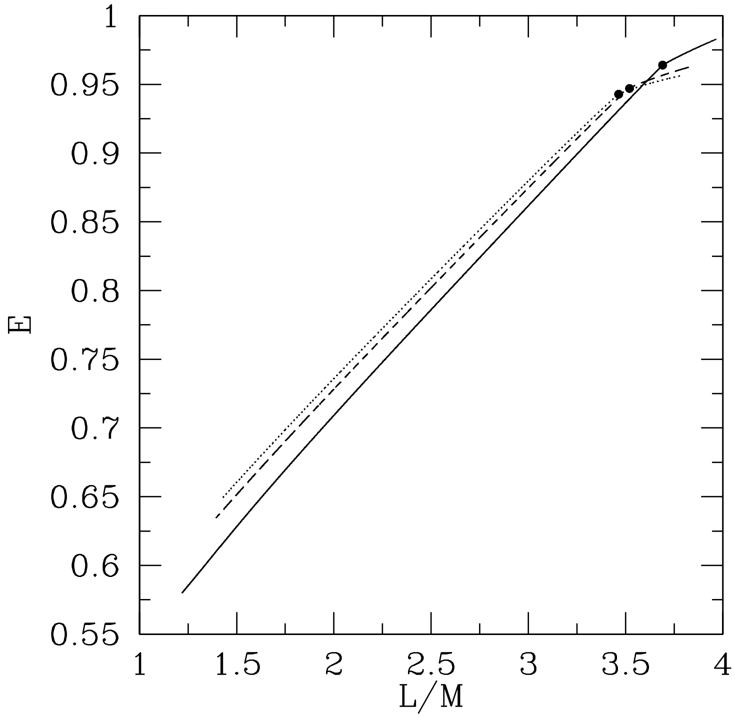

Figure 6.4: The phase-space trajectories $\{L(t), E(t)\}$ for the three EMRIs depicted in Fig. 6.3. As before, the solid, dashed, and dotted curves begin with eccentricities $e_0 = 0.8$, 0.4, and 0.0, respectively. The large black dots correspond to the points at which the COs plunge into the boson star. Note that the curves cross each other, implying that two different geodesics can be characterized by the *same* energy and angular momentum. This feature, unique to the boson-star case, follows from the existence of bound geodesics, both interior and exterior to the local maximum of the effective potential for certain energy and angular momenta. No such crossing appear in Fig. 6.3 as p and e do uniquely specify a given geodesic.

initially eccentric EMRIs appear qualitatively similar. Both the $e_0 = 0.4$ and $e_0 = 0.8$ EMRIs experience a two-step plunge as described qualitatively in §6.3. In the first stage, the pericenter migrates deep into the interior of the boson star. Only several orbits later does the apocenter finally move inward of the local maximum. While two-step plunge may not be immediately apparent from the real-space trajectory, it leaves a very distinctive signature in the gravitational waveform.

First we consider the waveform of the initially circular EMRI as shown in Fig. 6.7. The system is viewed at an inclination angle of 45° at a distance of 10 Mpc. The two

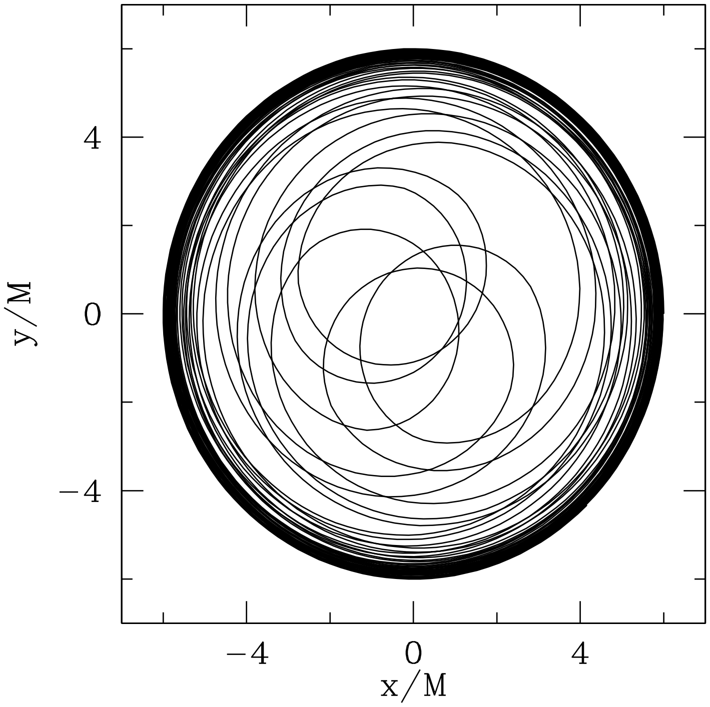

Figure 6.5: The real-space trajectory $\{x(t), y(t)\}$ for an initially circular EMRI. For purposes of clarity, only a period of approximately 150,000 s in the vicinity of the plunge is depicted.

polarizations provide similar information because the given decomposition depends on our arbitrary choice of basis tensors. Before the plunge, the waveform is identical to that of a circular inspiral into a Schwarzschild black hole. The fundamental oscillation has a period of about 500 s, one-half the azimuthal period over which ϕ changes by 2π radians. Since the orbit is circular, the amplitude only varies on the extremely long timescale on which energy and angular momentum are lost. After the plunge, the eccentricity increases sharply and we see an amplitude modulation with a period of about 10,000 s. This corresponds to radial oscillations about the inner minimum within the boson star. The amplitude creeps upwards as the CO approaches pericenter with increasing acceleration. Note, however, the peculiar feature at pericenter itself, where the amplitude should be greatest. Instead, the cycle at pericenter is suppressed because once the CO crosses the boson-star surface, it is only accelerated by the mass

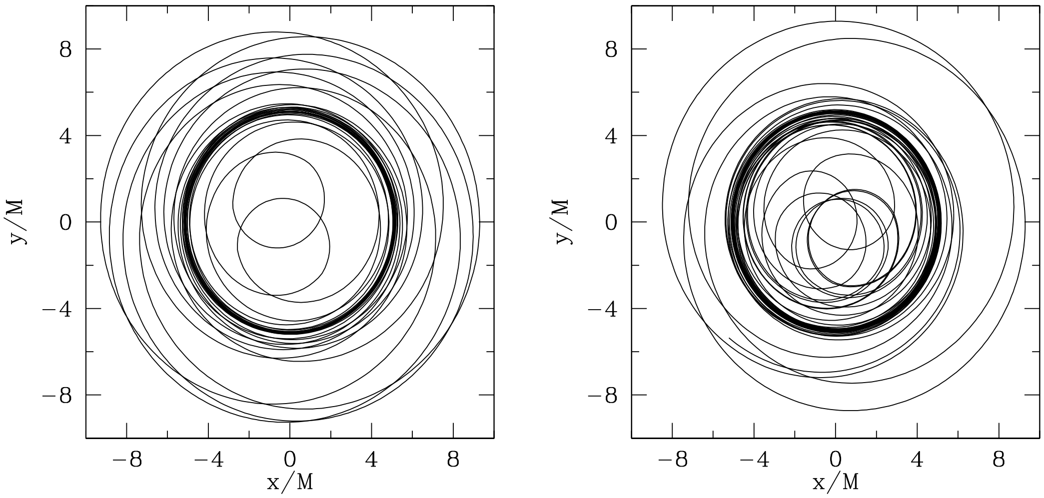

Figure 6.6: The real-space trajectory $\{x(t), y(t)\}$ for the EMRI with $e_0 = 0.4$, $p_0/M = 10.0$. The left panel shows 40,000 s including the first stage of the two-step plunge, while the right panel shows about 50,000 s including the second stage.

interior to its position. We present a close-up of this distinctive feature in the bottom panel of Fig. 6.7.

The waveform of the initially eccentric EMRI as shown in Fig. 6.8 is even more unique. All three phases of the EMRI can be easily distinguished by the naked eye. In the first phase, from the beginning of the waveform until the first-step of the plunge at $t = 1.4 \times 10^4$ s, the CO moves on ordinary eccentric geodesics of Schwarzschild. Amplitude modulation corresponding to oscillations about the outer minimum are seen, the crests near pericenter are smooth because there are as yet no close approaches to the boson star. In the second phase, between the two steps of the plunge, the CO moves on geodesics with *both* minima between pericenter and apocenter. These geodesics are extremely eccentric, leading to longer radial oscillations and an approximate 50% increase in the period of the amplitude modulation. The apocenter remains nearly constant during the plunge, marked by the continued preesnce of deep troughs in the amplitude modulation. The pericenter however has migrated deep into the interior of the boson star as evidenced by the spikes on each crest produced during close approaches. The peculiar feature at pericenter itself is barely visible on this

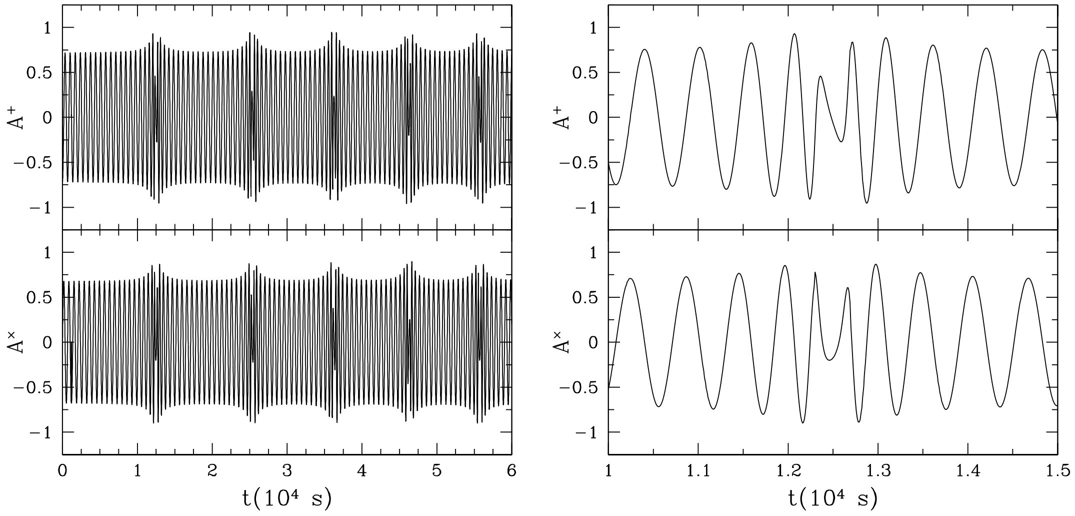

Figure 6.7: The GW amplitudes $\{A^{+}(t), A^{\times}(t)\}$ for the initially circular EMRI depicted in Fig. 6.5. The dimensionless amplitudes are given in units of 10^{-20}, while the time is in units of 10^4 s. The upper panel shows 60,000 s of the inspiral, including the plunge at $t = 0.12 \times 10^4$ s. The lower panel zooms in on the portion of the waveform produced in the interior of the boson star itself.

scale as a very narrow gap in these spikes. The second step of the plunge occurs at $t = 8.0 \times 10^4$ s, when the CO crosses over the local maximum for the final time. The second step is in some sense the opposite of the first, in that the pericenter remains constant while the apocenter plunges inwards. We see that the spikes marking pericenter remain unchanges between the second and third phases, but the deep troughs in the amplitude modulation from distant apocenters have vanished. Close-ups of the two steps of the plunge are shown in Fig. 6.9.

6.7 Discussion

Strong evidence supports the existence of supermassive compact objects at the centers of many, if not most, galaxies. Accretion onto these objects is presumed to power active galactic nuclei, and direct observations of velocity dispersions within galactic cusps reveal that a point-like mass dominates the dynamics within ∼ 1 pc from the center. While all this evidence is consistent with these objects being black holes, the

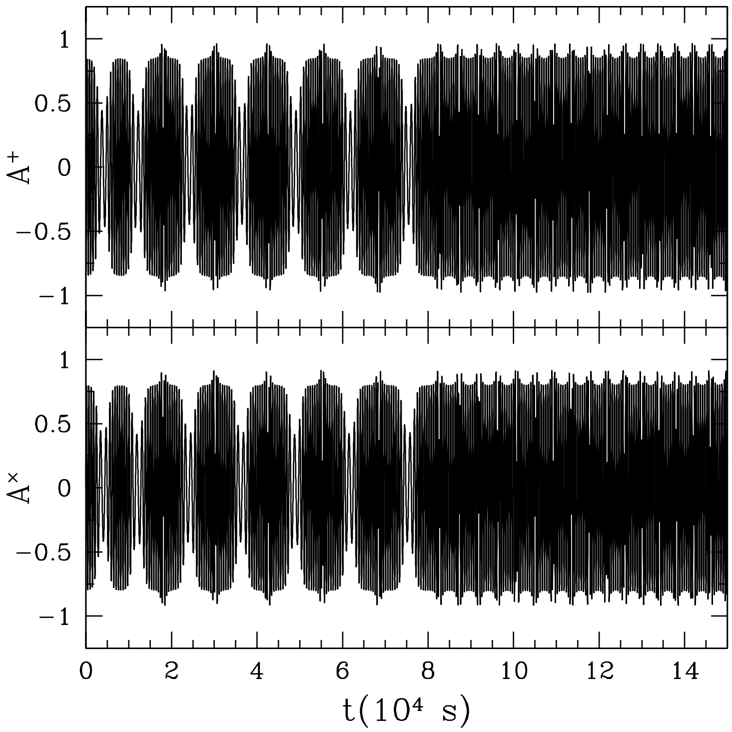

Figure 6.8: The GW amplitudes $\{A^+(t), A^\times(t)\}$ for the initially eccentric EMRI depicted in Fig. 6.6. As in Fig. 6.7, the dimensionless amplitudes are given in units of 10^{-20}, while the time is in units of 10^4 s. All three portions of the inspiral are clearly visible.

"smoking-gun" signature of an event horizon has yet to be observed. Until such a definitive determination is made, other candidates such as boson stars should continue to be considered. Previous work has suggested that accretion-induced X-ray bursts might be an optical signature of boson stars (Yuan et al., 2004). We rely here on distinguishing boson stars from black holes through their gravitational effects. One possibility is their different properties as a strong gravitational lens; a "shadow" induced by a black hole's event horizon may soon be observed in Sag A* at submillimeter wavelengths (Falcke et al., 2000). In this paper, we have investigated the possibility that the inspiral of a several solar-mass CO into a boson star may produce a distinctive spectrum of GWs. LISA is conservatively expected to observe several such inspirals each year with appreciable signal-to-noise ratios (Sigurdsson, 1997; Sigurdsson & Rees, 1997).

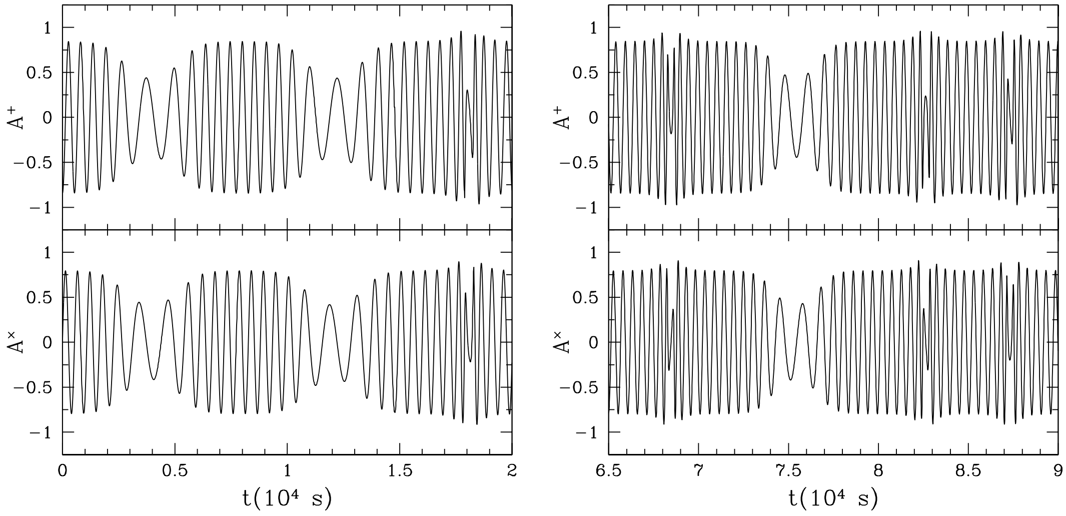

Figure 6.9: Close-ups of the waveform shown in Fig. 6.8 for the initially eccentric EMRI. The top panel depicts the first step of the plunge; note that only the right-most crest exhibits the peculiar feature produced in the boson-star interior. The bottom panel shows the second step of the plunge; note how different the final distant apocenter at $t = 7.5 \times 10^4$ s appears from the apocenter at $t = 8.5 \times 10^4$ s.

Previous work has examined the possibility of measuring the central object's multipole moments, and testing whether or not they satisfy the predictions of the "no-hair" theorem for black holes (Ryan, 1995). This formalism has been applied in particular to boson stars with anomalously large mass-quadrupole moments, where an appreciable effect on the GW-induced loss of energy was discovered (Ryan, 1997a). Here we have gone beyond such an analysis by calculating approximations to the waveforms that might actually be observed by LISA. The spherically symmetric boson stars we consider have Schwarzschild spacetimes outside their surfaces; a multipole analysis of GWs produced while the CO is still outside the boson star would not reveal any violations of the "no-hair" theorem. Solving Einstein's equations exactly for the gravitational radiation produced would be prohibitively expensive computationally. Instead, we rely on a series of analytic approximations: geodesic motion for the CO in the extreme mass-ratio limit, GW propagation unaffected by the boson-star spacetime, and direct indentification of Schwarzschild coordinates with flat-space spherical coordinates. These approximations, the “particle on a string” approach, allow a direct application of multipole-moment expansions for both the energy and angular momentum losses and the waveform itself. We then drop all but the mass quadrupole moment terms in these expansions. These approximations cannot be fully justified theoretically in the highly relativistic, strong-field regime, but direct comparison with rigorous TSN calculations in the black hole case shows much better agreement than one might have expected (Glampedakis et al., 2002; Babak et al., 2005). We then assume that the energy and angular momentum are changing adiabatically on timescales much longer than an orbital period. This approximation – valid in the extreme mass-ratio limit except during the plunge itself – allows us to apply a “hybrid” two-stage approach. In the first stage, differential equations for orbital elements p and e are obtained by relating these quantities to E and L with exact analytic derivatives. These equations are then integrated on a radiation-reaction timescale, with a timestep that is correspondingly long compared to an orbital period. In the second stage, the orbit of the CO, $r(t)$ and $\phi(t)$, is computed with small timesteps by solving the geodesic equations and linearly interpolating the trajectories $\{p(t), e(t)\}$ produced in the first stage. The hybrid approach allows us to calculate real-space trajectories and waveforms much more quickly than would be possible using the same short timestep for both stages.

As anticipated, the waveforms produced by this method exhibit distinctive features that allow them to be readily distinguished from those produced during EMRIs into black holes. In the model considered here, the boson-star inspiral is identical to a black-hole inspiral until the CO falls over the angular-momentum barrier. If LISA observed GWs from part of an inspiral including this plunge, the parameters of the exterior black-hole spacetime could be determined very accurately using black-hole EMRI templates from the part of the inspiral up until plunge. The “smoking gun” for a boson-star inspiral would be that GWs from the inspiral persist after the plunge. This persistence could be seen using, for instance, a time-frequency analysis of the LISA data stream. GWs from an event like this could not be mistaken for an inspiral into a black hole with different parameters, because the early stages of the inspiral are identical to the black-hole inspiral. If only the post-plunge stage of a boson-star inspiral were seen, it is not clear whether this could be mistaken for a black-hole EMRI without a proper Fisher-matrix analysis.

The waveforms produced in this paper are highly approximate, and far too crude to use in any attempts at matched filtering for LISA. Nonetheless, they may serve as substitutes for purposes of scoping out the data analysis and design specification until better waveforms are available. Improved waveforms might result from refining our method, for example, by incorporating higher-order multipole moments in our expansions. Press developed an improved formula for the GW field h_{ij}^{TT} that accounts for time delays in the source, as well as some of the relativistic effects provided by the higher moments (Press, 1977). Efforts are underway to adapt this formula to the "particle on a string" approach (Gair & Fang, 2004; Creighton et al., 2005). Analysis of higher harmonics of the waveform might also prove to be a useful way of distinguishing EMRIs into black holes and boson stars. These higher harmonics become comparable to the fundamental $m = 2$ mode at high eccentricities (Barack & Cutler, 2004). One might expect the sudden increase in eccentricity at plunge to manifest itself in frequency space as a sharp rise in the higher harmonics. Only a rigorous comparison of our waveforms with the LISA noise curve can truly determine whether LISA can differentiate between black-hole and boson-star EMRIs. We hope to conduct such a comparison in the near future using *Synthetic LISA*, an actual simulation of the LISA interferometer that models its response to incident gravitational waves (Vallisneri, 2005).

Going beyond these approximations to produce waveforms suitable for matched-filtering by LISA will involve the use of perturbation theory in the the TSN formalism. The resulting partial differential equations for a generic boson star should be separable in the spherical case, and we plan to consider this problem in the near future. Hopefully, these improvements or others will provide the tools to calculate accurate waveform templates in time for LISA's great attempt at "holiodesy", the mapping of spacetimes about the compact objects in galactic centers. Including exotic waveforms such as those produced by boson stars in our suite of templates will ensure that LISA

Table 6.1: Coefficients describing the contribution of higher-order multipole moments to the energy flux of the inspiraling CO.

l	A_l	B_l
2	6.400	1.067
3	6.458	4.571
2	14.966	1.250

will not miss the opportunity to discover something truly fundamental.

We wish to thank L. Lindblom and E. S. Phinney for useful conversations. Micah Solomon also provided help in preparing the Appendix. Kesden was supported by the NASA Graduate Student Research Program. JG's work was supported by NASA grants NAG5-12384 and NAG5-10707. This work was supported in part by DoE DE-FG03-92-ER40701 and NASA NAG5-9821.

6.8 Beyond the Quadrupole Approximation

A key approximation made in § 6.4 was to drop all but the mass quadrupole terms from the multipole-moment expansions of the GW field, energy flux, and momentum flux in Eqs. (6.17), (6.18), and (6.19). For circular orbits in the Schwarzschild metric, these expansions can be expressed as power series in M/r with coefficients that can be calculated analytically. To illustrate, the contribution of higher-order mass and current multipole moments to the energy flux is given by

$$\frac{dE}{d(t/M)} = \begin{cases} A_l \eta (r/M)^{-(l+3)}, & \text{mass multipoles} \\ B_l \eta (r/M)^{-(l+4)}, & \text{current multipoles} \end{cases} \tag{6.39}$$

where the dimensionless coefficients A_l and B_l are listed in Table 6.1. Neglecting the higher-order terms is clearly justified at large radii, but the expansion parameter $M/r = 1/6$ is not particularly small at the ISO. It is still small enough, however, so that the next largest terms which happen to be the $l = 3$ and $l = 4$ mass multipole moments provide only 16.8% and 6.5% corrections, respectively. We continue to drop the higher-order terms as the CO plunges past the ISO into the interior of the boson star. If the singular behavior of Eq. (6.39) held within the boson star, the expansion would become formally divergent at $r = M$, but fortunately, this is not the case. Outside the boson star, the angular frequency Ω is Keplerian, $M\Omega = (r/M)^{-3/2}$, but in the interior Ω is a monotonically increasing function of r/M. As such, the higher-order terms in the multipole-moment expansions which involve increasingly more time derivatives become increasingly steep functions of r/M. The multipole-moment expansion is therefore perfectly regular at the origin, and the quadrupole approximation again becomes more accurate at small radii. The errors associated with this approximation, while excessive for the purposes of data analysis, are acceptable for an initial search for qualitative features of a boson-star EMRI.

Bibliography

Abbott L. F., Wise M. B., 1984, Nuc. Phys. B, 244, 541

Afshordi N., Loh Y., Strauss M., 2004, Phys. Rev. D, 69, 083524

Albrecht A., Steinhardt P. J., 1982, Phys. Rev. Lett., 48, 1220

Alpher R. A., Bethe H. A., Gamow G., 1948, Phys. Rev., 73, 803

Amblard A., Vale C., White M., 2004, New Astron., 9, 687

Babak S., Fang H., Gair J. R., Glampedakis K., 2005, in preparation

Balbi A., Ade P., Bock J., Borrill J., Boscaleri A., de Bernardis P., Ferreira P. G., Hanany S., Hristov V. V., Jaffe A. H., Lee A. T., Oh S., Pascale E., Rabii B., Richards P. L., Smoot G. F., Stompor R., Winant C. D., Wu J. H. P., 2000, Astrophys. J., 545, L1

Barack L., Cutler C., 2004, Phys. Rev. D, 69, 082005

Bardeen J., Steinhardt P., Turner M., 1983, Phys. Rev. D, 28, 679

Bartelmann M., Schneider P., 1992, Astron. Astrophys., 259, 413

Benabed K., Bernardeau F., van Waerbeke L., 2001, Phys. Rev. D, 63, 043501

Bennett C. L., Banday A., Gorski K. M., Hinshaw G., Jackson P., Keegstra P., Kogut A., Smoot G. F., Wilkinson D. T., Wright E. L., 1996, Astrophys. J. Lett., 464, L1

Bennett C. L., Halpern M., Hinshaw G., Jarosik N., Kogut A., Limon M., Meyer S. S., Page L., Spergel D. N., Tucker G. S., Wollack E., Wright E. L., Barnes C., Greason M. R., Hill R. S., Komatsu E., Nolta M. R., Odegard N., Peirs H. V., Verde L., Weiland J. L., 2003, Astrophys. J. Suppl., 148, 1

Bernardeau F., 1997, Astron. Astrophys., 324, 15

Bernardeau F., 1998, Astron. Astrophys., 338, 767

Blanchard A., Douspis M., Rowan-Robinson M., Sarkar S., 2003, Astron. Astrophys., 412, 35

Blanchard A., Schneider J., 1987, Astron. Astrophys., 184, 1

Blandford R. D., Saust A. B., Brainerd T. G., Villumsen J. V., 1991, Mon. Not. R. Astron. Soc., 251, 600

Bond J. R., Efstathiou G., Tegmark M., 1997, Mon. Not. R. Astron. Soc., 291, L33

Bouchet F. R., Prunet S., Sethi S. K., 1999, Mon. Not. R. Astron. Soc., 302, 663

Boughn S. P., Crittenden R. G., 2003, A correlation of the cosmic microwave sky with large scale structure, astro-ph/0305001

Boughn S. P., Crittenden R. G., Turok N. G., 1998, New Astron., 3, 275

Bridle S. L., Lewis A. M., Weller J., Efstathiou G., 2003, Mon. Not. R. Astron. Soc., 342, L72

Cayon L., Martinez-Gonzalez E., Sanz J., 1993, Astrophys. J., 413, 10

Chandrasekhar S., 1992, The Mathematical Theory of Black Holes. Oxford University Press, New York

Cline J. M., Crotty P., Lesgourgues J., 2003, J. Cos. Astropart. Phys., 0309, 010

Collins N. A., Hughes S. A., 2004, Phys. Rev. D, 69, 124002

Colpi M., Shapiro S. L., Wasserman I., 1986, Phys. Rev. Lett., 57, 2485

Contaldi C. R., Peloso M., Kofman L., Linde A., 2003, J. Cos. Astropart. Phys., 0307, 002

Cooray A., 2001a, Phys. Rev. D, 64, 043516

Cooray A., 2001b, in Kursunoglu B., Perlmutter A., eds, 2001 Coral Gables Conference: Cosmology and Particle Physics After Acoustic Peaks: What's Next in CMB?

Cooray A., 2002, Phys. Rev. D, 65, 063512

Cooray A., Baumann D., 2003, Phys. Rev. D, 67, 063505

Cooray A., Hu W., 2000, Astrophys. J., 534, 533

Cooray A., Hu W., Tegmark M., 2000, Astrophys. J., 540, 1

Cooray A., Kesden M., 2003, New Astron., 8, 231

Cooray A., Sheth R. K., 2002, Phys. Rep., 372, 1

Creighton T., Gair J. R., Hughes S. A., Vallisneri M., 2005, in preparation

Crittenden R., Davis R., Steinhardt P., 1993, Astrophys. J., 417, L13

Crittenden R. G., Turok N., 1996, Phys. Rev. Lett., 76, 575

de Bernardis P., Ade P. A. R., Bock J. J., Bond J. R., Borrill J., Boscaleri A., Coble K., Crill B. P., Gasperis G. D., Farese P. C., Ferreira P. G., Ganga K., Giacometti M., Hivon E., Hristov V. V., Iacoangeli A., *et al.* A. H. J., 2000, Nature, 404, 955

DeDeo S., Caldwell R. R., Steinhardt P. J., 2003, Phys. Rev. D, 67, 103509

Dicke R. H., Peebles P. J. E., Roll P. G., Wilkinson D. T., 1965, Astrophys. J., 142, 414

Djouadi A., Gambino P., Kniehl B. A., 1998, Nucl. Phys. B, 523, 17

Dreyer O., Kelly B., Krishnan B., Finn L. S., Garrison D., Lopez-Aleman R., 2004, Class. Quantum Grav., 21, 787

Efstathiou G., 2003a, Mon. Not. R. Astron. Soc., 346, L26

Efstathiou G., 2003b, Mon. Not. R. Astron. Soc., 343, L95

Einstein A., 1918, Sb. Preuss. Akad. Wiss., p. 154

Eisenstein D. J., Hu W., Tegmark M., 1999, Astrophys. J., 518, 2

Falcke H., Melia F., Agol E., 2000, Astrophys. J., 528, 13

Fang H., Lovelace G., 2005, Tidal coupling of a Schwarzschild black hole and circularly orbiting moon, in preparation

Feng B., Zhang X., 2003, Phys. Lett. B, 570, 145

Finn L. S., Thorne K. S., 2000, Phys. Rev. D, 62, 124021

Fosalba P., Gaztanaga E., 2004, Mon. Not. R. Astron. Soc, 350, L37

Friedberg R., Lee T. D., Pang Y., 1987a, Phys. Rev. D, 35, 3640

Friedberg R., Lee T. D., Pang Y., 1987b, Phys. Rev. D, 35, 3658

Gair J. R., Fang H., 2004, Using the Press formula to generate numerical kludged waveforms, unpublished

Gair J. R., Glampedakis K., 2005, in preparation

Gaztanaga E., Wagg J., Multamaki T., Montana A., Hughes D. H., 2003, Mon. Not. R. Astron. Soc., 346, 47

Geroch R., 1970, J. Math. Phys., 11, 2580

Giardino G., Banday A. J., Grski K. M., Bennett K., Jonas J. L., Tauber J., 2002, Astron. Astrophys., 387, 82

Glampedakis K., Hughes S. A., Kennefick D., 2002, Phys. Rev. D, 66, 064005

Glampedakis K., Kennefick D., 2002, Phys. Rev. D, 66, 044002

Goldberg D. M., Spergel D. N., 1999, Phys. Rev. D, 59, 103002

Gunn J. E., 1967, Astrophys. J., 150, 737

Gurzadyan V. G., .Ade P. A. R., de Bernardis P., Bianco C. L., Bock J. J., Boscaleri A., Crill B. P., Troia G. D., Hivon E., Hristov V. V., Kashin A. L., Lange A. E., Masi S., Mauskopf P. D., Montroy T., Natoli P., *et al.* C. B. N., 2005, Mod. Phys. Lett. A, 20, 491

Guth A. H., 1981, Phys. Rev. D, 23, 347

Guth A. H., Pi S.-Y., 1982, Phys. Rev. Lett., 49, 1110

Halverson N. W., Leitch E. M., Pryke C., Kovac J., Carlstrom J. E., Holzapfel W. L., Dragovan M., Cartwright J. K., Mason B. S., Padin S., Pearson T. J., Readhead A. C. S., Shepherd M. C., 2002, Astrophys. J., 568, 38

Hanany S., Ade P., Balbi A., Bock J., Borrill J., Boscaleri A., de Bernardis P., Ferreira P. G., Hristov V. V., Jaffe A. H., Lange A. E., Lee A. T., Mauskopf P. D., Netterfield C. B., Oh S., Pascale E., Rabii B., Richards P. L., Smoot G. F., Stompor R., Winant C. D., Wu J. H. P., 2000, Astrophys. J. Lett., 545, L5

Hansen R. O., 1974, J. Math. Phys., 15, 46

Hawking S. W., 1982, Phys. Lett. B, 115, 295

Hinshaw G., Banday A. J., Bennett C. L., Górski K. M., Kogut A., 1995, Astrophys. J., 446, L67

Hinshaw G., Banday A. J., Bennett C. L., Górski K. M., Kogut A., Smoot G. F., Wright E. L., 1995, Astrophys. J., 464, L17

Hirata C., Seljak U., 2003a, Phys. Rev. D, 67, 043001

Hirata C., Seljak U., 2003b, Phys. Rev. D, 68, 083002

Hu W., 2000a, Phys. Rev. D, 62, 043007

Hu W., 2000b, Astrophys. J., 529, 12

Hu W., 2001a, Astrophys. J., 557, L79

Hu W., 2001b, Phys. Rev. D, 64, 083005

Hu W., Dodelson S., 2002, Ann. Rev. Astron. Astrophys., 40, 171

Hu W., Okamoto T., 2002, Astrophys. J., 574, 566

Hughes S. A., 2000, Phys. Rev. D, 61, 084004

Hughes S. A., 2001, Phys. Rev. D, 64, 064004

Hui L., Gaztanaga E., 1999, Astrophys. J., 519, 622

Jaffe A., Kamionkowski M., Wang L., 2000, Phys. Rev. D, 61, 083501

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

Juszkiewicz R., Weinberg D. H., Amsterdamski P., Chodorowski M., Bouchet F., 1995, Astrophys. J., 442, 39

Kaiser N., 1992, Astrophys. J., 388, 272

Kamionkowski M., 1996, Phys. Rev. D, 54, 4169

Kamionkowski M., Kosowsky A., 1998, Phys. Rev. D, 57, 685

Kamionkowski M., Kosowsky A., 1999, Ann. Rev. Nucl. Part. Sci., 49, 77

Kamionkowski M., Kosowsky A., Stebbins A., 1997a, Phys. Rev. Lett., 78, 2058

Kamionkowski M., Kosowsky A., Stebbins A., 1997b, Phys. Rev. D, 55, 7368

Kamionkowski M., Loeb A., 1997, Phys. Rev. D, 56, 4511

Kamionkowski M., Spergel D. N., 1994, Astrophys. J., 432, 7

Kashlinsky A., 1988, Astrophys. J., 331, L1

Kawasaki M., Takahashi F., 2003, Phys. Lett. B, 570, 151

Kendall M. G., Stuart A., 1969, The Advanced Theory of Statistics, Vol. II. Griffin, London

Kerr R. P., 1963, Phys. Rev. Lett., 11, 237

Kesden M., Cooray A., Kamionkowski M., 2002a, Phys. Rev. D, 66, 083007

Kesden M., Cooray A., Kamionkowski M., 2002b, Phys. Rev. Lett., 89, 011304

Kesden M., Cooray A., Kamionkowski M., 2003, Phys. Rev. D, 67, 123507

Kesden M., Gair J., Kamionkowski M., 2005, Phys. Rev. D, 71, 044015

Kesden M., Kamionkowski M., Cooray A., 2003, Phys. Rev. Lett., 91, 221302

Kinkhabwala A., Kamionkowski M., 1999, Phys. Rev. Lett., 82, 4172

Knox L., 1995, Phys. Rev. D, 52, 4307

Knox L., Song Y.-S., 2002, Phys. Rev. Lett., 89, 011303

Kormendy J., Gebhardt K., 2002, in Martel H., Wheeler J. C., eds, Proceedings of the 20th Texas Symposium on Relativistic Astrophysics Supermassive Black Holes in Nuclei of Galaxies. AIP

Kovac J. M., Leitch E. M., Pryke C., Carlstrom J. E., Halverston N. W., Holzapfel W. L., 2002, Nature, 420, 772

Kristian J., Sachs R. K., 1966, Astrophys. J., 143, 379

Lange A. E., Ade P. A., Bock J. J., Bond J. R., Borrill J., Boscaleri A., Coble K., Crill B. P., de Bernardis P., Farese P. C., Ferreira P. G., Ganga K., Giacometti M., Hivon E., Hristov V. V., Iacoangeli A., *et al.* A. H. J., 2001, Phys. Rev. D, 63, 042001

Lee T. D., Pang Y., 1992, Phys. Rep., 221, 251

Lewis A., Challinor A., Turok N., 2002, Phys. Rev. D, 65, 023505

Liddle A. R., Lyth D. H., 2000, Cosmological Inflation and Large-Scale Structure. Cambridge University Press, New York

Linde A. D., 1982a, Phys. Lett. B, 108, 389

Linde A. D., 1982b, Phys. Lett. B, 116, 335

Linder E. V., 1988, Astron. Astrophys., 206, 199

Lynds R., Petrosian V., 1989, Astrophys. J., 336, 1

Maoz E., 1995, Astrophys. J. Lett., 447, L91

Mason B. S., Pearson T. J., Readhead A. C. S., Shepherd M. C., Sievers J. L., Udomprasert P. S., Cartwright J. K., Farmer A. J., Padin S., Myers S. T., Bond J. R., Contaldi C. R., Pen U. L., Prunet S., Pogosyan D., Carlstrom J. E., Kovac J., Leitch E. M., Pryke C., Halverson N. W., *et al.* W. L. H., 2003, Astrophys. J., 591, 540

Miller A. D., Caldwell R., Devlin M. J., Dorwart W. B., Herbig T., Nolta M. R., Page L. A., Puchalla J., Torbet E., Tran H. T., 1999, Astrophys. J. Lett., 524, L1

Miralda-Escudé J., 1991, Astrophys. J., 380, 1

Misner C. W., Thorne K. S., Wheeler J. A., 1973, Gravitation. W. H. Freeman and Company, New York

Multamaki T., Manera M., Gaztanaga E., 2004, Phys. Rev. D, 69, 023004

Murray C. D., McDermott S. F., 1999, Solar System Dynamics. Cambridge University Press, New York

Newman E., Penrose R., 1962, J. Math. Phys., 3, 566

Nolta M. R., Wright E. L., Page L., Bennett C. L., Halpern M., Hinshaw G., Jarosik N., Kogut A., Limon M., Meyer S. S., Spergel D. N., Tucker G. S., Wollack E., 2004, Astrophys. J., 608, 10

Okamoto T., Hu W., 2003, Phys. Rev. D, 67, 083002

Peebles P. J. E., Yu J. T., 1970, Astrophys. J., 162, 815

Peiris H. V., Spergel D. N., 2000, Astrophys. J., 540, 605

Penzias A., Wilson R., 1965, Astrophys. J., 142, 419

Peters P. C., 1964, Phys. Rev., 136, B 1224

Peters P. C., Mathews J., 1963, Phys. Rev., 131, 435

Poisson E., 1993, Phys. Rev. D, 47, 1497

Poisson E., 2004, Living Rev. Relativity, 7, 6

Press W. H., 1977, Phys. Rev. D, 15, 965

Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, Numerical Recipes in C: The Art of Scientific Computing. Cambridge University Press, New York

Prunet S., Sethi S. K., Bouchet F. R., 2000, Mon. Not. R. Astron. Soc., 314, 348

Regge T., Wheeler J. A., 1957, Phys. Rev., 108, 1063

Richstone D., Ajhar E. A., Bender R., Bower G., Dressler A., Faber S. M., Filippenko A. V., Gebhardt K., Green R., Ho L. C., Kormendy J., Lauer T. R., Magorrian J., Tremaine S., 1998, Nature, 395, A14

Ryan F. D., 1995, Phys. Rev. D, 52, 5707

Ryan F. D., 1997a, Phys. Rev. D, 55, 6081

Ryan F. D., 1997b, Phys. Rev. D, 56, 1845

Rybicki G. B., Lightman A. P., 1979, Radiative Processes in Astrophysics. John Wiley & Sons, New York

Sachs R. K., 1961, Proc. R. Soc. A, 264, 309

Sachs R. K., Wolfe A. M., 1967, Astrophys. J., 147, 73

Sasaki M., Nakamura T., 1982, Prog. Theor. Phys., 67, 1788

Schunck F. E., Liddle A. R., 1997, Phys. Lett. B, 404, 25

Schwarz D., Starkman G., Huterer D., Copi C., 2004, Phys. Rev. Lett., 93, 221301

Scranton R., Connolly A. J., Nichol R. C., Stebbins A., Szapudi I., Eisenstein D. J., Afshordi N., Budavari T., Csabai I., Frieman J. A., Gunn J. E., Johnson D., Loh Y., Lupton R. H., Miller C. J., Sheldon E. S., Sheth R. S., Szalay A. S., Tegmark M., *et al.* Y. X., 2003, Physical Evidence for Dark Energy, astro-ph/0307335

Seljak U., 1996, Astrophys. J., 463, 1

Seljak U., 1997, Astrophys. J., 482, 6

Seljak U., Zaldarriaga M., 1997, Phys. Rev. Lett., 78, 2054

Seljak U., Zaldarriaga M., 1999a, Phys. Rev. Lett., 82, 2636

Seljak U., Zaldarriaga M., 1999b, Phys. Rev. D, 60, 043504

Sievers J. L., Pearson T. J., Readhead A. C. S., Shepherd M. C., Udomprasert P. S., Cartwright J. K., Mason B. S., Padin S., Myers S. T., Bond J. R., Contaldi C., Prunet S., Pogosyan D., 2001, in Bulletin of the American Astronomical Society Vol. 33, The CMB Power Spectrum from CBI Mosaics. p. 1357

Sigurdsson S., 1997, Class. Quantum Grav., 14, 1425

Sigurdsson S., Rees M. J., 1997, Mon. Not. R. Astron. Soc., 284, 318

Silk J., 1968, Astrophys. J., 151, 459

Smoot G. F., Bennett C. L., Kogut A., Wright E. L., Aymon J., Boggess N. W., Cheng E. S., de Amici G., Gulkis S., Hauser M. G., Hinshaw G., Jackson P. D., Janssen M., Kaita E., Kelsall T., Keegstra P., Lineweaver C., Loewenstein K., *et al.* P. L., 1992, Astrophys. J., 396, L1

Spergel D. N., Goldberg D. M., 1999, Phys. Rev. D, 59, 103001

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

Starobinskii A. A., 1985, Soviet Astron. Lett., 11, 133

Starobinsky A. A., 1982, Phys. Lett. B, 117, 175

Sunyaev R. A., Zel'dovich Y. B., 1970, Astrophys. Space Sci., 7, 3

Sunyaev R. A., Zel'dovich Y. B., 1980, Mon. Not. R. Astron. Soc., 190, 413

Tegmark M., de Oliveira-Costa A., Hamilton A. J. S., 2003, Phys. Rev. D, 68, 123523

Teukolsky S. A., 1973, Astrophys. J., 185, 635

Teukolsky S. A., Press W. H., 1974, Astrophys. J., 193, 443

Thorne K. S., 1980, Rev. Mod. Phys., 52, 299

Torres D. F., Capozziello S., Lambiase G., 2000, Phys. Rev. D, 62, 104012

Tucci M., Martinez-Gonzalez E., Vielva P., Delabrouille J., 2004, Limits on the detectability of the CMB B-mode polarization imposed by foregrounds, astro-ph/0411567

Turner M. S., 1991, Phys. Rev. D, 44, 3737

Tyson J. A., Valdes F., Wenk R., 1990, Astrophys. J., 349, L1

Vallisneri M., 2005, Phys. Rev. D, 71, 022001

Villumsen J. V., 2001, in Brainerd T. G., Kochanek C. S., eds, Gravitational Lensing: Recent Progress and Future Goals Vol. 237, Non-Gaussian Signatures of Gravitational Lensing in the Microwave Background. p. 349

Wald R. M., 1984, General Relativity. University of Chicago Press, Chicago

Weinberg N., Milosavljevic M., Ghez A., 2004, Stellar Dynamics at the Galactic Center with a Thirty Meter Telescope, astro-ph/0404407

Yuan Y., Narayan R., Rees M. J., 2004, Astrophys. J., 606, 1112

Zaldarriaga M., 1997, Phys. Rev. D, 55, 1822

Zaldarriaga M., 2000, Phys. Rev. D, 62, 063510

Zaldarriaga M., Seljak U., 1998, Phys. Rev. D, 58, 023003

Zaldarriaga M., Seljak U., 1999, Phys. Rev. D, 59, 123507

Zaldarriaga M., Spergel D. N., Seljak U., 1997, Astrophys. J., 488, 1

Zerilli F. J., 1970, Phys. Rev. D, 2, 2141