

Pixelation Effects in Weak Lensing

F. WILLIAM HIGH,^{1,2} JASON RHODES,^{2,3} RICHARD MASSEY,² AND RICHARD ELLIS²

Received 2007 March 17; accepted 2007 September 6; published 2007 December 3

ABSTRACT. Weak gravitational lensing can be used to investigate both dark matter and dark energy but requires accurate measurements of the shapes of faint, distant galaxies. Such measurements are hindered by the finite resolution and pixel scale of digital cameras. We investigate the optimum choice of pixel scale for a space-based mission, using the engineering model and survey strategy of the proposed *Supernova Acceleration Probe* as a baseline. We do this by simulating realistic astronomical images containing a known input shear signal and then attempting to recover the signal using the Rhodes, Refregier, & Groth algorithm. We find that the quality of shear measurement is always improved by smaller pixels. However, in practice, telescopes are usually limited to a finite number of pixels and operational life span, so the total area of a survey increases with pixel size. We therefore fix the survey lifetime and the number of pixels in the focal plane while varying the pixel scale, thereby effectively varying the survey size. In a pure trade-off for image resolution versus survey area, we find that measurements of the matter power spectrum would have minimum statistical error with a pixel scale of $0.09''$ for a $0.14''$ FWHM point-spread function (PSF). The pixel scale could be increased to $\sim 0.16''$ if images dithered by exactly half-pixel offsets were always available. Some of our results do depend on our adopted shape measurement method and should be regarded as an upper limit: future pipelines may require smaller pixels to overcome systematic floors not yet accessible, and, in certain circumstances, measuring the shape of the PSF might be more difficult than those of galaxies. However, the relative trends in our analysis are robust, especially those of the surface density of resolved galaxies. Our approach thus provides a snapshot of potential in available technology, and a practical counterpart to analytic studies of pixelation, which necessarily assume an idealized shape measurement method.

1. INTRODUCTION

Gravitational lensing is the deflection and distortion of light passing through the gravity potentials of massive celestial structures (cf. Blandford et al. 1991; Narayan & Bartelmann 1996). Rapid progress has been made in measuring weakly distorted shapes of galaxies in order to determine the distribution of dark matter in space and time (for reviews see Mellier 1999; Bartelmann & Schneider 2001; Refregier 2003b). The growth of large-scale structures is sensitive to both the gravitational attraction of dark matter and the repulsive effect of dark energy, so weak gravitational lensing is emerging as a promising cosmological probe (Dark Energy Task Force; Albrecht et al. 2006).

SNAP, the *Supernova Acceleration Probe* (Aldering et al. 2004), is one of several proposed space-based missions planning to use weak lensing as a major cosmological probe (e.g., *Dark Universe Explorer [DUNE]*; Refregier et al. 2006). A key design element is the detector pixel scale, θ_{ccd} (arcseconds). *SNAP* has nominally set $\theta_{\text{ccd}} = 0.10''$, which slightly under-

samples its $0.14''$ point-spread function (PSF). Weak-lensing studies require shape measurements of small and faint galaxies, so the question naturally arises as to whether a smaller θ_{ccd} would be better for *SNAP*'s weak-lensing analyses. This is the question that we explore in this paper.

We perturb the *SNAP* pixel scale about its baseline value, studying the effect that this has on the quality of the weak-lensing signal recovery and resulting cosmological measurements. All other survey parameters are fixed, such as aperture, survey lifetime, and number of detector pixels, as we only want to consider perturbations small enough not to affect any of these. Because we are fixing the number of pixels, the area of sky surveyed per unit time changes as θ_{ccd}^2 . Therefore, this can be recast effectively as a trade-off between image quality and survey area, or even weak-lensing errors and cosmological errors. Our principal tools are simulated, weakly lensed images at the different pixel scales, which we subject to a full, realistic weak-lensing analysis, ultimately arriving at estimated errors on the cosmological matter power spectrum.

The shapes of weakly lensed (i.e., “weakly sheared”) galaxies, which are usually faint and small, are inevitably degraded by the convolutions arising from a finite point-spread function and detector pixelation. Mathematically, the observed image I^{obs} at two-dimensional angular position θ relates to the true,

¹ Department of Physics, Harvard University, Cambridge, MA.

² California Institute of Technology, MS 105-24, Pasadena, CA.

³ Jet Propulsion Laboratory, California Institute of Technology, MS 169-506, Pasadena, CA.

preobserved image I^{tr} as (cf. Lauer 1999a, 1999b; Bernstein 2002)

$$I^{\text{obs}}(\theta) = I^{\text{tr}}(\theta) * \mathcal{P}(\theta) * \mathcal{R}(\theta), \quad (1)$$

where the asterisk denotes convolution, \mathcal{P} is the point-spread function, and \mathcal{R} is the pixel-response function (PRF) that describes the spatial sensitivity of a detector pixel. The ultimate image resolution is given by the effective PSF “ePSF,” which is equal to $\mathcal{P} * \mathcal{R}$.

It is possible to achieve the better resolution, in terms of ePSF size, by increasing the primary aperture size rather than decreasing the pixel scale, although this also changes the image signal-to-noise ratio (S/N) at fixed exposure time. We are only interested in the possibility of changing the detector configuration, however, so we fix the aperture.

The PSF of a well-designed optical space telescope will be dominated by diffraction arising from the system pupils, mirror struts, and other physical structures that perturb incident light. But pixelation further degrades the shapes. In the case of a CCD detector, a photon incident at a photosite, or pixel, photoelectrically liberates an electron, which is then bound near the CCD surface to be measured later. The photosite has finite cross-sectional area and counts photons that fall anywhere inside its perimeter, averaging over all subpixel-scale image features. Liberated electrons may also diffuse to the surface of adjacent photosites, an effect termed “charge diffusion.” Averaging and diffusion will erase features of a similar or smaller scale than a pixel, biasing the object’s shape. Pixelation is therefore a convolution, which is equivalent to the averaging described, followed by ideal sampling. Charge diffusion can be modeled as a PRF with spatial extent that is larger than the pixel size.

The net effect of diffraction, pixelation, and charge diffusion on imaging is fully described by the ePSF convolution. As this operation affects the weak shear signal, it must be deconvolved in some way. This is usually done by measuring the ePSF directly from images of stars, then either performing a full deconvolution or calculating the effect of the ePSF on the shear estimators, then subtracting it.

While CCD photosites smooth out small-scale image information, it is nonetheless possible to recover some information by dithering. Dithering involves taking multiple exposures of the same stationary objects so that their subpixel positions in the CCD are different each time. This can be done by slightly translating the camera between exposures. “Ideal interlacing” is the case in which the exposures are translated by steps of $(1/N)\theta_{\text{ccd}}$, where N is an integer. In general, nonideally dithered exposures can be combined before shape analysis using the Drizzle algorithm (Fruchter & Hook 2002) or the Fourier techniques of Lauer (1999a, 1999b), for example, or during shape analysis by somehow “averaging” the measurements. We do not make use of these methods, and instead we investigate the extent to which ideal interlacing with $N = 2$ can recover weak-lensing signals from undersampled data.

Our results are part of a wider investigation of the accuracy of weak-lensing measurements and the effects of various instrumental and algorithmic parameters. Our image simulation software is employed by the Shear Testing Program⁴ (STEP; Heymans et al. 2006; Massey et al. 2007b). The STEP project comprises independent blind analyses of simulated data of various types by international groups with a goal of verifying the limitations of extant weak-lensing measurement algorithms. In a similar study, Massey et al. (2004b) analyze the effect of varying the exposure time on weak shear recovery from a *SNAP*-like mission. Refregier et al. (2004) study semianalytically the trade-off between a wide and deep *SNAP* survey strategy by looking at the expected errors on cosmological parameters from weak-lensing data. Bernstein (2002) and G. Bernstein et al. (2008, in preparation) study pixelation effects, but starting from the opposite end of analytic first principles, evaluating the information loss during pixelation that is irreversible, even assuming a perfect shape measurement method.

This paper is organized as follows. Section 2 introduces our adopted PSF model (§ 2.1) and simulation software (§ 2.2) and discusses how we vary the pixel scale (§ 2.3), our simulated dither and stacking strategy (§ 2.4), and our shear measurement algorithm (§ 2.5). Results are presented in § 3. Properties of the simulated images are shown in § 3.1, and then the recovered weak-lensing signals are analyzed in terms of the surface density of usable galaxies (§ 3.2), the bias determination (§ 3.3), the standard error of shear estimation (§ 3.4), and the predicted error on the matter power spectrum (§ 3.5). We conclude with § 4.

2. METHOD

We have developed a software suite to simulate all relevant steps from the acquisition of imaging data, through its reduction, to the measurement of the weak shear signal. Our method is to

1. realistically simulate images at different pixel scales, containing galaxies with known shear signal;
2. redetect galaxies in the noisy images and measure their shapes;
3. detect stars in separate images to characterize the ePSF convolution kernel;
4. correct galaxy shapes for ePSF effects; and
5. estimate the weak-lensing errors and the errors on the matter power spectrum.

We have four sets of simulated images, which are summarized in Table 1. The following sections describe in detail the main components of our methods and data.

2.1. Adopted PSF and Charge Diffusion

We assume an ePSF that is spatially and temporally invariant. This isolates the problem of shear measurement in which we

⁴ See <http://www.physics.ubc.ca/~heyman/step.html>.

TABLE 1
SUMMARY OF THE SIMULATED IMAGES

Label ^a	Variable ^b	Plate Scale	σ_{ch}^c	Dithered?
PH	Photosite size in microns	$10'' \text{ mm}^{-1}$	$4 \mu\text{m} = 0.04''$	No
PH _d	Photosite size in microns	$10'' \text{ mm}^{-1}$	$4 \mu\text{m} = 0.04''$	Yes
FL	Focal length (plate scale)	$10'' \text{ mm}^{-1} (\theta_{\text{ccd}}/\theta_0)$	$4 \mu\text{m} = 0.04'' (\theta_{\text{ccd}}/\theta_0)$	No
FL _d	Focal length (plate scale)	$10'' \text{ mm}^{-1} (\theta_{\text{ccd}}/\theta_0)$	$4 \mu\text{m} = 0.04'' (\theta_{\text{ccd}}/\theta_0)$	Yes

^a The shorthand label assigned to a common set of simulated images.

^b What is perturbed in order to change θ_{ccd} .

^c Standard deviation of the Gaussian charge diffusion kernel as a function of perturbed CCD pixel scale.

are primarily interested, without distractions of ePSF interpolation, which is a separable problem that is being widely discussed elsewhere.

Our PSF model is based on an early engineering design, called TMA63, of the *Supernova Acceleration Probe*⁵ space telescope (Aldering et al. 2004). The simulated diffraction pattern comes from a ray-traced model of the *SNAP* f/11 optical system. It simulates light from an 820 nm wavelength point source incident on a 2 m primary mirror, a secondary supported in front of the primary by three struts, a folding flat, and a tertiary. The point source is taken to be off-axis such that its image appears on the focal plane 198 mm radially away from the optical axis, where the typical *SNAP* CCD lies. Figure 1 shows the diffraction profile.

SNAP plans to employ Lawrence Berkeley National Labo-

ratory's new high-resistivity CCDs, in which charge diffusion further spreads point-source light (Holland et al. 2003; Stover et al. 2000; Groom et al. 2000). We model charge diffusion as an additional Gaussian convolution with standard deviation $\sigma_{\text{ch}} = 4 \mu\text{m}$, the expected level for these CCDs, which have $10 \mu\text{m}$ wide photosites. The net ePSF at the baseline *SNAP* CCD pixel scale of $0.10''$ has a final FWHM of $0.14''$.

2.2. Simulated Images and Input Shear

The image simulation suite of Massey et al. (2004a), *Simage*, is our main tool for creating artificial astronomical images and applying an arbitrary weak shear signal. *Simage* uses Shapelets,⁶ a parameterization of galaxy morphologies as a weighted sum of a complete, orthonormal set of basis functions (Refregier 2003a).

Realistic morphologies are generated by empirically matching the measured properties of actual galaxies in the Hubble Deep Fields (HDFs; Williams et al. 1996, 1998). The HDF galaxy positions, orientations, and morphologies (plus sizes and magnitudes) are randomized when generating new images. In this way, realistic magnitude-morphology trends are produced, although no spatial clustering is imposed and no redshift information is encoded. Because the HDF catalog sources on which the simulations are based are already preconvolved with the HDF PSF, they have slightly larger intrinsic sizes than the true galaxy population. However, this does not affect the process of shear addition and measurement.

The analytic Shapelet models of galaxies are then sheared as described by Massey & Refregier (2005) including terms up to fourth order in the shear γ . Our input shears $\gamma = (\gamma_1, \gamma_2)$ range from -0.05 to $+0.05$ in steps of 0.025 in each component, while the other component is fixed at 0 , viz.,

$$\gamma_1 \in \{-0.05, -0.025, 0, 0.025, 0.05\} \text{ and } \gamma_2 = 0; \quad (2)$$

$$\gamma_1 = 0 \text{ and } \gamma_2 \in \{-0.05, -0.025, 0.025, 0.05\}. \quad (3)$$

A total of nine different input shears per pixel scale are thus produced. We make one image per input shear, uniformly distorting all galaxies in the field. The sheared galaxies are con-

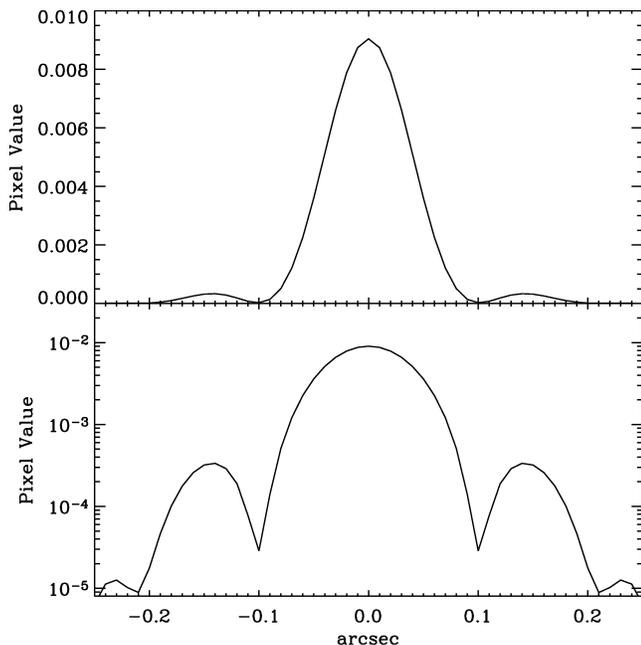


FIG. 1.—Cross sections through the peak of the diffraction pattern on linear (top) and log (bottom) scales. This comes from ray-tracing simulations of the *SNAP* telescope, finely pixelated at $0.01''$. The diffraction has FWHM $\sim 0.09''$. The y-axis scale is arbitrary but common to both panels.

⁶ See <http://www.astro.caltech.edu/~rjm/shapelets/> for a Shapelet analysis package.

volved with the PSF and transformed to real image space, integrating the basis functions analytically within each pixel. This integration is equivalent to convolution with a top-hat pixel-response function. Finally, photon noise and a realistic space background signal are added.

2.3. Perturbing the Pixel Scale

The key question motivating this study is what pixel scale most benefits weak-lensing analyses. To address this we simulate and analyze images with CCD pixel scales perturbed about a baseline value chosen to be $\theta_{\text{ccd}} = \theta_0 \equiv 0.10''$. In practice, pixel scales can be adjusted in one of two ways, and this defines two categories of simulations.

First, we simulate changing the CCD pixel scale θ_{ccd} by changing only the physical size of the photosites, labeling this data set ‘‘PH,’’ as the photosite size is the variable. Second, we simulate changing θ_{ccd} by adjusting only the focal length of the telescope, which changes the plate scale, calling this data set FL, as the focal length is the variable. While these each linearly perturb θ_{ccd} , their detailed effects on the final image are distinct and therefore each merits analysis.

The PH data set consists of different simulated images with detector pixel scales between $0.04''$ and $0.16''$ at $0.01''$ intervals. This ranges from the resolution of the processed, end-product HDF images to just larger than the baseline ePSF size. We adopt a baseline photosite size of $10 \mu\text{m}$ and then vary this between 4 and $16 \mu\text{m}$ in $1 \mu\text{m}$ steps. In practice, such changes would need to be made at the CCD manufacturing stage. Both the diffraction pattern and the baseline plate scale of $0.01'' \mu\text{m}^{-1}$ are unaffected. Assuming that the electron diffusion length is constantly $4 \mu\text{m}$ independent of the photosite size, the different photosite sizes cause no net effect in the apparent *angular size* of the charge diffusion. The final shape and angular size of the PSF, that is, before pixelation, are left unchanged at all θ_{ccd} . This scheme for pixel scale adjustment represents a relatively academic exercise into the effect of pixelation of a fixed PSF, as it is more costly and difficult to achieve in practice.

The FL data set emulates the effect of changing the focal length while keeping the physical size of the CCD photosite fixed, thereby changing the plate scale and θ_{ccd} from $0.04''$ to $0.16''$. At longer focal lengths the diffraction, which is constant in angular size, grows linearly in microns at the focal plane while the pixel scale decreases in arcseconds and the charge diffusion remains at $4 \mu\text{m}$. At these small pixel scales, the diffraction pattern dominates and the first Airy ring is readily apparent. At shorter focal lengths the diffraction size shrinks while the pixel scale increases. In this case the charge diffusion dominates, smoothing out the diffraction features and broadening the ePSF. This is perhaps a more practical engineering solution to adjusting the pixel scale.

In both cases, the $\theta_{\text{ccd}} = 0.04''$ images contain 4096×4096 pixels. These dimensions change linearly with θ_{ccd} such that the subtended solid angle per image is constant. This does not simulate how an instrument would change its CCD array

size with θ_{ccd} in practice; indeed, we assume in our final figure of merit (§ 3.5) that the number of pixels in a real instrument would be fixed. We add about 6000 galaxies and no stars to each sheared image, extending their distribution well below the intended detection threshold. Adding noise to fix the survey depth to $m \approx 27.7$ with 2000 s exposures, we reproduce the number of galaxies useful for weak shear estimation, $n_g \sim 100 \text{ arcmin}^{-2}$, found at the baseline pixel scale in earlier studies (Massey et al. 2004b). For each pixel scale, we also make one additional image containing only stars, simulating a weak-lensing survey that periodically points at stellar fields to characterize the ch- and ePSF.

2.4. Simulating and Combining Dithers

In addition to the two sets of single exposures introduced above, we also create two additional data sets, labeled PH_d and FL_d , to test how much pixel resolution can be recovered by combining dithered images. PH_d emulates PH in how θ_{ccd} varies by changing the size of CCD photosites, and FL_d emulates how FL changes the pixel scale by changing the plate scale. The new data consist of sets of four shallower images, dithered in a 2×2 interlacing pattern by half-pixel shifts in both directions. These resample the ePSF-convolved image from equation (1) at shifted intervals $\Delta\theta = (d_x, d_y)\theta_{\text{ccd}}$, where

$$\text{dither 1 : } (d_x, d_y) = (0, 0), \quad (4)$$

$$\text{dither 2 : } (d_x, d_y) = (0.5, 0), \quad (5)$$

$$\text{dither 3 : } (d_x, d_y) = (0.5, 0.5), \quad (6)$$

$$\text{dither 4 : } (d_x, d_y) = (0, 0.5). \quad (7)$$

We then perform an ideal interlacing, stacking the four dithers on a pixel grid that is twice as fine at $\theta_{\text{ccd}}/2$. We map each dither pixel to one resulting pixel with no overlap, such that the noise between adjacent pixels in the stack is uncorrelated. This represents the most optimistic case, although space telescopes should deliver sufficient pointing accuracy to achieve nearly ideal interlacing.

All dithered and stacked images subtend the same solid angle as in the other two studies. Each dither is a 500 s exposure, resulting in a 2000 s effective exposure after stacking. The CCD pixel scale θ_{ccd} varies from $0.08''$ to $0.24''$ in steps of $0.02''$, so the final stacked pixel sampling scales range from $0.04''$ to $0.12''$. Again, Table 1 summarizes the four simulation sets.

2.5. Weak-Lensing Measurement

We now turn to exploiting the simulated images in order to recover the weak-lensing signal and thereby determine the surface density of usable galaxies under various assumed pixel sizes and exposure strategies. We locate galaxies in the noisy images using Source Extractor (Bertin & Arnouts 1996) configured to detect all objects as near to the noise threshold as

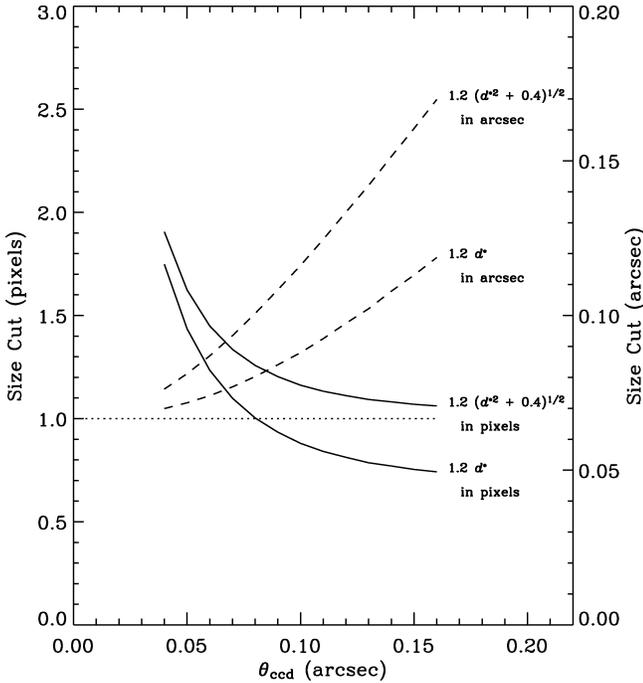


FIG. 2.—Two different lower size cuts d_{cut} in pixels (solid line, left axis) and arcseconds (dashed line, right axis) vs. θ_{ccd} . (Note that the left axis does not map horizontally to the right axis, since θ_{ccd} is made to vary.) Here d_* is defined as the median size of high-S/N stars at each pixel scale. We ultimately apply cuts with the 0.4 “floor” term, which has little effect at the smaller pixel scales but causes the lower size limit to tend toward about 1 pixel (dotted line) at large pixel scales. This ensures that galaxies nearly the size of the PSF are not used in the shear estimation. (Plotted here are cuts made in the FL study, with size cuts in other simulation sets behaving similarly.)

possible via a Gaussian detection kernel matched to the known size of the ePSF at each pixel scale. Source Extractor convolves the images with this detection kernel, blurring features smaller than the size of the kernel, such as pixel-to-pixel photon counting noise, to avoid spurious detections.

We adopted the shear measurement method by Rhodes et al. (2000; hereafter RRG) to measure the shapes of detected galaxies. RRG was specifically developed with space-based weak-lensing measurements in mind (Rhodes et al. 2000) and has undergone extensive tests on simulations (Leauthaud et al. 2007) and use on *Hubble Space Telescope* images (Rhodes et al. 2001, 2004; Refregier et al. 2002; Massey et al. 2007a) to constrain cosmological parameters including σ_8 , the normalization of the dark matter power spectrum.

RRG is a modification of the KSB+ (Kaiser et al. 1995; Hoekstra et al. 1998) method, which measures Gaussian-weighted multipole image moments,

$$J_{ij} = \int d^2 \theta w(\theta) I(\theta) \theta_i \theta_j, \quad (8)$$

where w is a Gaussian; $i, j \in \{x, y\}$, corresponding to the or-

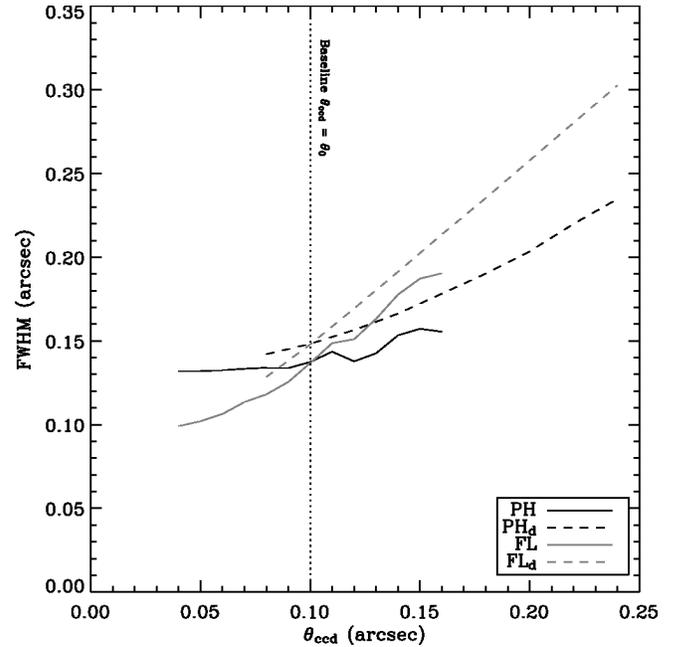


FIG. 3.—Stellar FWHM values (ePSF), measured from the SExtracted simulations at pixel scale θ_{ccd} . For the interlaced images, we measure the stellar FWHM from the stacked images, which have a final image sample rate that is $\frac{1}{2}\theta_{\text{ccd}}$. The two different methods of perturbing the pixel scale, labeled PH and FL, are expected to agree at the nominal $\theta_{\text{ccd}} = 0.10''$ by design, and this plot shows this to be the case for both the single exposures and the combined dithers.

thogonal image coordinates; and θ is chosen such that the weighted barycenter is zero. RRG corrects the galaxy image moments for ePSF effects using the moments of the measured ePSF. More advanced shape measurement algorithms that are being developed, including the Shapelets-based method of Massey et al. (2007), instead model the charge diffusion-convolved PSF. This requires higher resolution data but allows information known a priori about a regular pixel response function to be included analytically, whether that be a fixed square top hat or known pixel-to-pixel variations in the PRF.

RRG forms a measure of ellipticity (in contrast to KSB+), only at the final stage, after correction of individual shape moments,

$$e = \frac{(J_{xx} - J_{yy}, 2J_{xy})}{J_{xx} + J_{yy}}. \quad (9)$$

A shear estimator is then formed,

$$\hat{\gamma} = \frac{e}{G}, \quad (10)$$

where the shear susceptibility G is a scalar function of higher order moments of the ensemble of galaxies. Here G is averaged over the galaxies in a given image.

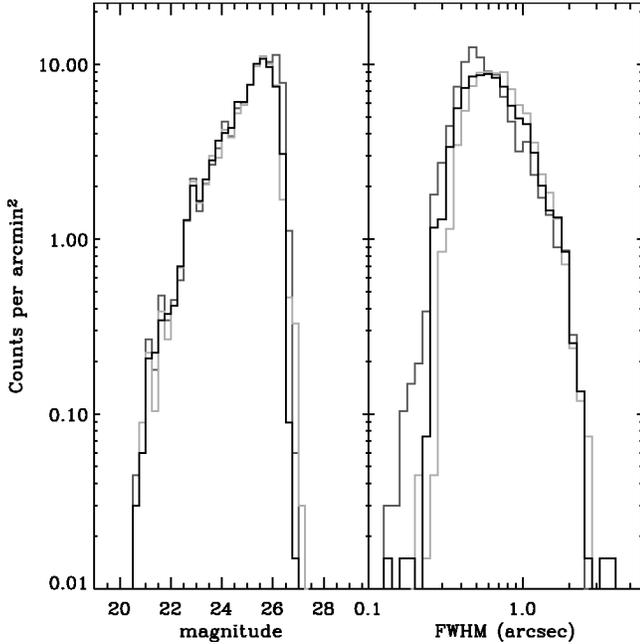


FIG. 4.—Surface density of galaxies as a function of magnitude (*left*) and size (*right*). These are the galaxies that we use to measure the shear, which survive various cuts as described in § 3.2, in PH images with $\theta_{\text{ccd}} = \theta_0$ (*black*) and $\theta_{\text{ccd}} = 0.4\theta_0$ (*dark gray*), and in PH_d images with $\theta_{\text{ccd}} = 2\theta_0$ (*light gray*). Magnitude bins are 0.25 wide, and size bins are 0.05 wide in \log_{10} space. Magnitude is taken from the Source Extractor `mag_auto` output, and FWHM from the `fwhm_image` output.

We then remove objects from the shear catalog in a similar manner to analyses of real data by Massey et al. (2007a) and Leauthaud et al. (2007) and shown by the dashed lines in Figure 5. We first eliminate the <1% of galaxies that Source Extractor misclassified as stars, and also use the Source Extractor S/N outputs `flux_auto` and `fluxerr_auto` to cut galaxies with $S/N < 10$.

We then calculate the rms size measure, from the combination of uncorrected quadrupole image moments

$$d_{\text{rrg}} = \sqrt{\frac{J_{xx} + J_{yy}}{2}}. \quad (11)$$

We remove large galaxies with $d_{\text{rrg}} > 2000$, which eliminates fewer than 1 detected object per arcmin^2 in our simulations. And via

$$d_{\text{rrg}} < 1.2\sqrt{d_*^2 + 0.4} \equiv d_{\text{cut}} \quad (12)$$

we eliminate galaxies whose size is nearly equal to the measured ePSF. The 0.4 term also eliminates galaxies that are only a few pixels across, which have $d_{\text{rrg}} \approx 1$. This term is negligible at small pixel scales, as Figure 2 shows, because the ePSF is large in pixel units. At large θ_{ccd} , where the ePSF size is nearly

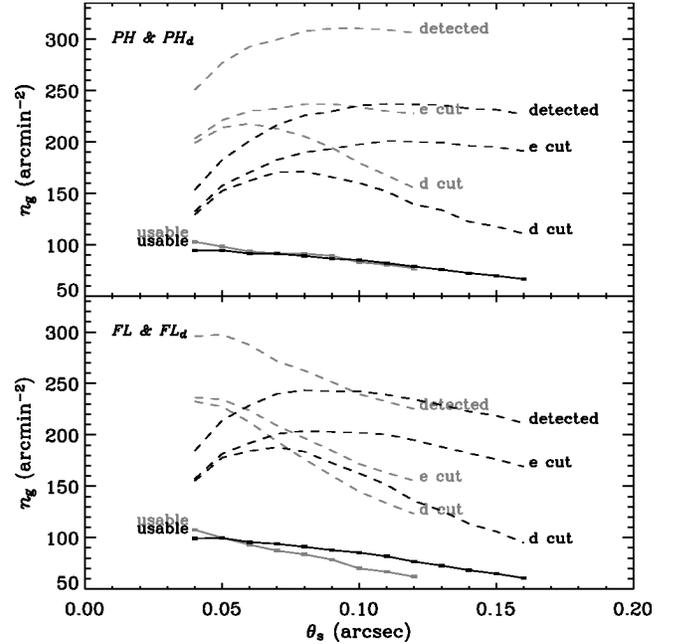


FIG. 5.—Surface density of galaxies detected and usable for weak lensing, n_g , along with counts after different cuts are applied to the data in turn, vs. the pixel scale of the reduced images, θ_* . The top panel plots data from PH (*black lines*) and PH_d (*gray lines*), and the bottom panel similarly for FL and FL_d. Galaxies are detected with Source Extractor, and cuts are made based on measured ellipticity (“e cut”), size (“d cut”; see eqs. [11] and [12]), and S/N (final cut, giving “usable” galaxies). Error bars are calculated as the sample variance of n_g between the different images at a given pixel scale.

1 pixel, the 0.4 term causes $d_{\text{cut}} \rightarrow 1$ pixel, also seen in Figure 2. This cut ensures that we measure shear on galaxies that are somewhat larger than both the PSF and the pixel size, which is the regime where shape measurement is most reliable. The effect of the lower size cut on the data is shown with the dashed line below the ellipticity cut in Figure 5. We also test applying a lower size limit of $d_{\text{cut}} = 1.2d_*$ instead of equation (12). This leaves $n_g(\theta_{\text{ccd}} \lesssim \theta_0)$ unchanged within a few galaxies per arcmin^2 at most, but $n_g(\theta_{\text{ccd}} > \theta_0)$ increases by up to $\sim 10 \text{ arcmin}^{-2}$. The result that smaller pixel scales always yield higher galaxy surface densities, however, is robust.

3. RESULTS

After discussing some properties of the simulated images, we explore as a function of pixel scale the number of detected and “useful” galaxies for weak-lensing measurements n_g , the performance of shear recovery, the sample variance of shear estimators, and the predicted errors on the dark matter power spectrum. The number of galaxies for which shear can be measured is mainly a function of the angular size of galaxies relative to the size of the PSF. These results are therefore likely to be independent of our adopted shear measurement method. However, the performance of shear recovery and the error on the

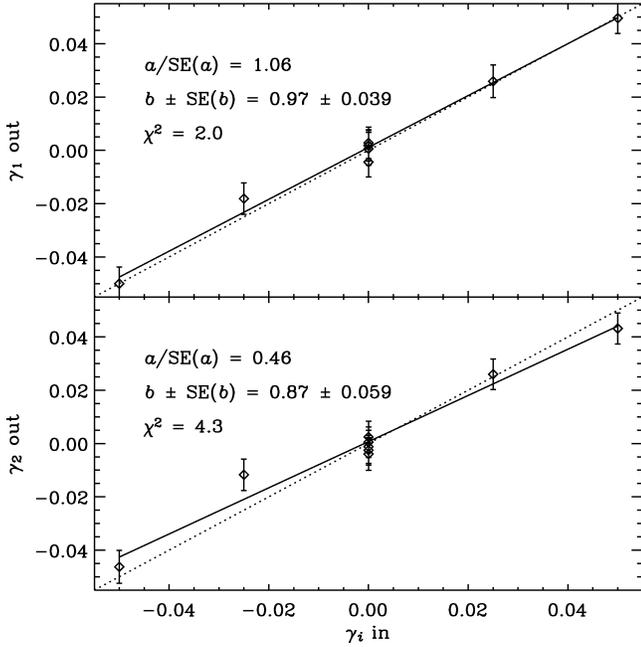


FIG. 6.—Example of known, input shear vs. measured shear, for one set of simulated images. Points are $\langle \hat{\gamma}_i \rangle$, measured from the data. Error bars are the standard error of shear estimators in each image. Here $\theta_{\text{px}} = 0.10''$. The dotted lines have slope 1, representing perfect shear recovery, and the solid lines are linear least-squares fits to the data points. Also shown are the y-intercept a divided by the standard error of a , $\text{SE}(a)$, and the slope b plus or minus the standard error of b , $\text{SE}(b)$ —all outputs from the least-squares algorithm. Also shown are the χ^2 values of the fit, which has 7 degrees of freedom.

power spectrum will depend on the sophistication of the chosen method, so those results are likely to improve before the launch of any future space-based mission.

3.1. Image Properties

Figure 3 shows the measured ePSF sizes in each simulation set. Because we are perturbing about a baseline of $\theta_{\text{ccd}} = 0.10''$, we expect results to agree there, which is shown to be the case. Charge diffusion dominates the FL and FL_d ePSFs at larger pixel scales, and diffraction dominates at smaller pixel scales. Because the charge diffusion is constant at all pixel scales in PH and PH_d, the only variable source of resolution degradation is the top-hat pixel kernel. Charge diffusion is clearly a detrimental effect in terms of ePSF size.

As an illustration, Figure 4 shows the size and magnitude distributions of galaxies useful for lensing measurements in images from PH at the baseline pixel scale. The ordinate axis units are counts per arcmin², and the data are measured on galaxies after we make cuts, as described in § 3.2. Our images have a depth of about magnitude 27.7, but the faintest galaxies are not ultimately used.

3.2. Surface Density of Sheared Galaxies

The surface density of galaxies that survive the various cuts discussed above is plotted as the solid line in Figure 5. The uppermost dashed lines in each simulation set show the raw galaxy detections using Source Extractor,⁷ averaged from all nine sheared images at each pixel scale in each simulation set. The surface density of galaxies that survive this cut is plotted in the next highest dashed line for each simulation set in Figure 5.

As expected, the usable n_g decreases monotonically in all simulation sets, indicating that smaller θ_{ccd} always helps decrease statistical error in mean shear measurement for a fixed survey area. Stacking subpixel dithers fully recovers the surface density of galaxies from PH, despite the larger underlying CCD pixel convolution. Some galaxies are lost in the FL_d data. Here the constituent dithers, unlike the PH and PH_d sets, have an even larger underlying ePSF than the nonstacked counterparts due to the dominant charge diffusion. Taken together, these results indicate that n_g worsens with increased charge diffusion but is independent of the CCD pixel spacing θ_{ccd} .

We estimate errors on the surface density of usable galaxies, plotted with error bars in Figure 5 (errors are quite small compared to the plot line width), to be the sample variance of n_g measurements between each of the nine simulations per pixel scale. These errors are of the order of one galaxy per arcmin² for all data, which is therefore negligible in our analysis. However, it is worth noting that the number counts of galaxies in the simulated images were calibrated against those in the Hubble Deep Field, which may not be a representative sample of the universe (van Waerbeke et al. 2006).

3.3. Accuracy of Weak Shear Recovery

We define “accuracy” to be the closeness of the measured value to the input value. Figure 6 illustrates the accuracy with which we recover shear measurement with RRG for one typical set of simulated images. In all sets, the shear recovery is well fitted by a linear model, thereby justifying the catalog cuts described in § 2.5. Relaxing the cuts and using smaller or fainter objects introduces systematic effects or a nonlinear shear response. Again, in all cases, the y-intercepts (“additive shear residual”) of the shear recovery are consistent with zero and accordingly set explicitly so in subsequent analyses.

On the other hand, the best-fit slopes (“multiplicative shear bias”) are systematically smaller than unity by roughly 20%. This effect has been known for some time as a limitation of KSB+ methods (e.g., Bacon et al. 2001), and it has been speculated to arise as a result of pixelation (Massey et al. 2007b). The population bias has appeared to be robust to effects such as galaxy morphological type, and we assume in all subsequent sections that it could, in practice, be determined to arbitrary

⁷ See http://terapix.iap.fr/rubrique.php?id_rubrique=91/.

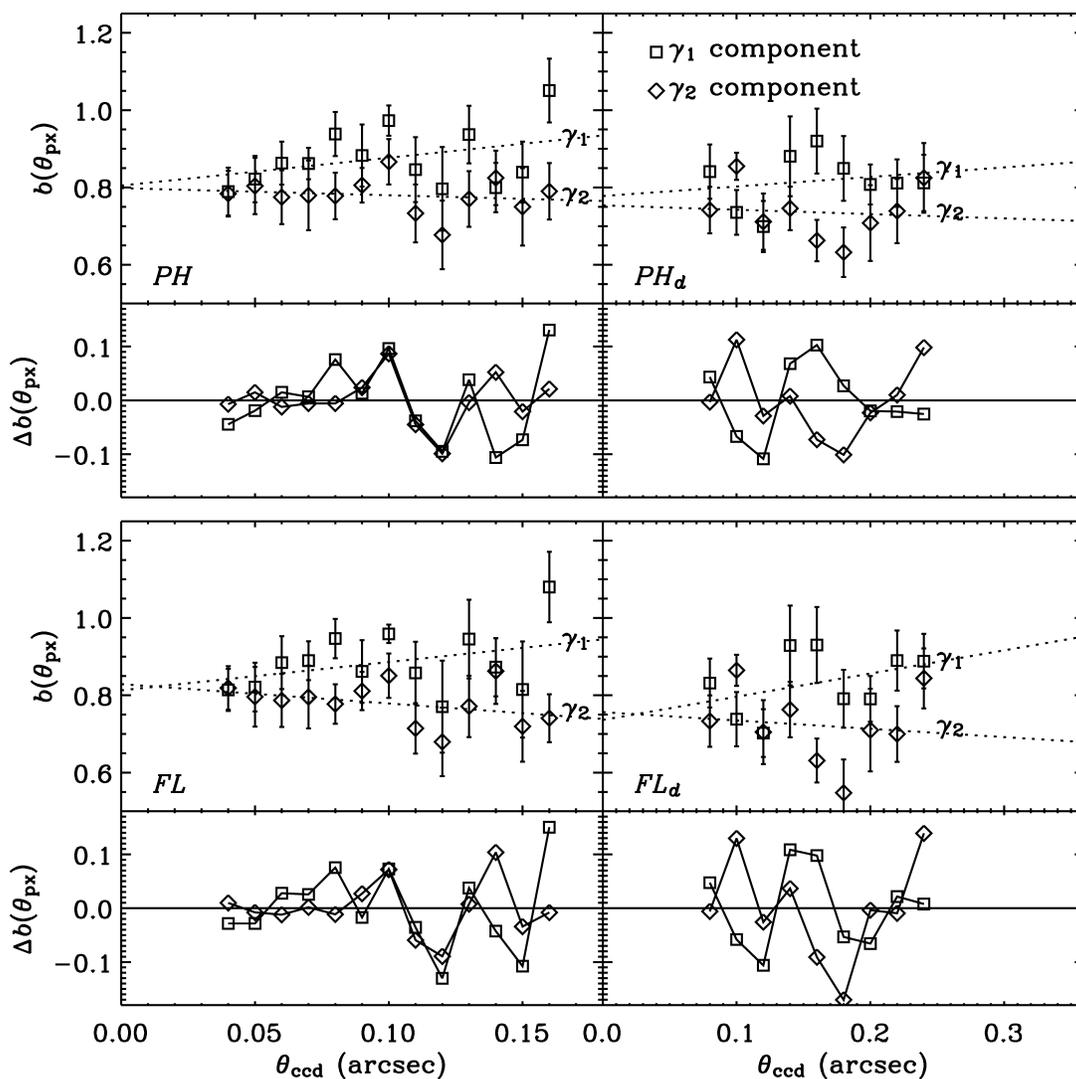


FIG. 7.—RRG shear estimator bias b , and deviation from the least-squares line fitted to each component Δb , as a function of pixel scale θ_{ccd} . Squares indicate the γ_1 component of shear, relative to the pixel grid, and diamonds the γ_2 component. Error bars show the standard error on slope measurement obtained from the least-squares fits to data like those shown in Fig. 6. The bias in each component of shear is different and varies with pixel scales. Extrapolating the pixel scale to zero appears to give a bias of about 0.8. Interlacing to smaller pixel scales does not remove the bias nor the trend toward 0.8.

accuracy using simulations. We therefore take

$$b \equiv \frac{\text{bias}}{\gamma} + 1, \quad (13)$$

where, conventionally,

$$\text{bias} \equiv \langle \hat{\gamma} \rangle - \gamma, \quad (14)$$

to be the true value with negligible error and correct for the bias using values from the line fitted as a function of pixel scale in Figure 7. We call b the bias for brevity.

In general, images that undersample the PSF, as our undithered data do at pixel scales larger than roughly our nominal

one, should suffer from aliasing. Indeed, in poorly sampled images there is great ambiguity as to the shapes of objects that are about the size of a pixel. However, we do not expect this to be a problem in our study because we always average shapes of many stars and galaxies that are randomly positioned throughout our images. In practice, as our plots show, we observe no significant or catastrophic bias or error as the pixel scales vary from under- to oversampled.

3.4. Precision of Weak Shear Recovery

We define “precision” to be the scatter in measured shear compared to the input signal. Weak lensing induces changes of only a few percent in galaxy ellipticities, but the rms of the

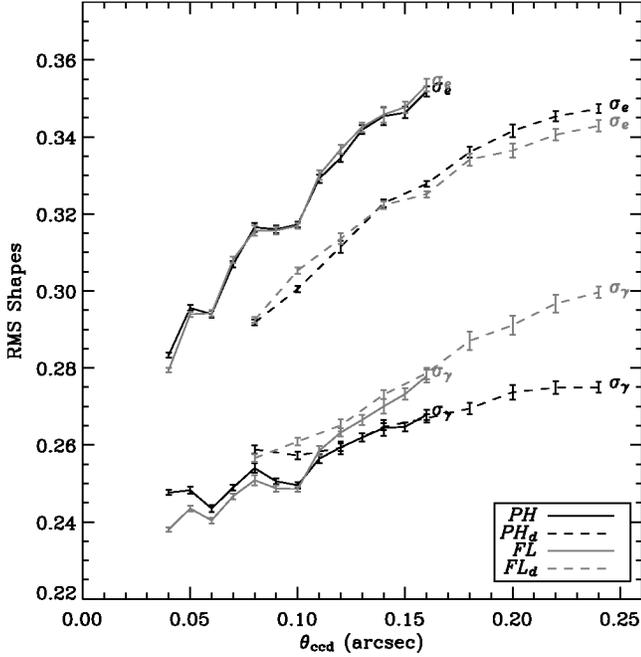


FIG. 8.—The rms of galaxy ellipticity and shear estimators in simulated images with single exposures (*black lines*) and stacked dithered (*gray lines*). The rms of the shear is calculated as $\sigma_\gamma = (\sigma_{\gamma_1}^2 + \sigma_{\gamma_2}^2)^{1/2}$, where σ_{γ_i} is the rms of the bias-corrected shear estimators. Errors are calculated as the sample variance of the rms shapes between different images at a given pixel scale. The rms of the galaxy ellipticities are calculated similarly from uncorrected ellipticities.

ellipticity distribution of faint galaxy populations are about 30% (e.g., Leauthaud et al. 2007). Figure 8 shows the standard deviation of ellipticities σ_e and shear estimators σ_γ (after correction for multiplicative bias), where σ_e is defined analogously to

$$\sigma_\gamma^2 = \sigma_{\gamma_1}^2 + \sigma_{\gamma_2}^2, \quad (15)$$

where

$$\sigma_{\gamma_i}^2 = \langle (\hat{\gamma}_i - \langle \hat{\gamma}_i \rangle)^2 \rangle. \quad (16)$$

As seen in Figure 8, the measured rms ellipticity increases monotonically with θ_{ccd} , meaning that larger pixels always make shape measurements noisier, an undesirable effect for weak-lensing analysis. The ePSF “dilutes” the shear signal in galaxies, and RRG is shown here to reverse this dilution at all pixel scales by decreasing the rms of shapes, as expected. The near coincidence of the nondithered and dithered σ_γ lines suggests that the rms shear depends most on the CCD pixel spacing θ_{ccd} and whether or not (ideal) interlacing is performed. This is, in fact, the ideal situation: the best image resolution that we can have is fundamentally limited by (in addition to the diffraction) the CCD pixel response—that is, the ePSF. Ideal in-

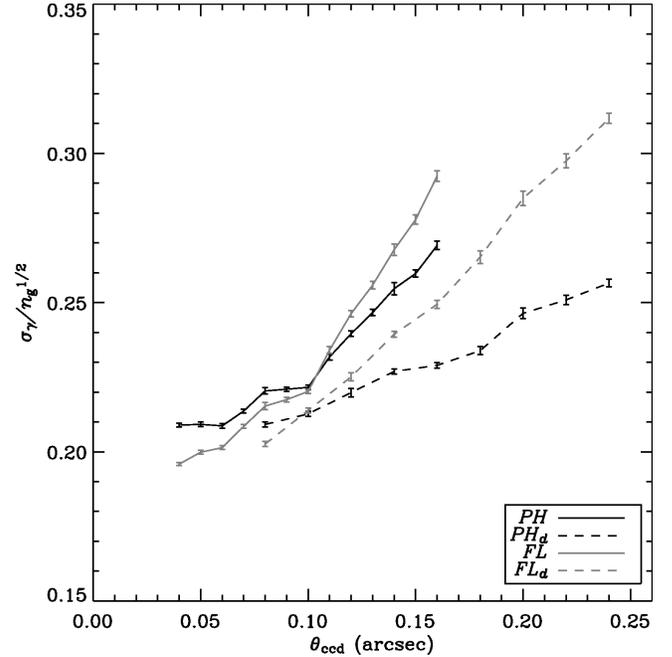


FIG. 9.—Standard error of the shear, $\sigma_\gamma / \sqrt{n_g}^{1/2}$, in the single-exposure (*black lines*) and dithered (*gray lines*) simulations. Error bars are calculated as the sample variance of the standard error between the different images at a given pixel scale. This quantity combines the method-independent n_g and the more method-dependent σ_γ into an estimate of the precision of our shear estimation. This is the figure of merit showing how precisely we can measure shear as a function of θ_{ccd} if the survey area is fixed.

terlacing increases the sample rate of a given object, which is why σ_e decreases when dithering. Figure 8 shows that RRG recovers the underlying shape information at large pixel scales *from the ensemble* even without dithering.

By assuming uncorrelated shapes in galaxy populations (for a discussion, see Hirata & Seljak 2004), measurement errors can be reduced by measuring many galaxies and applying Poisson statistics. An ability to use more galaxies per arcmin², n_g , improves the precision of shear recovery, which is best quantified by the sample variance of shear estimators, or its square root $\sigma_\gamma / \sqrt{n_g}$. This is shown as a function of pixel scale in Figure 9. This confirms our result that smaller pixel scales always improve the precision of weak shear estimation. However, smaller pixel scales face diminishing returns: the gains in precision from slightly decreasing the pixel scale about θ_0 appear to flatten out. Dithering decreases the error somewhat, thanks entirely to the behavior of n_g with θ_{ccd} . The shear errors from the FL and FL_d sets degrade faster with increased θ_{ccd} . This indicates that charge diffusion, which dominates larger θ_{ccd} in the FL and FL_d sets, should be minimized in addition to θ_{ccd} .

One potential concern is the observation by Kaiser (2000) that the distribution of shear estimators in practice is not Gaussian and that extended wings may even make the second moments of this distribution infinite. This is not surprising because

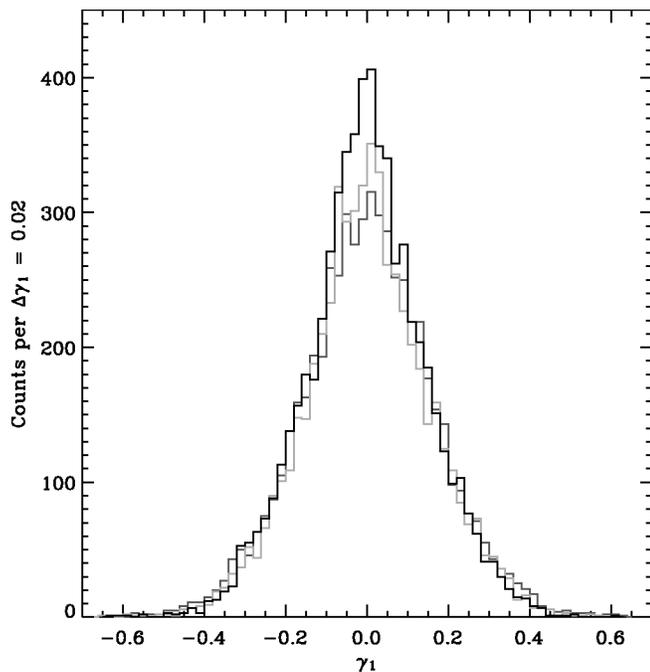


FIG. 10.—Distribution of shear estimators $\hat{\gamma}_1$ obtained from PH images with $\theta_{\text{ccd}} = 0.10''$ (black histogram) and $\theta_{\text{ccd}} = 0.04''$ (dark gray histogram), and PH_d images with $\theta_{\text{ccd}} = 0.20''$ (light gray histogram) pixel scales. Each sample comes from all images at the quoted pixel scales, which have different input shears applied. The input shears perturb this distribution very slightly but are all small and have a mean of zero.

the shear estimator involves a ratio of two noisy quantities, which produces a more general Cauchy distribution. Our data suggest this in fact to be the case. As shown in Figure 10, the distribution of RRG shear estimators does indeed have a pronounced peak and wings. The kurtosis as a function of pixel scale is shown explicitly in Figure 11 (the rms is shown in Fig. 8). The kurtosis shows that the distribution has a sharper peak and larger wings than a Gaussian. There does not appear to be any significant increase in the kurtosis at larger pixel scales. The skewness of the distributions is also consistent with zero at all pixel scales, suggesting that the distribution is symmetric, as expected.

3.5. Cosmological Implications

As we have already seen, the *quality* of shear measurements is always improved with small pixel scales. However, for a mission with a fixed lifetime, larger pixels would allow a linear increase in the total survey volume and a corresponding decrease in sample (or “cosmic”) variance errors. These two effects combine in a measurement of the dark matter power spectrum from cosmic shear, which would have a total statistical

error (cf. Refregier et al. 2004)

$$\Delta C_l = \sqrt{\frac{2}{(2l+1)f_{\text{sky}}}} \left(C_l + \frac{\sigma_\gamma^2}{2n_g} \right), \quad (17)$$

where C_l is the power and l is a multipole. The fraction of sky surveyed is

$$f_{\text{sky}}(\theta_{\text{ccd}}) = f_0 \frac{\theta_{\text{ccd}}^2}{\theta_0^2}, \quad (18)$$

where f_0 is the baseline fraction, θ_0 the baseline pixel scale, and f_{sky} at the largest simulated pixel scale is assumed to be still smaller than the observable sky. If the sole purpose of a mission is to produce a standalone measurement of cosmological parameters, this introduces a tension between weak-lensing precision and statistical errors. In this analysis, we use the overall error ΔC_l as our figure of merit.

The measured values of the second term (i.e., in the limit of small C_l) are shown in Figure 12. The ordinate axis uses a logarithmic scale, in units of the baseline survey because the absolute values are heavily dependent on the survey area, multipole l , number of galaxies, etc. Cosmic variance dominates errors at small θ_{ccd} , and increased sample variance in shear estimators takes effect at large θ_{ccd} . The error flattens out at larger pixel scales, affording some freedom in choosing the pixel scale, in which case the smaller pixel scales are clearly preferred for reasons of caution.

Our idealized interlacing decreases the error, thanks entirely to the behavior of n_g with θ_{ccd} , indicating that larger CCD pixel scales are acceptable if perfect interlacing is possible and charge diffusion remains fixed. It is important to note that larger θ_{ccd} degrades n_g , rms shapes, and ultimately shear errors.

Increasing the survey lifetime while fixing the depth would increase the nominal survey area f_0 , changing the normalization but not the functional form of the curve in Figure 12; however, this may increase the range of angular scales probed (l) and thus change the overall behavior of equation (17) with θ_{ccd} .

4. DISCUSSION

To explore the effects of pixelation on galaxy shape measurement, we have realistically simulated weakly sheared galaxies at a range of pixel scales. We have examined two different ways in which a future weak-lensing space mission could adjust the pixel scale, and both nondithered and dithered exposure strategies, for a total of four different simulation sets. We have then applied a full shear measurement pipeline as would be used on real data. We have finally combined the observed surface density of galaxies, statistics of the shear estimators, and our privileged knowledge of the true input signal, to arrive at figures of merit for the pixel scale.

We find that, as expected, smaller pixel scales, θ_{ccd} , consis-

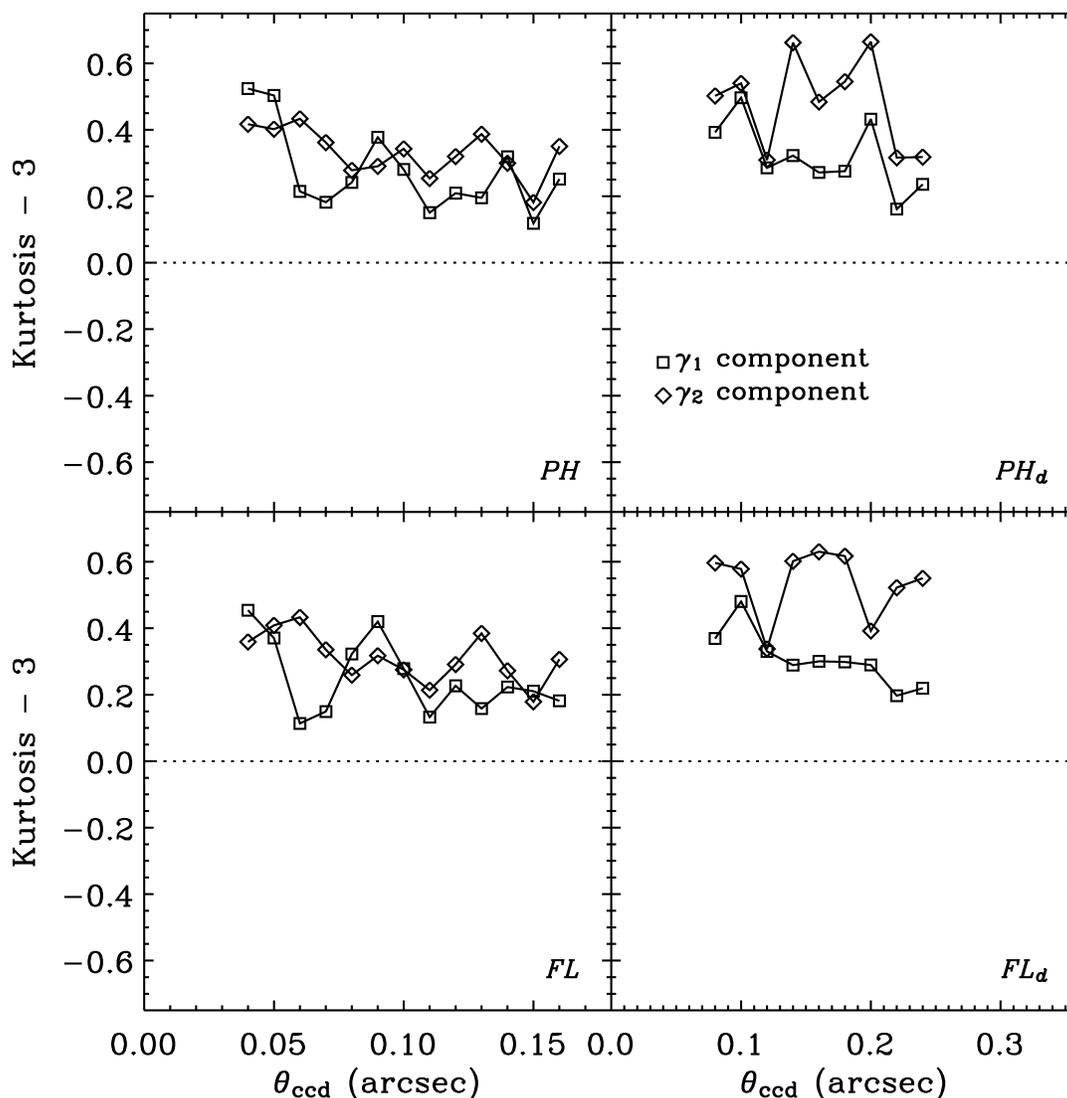


FIG. 11.—Kurtosis of shear estimators minus 3 as a function of pixel scale. The kurtosis is calculated as $\langle (\hat{\gamma}_i - \langle \hat{\gamma}_i \rangle)^4 \rangle / \sigma_{\hat{\gamma}_i}^2$, where $\sigma_{\hat{\gamma}_i}$ is the variance of $\hat{\gamma}_i$. The $\hat{\gamma}_i$ are measured from all nine images at each θ_{ccd} . Here 3 is the kurtosis of a Gaussian distribution. All data here are greater than 0, showing that the distribution of measured shear estimators has a sharper peak and more extended wings than a Gaussian.

tently improve the *quality* of shear recovery and tighten constraints on the dark matter power spectrum for a hypothetical survey of fixed area. Ideal interlacing of dithered images can provide further improvement. On the other hand, if the survey area changes with the pixel scale, larger pixel scales minimize statistical errors on the measured matter power spectrum because the survey area increases as θ_{ccd}^2 . For the baseline *SNAP* design, we find that these errors are relatively flat above $\theta_{\text{ccd}} = 0.09''$, or about 0.75–0.80 the FWHM of the PSF size. Considering situations of both fixed and variable survey area, this scale would therefore appear to be the optimal value.

We have not found a satisfactory explanation for the relative bias between measurements of γ_1 and γ_2 discussed by Massey et al. (2007b). We do see a bias, and it does vary with pixel

scale. However, the behavior is not consistent with that expected from a pure pixelation effect. The bias may arise from a combination of effects due to pixelation and the fitting of the shear susceptibility; our results are too noisy, and our full simulation pipeline too complex to distinguish these. For example, images with smaller pixel scales allow the measurement of additional, small galaxies—but these also tend to be fainter and their higher order moments in particular are noisier. The plan by the STEP collaboration,⁸ to return to basics, with analyses of a large number of identical galaxies at different pixel scales, is needed.

⁸ See <http://www.physics.ubc.ca/~heyman/step.html>.

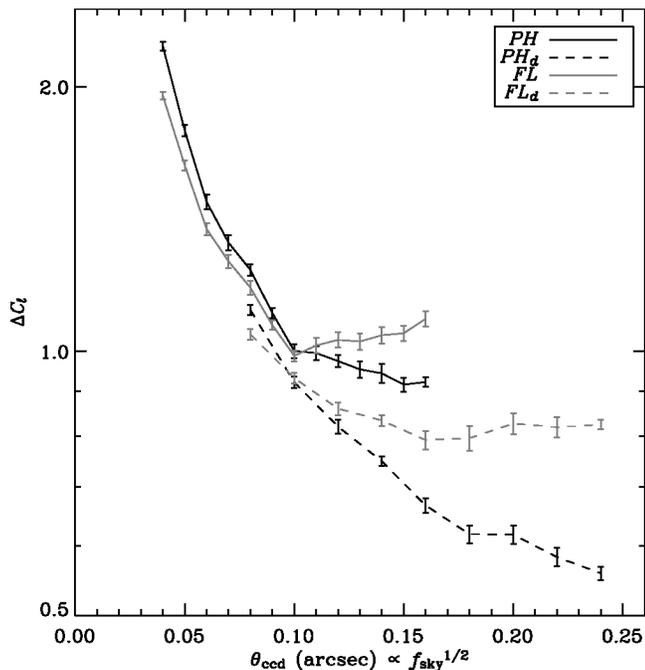


FIG. 12.—Contribution of shear sample variance to the predicted error on the matter power spectrum in single exposure (*black lines*) and dithered (*gray lines*) simulations. This assumes that the number of pixels in the focal plane is fixed when perturbing θ_{cdd} , such that the survey area scales as in eq. (18). The ordinate axis scale is \log_{10} , normalized to the PH error value at the baseline pixel scale, $\theta_{\text{cdd}} = 0.10''$.

We used a current-generation shear measurement method as the basis of our study: a snapshot of available technology. Better methods are certainly needed to fully exploit the ambitious future surveys now being planned. We can speculate that these may either be better able to cope with poor resolution or (more likely, since information is irrevocably lost during pixelation) require smaller pixels to overcome systematic floors revealed by the lowering of statistical errors. This is being suggested by the variable bias in current results, which is now well documented but poorly understood. We simply calibrate away this bias, since we know the true input shear. It even remains to be argued whether the best use of an expensive space mission would be to minimize statistical errors on a self-contained measurement of the matter power spectrum. By imaging a smaller region, but at greater resolution or depth—and with a higher density of usable galaxies—a mission could alternatively be used to obtain the phase information needed for mass maps and higher order correlation functions or even to calibrate the shape measurement of larger, ground-based surveys. As shown by our results, this approach would prefer smaller pixel scales.

Two important simplifications were imposed on our pipeline. First, we allowed no temporal or spatial variation in the PSF, and we created a comfortably large number of fake stars to characterize the PSF shape. As demonstrated in Rhodes et al. (2007), pixelation especially adds noise to peaky objects such as a diffraction-limited PSF. A typical survey region is likely to lie at high Galactic latitude. If sufficiently many bright stars cannot be imaged within the time taken for the PSF to vary, noise in the measurement of PSF shapes (which we have beaten down by simply using a large number of stars) could potentially dominate that in galaxy shapes. S. Paulin-Henriksson et al. (2008, in preparation) consider a case in which the measurement of stars rather than galaxies drives mission design. Second, real dither strategies never provide perfect interlacing. For example, optical distortions differentially alter the spacing of the pixel grid in different places, effectively causing θ_{cdd} to be a function of position in the focal plane. Consequently, Drizzling real images correlates noise between adjacent pixels because it must average nearby pixel values. Correlated noise hinders both object detection and shape measurement, so Drizzling is detrimental to weak lensing. The dithering implemented here is therefore idealized and provides a best-case scenario.

With these caveats in mind, our approach has provided a practical analysis that is achievable with existing methodology. It is complementary to studies starting from analytic first principles and assuming the existence of an idealized shape measurement method. In practice, our result on the optimum pixel scale should sensibly be considered as an upper limit, pending future developments in shape measurement methodology.

We are grateful to Steve Kent for providing ray-tracing software configured for the *SNAP* design. We also thank Alexandre Refregier, David Johnston, Matt Ferry, Gary Bernstein, Mike Jarvis, Molly Peeples, Chris Stubbs, and Adam Amara for useful discussions. The Parallel Distributed Systems Facility,⁹ a Linux cluster run by the Department of Energy's National Energy Research Scientific Computing Center, made our large-scale simulations and analysis possible: we particularly thank Iwona Sakrejda for consultation. The research described in this paper was carried out in part at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration; as well as at the California Institute of Technology main campus. This research was supported in part by a 2004 Caltech Summer Undergraduate Research Fellowship and Department of Energy grant DE-FG02-04ER41316.

⁹ <http://www.nersc.gov/nusers/resources/PDSF/>.

REFERENCES

- Albrecht, A., et al. 2006, preprint (astro-ph/0609591)
- Aldering, G., et al. 2004, *PASP*, submitted (astro-ph/0405232)
- Bacon, D. J., Refregier, A., Clowe, D., & Ellis, R. S. 2001, *MNRAS*, 325, 1065
- Bartelmann, M., & Schneider, P. 2001, *Phys. Rep.*, 340, 291
- Bernstein, G. 2002, *PASP*, 114, 98
- Bertin, E., & Arnouts, S. 1996, *A&AS*, 117, 393
- Blandford, R. D., Saust, A. B., Brainerd, T. G., & Villumsen, J. V. 1991, *MNRAS*, 251, 600
- Fruchter, A. S., & Hook, R. N. 2002, *PASP*, 114, 144
- Groom, D. E., Holland, S. E., Levi, M. E., Palaio, N. P., Perlmutter, S., Stover, R. J., & Wei, M. 2000, *Nucl. Instrum. Methods Phys. Res. A*, 442, 216
- Heymans, C., et al. 2006, *MNRAS*, 368, 1323
- Hirata, C. M., & Seljak, U. 2004, *Phys. Rev. D*, 70, 063526
- Hoekstra, H., Franx, M., Kuijken, K., & Squires, G. 1998, *ApJ*, 504, 636
- Holland, S. E., Groom, D. E., Palaio, N. P., Stover, R. J., & Wei, M. 2003, *IEEE Trans. Electron Devices*, 50, 225
- Kaiser, N. 2000, *ApJ*, 537, 555
- Kaiser, N., Squires, G., & Broadhurst, T. 1995, *ApJ*, 449, 460
- Lauer, T. R. 1999a, *PASP*, 111, 227
- . 1999b, *PASP*, 111, 1434
- Leauthaud, A., et al. 2007, *ApJS*, 172, 219
- Massey, R., & Refregier, A. 2005, *MNRAS*, 363, 197
- Massey, R., Rowe, B., Refregier, A., Bacon, D. J., & Berge, J. 2007, *MNRAS*, 380, 229
- Massey, R., et al. 2004a, *MNRAS*, 348, 214
- . 2004b, *AJ*, 127, 3089
- . 2007a, *ApJS*, 172, 239
- . 2007b, *MNRAS*, 376, 13
- Mellier, Y. 1999, *ARA&A*, 37, 127
- Narayan, R., & Bartelmann, M. 1996, preprint (astro-ph/9606001)
- Refregier, A. 2003a, *MNRAS*, 338, 35
- . 2003b, *ARA&A*, 41, 645
- Refregier, A., Rhodes, J., & Groth, E. J. 2002, *ApJ*, 572, L131
- Refregier, A., et al. 2004, *AJ*, 127, 3102
- . 2006, *Proc. SPIE*, 6265, 62651Y
- Rhodes, J., Refregier, A., & Groth, E. J. 2000, *ApJ*, 536, 79
- . 2001, *ApJ*, 552, L85
- Rhodes, J., et al. 2004, *ApJ*, 605, 29
- . 2007, *ApJS*, 172, 203
- Stover, R. J., et al. 2000, in *Optical Detectors for Astronomy II: State-of-the-Art at the Turn of the Millenium*, ed. P. Amico & J. W. Beletic (Dordrecht: Kluwer), 239
- van Waerbeke, L., White, M., Hoekstra, H., & Heymans, C. 2006, *Astropart. Phys.*, 26, 91
- Williams, R. E., et al. 1996, *AJ*, 112, 1335
- . 1998, *BAAS*, 30, 1366