Cosmological Parameters from Eigenmode Analysis of Sloan Digital Sky Survey Galaxy Redshifts

Alexander Szalay,[1] Adrian Pope,[1] Takahiko Matsubara,[2] Michael R. Blanton,[3] Daniel J. Eisenstein,[4] Jim Gray,[5] and Bhuvnesh Jain[6] for the SDSS Collaboration

[1] *Dept. of Physics and Astronomy, The Johns Hopkins University, Baltimore, MD 21218, USA*
[2] *Dept. of Physics and Astrophysics, Nagoya University, Chikusa, Nagoya 464-8602, Japan*
[3] *Center for Cosmology and Particle Physics, Dept. of Physics, New York University, New York, NY 10003, USA*
[4] *Steward Observatory, University of Arizona, Tucson, AZ 85721, USA*
[5] *Microsoft Bay Area Research Center, San Francisco, CA 94105, USA*
[6] *Dept. of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104, USA*

Abstract. We present estimates of cosmological parameters from the application of the Karhunen-Loève transform to the analysis of the 3D power spectrum of density fluctuations using Sloan Digital Sky Survey galaxy redshifts. We use $\Omega_m h$ and $f_b = \Omega_b/\Omega_m$ to describe the shape of the power spectrum, σ_{8g}^L for the (linearly extrapolated) normalization, and β to parametrize linear theory redshift space distortions. On scales $k \leq 0.16 h \mathrm{Mpc}^{-1}$, our maximum likelihood values are $\Omega_m h = 0.264 \pm 0.043$, $f_b = 0.286 \pm 0.065$, $\sigma_{8g}^L = 0.966 \pm 0.048$, and $\beta = 0.45 \pm 0.12$. When we take a prior on Ω_b from WMAP, we find $\Omega_m h = 0.207 \pm 0.030$, which is in excellent agreement with WMAP and 2dF. This indicates that we have reasonably measured the gross shape of the power spectrum but we have difficulty breaking the degeneracy between $\Omega_m h$ and f_b because the baryon oscillations are not resolved in the current spectroscopic survey window function.

1. Introduction

Redshift surveys are an extremely useful tool to study the large scale distribution of galaxies. Of the many possible statistical estimators, the power spectrum of the density fluctuations has emerged as one of the easiest to connect to theories of structure formation in the Universe, especially in the limit of Gaussian fluctuations where the power spectrum is the complete statistical description. There are several ways to measure the power spectrum (for a comparison of techniques see Tegmark et al. 1998). Over the last few years, the Karhunen-Loève method (Vogeley & Szalay 1996; hereafter VS96) has been recognized as the optimal way to build an orthogonal basis set for likelihood analysis, even if the underlying survey has a very irregular footprint on the sky. A variant of the same technique is used for the analysis of CMB fluctuations (Bond et al. 2000).

The shape of the power spectrum is well described by a small set of parameters (Eisenstein & Hu 1998). For redshift surveys, it is of particular importance to consider the large-scale anisotropies caused by infall (Kaiser 1987). Using a forward technique that compares models directly to the data, like the KL-transform, enables us to easily consider these anisotropies in full detail. Here we present results of a parametric analysis of the shape of the fluctuation spectrum for the SDSS galaxy catalog.

2.   Data

2.1.   Sloan Digital Sky Survey

The Sloan Digital Sky Survey (SDSS; York et al. 2000, Stoughton et al. 2002) plans to map nearly one quarter of the sky using a dedicated 2.5 meter telescope at Apache Point Observatory in New Mexico. A drift-scanning CCD camera (Gunn et al. 1998) is used to image the sky with a custom set of 5 filters ($ugriz$; Fukugita et al. 1996 Smith et al. 2002) to a limiting Petrosian (1976) magnitude of $m_r \sim 22.5$. Observations are calibrated using a 0.5 meter photometric telescope (Hogg et al. 2001). After a stripe of sky has been imaged, reduced, and astrometrically calibrated (Pier et al. 2003), additional automated software selects potential targets for spectroscopy. These targets are assigned to 3° diameter (possibly overlapping) circles on the sky called tiles (Blanton et al. 2003a). Aluminum plates drilled from the tile patterns hold optical fibers that feed into the SDSS spectrographs (Uomoto et al. 1999). The SDSS Main Galaxy Sample (MGS; Strauss et al. 2002) will consist of spectra of nearly one million low redshift ($\langle z \rangle \sim 0.1$) galaxies creating a three dimensional map of local large scale structure.

2.2.   Large Scale Structure Sample

Considerable effort has been invested in preparing SDSS MGS redshift data for large scale structure studies. The first task is to correct for fiber collisions. The minimum separation between optical fibers is 55″ which causes a correlated loss of redshifts in areas covered by a single plate. Galaxy targets that were not observed due to collisions are assigned the redshift of their nearest neighbor. Next the sky is divided into unique regions of overlapping spectroscopic plates called sectors. The angular completeness is calculated for each sector as if the collided galaxies had been successfully measured. Galaxy magnitudes are extinction-corrected with the (Schlegel et al. 1998) dust maps, then k-corrections are applied and rest frame colors and luminosities are calculated (Blanton et al. 2003b). Subsamples are created by making appropriate cuts in luminosity, color, and/or flux. A luminosity function is then calculated for each subsample (Blanton et al. 2003c) and used to create a radial selection function assuming $\Omega_m = 0.3$ and $\Omega_\Lambda = 0.7$ cosmology.

This analysis considers two samples of SDSS data, which we will label sample 10 and sample 12. Both samples were prepared in similar manners, although using different versions of software. Sample 12 represents a later state of the survey and the sample 10 area is contained in sample 12. Sample 10 represents 1983.39 completeness-weighted square degrees of spectroscopically ob-

served SDSS data and 165,812 MGS redshifts. Sample 12 has 205,484 redshifts over 2406.74 square degrees. Both samples are larger than the 1360 square degrees of spectroscopy in data release 1 (DR1; Abazajian et al. 2003) of the SDSS. The geometry of the samples and DR1 are qualitatively similar, consisting of two thick slices in the northern cap of the survey and three thin stripes in the south. The samples used have a luminosity cut of $-19 \geq M_r \geq -22$, where $h = 1.0$ and $M_* = -20.44$ (Blanton et al. 2003c). Rest frame quantities (ie absolute magnitudes) are given for the SDSS filters at z=0.1, the median depth of the MGS. In a study of the two point correlation function of SDSS galaxy redshifts, Zehavi et al. (2002) found that the bias relative to M_* galaxies varies from 0.8 for galaxies with $M = M_* + 1.5$ to 1.2 for galaxies with $M = M_* - 1.5$. Norberg et al. (2001) found similar results for the 2dF, with the trend becoming more pronounced at luminosities significantly greater than L_*. The dependence of clustering strength on luminosity could induce an extra tilt in the power spectrum because more luminous galaxies contribute more at large scales and less luminous galaxies contribute more at small scales due to the number of available baselines. We minimize this effect by staying within $M = M_* \pm \sim 1.5$. A uniform flux limit of $m_r \leq 17.5$ was applied, leaving 110,345 redshifts for sample 10 and 134,141 for sample 12. Although there are luminosity limits for this sample, it is essentially a flux-limited sample with a (slowly) varying selection function. We used galaxies in the redshift range $0.05 \leq z \leq 0.17$.

3. Algorithm

3.1. The Karhunen-Loève Eigenbasis

Following the strategy described in VS96, the first step in a Karhunen-Loève (KL) eigenmode analysis of a redshift survey is to divide the survey volume into cells and use the vector of galaxy counts within the cells as our data. This allows a large compression in the size of the dataset without a loss of information on large scales. Our data vector of fluctuations $\mathbf{d}$ is defined as

$$d_i = c_i/n_i - 1 \tag{1}$$

where c_i is the observed number of galaxies in the i^{th} cell and $n_i = \langle c_i \rangle$ is its expected value, calculated from the angular completeness and radial selection function. The data is "whitened" by the factor $1/n_i$ to control shot noise properties in the transform (VS96). We call this the "overdensity" convention.

The KL modes are the solutions to the eigenvalue problem $\mathbf{R}\boldsymbol{\Psi}_n = \lambda_n \boldsymbol{\Psi}_n$ with the correlation matrix of the data given by

$$R_{ij} = \langle d_i d_j \rangle = \xi_{ij} + \delta_{ij}/n_i + \eta_{ij}/(n_i n_j) \tag{2}$$

where ξ_{ij} is the cell-averaged correlation matrix, δ_{ij}/n_i is the shot noise term, and $\eta_{ij}/(n_i n_j)$ can be used to account for correlated noise (not used in this analysis). The most obvious source of correlated noise in the MGS would be differences in photometric zero points between different SDSS imaging runs, which would result in "zebra stripe" patterns of density fluctuations. The MGS selection has a magnitude limit, but no color selection terms, so the variation in target density depends only linearly on the photometric calibration. The r

band zero point variation is 0.02 mag rms (Abazajian et al. 2003), indicating that the density variation should be $\leq 2\%$. The transformed data vector $\mathbf{B}$ is the expansion of $\mathbf{d}$ over the KL modes $\mathbf{\Psi}_n$:

$$\mathbf{d} = \sum_n B_n \mathbf{\Psi}_n. \tag{3}$$

The KL basis is defined by two properties: orthonormality of the basis vectors, $\mathbf{\Psi}_m \cdot \mathbf{\Psi}_n = \delta_{mn}$, and statistical orthogonality of the transformed data, $\langle B_m B_n \rangle = \langle B_n^2 \rangle \delta_{mn}$.

3.2.   The Correlation Function in Redshift Space

In order to directly compare cosmological models to our redshift data using a two point statistic we must calculate the redshift space correlation function $\xi^{(s)}(\mathbf{r}_i, \mathbf{r}_j)$, where $\mathbf{r}_i$ and $\mathbf{r}_j$ describe positions in the observable angles and redshift. The infall onto large scale structures affects the velocities of galaxies leading to an anisotropy in redshift space for a power spectrum that is isotropic in real space (Kaiser 1987). Szalay et al. (1998) derived an expansion of the correlation function that accounts for this anisotropy in linear theory for arbitrary angles. The expansion is

$$\xi^{(s)}(\mathbf{r}_i, \mathbf{r}_j) = c_{00}\xi_0^{(0)} + c_{02}\xi_2^{(0)} + c_{04}\xi_4^{(0)} + ..., \tag{4}$$

$$\xi_L^{(n)}(r) = \frac{1}{2\pi^2} \int dk k^2 k^{-n} j_L(kr) P(k) \tag{5}$$

where the c_{nL} coefficients are polynomials of β and functions of the relative geometry of the two points. The quantity β relates infall velocity to matter density and is well approximated by the fitting formula $\beta = \Omega_m^{0.6}/b$ where b is the bias parameter. Further terms in Eq. (4) are negligible as long as $2 + \partial \ln \phi(r)/\partial \ln r$ (where r is the distance to the cell and $\phi(r)$ is the radial selection function) does not significantly differ (i.e. orders of magnitude) from unity. For the redshift range considered in this analysis $|2 + \partial \ln \phi(r)/\partial \ln r| \leq 4$. When using counts-in-cells, we must calculate the cell-averaged correlation matrix

$$\xi_{ij} = \int d^3\mathbf{r}_1 \int d^3\mathbf{r}_2 \xi^{(s)}(\mathbf{r}_1, \mathbf{r}_2) W_i(\mathbf{x}_i - \mathbf{r}_1) W_j(\mathbf{x}_j - \mathbf{r}_2) \tag{6}$$

where $W_i(\mathbf{y})$ is the cell window function and $\mathbf{x}_i$ is the position of the i^{th} cell. To be precise, $W_i(\mathbf{y})$ should describe the shape of the cell in redshift space. Numerical calculation of this multi-dimensional integral can be computationally expensive. However, for the case of spherically symmetric cells we can change the order of integration and perform the redshift space integrals in Eq. (6) analytically before the k-space integral in Eq. (5). If both cells have the same window function, we can use Eq. (4) as our cell-averaged correlation function (with $\mathbf{r}_i$ and $\mathbf{r}_j$ indicating the cell positions) if we replace $P(k)$ with $P(k)\tilde{W}^2(k)$ in Eq. (5) where $\tilde{W}(k)$ is the Fourier transform of the cell window function. This results in a one dimensional numerical integral. The full technical details of our method will be presented in Matsubara et al. (2004).

We used hard spheres as our cell shape and placed them in a hexagonal closest packed (the most efficient 3D packing, with a 74% space-filling factor) arrangement. The current slice-like survey geometry and packing arrangement causes some spheres to partially protrude outside the survey. The effective fraction of the sphere that is sampled is also affected by the angular completeness of our survey (which averages $\sim 97\%$). We calculate our expected counts as if the sphere was entirely filled and multiply the observed galaxy counts by $1/f_i$ where f_i is the fraction of the ith sphere's volume that was effectively sampled. This sparser sampling also increases the shot noise by a factor of $1/f_i$. Cells with $f_i < 0.65$ were rejected as too incomplete. We found that a $6h^{-1}\mathrm{Mpc}$ sphere radius allowed us to fill the survey volume with a computationally feasible number of cells without the spheres protruding too much out of the survey, while smoothing on sufficiently small length scales so that we do not lose information in the linear regime ($2\pi/k \geq 40h^{-1}$ Mpc). We used 14,194 cells for sample 10 and 16,924 for sample 12.

The calculation of the sampling fraction for each cell is difficult due to the complicated shapes of the sectors (§2.2.). We created a high resolution angular completeness map in a SQL Server database using 10^7 random angular points over the entire sky. Each point was assigned a completeness weighting by finding which sector contained the point or setting the completeness to zero for points outside the survey area. We used a Hierarchical Triangular Mesh (HTM; Kunszt et al. 2001) spatial indexing scheme to find all points in the completeness map that pierce a cell and calculate the volume weighted completeness for that cell.

3.3. Eigenmode Selection

The KL transform is linear, so there is no loss of information if we use all of the eigenmodes. However, if we perform a truncated expansion we can use the KL transform for compression and filtering. The difference between the original data vector and a truncated reconstruction, $\hat{\mathbf{d}} = \sum_{i=1}^{M<N} B_i \boldsymbol{\Psi}_i$, where we use only M out of a possible N modes can be related to the eigenvalues of the excluded modes by $(\mathbf{d} - \hat{\mathbf{d}})^2 = \sum_{i=M+1}^{N} \lambda_i$. The error is minimized (in a squared sense) when we retain modes with larger eigenvalues and drop modes with smaller eigenvalues, which is sometimes called optimal subspace filtering (Therrien 1992).

The eigenvalue of a KL mode is also related to the range in k-space sampled by that mode. Our models assume that linear theory is a good approximation, which is only valid on larger scales. Consequently we only wish to use KL modes that fall inside a "Fermi sphere" whose radius is set by our cutoff wavenumber k_f. If we sort modes by decreasing eigenvalue, they will densely pack k-space starting from the origin. The modes resist overlapping in k-space due to orthogonality. The shape of a KL mode in k-space resembles the Fourier transform of the survey window function. This means that the number of KL modes within the "Fermi sphere" depends mostly on the survey window function and does not drastically change if we change the size of our cells, as long as we have significantly more cells than modes (which means that our cells must be smaller than the cutoff wavelength). In a fully three dimensional survey the modes would fill k-space roughly spherically and $M \propto k_f^3$. However, the current SDSS geometry resembles several two dimensional slices, resulting in KL modes that resemble cigars in k-

space. These modes pack layer-by-layer into spherical shells whose diameters are integer multiples of the long axis of the mode. See Figure 5 in (Szalay et al. 2003) for a visualization. This results in a scaling more like $M \propto k_f^2$.

In choosing the number of KL modes to use in our analysis we try to keep as many modes as possible for better constraints on our parameter values while requiring that our modes are consistent with linear theory. We have developed a convenient method for determing the range in k-space probed by each KL mode. We separate the integral in Eq. (5) into bandpowers in k. This allows us to determine how strongly each mode couples to each bandpower, which shows a coarse picture of the spherically averaged position of the mode in k-space. Figure 1 shows a grayscale image of how the modes couple to the bandpowers. Once we choose a value for the cutoff wavenumber k_f, we truncate our expansion at the mode where wavenumbers larger than k_f start to dominate.
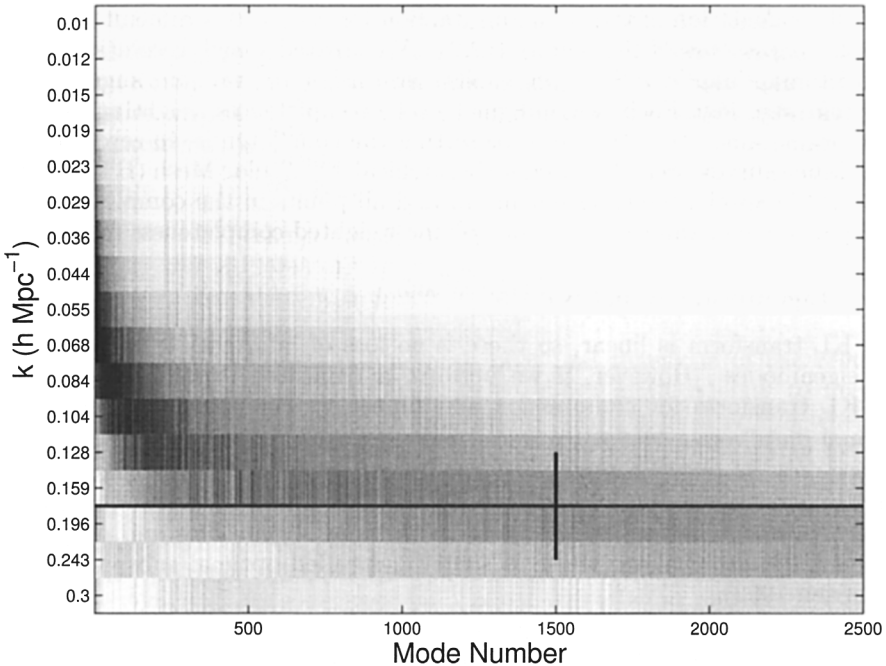

Figure 1.   Grayscale image of wave number vs. mode number. The horizontal line indicates $k_f = 0.16h$ Mpc^{-1}. The vertical line indicates the truncated number of modes used for likelihood analysis.

We can use the statistical properties of the transformed data to check that we are avoiding non-linearities. A rescaled version of the KL coefficients $b_n = B_n/\sqrt{\Psi_n}$ should be normally distributed. Non-linear effects would cause skewness and kurtosis in the distribution of b_n. We do not see evidence of non-linear effects when we use $k_f \leq 0.16h$ Mpc^{-1} (corresponding to length scales $2\pi/k_f \geq 40h^{-1}$ Mpc). This value for the cutoff wavenumber leaves us with 1500 modes for sample 10 and 1850 modes for sample 12.

3.4. Model Testing

We estimate cosmological parameters by performing maximum likelihood analysis in KL space. The likelihood of the observed data given a model m is

$$\mathcal{L}(\mathbf{B}|m) = (2\pi)^{-M/2}|\mathbf{C}_m|^{-1/2} \exp\left[-\frac{1}{2}\mathbf{B}^T\mathbf{C}_m^{-1}\mathbf{B}\right] \qquad (7)$$

where $\mathbf{C}_m$ is the covariance matrix and can be calculated as the projected model correlation matrix,

$$(\mathbf{C}_m)_{ij} = \langle B_i B_j \rangle_m = \mathbf{\Psi}_i^T \mathbf{R}_m \mathbf{\Psi}_j. \qquad (8)$$

Our method is based upon a linear comparison of models to data, thus the $\mathbf{R}_m$ (and $\mathbf{C}_m$) model matrices only contain second moments of the density field. This linear estimator is computationally more expensive than quadratic or higher order estimators, but the results are less sensitive to non-linearities. For a comparison of different estimation methods, see Tegmark et al. (1998).

In practice we must decide on an explicit parametrization. We construct a power spectrum assuming a primordial spectrum of fluctuations with a spectral index $n_s = 1$. We use a fitting formula from Eisenstein & Hu (1998) to characterize the transfer function, including the baryon oscillations. We fit for $\Omega_m h$ and $f_b = \Omega_b/\Omega_m$ while taking a prior of $H_0 = 72 \pm 8$ km s^{-1} from the Hubble key project (Freedman et al. 2001) and fixing $T_{\text{CMB}} = 2.728K$ (Fixsen et al. 1996). We fit the linearly extrapolated σ_{8g}^L for normalization, where $\sigma_{8g}^L = b\sigma_{8m}$ and b is the bias. Linear theory redshift-space distortions are characterized by β (see §3.2.).

In order to search an appreciable portion of parameter space we have developed efficient methods to calculate the model covariance matrices $\mathbf{C}_m$. The straightforward approach would be to calculate the model correlation matrix for a set of parameters and then project into the KL basis and calculate a likelihood, but this is computationally expensive. The covariance matrix can easily be written as a linear combination of matrices and powers of σ_{8g}^L and β (see §3.2.), so we can project pieces of the correlation matrix and add them in the appropriate proportions for those parameters. However, the shape of the power spectrum depends on Ω_m, f_b, and H_0 in a non-trivial way. We project each bandpower of the correlation matrix (see §3.3.) separately and add the pieces of the covariance matrix together with appropriate weighting to represent different power spectrum shapes. This alleviates the need for further projections. We must be careful when choosing our bandpowers so that we retain sufficient resolution to accurately mimic power spectrum shapes (especially baryon oscillations), but we must also be careful that our k ranges are large enough that the integrals converge correctly.

Note that a non-optimal choice of fiducial parameters does not bias our results, but it can result in non-minimal error bars. This procedure can be iterated if necessary.

4.   Results and Discussion

Our best-fit maximum-likelihood parameter values for samples 10 and 12 are presented in Table 1. Results are given for the priors described in §3.4. and also when using the additional prior $\Omega_b = 0.047 \pm 0.006$ from WMAP (Spergel et al. 2003). We show the results of sample 10 and 12 to give some indication of sample variance, although sample 10 is a subset of sample 12.

   The middle column of Figure 2 shows the marginalized one-dimensional and two-dimensional confidence regions for the power spectrum shape parameters $\Omega_m h$ and f_b for sample 10 without the additional prior on Ω_b. There is a strong correlation between $\Omega_m h$ and f_b. The gross shape of the power spectrum (i.e. ignoring the baryon oscillations and concentrating on the position of the peak and slope of the tail) is nearly constant along the ridge of this correlation due to a degeneracy between shifting the position of the peak with $\Omega_m h$ and adding power to the peak with f_b. However, the strength of the baryon oscillations varies significantly over this range. Table 1 shows that our estimates of $\Omega_m h$ agree well with the WMAP value of 0.194 ± 0.04 (Spergel et al. 2003) and the 2dF value of 0.20 ± 0.03 (Percival et al. 2001) when we use the additional prior on Ω_b, and the associated confidence regions are shown in the left column of Figure 2. The results with the Ω_b prior indicate that the gross shape of the power spectrum we measure is consistent with WMAP and 2dF, as can be seen in Figure 3 which shows the (isotropic) real-space power spectra inferred from the cosmological parameter estimates from the three surveys. However, the results without the Ω_b prior show that we have difficulty breaking the degeneracy between $\Omega_m h$ and f_b because the baryon oscillations are not resolved due to the current state of the SDSS window function.

Table 1.   Maximum likelihood parameter values and 68% confidences (marginalized over all other parameters). Ω_b indicates that a WMAP prior was used.

Param.	10	$10 + \Omega_b$	12	$12 + \Omega_b$
$\Omega_m h$	0.264 ± 0.043	0.207 ± 0.030	0.270 ± 0.057	0.229 ± 0.029
f_b	0.286 ± 0.065	0.163 ± 0.031	0.233 ± 0.088	0.149 ± 0.026
σ_{8g}^L	0.966 ± 0.048	0.971 ± 0.049	0.978 ± 0.043	0.980 ± 0.043
β	0.45 ± 0.12	0.44 ± 0.12	0.44 ± 0.11	0.43 ± 0.11

   The right column of Figure 2 shows the marginalized one-dimensional and two-dimensional confidence regions for σ_{8g}^L (normalization) and β (distortions) for sample 10. Again there is a strong correlation between these parameters, which is expected from their dependence on b. Our constraint on σ_{8g}^L is strong, but we can only measure β to $\sim 20\%$ which limits our ability to perform an independent estimate of b. We can compare our results to WMAP by examining the combination of parameters $\sigma_{8g}^L \beta = \sigma_{8m}^L \Omega_m^{0.6}$, for which we obtain the value 0.44 ± 0.12, in excellent agreement with the WMAP result of 0.44 ± 0.10 (Spergel et al. 2003). By combining our measurements with WMAP results we find $b = 1.07 \pm 0.13$ for our galaxy sample, but this compares information dominated by
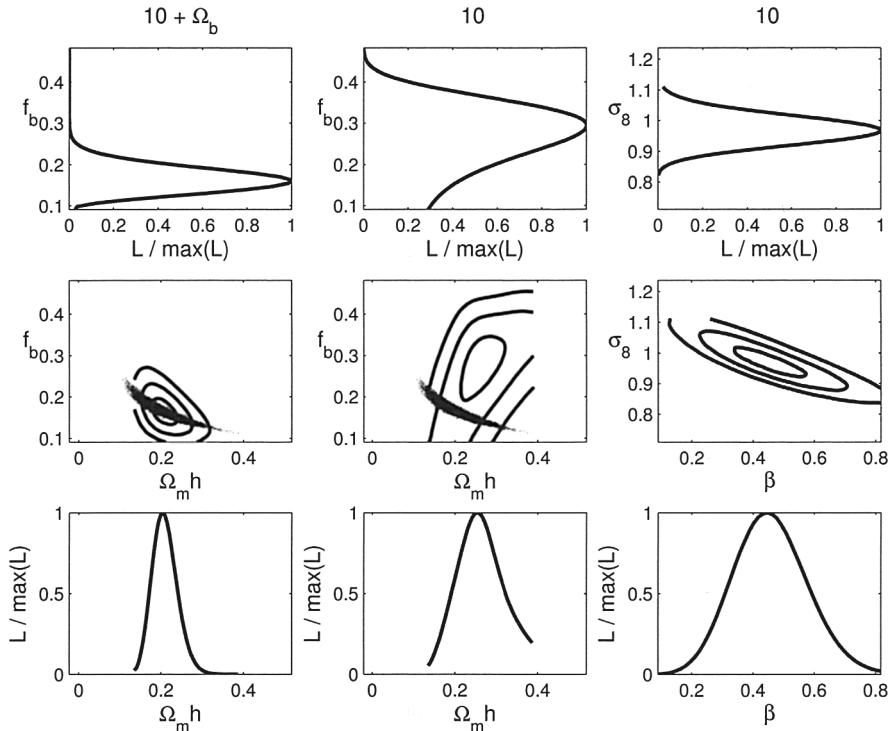
Figure 2. Likelihoods for parameters using sample 10. The left column shows the power spectrum shape parameters with an Ω_b prior. The middle column shows the power spectrum shape parameters without an Ω_b prior. The right column shows normalization and distortion parameters. The contours in the joint parameter plots are the two-dimensional 1, 2, and 3 σ contours. The points in the f_b vs. $\Omega_m h$ plots are MCMC points from WMAP (alone). Parameter combinations not plotted are nearly uncorrelated.

galaxies with redshifts $0.1 \leq z \leq 0.15$ to present-day matter. If we use a ΛCDM model to extrapolate to the present, we would find $b \approx 1.16$. Our galaxies cover a range of luminosities but our signal is dominated by the more luminous galaxies (brighter than L_*) because there are more long baselines available for the more distant galaxies. This must be kept in mind when comparing our measurement of σ_{8g}^L with other estimates using SDSS data which focus on L_* galaxies (Szalay et al. 2003, Tegmark et al. 2004).

This analysis used less than one third of the data that will comprise the completed SDSS survey. Our ability to measure cosmological parameters will increase as the survey area increases, but we should also gain leverage in resolving features in the power spectrum as our survey window function becomes cleaner. The thickest slice of data from the samples used was roughly 10°, implying a thickness of $\sim 50h^{-1}$ Mpc at $z \sim 0.1$. As the slices become thicker, the KL modes will become much more compact in that direction in k-space. Thus we
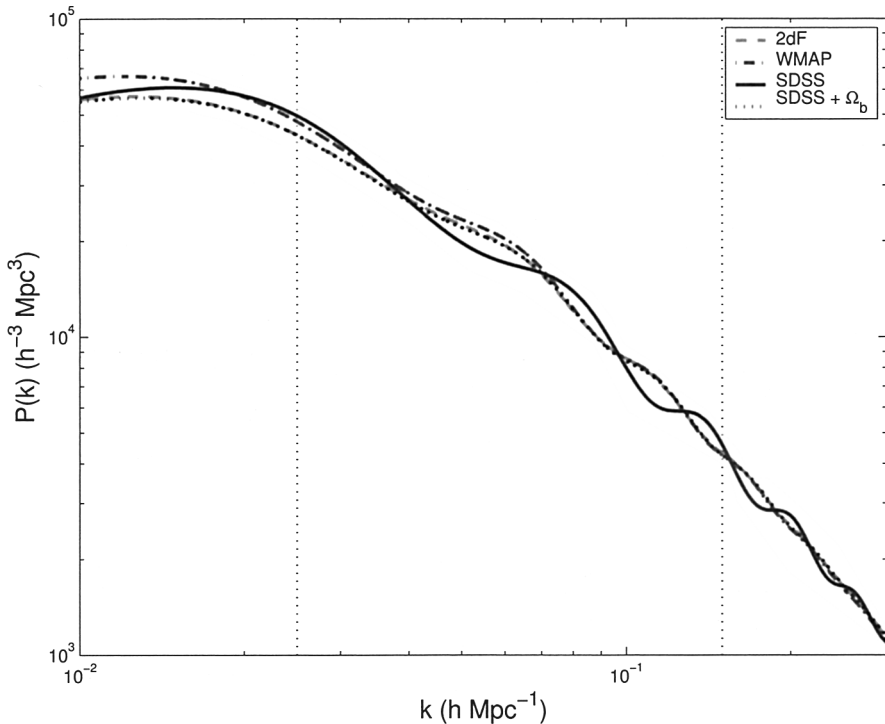
Figure 3.   Plots of the real space P(k) from best-fit model parameters for SDSS (sample 10 with and without the Ω_b prior), WMAP, and 2dF. All use σ_{8g}^L from the SDSS for normalization. The vertical dotted lines indicate the range in k used in the SDSS analysis.

will benefit from the change in the survey aspect ratio in addition to the increase in survey area.

Astrophysics (MPA), New Mexico State University, University of Pittsburgh, Princeton University, the United States Naval Observatory, and the University of Washington.

References

Abazajian, K., et al. 2003, AJ, 126, 2081
Blanton, M. R., et al. 2003a, AJ, 125, 2276
Blanton, M. R., et al. 2003b, AJ, 125, 2348
Blanton, M. R., et al. 2003c, ApJ, 592, 819
Bond, J. R., Jaffe, A. H., & Knox, L. 2000, ApJ, 533, 19
Eisenstein, D. J., & Hu, W. 1998, ApJ, 496, 605
Fixsen, D. J., et al. 1996, ApJ, 473, 576
Freedman, W. L., et al. 2001, ApJ, 553, 47
Fukugita, M., et al. 1996, AJ, 111, 1748
Gunn, J. E., et al. 1998, AJ, 116, 3040
Hogg, D. W., Finkbeiner, D. P., Schlegel, D. J., & Gunn, J. E. 2001, AJ, 122, 2129
Kaiser, N. 1987, MNRAS, 227, 1
Kunszt, P. Z., Szalay, A. S., & Thakar, A. R. 2001, Mining the Sky, 631
Matsubara, T., Szalay, A. S., & Pope, A. C. 2004, ApJ, 606, 1
Norberg, P., et al. 2001, MNRAS, 328, 64
Percival, W. J., et al. 2001, MNRAS, 327, 1297
Petrosian, V. 1976, ApJ, 209, L1
Pier, J. R., et al. 2003, AJ, 125, 1559
Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, ApJ, 500, 525
Smith, J. A., et al. 2002, AJ, 123, 2121
Spergel, D. N., et al. 2003, ApJS, 148, 175
Stoughton, C., et al. 2002, AJ, 123, 485
Strauss, M. A., et al. 2002, AJ, 124, 1810
Szalay, A. S., Matsubara, T., & Landy, S. D. 1998, ApJ, 498, L1
Szalay, A. S., et al. 2003, ApJ, 591, 1
Tegmark, M., Hamilton, A. J. S., Strauss, M. A., Vogeley, M. S., & Szalay, A. S. 1998, ApJ, 499, 555
Tegmark, M., et al. 2004, AJ, 606, 702
Therrien, C. W. 1992, Discrete Random Signals and Statistical Signal Processing (New Jersey: Prentice-Hall)
Uomoto, A., et al. 1999, BAAS, 31, 1501
Vogeley, M. S., & Szalay, A. S. 1996, ApJ, 465, 34
York, D. G., et al. 2000, AJ, 120, 1579
Zehavi, I., et al. 2002, ApJ, 571, 172