WEAK GRAVITATIONAL LENSING WITH COSMOS: GALAXY SELECTION AND SHAPE MEASUREMENTS

ALEXIE LEAUTHAUD,[1] RICHARD MASSEY,[2] JEAN-PAUL KNEIB,[1,2] JASON RHODES,[2,3] DAVID E. JOHNSTON,[3] PETER CAPAK,[2]
CATHERINE HEYMANS,[4] RICHARD S. ELLIS,[2] ANTON M. KOEKEMOER,[5] OLIVER LE FÈVRE,[1] YANNICK MELLIER,[6,7]
ALEXANDRE RÉFRÉGIER,[8] ANNIE C. ROBIN,[9] NICK SCOVILLE,[2] LIDIA TASCA,[1]
JAMES E. TAYLOR,[10] AND LUDOVIC VAN WAERBEKE[3]

Received 2006 September 21; accepted 2007 February 11

ABSTRACT

With a primary goal of conducting precision weak-lensing measurements from space, the COSMOS survey has imaged the largest contiguous area observed by *Hubble Space Telescope* to date, using the Advanced Camera for Surveys (ACS). This is the first paper in a series in which we describe our strategy for addressing the various technical challenges in the production of weak-lensing measurements from COSMOS data. We first construct a source catalog from 575 ACS/WFC tiles (1.64 deg^2) subsampled at a pixel scale of 0.03″. Defects and diffraction spikes are carefully removed, leaving a total of 1.2×10^6 objects to a limiting magnitude of F814W = 26.5. This catalog is made publicly available. Multiwavelength follow-up observations of the COSMOS field provide photometric redshifts for 73% of the source galaxies in the lensing catalog. We analyze and discuss the COSMOS redshift distribution and show broad agreement with other surveys to $z \sim 1$. Our next step is to measure the shapes of galaxies and correct them for the distortion induced by the time-varying ACS point-spread function and for charge transfer efficiency (CTE) effects. Simulated images are used to derive the shear susceptibility factors that are necessary in transforming shape measurements into unbiased shear estimators. For every galaxy we derive a shape measurement error and utilize this quantity to extract the intrinsic shape noise of the galaxy sample. Interestingly, our results indicate that intrinsic shape noise varies little with size, magnitude, or redshift. Representing a number density of 66 galaxies per arcmin^2, the final COSMOS weak-lensing catalog contains 3.9×10^5 galaxies with accurate shape measurements. The properties of the COSMOS weak-lensing catalog described throughout this paper will provide key input numbers for the preparation and design of next-generation wide field space missions.

Subject headings: cosmology: observations — gravitational lensing — large-scale structure of universe

Online material: color figures

1. INTRODUCTION

As we look toward distant galaxies, fluctuations in the intervening mass distribution cause a slight, coherent distortion of their intrinsic shapes. This effect, known as weak gravitational lensing, has been used for more than a decade to probe the cosmography and the growth of structure (for a review, see Bartelmann & Schneider 2001). Although technically challenging because the weak-lensing signal is minuscule and buried in a considerable amount of noise, this field has shown substantial progress due to the advent of high-resolution space-based imaging, the proliferation of wide-field multicolor surveys, and a determined effort to improve image analysis methods and minimize systematic errors through the Shear Testing Program (STEP; Heymans et al. 2006; Massey et al. 2007b).

Historically first observed only around cluster cores (Tyson et al. 1990), weak lensing has emerged as a versatile and effective technique to probe the mass distribution of clusters (e.g., Kneib et al. 2003), to measure the clustering of dark matter around galaxies ensembles (e.g., Natarajan et al. 1998; Hoekstra et al. 2004; Sheldon et al. 2004; Mandelbaum et al. 2006), and to put constraints on the matter density parameter Ω_m and the amplitude of the matter power spectrum σ_8 (e.g., Hoekstra et al. 2006; Semboloni et al. 2006; Schrabback et al. 2006; Hetterscheidt et al. 2006; Jarvis et al. 2006). Most applications currently only use the first-order deformation induced by the mass distribution, but novel techniques are under development to take into account second-order deformations (also called flexion; Goldberg & Bacon 2005; Bacon et al. 2006; Okura et al. 2006; Goldberg & Leonard 2006) and may prove to be more efficient probes of compact structures such as galaxies and groups of galaxies. Weak-lensing measurements are particularly powerful when combined with the knowledge of the three-dimensional (3D) galaxy distribution. Sophisticated lensing tomography techniques that utilize redshifts to analyze the 3D shear field are a sensitive probe of the growth of structure and the equation of state of dark energy (Jain & Taylor 2003; Bernstein & Jain 2004; Bacon et al. 2005). Applied in many different ways, weak-lensing techniques unravel the mass distribution of structures and their evolution in the universe.

High-quality measurements of weak shear depend on the accurate determination of the shapes and redshifts of distant, faint

[1] Laboratoire d'Astrophysique de Marseille, BP 8, Traverse du Siphon, 13376 Marseille Cedex 12, France; alexie.leauthaud@oamp.fr.

[2] California Institute of Technology, MC 105-24, 1200 East California Boulevard, Pasadena, CA 91125.

[3] Jet Propulsion Laboratory, Pasadena, CA 91109.

[4] Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada.

[5] Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218.

[6] Institut d'Astrophysique de Paris, UMR7095 CNRS, Université Pierre et Marie Curie-Paris, 98 bis bd Arago, 75014 Paris, France.

[7] Observatoire de Paris, LERMA, 61, avenue de l'Observatoire, 75014 Paris, France.

[8] Service d'Astrophysique, CEA/Saclay, 91191 Gif-sur-Yvette, France.

[9] Observatoire de Besançon, BP1615, 25010 Besançon Cedex, France.

[10] Department of Physics and Astronomy, University of Waterloo, 200 University Avenue West, Waterloo, ON N2L 3G1, Canada.

galaxies. The COSMOS[11] program has imaged the largest contiguous area (1.64 deg^2) with the *Hubble Space Telescope* (*HST*) to date using the Advanced Camera for Surveys (ACS) Wide Field Channel (WFC). The full width half-maximum (FWHM) of the point-spread function (PSF) of the ACS/WFC is 0.12″ at the detector,[12] yielding a much better resolution of small galaxies than ground-based surveys, which are typically limited by the seeing to a PSF of FWHM ~1″. Shape measurements also benefit from ACS/WFC imaging compared to ground-based observations, because smaller corrections are required for the PSF and the shear measurements are less diluted by PSF smearing. The imaging quality and unprecedented area of the COSMOS ACS/WFC data, combined with extensive follow-up observations at other wavelengths (Scoville et al. 2007; Koekemoer 2007) to provide accurate photometric redshifts (Mobasher et al. 2007), make COSMOS a unique data set for weak-lensing studies.

To exploit the weak-lensing potential of the COSMOS ACS/WFC data, a carefully designed catalog of resolved galaxies with shape measurements must be extracted from the imaging data. The challenges and requirements of such a catalog are the following. First, the large survey size makes a robust automation of catalog generation essential. Second, the lensing sensitivity increases with the number density of resolved faint galaxies. Thus, it is important to detect all galaxies to faint magnitudes while taking care to minimize spurious detections that will add noise to weak-lensing measurements. Third, the high spatial resolution of the ACS/WFC allows for an excellent separation of close pairs. Given accurate deblending and a high number density of galaxies, one can expect to measure shear statistically on subarcminute scales, where baryonic physics may begin to influence the dark matter distribution. As the COSMOS ACS/WFC data set is likely to be the only large space-based lensing survey until the launch of next-generation wide field space missions such as *SNAP*[13] or *DUNE* (Réfrégier et al. 2006), the knowledge acquired through the COSMOS data will be unique and of crucial importance for the preparation and design of these future missions.

This is the first paper of a series describing the galaxy selection, galaxy shape measurement, lensing analysis, and cosmological interpretation of COSMOS data. Details regarding PSF corrections as well as tests for systematic effects are presented in the second paper of this series (Rhodes et al. 2007). A 3D cosmic shear analysis is presented in the third paper of this series (Massey et al. 2007c). Finally, a fourth paper presents high-resolution dark matter mass maps of the COSMOS field (Massey et al. 2007a).

In this paper, we describe our methods for constructing a galaxy catalog from the ACS/WFC data, to be used in subsequent weak-lensing work with COSMOS. Our goal is to produce a catalog of galaxies with photometric redshifts and PSF-corrected shape measurements, free of contaminating stars, cosmic rays,

Fig. 1.—Date of observation for each of the survey pointings. The PSF of the ACS/WFC varies on timescales that are much shorter than the period over which COSMOS was observed. [*See the electronic edition of the Supplement for a color version of this figure.*]

diffraction spikes, and other artifacts. The paper is organized as follows. In § 2 we present the data. In § 3 we describe the pipeline that locates and measures the properties of all detected objects. In § 4 we assess the quality of the data and analyze the COSMOS redshift distribution. In § 5 and in § 6 we present the PSF and charge transfer efficiency (CTE) correction schemes, the shape and shear measurement methods, and our selection criteria for the final lensing catalog. In § 7 we extract the intrinsic shape noise of the galaxy sample as a function of redshift and discuss the implications for future weak-lensing surveys. Where necessary, we assume a standard cosmological model with $\Omega_M = 0.3$, $\Omega_\Lambda = 0.7$, $H_0 = 100\ h$ km s^{-1} Mpc^{-1} and $h = 0.7$.

2. THE COSMOS ACS DATA

The COSMOS *HST* ACS field (Scoville et al. 2007; Koekemoer 2007) is a contiguous 1.64 deg^2, centered at $10^h00^m28.6^s$, $+02°12'21.0''$ (J2000.0). Between 2003 October and 2005 June (*HST* Cycles 12 and 13), the region was completely tiled by 575 adjacent and slightly overlapping pointings of the ACS/WFC (see Fig. 1). Images were taken through the wide F814W filter (broad *I*). The camera has a 203″ × 203″ field of view, covered by two 4096 × 2048 CCD chips with a native pixel scale of 0.05″ (Ford et al. 2003). The median exposure depth across the field is 2028 s (one *HST* orbit). At each pointing, four 507 s exposures were taken, each dithered by 0.25″ in the *x* direction and 3.08″ in the *y* direction from the previous position. This strategy ensures that the 3″ gap between the two chips is covered by at least three exposures and facilitates the removal of cosmic rays. Pointings were taken with two approximately 180° opposed orientation angles (PA_V3 = 100° ± 10° and 290° ± 10°). In this paper we use the "unrotated" images (as opposed to north up) to avoid rotating the original frame of the PSF. By keeping the images in the default unrotated detector frame, they can be stacked to map out the observed PSF patterns. For similar reasons, we perform detection in
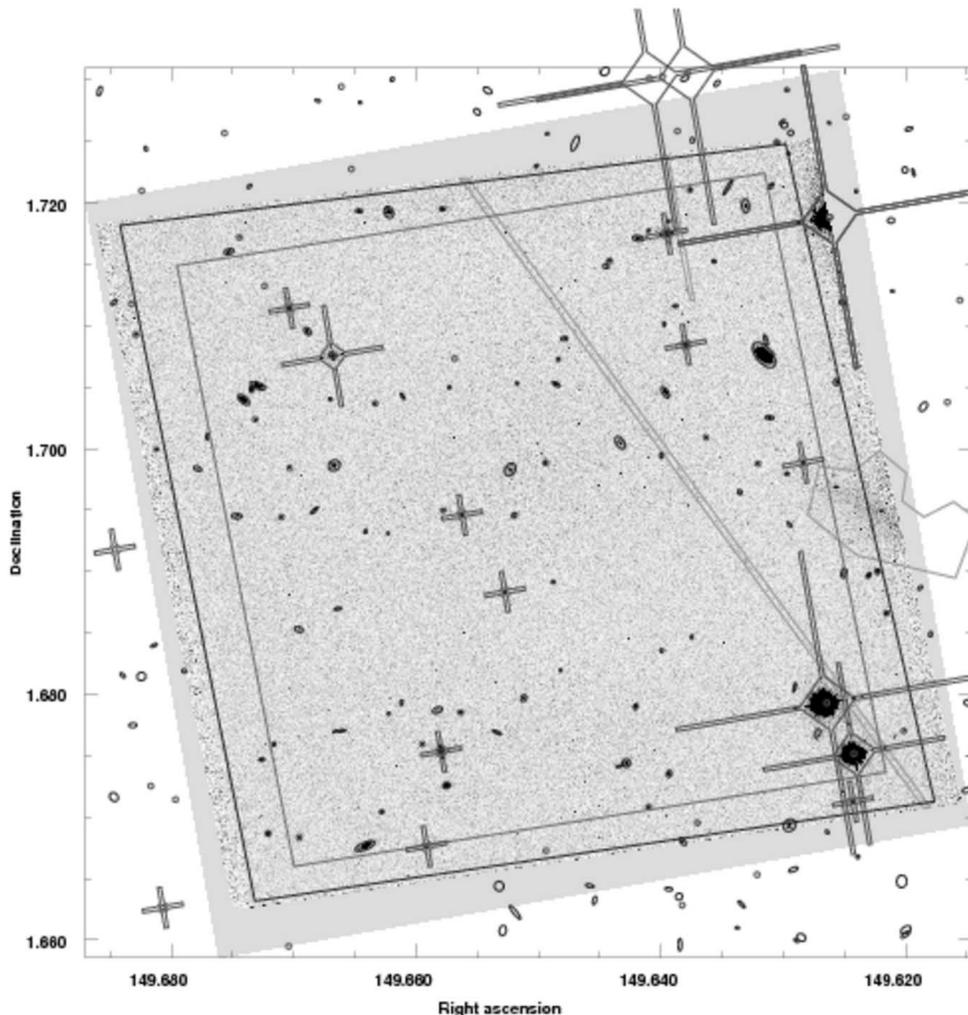
Fɪɢ. 2.—COSMOS pointing acs_I_095836+0141_unrot_sci_12.fits with bright detections, masking, and edge definition indicated. Adjacent images have a sizable overlap (*smaller box*), which allows us to discard detections on the boundaries of each tile (*larger box*) without losing any objects in the final concatenated catalog. The automated masking of the diffraction spikes around bright stars (F814W < 23) is pictured here by the polygons. The basic shape of the star masks is predefined and then scaled with the magnitude of the star. The rectangles correspond to the manual masking of asteroid trails and various other image defects. Toward the right edge of this image, the dwarf galaxy L1-099 (Impey et al. 1996) is identified and flagged in a special category. Bright galaxies with F814W < 23 are depicted by ellipses and bright stars with F814W < 23 by circles. [*See the electronic edition of the Supplement for a color version of this figure.*]

individual ACS/WFC tiles instead of on a larger mosaic (where the orientation of the PSF frame would be unknown). Figure 2 shows a COSMOS ACS pointing with the bright detections (F814W < 23) and the masking of stars, asteroid trails, and image defects (see § 3.5).

To build our catalog, we use version 1.3 of the unrotated ACS/ WFC data, which has been specially reduced for lensing purposes (see Koekemoer 2007 for technical details). Image registration, geometric distortion, sky subtraction, cosmic-ray rejection, and the final combination of the dithered images were performed by the MultiDrizzle algorithm (Koekemoer et al. 2002). As described in (Rhodes et al. 2007), the MultiDrizzle parameters have been chosen for precise galaxy shape measurement in the co-added images. In particular, a finer pixel scale of $0.03''$ pixel^{-1} was used for the final co-added images (7000×7000 pixels), even though this implies more strongly correlated pixel noise (see § 3.8). Hereafter, when we refer to pixels, we will assume a pixel scale of $0.03''$ pixel^{-1}. Pixelization acts as a convolution followed by a resampling, and although current shear measurement methods can successfully correct for convolution, the formalism to properly treat resampling is still under development for the next generation of

methods. Again following the recommendations of Rhodes et al. (2007), a Gaussian and isotropic MultiDrizzle convolution kernel was used, with scale = 0.6 and pixfrac = 0.8, small enough to avoid smearing the object unnecessarily while large enough to guarantee that the convolution dominates the resampling. This process is then properly corrected by existing shear measurement methods.

The ACS/WFC CCDs also suffer from imperfect charge transfer efficiency (CTE) during readout. As charges are transferred during the readout process, a certain fraction are retained by charge traps (created by cosmic-ray hits) in the pixels. This causes flux to be trailed behind objects as the traps gradually release their charge, spuriously elongating them in a coherent direction that mimics a lensing signal. Since this effect is produced by a fixed number of charge traps within the CCD substrate, it affects faint sources (with a larger fraction of their flux being trailed) more than bright ones. This is an insidious effect that mimics an increasing shear signal as a function of redshift, and prevents the use of the traditional way of dealing with the calibration of faint galaxies in a lensing analysis by looking at bright stars. Ideally this effect would be corrected for on a pixel-by-pixel basis in the raw images; unfortunately our

current physical understanding of this effect is insufficient and a more in-depth analysis is still underway. The CTE effect can be quantified sufficiently well, however, that in a first step we can adopt a postprocessing correction scheme based on an object's position, flux, and date of observation. Further details regarding this model can be found in Rhodes et al. (2007).

3. THE COSMOS ACS GALAXY CATALOG

In this section, we discuss the construction of the COSMOS ACS/WFC source catalog. This catalog is carefully cleaned of defects and artifacts and is made publicly available through the Infrared Science Archive (IRSA) database.[14]

3.1. Detection Strategy

We use version 2.4.3 of the SExtractor photometry package (Bertin & Arnouts 1996) to extract a source catalog of positions and various photometric parameters. In the construction of this catalog, our main concern is to pick out the small faint objects that contain most of the lensing signal. The detection strategy that we therefore adopt is to configure SExtractor with very low thresholds (even if this leads to more false detections in the catalog) and to control our sample selection via subsequent "lensing cuts" (see § 6). We hope to thus reduce unknown selection biases introduced by the SExtractor detection algorithm. When configured with low detection thresholds however, SExtractor also inevitably (1) overdeblends low surface brightness spirals and patchy irregulars, (2) deblends the outer features of bright galaxies, (3) detects spurious objects in the scattered light around bright objects, and (4) underdeblends close pairs.

The correct detection of close pairs enables lensing measurements on very small scales. However, overdeblending and spurious detections add noise to these measurements. In particular, false detections around bright objects can have quite high signal-to-noise ratio (S/N) values and are not trivial to remove with lensing cuts (see § 6). The method presented here is a partial solution to (2), (3), and (4). The overdeblending of low surface brightness and patchy galaxies remains a difficult problem, however, especially for high-resolution imaging, and calls for an improvement of existing detection algorithms. With the advent of high-resolution multiwavelength surveys, a possible solution would be to incorporate color and morphological information into the detection process (e.g., Lupton et al. 2001).

While this overdeblending problem persists in our catalog, it affects less than 1% of the objects. This problem is furthermore mitigated by the centroiding process during the shape measurement stage. Indeed, objects for which the centroid algorithm fails to converge, which will often be the case for overdeblended features, are discarded from the catalog. To remedy the remaining problems (2), (3), and (4), we adopt and improve the method (known as the "Hot-Cold" method) employed in Rix et al. (2004). In this method, we run SExtractor twice, once with a configuration optimized for the detection of only the brightest objects ("cold" step) and then again with a configuration optimized for the faint objects ("hot" step). This double extraction helps improve the detection of close pairs. The two samples are then merged together to form the final catalog, and masks are created around the bright detections, minimizing the effects of (3) and (4).

For the hot and cold steps, we vary four main parameters to optimize the detection: (1) DETECT_THRESHOLD, the minimum S/N per pixel above the background level; (2) MIN_AREA, the number of contiguous pixels exceeding this threshold; (3) BACK_

TABLE 1
SExtractor Configuration Parameters

Parameter	Bright Objects	Faint Objects
DETECT_MINAREA	140	18
DETECT_THRESH	2.2[a]	1.0
DEBLEND_NTHRESH	64	64
DEBLEND_MINCONT	0.04	0.065
CLEAN_PARAM	1.0	1.0
BACK_SIZE	400	100
BACK_FILTERSIZE....................	5	3
BACKPHOTO_TYPE	Local	Local
BACKPHOTO_THICK	200	200

[a] Because of correlated noise (see § 3.8), the effective threshold levels are DETECT_THRESH ∼ 1.25 for bright objects and DETECT_THRESH ∼ 0.57 for faint objects.

SIZE, the mesh size of the background map; and (4) DEBLEND_NTHRES and DEBLEND_MINCONT, the parameters regulating deblending. In both cases, the data are filtered prior to detection by a 5 pixel (0.15″) Gaussian filtering kernel. Our choice of parameters for both steps is provided in Table 1.

The two-step method also allows one to adjust the estimation of the background map according to the typical size of objects one expects to detect, improving detections with SExtractor. The background map is constructed by computing an estimator for the local background on a grid of mesh size BACK_SIZE. We adjust BACK_SIZE so as to capture the small-scale variations of the background noise while keeping it large enough not to be affected by the presence of objects.

For each exposure, a weight map is produced by MultiDrizzle, describing the combined noise properties of the readout, the dark current, and the sky background (Koekemoer 2007). These maps describe the noise intensity at each pixel and are used to account for the spatial-dependent noise pattern in the co-added image with the SExtractor WEIGHT_IMAGE option set to WEIGHT MAP.

Each ACS/WFC pointing consists of four slightly offset, dithered exposures, making cosmic-ray rejection more difficult and detection more unreliable on the boundaries of each tile, where there are fewer than four input exposures. Because adjacent images overlap sufficiently, we can trim the edges of the images without actually removing data (see Fig. 2).

3.2. Bright Object Detection

In the first step, we detect only the brightest and largest objects in the image, with 140 or more contiguous pixels (corresponding to a diameter of 0.4″ for a circular object) rising more than 2.2 σ per pixel above the background level. The BACK_SIZE parameter is set to 12″, or 30 times the diameter of the smallest objects detected. The detection threshold and the deblending parameters DEBLEND_NTHRES and DEBLEND_MINCONT are calibrated heuristically on several images to separate close pairs as much as possible without deblending patchy, extended spiral galaxies. Because faint objects are captured in a second run, we are free to choose the value of the detection threshold during this step. We found that this flexibility greatly helped to calibrate the parameters that optimize the deblending. Indeed, if DETECT_THRESHOLD is set to a low value (for example, to detect faint objects), close pairs will be detected as a single object and are difficult to deblend. Figure 3 illustrates the improvements of this two-step method compared to a single-step method. During this first step, all pixels associated with a detection are recorded by SExtractor in an image called a segmentation map. These segmentation maps are used at a
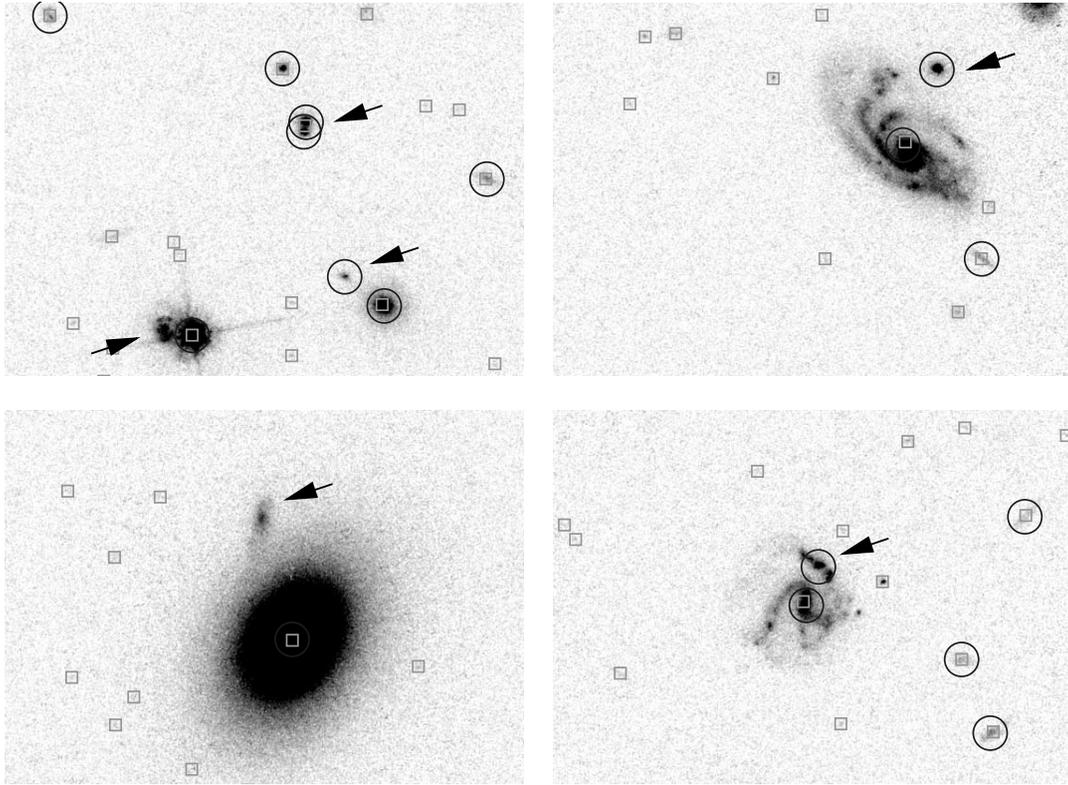
Fɪɢ. 3.—This figure illustrates the difficulty of correctly deblending close pairs while keeping patchy spirals with strong star-forming regions intact. Squares indicate detections from the faint step and circles indicate detections from the bright step. The top two panels show three objects that are not detected or incorrectly deblended by the faint step but that are picked by two-step method. The bottom panel and the arrow toward the star in the upper left panel show that even with this method, a perfect configuration is still difficult to reach. [*See the electronic edition of the Supplement for a color version of this figure.*]

later stage to merge the bright and the faint catalogs (see § 3.4). This first catalog of bright objects is referred to as $\mathcal{C}_{\mathrm{cold}}$.

3.3. *Faint Object Detection*

In the second step, we configure SExtractor to pick up the small, faint objects, taking care to choose the detection parameters to be less conservative than any subsequent lensing cuts (see § 6). The MIN_AREA parameter is set to 18 pixels (corresponding to a diameter of 1.2 times the FWHM of the PSF) and the detection threshold is 1 σ above the background level. As objects detected at this step are smaller, the background estimation can be improved by refining the mesh size of the background map and setting BACK_SIZE to 100 pixels, or 20 times the diameter of the smallest objects detected. This second catalog is referred to hereafter as $\mathcal{C}_{\mathrm{hot}}$.

3.4. *Merging the Two Samples*

The final catalog is obtained by merging the detections from $\mathcal{C}_{\mathrm{cold}}$ and $\mathcal{C}_{\mathrm{hot}}$, keeping all objects in $\mathcal{C}_{\mathrm{cold}}$ and only the objects from $\mathcal{C}_{\mathrm{hot}}$ not detected in $\mathcal{C}_{\mathrm{cold}}$. To determine which objects to discard from $\mathcal{C}_{\mathrm{hot}}$, we use the segmentation maps created during the bright detection step. To begin with, we enlarge the flagged areas in these segmentation maps by approximately 20 pixels (0.6″). We then discard all objects from $\mathcal{C}_{\mathrm{hot}}$ for which the central pixel lies within a flagged area of these maps. Thus, we remove duplicate detections and create a mask around all bright objects, immediately cleaning the catalog of a certain number of spurious detections. By a visual inspection of the data, we estimate that this method solves about half of the deblending problems that we observe (excluding the low surface brightness galaxies). The final catalog

of raw SExtractor detections $\mathcal{C}_1$ contains 1.8 million objects in total (see Table 2).

3.5. *Cleaning the Catalog*

Great care was taken to mask unreliable regions within images and to remove false detections from the catalog, especially those that can mimic a lensing signal. As illustrated in Figure 2, an automatic algorithm was developed to define polygonal-shaped masks around stars with F814W < 19 (the limit at which stars saturate in the COSMOS images), with a size scaled by the magnitude of the star. Objects near bright stars or saturated pixels were masked to avoid shape biases due to any background gradient. All the images were then visually inspected. In a few cases the automatic algorithm failed (very saturated stars for which the centroid of the star is widely offset) and the stellar masks were corrected by hand. Other contaminated regions of the images were also masked out, including reflection ghosts, asteroids, and satellite trails. Astronomical sources such as H ɪɪ regions around bright galaxies, stellar clusters, and nearby dwarf galaxies, were also flagged and removed from the lensing catalog.

Objects with double entries in the catalog (from the overlap between adjacent images) are identified and the counterpart with the highest SExtractor flag (indicating a poor detection) is discarded, leaving a catalog of unique objects. However, the duplicated objects from overlapping regions are a valuable asset for consistency checks and are used, for example, to check the galaxy shape measurement error (see § 7 and Fig. 17).

The final clean catalog ($\mathcal{C}_2$) is free of spurious or duplicate detections and contains 1.2 million sources in total (see Table 2).

TABLE 2
SUMMARY OF THE CONSTRUCTION OF THE ACS LENSING CATALOG

Parameter	Number	Percent of $\mathcal{C}_x$
Catalog of Raw SExtractor Detections: $\mathcal{C}_1$		
Total number of objects in $\mathcal{C}_1$	1.8×10^6	...
Number of Hot (faint) detections...	1.6×10^6	88
Number of Cold (bright) detections...	2.2×10^5	12
Details of the Cleaning Process		
Number of objects within the noisy border of a tile..	3.2×10^5	17
Number of Hot detections with central pixel in Cold segmentation map	2.0×10^5	11
Number of objects within automatically defined star masks..	2.4×10^4	1
Number of objects within manually defined masks	4.1×10^4	2
Number of objects detected more than once in adjacent tiles..	6.6×10^4	4
Catalog Cleaned of Image Defects: $\mathcal{C}_2$		
Total number of objects in $\mathcal{C}_2$	1.2×10^6	...
Number of galaxies (ID = 1)...	1.1×10^6	96
Number of point sources (ID = 2)...	2.8×10^4	2
Number of fake detections (ID = 3)...	1.7×10^4	2
ACS Galaxies from $\mathcal{C}_2$ with F814W$_{AB}$ < 26.5: $\mathcal{C}_3$		
Total number of objects in $\mathcal{C}_3$	7.0×10^5	...
Number of galaxies with a counterpart in the photometric catalog...............................	6.0×10^5	85
Number of galaxies that have been matched but that are in ground-based masks	8.3×10^4	12
Number of galaxies for which the redshift code did not converge..............................	1.1×10^4	1.7
Total number of galaxies with accurate photometric redshifts, $\mathcal{C}_3$................................	5.0×10^5	71
Final COSMOS ACS Lensing Catalog: $\mathcal{C}_4$		Galaxy Number Density
Total number of galaxies in final lensing catalog...	3.9×10^5	66 arcmin^{-2}
Number of galaxies with accurate photometric redshifts	2.8×10^5	48 arcmin^{-2}

3.6. Star-Galaxy Classification

The correct identification of stars has two implications for the lensing analysis. First, bright stars are useful for PSF modeling and second, stars must be correctly identified in order to apply our automatic masking algorithm of diffraction spikes. A robust star-galaxy classification is thus necessary.

SExtractor produces a continuous stellar classification index parameter ranging from 0 (extended sources) to 1 (point sources). This index has two drawbacks: first, the definition of the dividing line is ambiguous, and second, the neural-network classifier used by SExtractor is trained with ground-based images and is therefore only valid for a sample of profiles similar to the original training set. With space-based images, this index becomes difficult to interpret, as illustrated in Figure 4, which depicts our star selection (described below) within the CLASS_STAR/MAG_AUTO plane.

We therefore test two alternative methods to classify point sources and galaxies, one based on the SExtractor parameter MU_MAX (peak surface brightness above the background level) and the other based on the half-light radius, RHL (e.g., Peterson et al. 1979; Bardeau et al. 2005). Both methods are motivated by the fact that the light distribution of a point source scales with magnitude. Point sources therefore occupy a well-defined locus in a MU_MAX/MAG_AUTO or a RHL/MAG_AUTO plane. Figure 5 shows how we can use this property to define stars (ID = 2) and galaxies (ID = 1) reliably up to a magnitude of F814W ~ 25. At fainter levels, the classification begins to break down and the
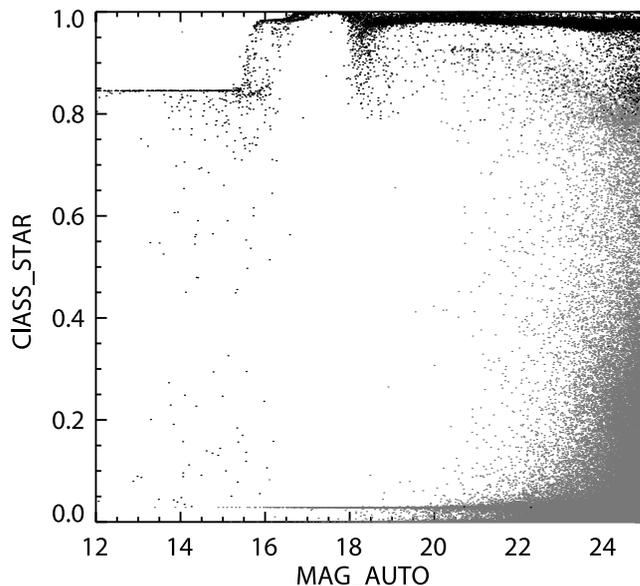
FIG. 4.—SExtractor stellar index (CLASS_STAR) for our point-source selection based on the peak surface brightness of objects. Gray points show the corresponding galaxy sample. A point-source selection of the form CLASS_STAR > 0.8 for example, will miss a certain number of bright stars (F814W < 17) and will grossly misclassify compact galaxies at 22 < F814W < 25.
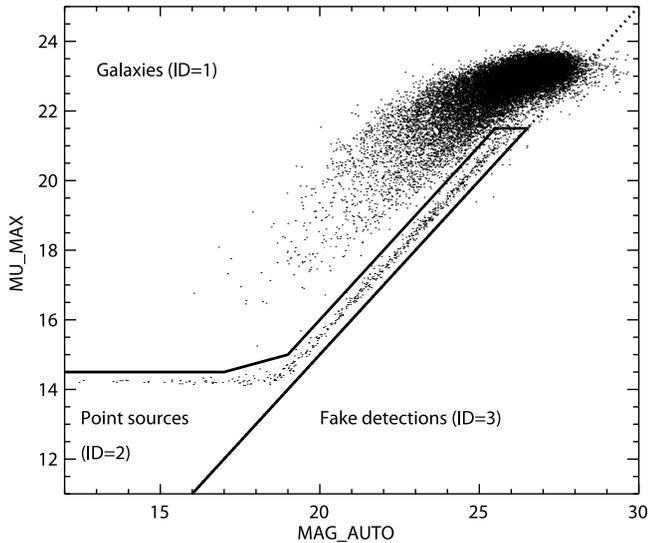
Fig. 5.—Classification of point sources, galaxies, and artifacts within the MU_MAX/MAG_AUTO plane. Point sources follow the PSF and are delimited by the solid region. Objects that are more sharply peaked than the PSF are contained in the dashed region and are considered to be artifacts. For clarity, only a 2% random selection of all objects are included in this plot.

point sources become indistinguishable from the small galaxies. We find that the two methods agree very well, within 1% at magnitudes less than F814W = 24 and within 2% at magnitudes less than F814W = 25. The small difference arises mainly from a misclassification of objects by the RHL method because of the presence of a close pair that distorts the estimation of RHL. Overall, the MU_MAX method proved to be more robust and has the advantage of a tighter correlation of the stellar locus and a clear break indicating the magnitude at which the stars saturate (MAG_AUTO ∼ 19). Moreover, with this method, a surface brightness cut at the faint end of the stellar sample is trivial to implement (at faint magnitudes, the catalog is surface brightness limited). The performance of this star galaxy separation scheme will be analyzed in more detail in § 4.1.

Using the MU_MAX method, we also define a set of objects that are more sharply peaked than the PSF, which is obviously nonphysical. A visual inspection finds that these objects are mainly artifacts, hot pixels, and residual cosmic rays. We flag these spurious objects in our catalog (ID = 3) and remove them for the lensing analysis.

Averaging over the COSMOS field, we find ∼15 stars per pointing with $19 \leq$ MAG_AUTO ≤ 23. This is an insufficient number to model the PSF in individual images using standard interpolation techniques. However, it is a sufficient number to identify the PSF pattern of each exposure given a finite set of recurring patterns (see § 5.1 and Rhodes et al. [2007] for further details).

3.7. Photometric Redshifts

In addition to the ACS/WFC (F814W) imaging, the COSMOS field has been imaged with Subaru Suprime-Cam (B_j, V_j, g^+, r^+, i^+, z^+, NB816), the Canada-French Hawaii Telescope (CFHT; u^*, i^*), and the KPNO/CTIO (K_s). Details of the ground-based observations and the data reduction are presented in Capak et al. (2007) and Taniguchi et al. (2007). Other observations were taken in the UV with *GALEX*, in the X-ray with *XMM-Newton*, and in the radio with VLA, CSO, and IRAM. Yet more observations are

underway, including intermediate and narrowband imaging with Subaru Suprime-Cam, deep Infrared imaging covering 1.0–2.2 μm (CFHT WIRCam, UKIRT WFCAM, and UH2.2 ULBcam), and observations with space-based facilities including *Chandra X-Ray Observatory* and *Spitzer Space Telescope*. This extensive multi-wavelength data set is a key component to COSMOS weak-lensing measurements because it allows us to accurately measure the COSMOS redshift distribution, to separate foreground and background structures, and to remove contamination from intrinsic galaxy alignments (Heymans & Heavens 2003) and shear-ellipticity correlations (King 2005).

Photometric redshifts were determined by the COSMOS photometric redshift code with a Bayesian prior based on luminosity functions and allowing for internal extinction (Mobasher et al. 2007). For each galaxy, the entire probability distribution $P(z)$, the most likely redshift, and a confidence level for that redshift are calculated. The knowledge of the full $P(z)$ allows us to apply weight to galaxies in weak-lensing measurements according to the uncertainty in their measured redshift. The accuracy of the photometric redshifts are estimated based on extensive simulations and by comparison with a sample of 958 galaxies with spectroscopic redshifts measured by the ESO VIMOS instrument as part as the zCOSMOS program (Lilly et al. 2007). Both the luminosity prior and the extinction corrections have been shown to improve the accuracy of the photometric redshifts when compared to the spectroscopic sample. The accuracy of the current photometric redshifts down to F814W ∼ 22.5 is

$$\sigma_{\Delta z}/(1 + z_s) = 0.031, \qquad (1)$$

with $\eta = 1.0\%$ of catastrophic errors, defined as $\Delta z/(1 + z_s) > 0.15$. This relation scales with magnitude in a similar fashion to Wolf et al. (2004). The accuracy of the photometric redshifts will continue to improve as more data become available (in particular the deeper u^*, J, K, *Spitzer*, and Subaru narrowband data).

The COSMOS optical and near-infrared catalog (Capak et al. 2007) provides multiband photometry for 89% of the COSMOS ACS/WFC galaxies. As demonstrated by Figure 6, the remaining 11% of the galaxies for which we lack multiwavelength information (and therefore photometric redshifts), are the small galaxies that cannot be detected by the ground-based imaging. In effect, although the Subaru data is deeper for sources 1″ in diameter, the ACS/WFC imaging will do a better job at detecting anything smaller. Because we apply a size cut to the final lensing catalog, however (see § 6), many of these small galaxies will be discarded from the final analysis. In total, after removing the galaxies with unreliable multiband photometry (because they are masked out in the ground-based data) as well as those for which the photometric redshift code failed to converge, 76% of the galaxies in the COSMOS ACS/WFC lensing catalog (F814W < 26) have photometric redshifts (see Fig. 6). For further discussions on the photometric redshifts and the COSMOS redshift distribution, see § 4.3.

3.8. Noise Properties

The drizzling process introduces pattern-dependent correlations between neighboring pixels and can artificially reduce the noise levels in co-added images. Noise and error estimates derived from drizzled images will thus tend to underestimate the true noise levels of the image. One should in theory take into account the exact covariance matrix of the noise in order to derive error estimates for drizzled images. For our purposes, however, a simple scaling

Fig. 7.—Significance of COSMOS detections defined as FLUX_AUTO/FLUXERR_AUTO, where FLUXERR_AUTO has been corrected for correlated noise.

As implemented by SExtractor, the formulas for the flux and magnitude uncertainties (for both AUTO and the ISO quantities) are given by

$$\text{FLUX ERR} = \sqrt{A\sigma^2 + F/g}, \tag{3}$$

$$\text{MAG ERR} = \frac{2.5}{\ln 10} \frac{\text{FLUX ERR}}{F}, \tag{4}$$

where A is the area (in pixels) over which the flux F (in ADU) is summed and g is the detector gain. To correct the magnitude and flux errors reported by SExtractor for the correlated noise, we replace σ, the standard deviation of the noise (in ADU) estimated by SExtractor, by $\sigma/\sqrt{F_A}$ within the above equations. The significance of a COSMOS detection, after this correction is applied, is defined as S/N = FLUX_AUTO/FLUXERR_AUTO (see Fig. 7).

4. QUALITY ASSESSMENT OF THE ACS CATALOG

4.1. *Galaxy and Stellar Counts*

The same MU_MAX parameter used to classify stars and galaxies can also be used as an indication of the background level and the depth of the data. For each image i we calculate the mode m_i of the MU_MAX parameter. We then divide the m_i into two bins according to the angle of the telescope with the Sun at the moment of the pointing. The histogram of m_i (Fig. 8) for these two bins reveals that the depth of the data is bimodal depending on whether the angle of the Sun is less or greater than a critical angle of $70°$. This is visually evident when we inspect the density map of the very faint objects (Fig. 8). There are 96 pointings out of 575 that have a Sun angle less than the critical value, making them slightly shallower than the average.

The number counts serve as a check of the approximate photometric calibration and the depth of the data. Figure 9 shows the counts for galaxies and stars compared to the reported HDF F814W counts (Williams et al. 1996). The magnitudes are given in the AB system. Stars have been subtracted from the galaxy counts up to F814W = 25. At fainter magnitudes their contribution is negligible. We plot raw number counts only, i.e., we do

Fig. 6.—*Top*: Number counts of galaxies that have been correctly matched to the ground-based catalog (*dotted line*) compared to the total number counts (*solid line*). The dashed line indicates galaxies for which we consider the photometric redshift to be reliable. The difference between these two curves is primarily due to larger masked areas in the ground-based data than in the ACS imaging. *Bottom*: Sizes of galaxies within the ACS catalog (*solid line*) compared to the sizes of those that have been matched with the ground-based catalog (*dashed line*). The objects for which we do not have multiband photometry are small galaxies that are detected with ACS but not with ground-based imaging (seeing $\sim 1''$). The vertical dotted line shows the approximate size cut that we make in the final lensing catalog.

of the noise level in each pixel by the same constant factor is sufficient. The scaling factor that we adopt, F_A, has been derived for MultiDrizzled images by Casertano et al. (2000). In principle, F_A is size dependent but converges rapidly with increasing size toward an asymptotic value given by

$$\sqrt{F_A} = \begin{cases} (s/p)[1 - s/(3p)] & \text{if } (s < p), \\ [1 - p/(3s)] & \text{if } (p < s), \end{cases} \tag{2}$$

where p and s are respectively the `pixfrac` and the `scale` configuration parameters of MultiDrizzle. For the purpose of this paper, we assume F_A to be constant and equal to $F_A \sim 0.316$ ($p = 0.8$ and $s = 0.6$). By assuming a constant corrective factor regardless of size, we make less than a 10% error on the noise estimation of the smallest objects detected.

FIG. 8.—*Top*: Histogram of m_i (mode of the MU_MAX parameter for each image) demonstrating that images with a Sun angle less than 70° are not quite as deep as images with a Sun angle greater than 70°. *Bottom*: Density of faint objects (27 < F814W < 28) within the COSMOS field. The circles indicate the pointings for which the angle with the Sun is less than 70°. [*See the electronic edition of the Supplement for a color version of this figure.*]

not correct for incompleteness at the faint end. To facilitate comparisons with other surveys, we fit the galaxy counts between F814W = 20 and F814W = 26 to an exponential of the form $N = B \times 10^{A \times \mathrm{mag}}$, where N has units of number deg^{-2} 0.5 mag^{-1}. For the deeper set of images (Sun angle > 70°), we find $A = 0.332$ and $\log(B) = -3.543$. The raw galaxy and stellar number counts are provided in Table 3.

We also fit the stellar counts to models as shown in Figure 9. The star count predictions have been done using the Besançon

model of the Galaxy (Robin et al. 2003, 2004) and are described in detail in Robin et al. (2007). By extrapolating the stellar number counts, we estimate that the galaxy catalog has less than a 3% contamination from stars at magnitudes greater than 25. The fit to stellar models is excellent between F814W = 20 and F814W = 25, and at magnitudes less than 19 we visually inspect the catalog to check that the star selection is correct to within 0.5%. This small error arises mainly from false detections by SExtractor of the diffraction spikes of bright, saturated stars.
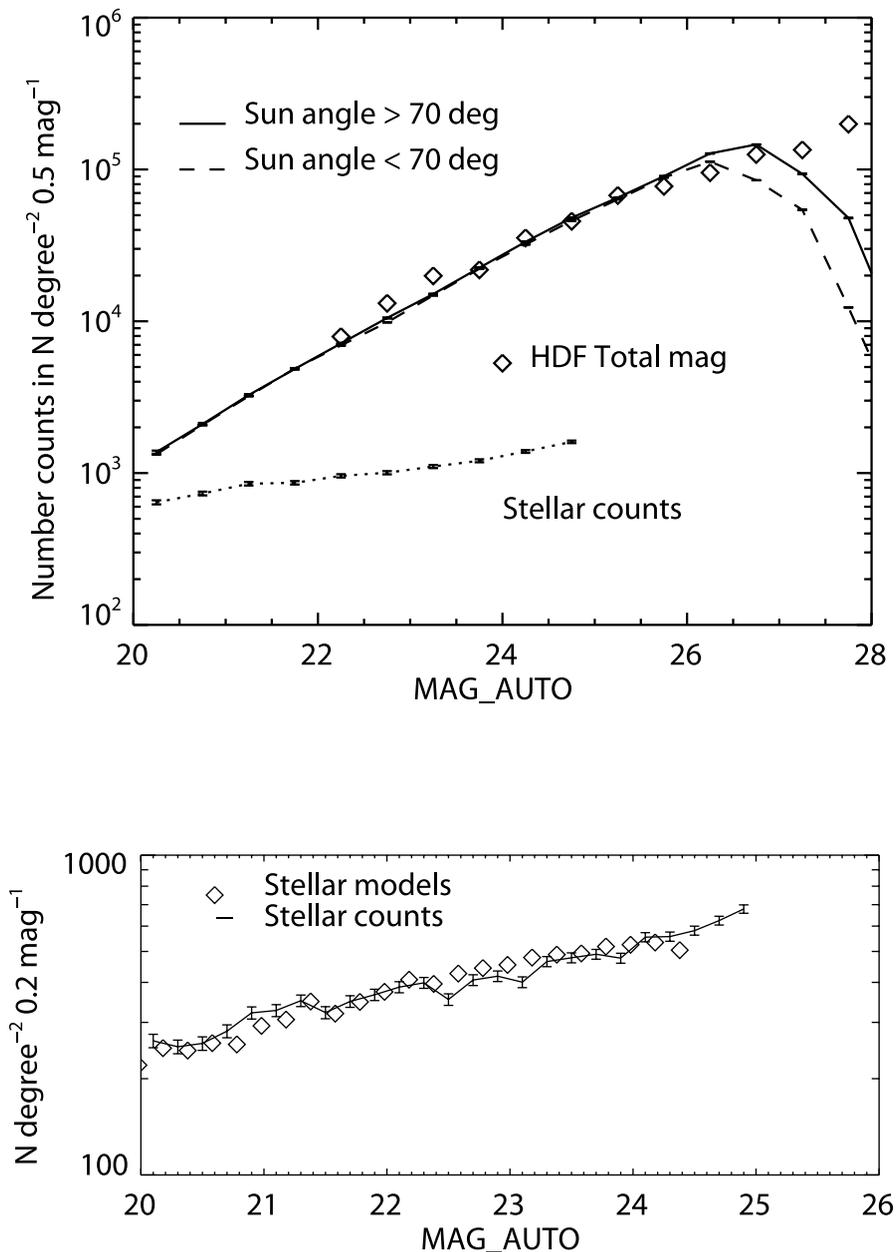
Fig. 9.—*Top*: Galaxy and stellar number counts as compared to the HDF. The dashed curve corresponds to images with a Sun angle of less than 70° and the solid curve corresponds to images with a Sun angle greater than 70°. Poisson error bars are also indicated but are very small. *Bottom*: Point-source selection for the catalog compared to stellar models computed from evolutionary tracks and constrained by local *Hipparcos* data.

4.2. *Completeness*

The probability that a galaxy enters our catalog will depend on its size and surface brightness profile. To quantify the completeness and detection limits of our SExtractor configuration, we insert fake objects with a Gaussian profile of varying FWHM and total magnitude into empty regions of an ACS image and test how well these objects can be recovered with our pipeline. Each artificial source was considered to be correctly detected if its centroid was within 10 pixels and the MAG_AUTO parameter was within 0.5 mag of the input value. From this analysis, we determine our completeness as a function of magnitude and FWHM and the results are shown in Figure 10. The completeness is about 90% for objects with a FWHM of 0.2″ at F814W = 26.6. These values should only be used as a rough estimate, however, as we do not actually model galaxies, but use a simple Gaussian profile for artificial objects.

4.3. *The COSMOS Redshift Distribution*

The estimation of the redshifts of galaxies is the major astrophysical uncertainty inherent to weak-lensing methods. To first approximation, cosmic shear and tomography are mainly sensitive to the median redshift of the sources, while galaxy-galaxy lensing benefits greatly from the knowledge of precise spectroscopic redshifts for the foreground lenses (Kleinheinrich et al. 2005). With the depth and area coverage of COSMOS, we surpass the current capability for complete spectroscopic followup. For most forthcoming weak-lensing surveys, this will also be the case, hence the importance of the photometric redshift technique to measure redshifts for a majority of the galaxies, to and beyond today's spectroscopic limits (Ilbert et al. 2006). COSMOS presents a unique advantage, in terms of current weak-lensing surveys, of a prodigious multiwavelength followup combined with the planned

TABLE 3
Cosmos F814W Galaxy and Stellar Number Counts

F814W	Galaxy Density [$\log(n)$ deg^{-2} 0.5 mag^{-1}]	Stellar Density [$\log(n)$ deg^{-2} 0.5 mag^{-1}]
20.25....................	3.138 ± 1.503	2.807 ± 1.297
20.75....................	3.323 ± 1.600	2.865 ± 1.329
21.25....................	3.514 ± 1.691	2.930 ± 1.361
21.75....................	3.686 ± 1.777	2.936 ± 1.365
22.25....................	3.853 ± 1.860	2.981 ± 1.385
22.75....................	4.022 ± 1.945	3.003 ± 1.401
23.25....................	4.180 ± 2.023	3.043 ± 1.414
23.75....................	4.352 ± 2.110	3.081 ± 1.437
24.25....................	4.523 ± 2.196	3.143 ± 1.463
24.75....................	4.682 ± 2.275	3.206 ± 1.493
25.25....................	4.814 ± 2.340	$\dots$
25.75....................	4.956 ± 2.412	$\dots$

Notes.—Galaxy counts are derived for the 479 images with a Sun angle greater than 70°. Magnitudes are the SExtractor MAG_AUTO.

measurement of ~50,000 spectroscopic redshifts by the ongoing zCOSMOS program (Lilly et al. 2007). On completion, this data set will provide the COSMOS lensing catalog with accurate photometric redshifts and will be vital for refining and improving the photometric technique in preparation for forthcoming weak-lensing surveys. We present here a first analysis of the COSMOS redshift distribution. A more detailed study of the systematic trends in the photometric redshifts and of their effects on the redshift distribution is beyond the scope of this paper and will be addressed elsewhere when more data become available. For the purposes of this paper, we adopt the magnitude-dependent parameterization of the redshift distribution, common to many other weak-lensing studies as given by Baugh & Efstathiou (1993)

$$\begin{cases} n(z, \text{mag}) \propto z^2 \exp\left[-\frac{z}{z_0(\text{mag})} \right] 1.5 \\ z_0(\text{mag}) = \frac{z_m(\text{mag})}{1.412}, \end{cases} \quad (5)$$

where z_m is the median redshift of the survey as a function of magnitude. We calculate z_m for the COSMOS ACS data by bins of ΔF814W $= 0.25$ for $20 <$ F814W < 24 and derive the best linear fit to z_m, given by

$$z_m = (0.18 \pm 0.01)\text{F814W} - (3.3 \pm 0.2). \quad (6)$$

The majority of the galaxies for which we have no redshift estimate at F814W < 24 are those in masked regions, and their exclusion from this derivation does not affect these results. At F814W < 24, our redshift incompleteness is less than 4%, and the dominant source of error is the photometric redshift uncertainty expressed previously in equation (1). We note that a significant number of galaxies fainter then this limit do have photometric redshifts. The fact that these galaxies represent a statistically incomplete sample does not matter for some applications. For example, these additional galaxies are used in the cosmic shear measurement by Massey et al. (2007c).

In Figure 11, we compare the median redshift of COSMOS to the UDF survey (Coe et al. 2006), the CFHTLS survey (Ilbert et al. 2006), the H-HDF-N survey (Capak et al. 2004), and SSA22 (Hu et al. 2004; Capak et al. 2004). Table 4 is a summary of the data and the methods employed by these different photometric redshift surveys. All photometric redshifts have been computed with a Bayesian prior based on luminosity functions. The agree-
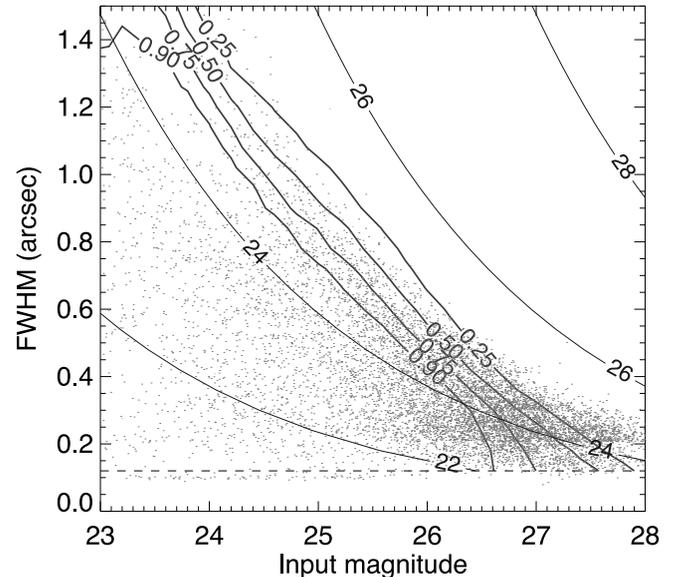
Fig. 10.—Completeness of the COSMOS F814W catalog as a function of total magnitude and FWHM determined by inserting fake objects into an ACS image. The thick contours show the percentage of fake objects recovered by SExtractor. The thin contours are the lines of constant surface brightness, in units of mag arcsec^{-2}, assuming a Gaussian profile. The gray points represent a random sample of objects from the COSMOS catalog plotted as a function of MAG_AUTO and FWHM_IMAGE. The dashed horizontal line indicates the size of the ACS PSF. Note that the simulations only consider objects with Gaussian profiles, whereas in reality the COSMOS objects exhibit a wide variety of profiles. [*See the electronic edition of the Supplement for a color version of this figure.*]

ment that we see between the various surveys at $z < 1$ is quite remarkable. At $z > 1$, however, the scatter in Figure 11 indicates the limits of current photometric techniques. Further simulations are clearly necessary in order to understand the biases introduced in the redshift distribution at $z > 1$. Although we reserve a full discussion for a future paper, we can already highlight some of the issues at hand.

The photometric redshift technique relies on detecting and measuring the strength of broad spectral features. These same features are used by color selection techniques to select objects at specific redshifts. The key features are the 4000 Å break, the Lyman break at ~912 Å, Lyman absorption at ~1216 Å, and coronal line absorption between 1500 and 2500 Å. Photometric redshifts are very robust if one or more of these features are detectable in the available data. However, at faint magnitudes the photometric errors and detection limits are often too large to constrain these features.

For example, problems arise for COSMOS between $1.5 < z < 3.2$, where the measurable features, the 4000 Å break and the coronal line absorption features, are difficult to detect. Indeed, at these redshifts, the 4000 Å break is well into the IR, where it is difficult to obtain deep data. A typical object at $z \simeq 2$ will be ~1.4 mag fainter at I than at K. This means objects fainter than F814W > 23.5 are not constrained by the present K-band data. At the other end of the spectrum the coronal absorption feature has a typical strength of ~0.15 mag, which requires a 25 σ detection in u^* to accurately differentiate from a similar break in $z < 0.5$ galaxies. With the present u^* data, this corresponds to objects brighter than F814W < 24.5 for typical galaxies.

In conclusion, the high-redshift tail of the COSMOS redshift distribution will be more accurately determined with the forthcoming near-IR data and the future deep zCOSMOS spectroscopy which specifically targets the $1.5 < z < 3$ region and will allow proper calibration down to F814W ~ 24. A more detailed
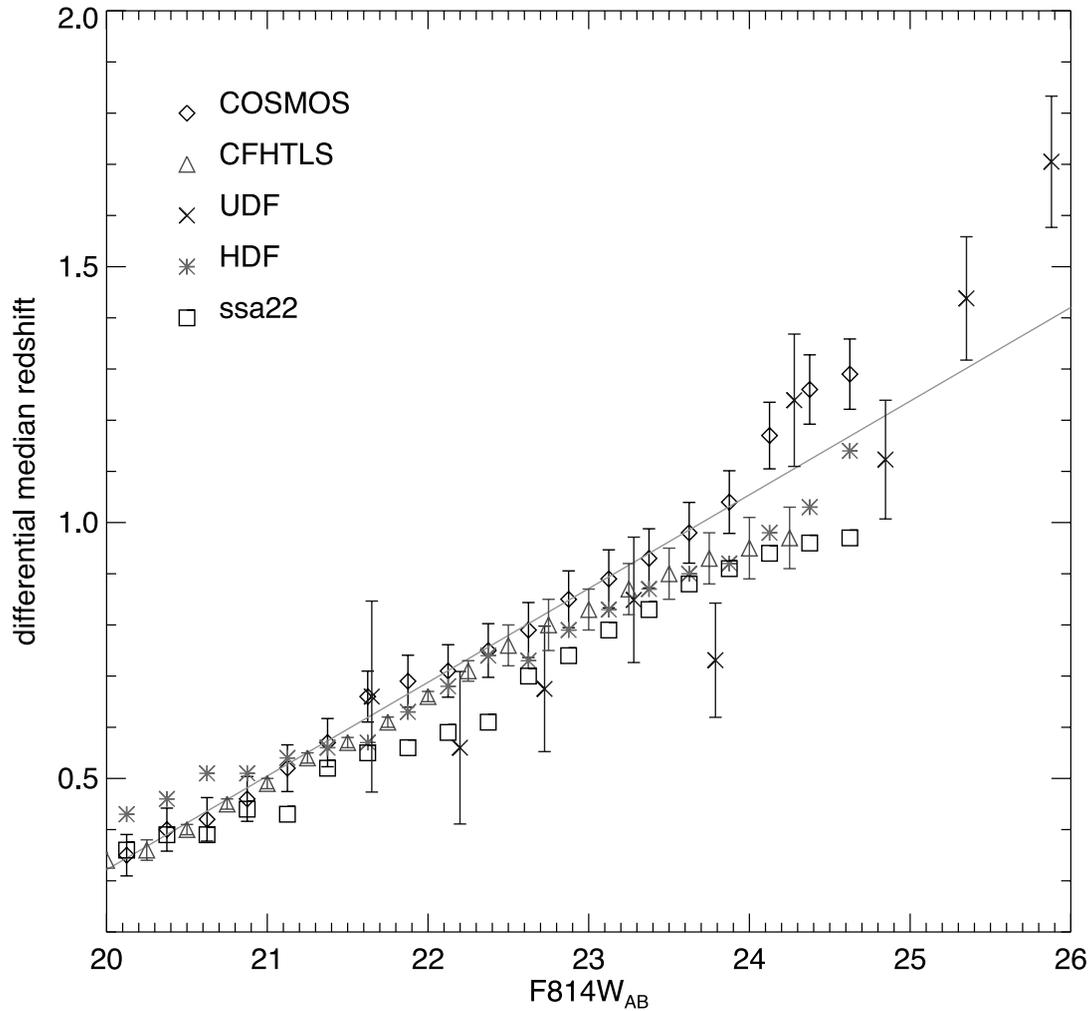
FIG. 11.—Median redshift of COSMOS compared to various photometric redshift surveys. For clarity, error bars are only shown for the COSMOS, CFHTLS, and UDF surveys. [*See the electronic edition of the Supplement for a color version of this figure.*]

analysis of the photometric redshifts will be conducted once the new data are available.

5. PSF CORRECTION AND SHEAR MEASUREMENT

In this section, we measure the shapes of galaxies and correct them for the convolution with the telescope's PSF and for other instrumental effects. For each galaxy, we construct an unbiased local estimator of the shear and derive the associated measurement error.

5.1. *PSF Modeling*

The ACS/WFC PSF is not as stable as one might naively hope from a space-based camera. As shown in Rhodes et al. (2007),

gradual changes to both the size and the ellipticity pattern of the PSF due to telescope "breathing" causes the PSF to change considerably on timescales of weeks. The long period of time over which the COSMOS field was observed forces us to take these variations into account (see Fig. 1). Although other strategies have been demonstrated successfully for observations conducted on a shorter time span, it would be inappropriate for us to assume, as in Lombardi et al. (2005) or Jee et al. (2005), that the PSF is constant or even, as in Heymans et al. (2005), that the focus is piecewise constant.

Fortunately, most of the PSF variation can be ascribed to a single physical parameter. Thermal expansions and contractions of *HST* alter the distance between the primary and secondary

TABLE 4
PRESENT DATE PHOTOMETRIC REDSHIFT SURVEYS

Survey	Area	Imaging Data	Calibration Spectra	Technique
COSMOS.................	1.67 deg^2	$B_j, V_j, g^+, r^+, i^+, z^+$, NB816, u^*, i^*, K_s	958	COSMOS/BPZ
UDF........................	11.97 arcmin^2	$B, V, i', z', J^a H^b$	76	BPZ
CFHTLS..................	3.2 deg^2	$u^*, g', r', i', z', J, K$	2867	Le Phare
H-HDF-N	0.2 deg^2	$U_j, B_j, V_j, R_C, I_C, Z^+, HK'$	2149	BPZ
SSA22	0.2 deg^2	$U^*_j, B_j, V_j, R_C, I_C, z^+, J, H, K, HK'$	452	BPZ

[a] 5.76 arcsec^2.
[b] 160 arcmin^2.
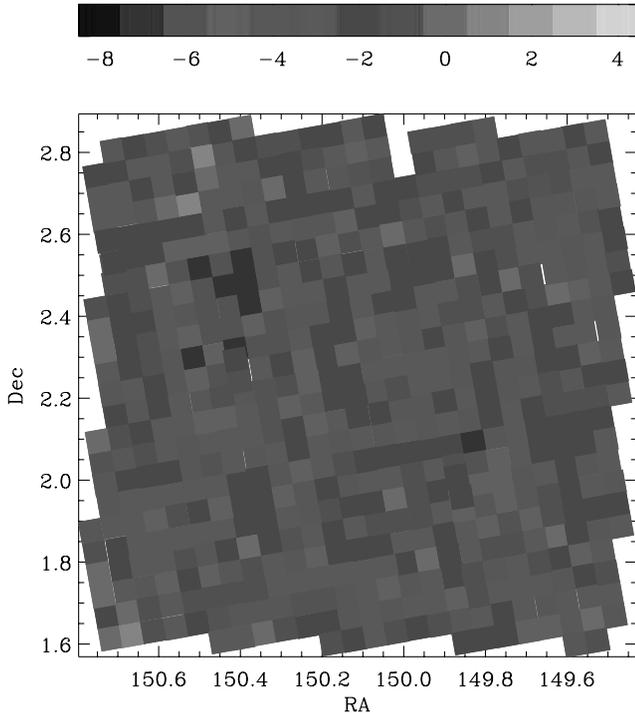
Fig. 12.—Adopted PSF model across the survey. The colors correspond to deviations in the apparent focus of the telescope away from nominal (μm). These are caused by expansion and contraction of the *HST* due to thermal variations. Note that the focus values are clustered and not randomly distributed. See Rhodes et al. (2007) for more details about the PSF pattern at particular focus positions. [*See the electronic edition of the Supplement for a color version of this figure.*]

mirrors. As the effective focus deviates from nominal, the PSF becomes larger and more elliptical, with the direction of elongation depending on the position above or below nominal focus (cf. Krist 2005). The thermal load on *HST* is constantly changing in a complicated way as it passes in and out of the shadow of the Earth and is rotated to different pointings.

As described in Rhodes et al. (2007), we have modified version 6.3 of the Tiny Tim ray-tracing program (Krist & Hook 2004) to create a grid of model PSF images at varying focus offsets. By comparing the ellipticity of ∼20 stars in each image to these, we can determine the image's effective focus. Tests of this algorithm on ACS/WFC images of dense stellar fields confirm that the best-fit effective focus can be repeatably determined from a random sample of 10 stars brighter than F814W = 23 with an rms error less than 1 μm. The effective focus of the COSMOS images are shown in Figure 12. An alternative correction scheme based on PSF models constructed from dense stellar fields has also been suggested by Schrabback et al. (2006).

Once images have been grouped by their effective focus position, we can combine the few stars in each image into one large catalog. We interpolate the PSF model parameters using a polynomial fit (of order $3 \times 2 \times 2$ in each CCD separately), in the usual weak-lensing fashion (cf. Massey et al. 2002). See Rhodes et al. (2007) for more details concerning the PSF modeling scheme.

5.2. Galaxy Shape Measurement

We use the shape measurement method developed for space-based imaging by Rhodes et al. (2000, hereafter RRG). The RRG method has been optimized for space-based images with small PSFs and has previously been used on weak-lensing analyses of WFPC2 and STIS data (Rhodes et al. 2001, 2004; Refregier

et al. 2002). In a manner similar to the common "KSB" method (Kaiser et al. 1995), RRG measures the second- and fourth-order Gaussian-weighted moments of each galaxy:

$$I_{ij} = \frac{\sum wIx_ix_j}{\sum wI}, \tag{7}$$

$$I_{ijkl} = \frac{\sum wIx_ix_jx_kx_l}{\sum wI}. \tag{8}$$

The sum is over all pixels, w is the size of the Gaussian weight function, I is the pixel intensity, and the x_i coordinates are measured in pixels. The Gaussian weight function is necessary to suppress divergent sky noise contributions in the measurement of the quadrupole moments. The RRG method is well suited to the small, diffraction-limited PSF obtained from space, because it decreases the noise on the shear estimators by correcting each moment for the PSF linearly and only dividing them to form an ellipticity at the last possible moment.

After the moments have been corrected for the PSF, an ellipticity $\varepsilon = (e_1, e_2)$ and size measure, d, are calculated for each galaxy:

$$e_1 = \frac{I_{xx} - I_{yy}}{I_{xx} + I_{yy}}, \tag{9}$$

$$e_2 = \frac{2I_{xy}}{I_{xx} + I_{yy}}, \tag{10}$$

$$d = \sqrt{\frac{I_{xx} + I_{yy}}{2}}. \tag{11}$$

Note that the d parameter is a measure of galaxy size but that its value will depend on the choice of the width of the Gaussian weight function, w.

5.3. Shear Measurement

The estimator $\varepsilon = (e_1, e_2)$ is not yet a shear estimator, because it does not respond linearly to changes in shear. It must first be normalized by a shear susceptibility factor (also known as the shear polarizability),

$$\tilde{\gamma} = \{\tilde{\gamma}_1, \tilde{\gamma}_2\} = \frac{\varepsilon}{G}, \tag{12}$$

where the shear susceptibility factor, G, is measured from moments of the global distribution of ε and other, higher order shape parameters (see eq. [28] in RRG). The RRG formalism does not allow for G to be calculated for any individual galaxy. However, G can be calculated for an ensemble of galaxies by averaging over a population's shape moments. Previous incarnations of RRG that have been used to measure cosmic shear (Rhodes et al. 2001, 2004; Refregier et al. 2002), have made use of a single value of G for the entire survey. Adopting this approach for the COSMOS data would yield a value of $G = 1.13$. However, STEP (Massey et al. 2007c) showed that G can vary significantly as a function of object flux, whether this be due to evolution in galaxy morphologies as a function of redshift, or noise in the wings of faint galaxies that simply impedes the measurement of their radial profiles and higher order moments. An increase in shear susceptibility with object S/N has also been seen in KSB-type analyses (Massey et al. 2004a).

Our tests have confirmed that a constant value would be insufficiently precise for a survey the size of COSMOS, and would particularly affect the kind of 3D analysis for which COSMOS is so well suited. We therefore calculate G from the COSMOS data
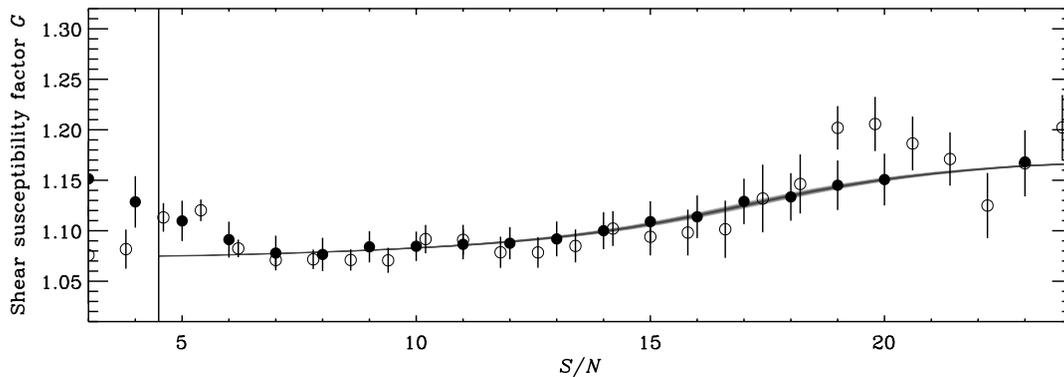
FIG. 13.—Interpolation of the shear susceptibility factor G. The filled circles show G, calculated in bins of S/N, for the COSMOS data. The open circles show the same for the simulated COSMOS images. The solid line shows the shear susceptibility model adopted for the data.

in bins of S/N (see Fig. 13). Simulated images of the COSMOS data are created using the shapelets-based method of Massey et al. (2004b) (see § 5.4). The variations of G as a function of S/N are apparent in the COSMOS data are well reproduced by the simulated data. For COSMOS galaxies, we find that variations in G as a function of S/N are well fit by

$$G = 1.125 + 0.04 \arctan\left(\frac{\text{S/N} - 17}{4}\right). \qquad (13)$$

Adopting the above model, we derive G for each galaxy as a function of S/N.

5.4. *Calibration via Simulated Images*

Using the shapelets-based method of Massey et al. (2004b), we have created simulated images with the same depth, noise properties, PSF, and galaxy morphology distribution as the real COSMOS data. A known shear signal was applied to the images, which we then attempted to measure using the same pipeline as the data. This exercise is similar to STEP (Heymans et al. 2006; Massey et al. 2007b) but tailored exclusively to COSMOS.

The simulated COSMOS images are each $4' \times 4'$, and contain ~500 galaxies after applying the same catalog cuts that were applied to the real data (see §6). Simulated galaxy morphologies are based on those observed in the Hubble Deep Fields (Williams et al. 1996, 1998), parameterized as shapelets and randomly rotated/flipped before being sheared. Different input galaxies were used in each simulated image, as if they were pointing to different patches of the sky, to keep them independent. To simplify later analysis, all of the galaxies within an image were sheared by the same amount. A total of 41 images were made, with shears applied in integer steps from -10% to $+10\%$ in the γ_1 component (while γ_2 was fixed at zero) and similarly for the γ_2 component. The images were then convolved with a model ACS PSF. Again to simplify the analysis, this was a constant PSF obtained from Tiny Tim. Its (e_1, e_2) ellipticity is $(-0.21\%, -2.07\%) \pm (0.14, 0.10)$. No stars were included in the simulated images; a separate star field was created, from which the PSF moments could be measured. Noise was added to all of these images, to the same depth as the COSMOS observations, and with a similar (but isotropic) correlation between adjacent pixels to mimic the effects of MultiDrizzle and unresolved background sources.

The recovered shear measurement from the simulated data is presented in Figure 14. We find that, in order to correctly measure the input shear on COSMOS-like images, the RRG method

requires an overall calibration factor of $C = (0.86^{+0.07}_{-0.05})^{-1}$, so that

$$\tilde{\gamma} = C \frac{\varepsilon}{G}. \qquad (14)$$

The necessity for such calibration factors has long been known in the field (e.g., Bacon et al. 2001; Erben et al. 2001), and is in accord with results from STEP1. STEP2 and F. W. High et al. (2007, in preparation) suggest that this may be intrinsic with KSB-related methods and, furthermore, that the calibration can vary for the two components of shear. For this reason, we fit each component separately, and use final calibration factors of $C_1 = (0.80)^{-1}$ for $\tilde{\gamma}_1$ and $C_2 = (0.92)^{-1}$ for $\tilde{\gamma}_2$. After this recalibration, Table 5 shows STEP-like estimates of the additive bias $\langle c \rangle$ and the multiplicative bias $\langle m \rangle$ obtained by fitting deviations of the recovered shear from the input shear. Both of these are consistent with ideal shear recovery, although the error on these estimates will be propagated through subsequent analyses.
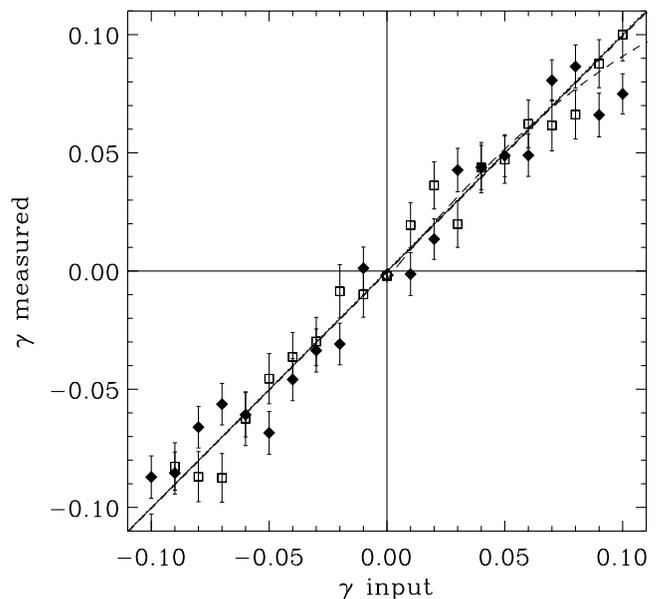
FIG. 14.—Calibration of the RRG shear measurement method from simulated COSMOS images containing a known input shear. The measured shear on the y-axis includes the shear calibration factor C. *Squares*: Measurements of γ_1. *Diamonds*: Measurements of γ_2. *Solid line*: Linear fit to deviations from the ideal case of $\gamma_{\text{measured}} = \gamma_{\text{input}}$ for all points. *Dashed line*: Quadratic fit demonstrating that the curvature terms are negligible.

TABLE 5
STEP-like Estimates of Additive and Multiplicative Bias

Bias	Value
Linear STEP fit	
$\langle c \rangle$	$(2.1 \pm 40.0) \times 10^{-4}$
c_1	$(5.6 \pm 28.5) \times 10^{-4}$
c_2	$(-1.3 \pm 28.0) \times 10^{-4}$
$\langle m \rangle$	$(-0.3 \pm 9.4) \times 10^{-2}$
m_1	$(-12.2 \pm 6.5) \times 10^{-2}$
m_2	$(11.6 \pm 6.3) \times 10^{-2}$
Quadratic STEP fit	
$\langle c \rangle$	$(-23.3 \pm 121.8) \times 10^{-4}$
$\langle m \rangle$	$(20.1 \pm 51.5) \times 10^{-2}$
$\langle q \rangle$	-2.84 ± 4.59

Notes.—Calibration of the RRG shear measurement method on simulated COSMOS images containing a known shear, described using STEP parameters. Figures are supplied after the application of the shear calibration factor.

For the kind of 3D shear analysis for which COSMOS is so well suited, the simultaneous calibration of shears from an entire population of galaxies is insufficient. Since our companion paper, Massey et al. (2007c), is concerned with the growth of the shear signal as a function of redshift, it is crucial that the shear calibration be equally precise for both distant and relatively nearby galaxies. Given that more distant galaxies are fainter and smaller and that the details of the shear measurement depend on a fixed PSF size, pixel size, and noise level, this requirement is not trivial. We have therefore split the simulated galaxy catalog in half by magnitude (at F814W = 25.4) and by size (at $d = 5.0$ pixels), and repeated the analysis. We find that our shear calibration $\langle m \rangle$ is robust for galaxies of different fluxes within 1% and of different sizes within 4%. Redshifts were not available for the simulated galaxies, so a direct split in redshift was not possible. Although this effect is clearly small in the regime of our current measurements, it will be significant in future weak-lensing surveys where the error budget will be dominated by systematic uncertainties.

5.5. *Error on the Shear Estimator*

For each galaxy, the error on the measured shear is estimated using the same method as implemented in the Photo pipeline (Lupton et al. 2001) to analyze data from the Sloan Digital Sky Survey. For each object, we assume that the optimal moments are the same as the moments corresponding to a best-fit Gaussian. This formulation allows us to determine the covariance matrix of the moments in the usual way of nonlinear least squares. To be precise, we model the image with a two-dimensional elliptical Gaussian model

$$M(\boldsymbol{x}) = \frac{f}{2\pi |Q|^{1/2}} \exp\left[-\frac{1}{2}(\boldsymbol{x} - \mu)^T Q^{-1}(\boldsymbol{x} - \mu)\right]. \quad (15)$$

This model has six parameters: the flux, f, the two centroids, μ, and the three moments which form the elements of the symmetric matrix Q. These parameters are noted p_l. Next, we derive the χ^2 in the usual way,

$$\chi^2 \equiv \frac{1}{\sigma^2} \sum_{ij} \left[M(\boldsymbol{x}_{ij}) - I_{ij}\right]^2, \quad (16)$$

where σ is the sky noise level. We can now compute the 6×6 Fisher matrix which is the matrix of the second derivatives of the χ^2

$$F_{kl} \equiv \frac{1}{2} \frac{\partial^2 \chi^2}{\partial p_k \partial p_l}$$

$$= \frac{1}{\sigma^2} \sum_{ij} \frac{\partial M(\boldsymbol{x}_{ij})}{\partial p_k} \frac{\partial M(\boldsymbol{x}_{ij})}{\partial p_l}$$

$$- \frac{1}{\sigma^2} \sum_{ij} \left[M(\boldsymbol{x}_{ij}) - I_{ij}\right] \frac{\partial^2 M(\boldsymbol{x}_{ij})}{\partial p_k \partial p_l}.$$

As is customary, we drop the second term which is proportional to the residuals. This term is usually very small compared to the first. The covariance matrix of the Gaussian parameters is then the inverse of the Fisher matrix. The 3×3 block of the covariance matrix corresponding to the second moments can then be extracted. Because the whole 6×6 matrix was inverted, this correctly marginalizes over centroid errors and other model parameter degeneracies. In this way, it differs from formulas which assume a constant (nonadaptive) weighting function and perfect centroiding.

This Fisher matrix does not depend explicitly on the data; it only depends on the best-fit parameters and can be computed analytically. It is simply a function of four numbers: the three second-order moments defined in equation (7) and the S/N (f/σ), which can be parameterized by the magnitude error (for example, MAGERR_AUTO). Since the flux computed with SExtractor is not exactly equal to f (which would be the best-fit Gaussian amplitude) we allow for a single calibration factor and multiply the covariance matrix by this. We calibrate this factor with image simulations and verify that the errors are correctly predicted.

Since the ellipticity components are computed from the moments, the variances of the ellipticity components can be computed by linearly propagating the covariance matrix of the moments. Finally, the two ellipticity components can be shown to be uncorrelated with each other.

6. FINAL GALAXY SELECTION

6.1. *Lensing Cuts*

Estimations of the gravitational shear will be improved by averaging only those galaxies with precise shape measurements on the condition that no ellipticity selection bias is introduced by the "lensing cuts." We apply strict cuts to the $\mathcal{C}_2$ catalog that are designed to extract a sample of resolved galaxies with reliable shape measurements. The resulting catalog is referred to as $\mathcal{C}_4$. Our lensing cuts are summarized in Table 6 and are based on the four following parameters:

1. The estimated significance of each galaxy detection, where the significance is defined as S/N = FLUX AUTO/FLUXERR AUTO.
2. The first order moments, I_{xx} and I_{yy}.
3. The total ellipticity, $e = (e_1^2 + e_2^2)^{1/2}$.
4. The galaxy size as defined by the RRG d parameter (see § 5.2).

The final size cut is designed to select galaxies with well-resolved shapes. Indeed, PSF corrections become increasingly significant as the size of a galaxy approaches that of the PSF and the intrinsic shape of a galaxy becomes more difficult to measure. In COSMOS images, the typical size (as defined by d) of a star is about $d_\star = 2.2$ pixels (0.066″). Our size cut is thus equivalent to selecting galaxies with $d_g > 1.6 d_\star$. Note that in this section, as well as in all following sections, d has not been corrected for the PSF.

Parameter	Galaxies Retained in $\mathcal{C}_4$
I_{xx} and I_{yy}.........................	Finite[a] I_{xx} and I_{yy}
RRG size parameter[b].......	$d > 3.6$ pixels
Significance......................	S/N > 4.5
Total ellipticity[c]...............	$e < 2$

[a] Indicating that the RRG code converged.
[b] Uncorrected for the PSF.
[c] Corrected for the PSF.

The ellipticity cut at $e < 2$ may be surprising given that, by definition, ellipticities are restricted to $e \leq 1$. In reality, however, because of noise, it is possible to measure an ellipticity of $e > 1$. Selecting galaxies with $e < 1$ could introduce an unwanted ellipticity bias, but because we are only interested in ensemble averages an acceptable solution is to cut out only a small number of large outliers ($e > 2$).

Note that further cuts may be required for some applications (for example, to isolate only those objects with well-measured photometric redshifts). In addition to these cuts, a galaxy-by-galaxy weighting scheme may also be used to minimize the impact of shape measurement noise.

6.2. Effective Galaxy Number Density

Once the lensing cuts have been applied, the $\mathcal{C}_4$ catalog contains only those galaxies which are useful for lensing analyses with the COSMOS data. Predictions for future weak-lensing surveys based on COSMOS results must consider the effective number of galaxies that are actually useful for lensing purposes (and not the number of raw detections). With these considerations in mind, we define the effective galaxy number density, $N_g(z)$, as the total number of galaxies within $\mathcal{C}_4$, per unit area, and with redshifts below a given redshift, z. The equivalent quantities as a function of galaxy magnitude and size are $N_g(m)$ and $N_g(d)$. For each galaxy, the SExtractor MAG_AUTO parameter is used to estimate the magnitude and the RRG d parameter is used as an estimate of the size. We also consider the derivatives of $N_g(z)$, $N_g(m)$, and $N_g(d)$ so that

$$N_g(z) = \int_0^z n_g(z')dz', \qquad (17)$$

$$N_g(m) = \int_{20}^m n_g(m')dm', \qquad (18)$$

$$N_g(d) = \int_d^\infty n_g(d')dd'. \qquad (19)$$

The total number density of galaxies in $\mathcal{C}_4$ is noted as N_g and is significantly lower than the number density of *detected* galaxies. In total, the final lensing catalog $\mathcal{C}_4$ contains 3.9×10^5 galaxies with accurate shape measurements and 2.8×10^5 galaxies with both shape and photometric redshift measurements. Table 2 shows a summary of the different steps leading to this catalog. The surveyed area of COSMOS is 1.64 deg^2, leading to an overall number density of $N_g \sim 66$ galaxies per arcmin^2 for the first sample and ~ 50 for the second. These numbers can be contrasted to a more sparse 15–25 galaxies per arcmin^2 typically resolved with deep, ground-based surveys. In Figures 15 and 16 we show the effective densities defined above, as well as their corresponding derivatives. From these figures, we can draw the following conclusions:
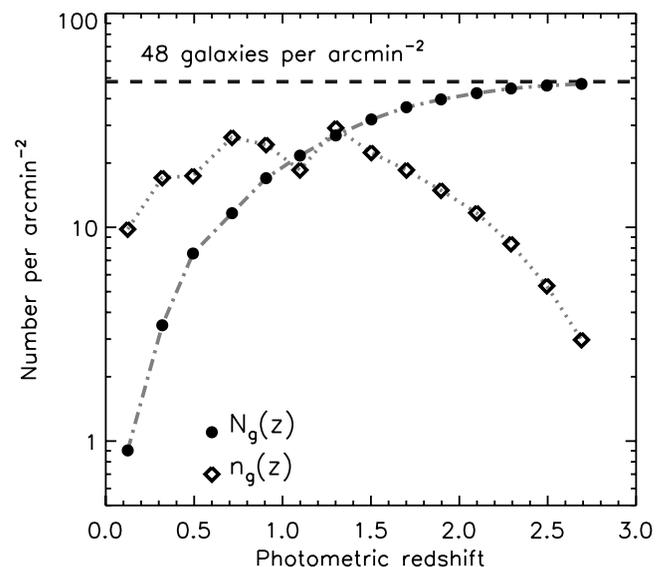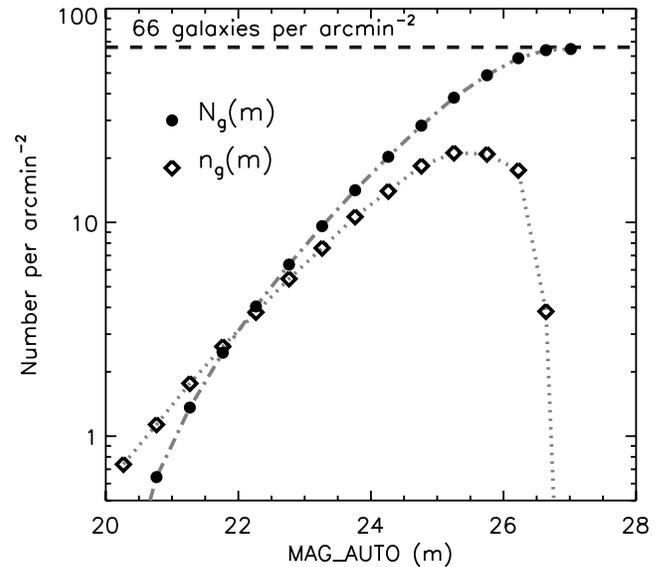
Fig. 15.—*Top*: Effective number density of galaxies as a function of F814W magnitude. Only resolved galaxies with precise shape measurement are included in these counts. *Bottom*: Effective number density of galaxies as a function of redshift. The effective density still evolves sharply after $z > 1$, demonstrating that the COSMOS lensing data is a powerful probe of structures at $z < 1$.

1. Over 60% of the COSMOS source galaxies are at redshifts higher than $z = 1$. The COSMOS weak-lensing data is therefore a powerful probe of the dark matter distribution from $z \sim 1$ to the present day.

2. About 18 galaxies per arcmin^{-2} (73%) are discarded when we select only those galaxies with accurate photometric redshifts. The primary cause of this loss is the larger areas masked out in the ground-based data as compared to the ACS data. Future space-based weak-lensing surveys could recover the remaining 27% by using space-based, multiwavelength imaging to derive photometric redshifts.

3. The effective number density rises very steeply with decreasing galaxy size; over 50% of our total number of sources have $d < 5$ pixels ($0.15''$). The current size cut for the COSMOS data is $d = 3.6$ pixels, and the size of the ACS PSF is $d_\star = 2.2$ pixels ($0.066''$). By pushing this size barrier to even smaller values, future surveys could very quickly obtain much higher effective densities.
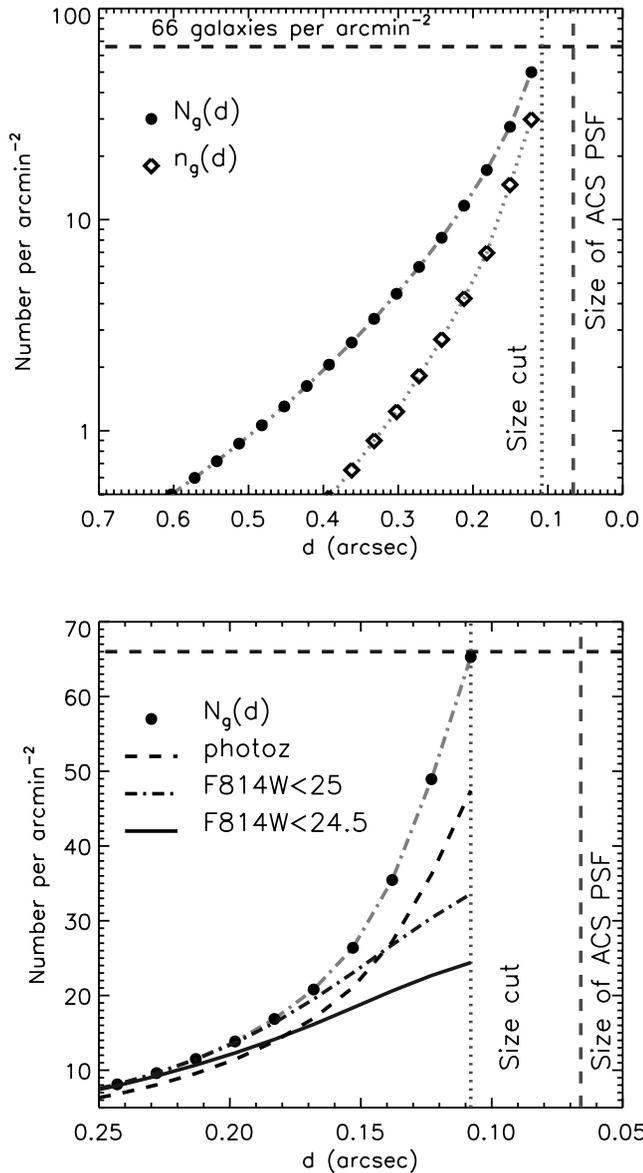
FIG. 16.—*Top*: Effective number density of galaxies as a function of size. The dotted vertical line indicates the size cut that we make in order to extract galaxies with precise shape measurements. This size cut is such that $d_g > 1.6d_\star$. Here N_g rises very sharply as a function of decreasing d, demonstrating that small galaxies make up the majority of our lensing sources. *Bottom*: N_g for COSMOS galaxies with accurate photometric redshifts and for magnitude cuts of F814W < 25 and F814W < 24.5.

Finally, the cuts that we have applied to the COSMOS data (in particular the size cut) are stricter than would be necessary without significant CTE effects that degrade the shape measurements of the faintest galaxies. Next-generation space-based missions designed to avoid the problems encountered and identified in COSMOS, as well as future implementations of the COSMOS catalog, will undoubtedly achieve higher effective number densities.

7. INTRINSIC SHAPE NOISE: A FUNDAMENTAL LIMIT TO WEAK-LENSING MEASUREMENTS

Under the assumption of weak gravitational lensing, a source galaxy with intrinsic shape ε_{int} and observed ellipticity ε_{obs} is related to the gravitational lensing induced shear γ according to

$$\varepsilon_{obs} = \varepsilon_{int} + \gamma. \tag{20}$$

Throughout this paper, the gravitational shear is noted as γ, whereas $\tilde{\gamma}$ represents our estimator of γ. The above relationship indicates that galaxies would be ideal tracers of the distortions caused by gravitational lensing if the intrinsic shape ε_{int} of each source galaxy were known a priori. A quick glance at an ACS image, however, reveals that galaxies display a very wide variety of shapes, which unfortunately prevents the extraction of γ for any single galaxy. Lensing measurements thus exhibit an intrinsic limitation encoded in the width of the ellipticity distribution of the galaxy population, noted here as σ_{int}, and often referred to as the intrinsic shape noise. Because the shape noise (of order $\sigma_{int} \sim 0.26$) is significantly larger than weak shear (typically $\gamma \sim 0.01$ for cosmic shear), γ must be estimated by averaging over a large number of galaxies. In this case equation (20) simplifies to

$$\langle \varepsilon_{obs} \rangle = \langle \gamma \rangle. \tag{21}$$

The uncertainty in the shear estimator, $\sigma_{\tilde{\gamma}}$, arises from a combination of unavoidable intrinsic shape noise, $\sigma_{int}^2 = \langle \varepsilon_{int}^2 \rangle$ and the measurement error of galaxy shapes σ_{meas}^2:

$$\sigma_{\tilde{\gamma}}^2 = \sigma_{int}^2 + \sigma_{meas}^2. \tag{22}$$

In the following analysis, $\sigma_{\tilde{\gamma}}$ will be referred to as the *shape noise* and σ_{int} will be called the *intrinsic shape noise*. The former includes the shape measurement error, σ_{meas}, and hence will vary according to the data set as well as the shape measurement method that is employed. Note that the uncertainty contributions from photon noise, PSF correction, CTE calibration, and the shape measurement method are all included in our definition of σ_{meas}^2. The weak-lensing distortions averaged over the whole COSMOS field are small, and represent a negligible perturbation to equation (22).

Instead of the simple arithmetic mean of equation (21), many lensing practitioners adopt in some form or another an optimized weighting scheme in order to estimate γ, which often incorporates both the measurement error and the shape noise (for, e.g., Bernstein & Jarvis 2002). A constant value (of order 0.3) is often assumed for σ_{int}. However, it would not be surprising that the same processes that shape galaxy formation also lead to a variation of σ_{int} as a function of magnitude, galaxy type, or redshift. Furthermore, as weak-lensing surveys increase in both scale and depth, it is of intense interest to obtain accurate estimates of the intrinsic shape noise floors that these surveys must confront. For these reasons, we undertake a measurement of the shape noise, $\sigma_{\tilde{\gamma}}$, as well as the intrinsic shape noise, σ_{int}, as a function of magnitude, size, and redshift. A more detailed analysis of the intrinsic shape noise as a function of galaxy morphology will be the subject of a future paper.

First, we estimate the shape noise $\sigma_{\tilde{\gamma}}$ directly from the COSMOS data as a function of size and F814W magnitude. To derive $\sigma_{\tilde{\gamma}}$ we consider (1) the mean variance of both shear components (including correction factors), $\sigma_{\tilde{\gamma}} = (\sigma_{\tilde{\gamma}1} + \sigma_{\tilde{\gamma}2})/2$, and (2) galaxies with well measured shapes (see § 6).

Second, we derive an empirical estimation of the shape measurement error, σ_{meas}, using a sample of 27,000 galaxies that belong to overlapping regions of adjacent pointings. Each of these galaxies provides us with two independent shape measurements. Using these overlaps, we find that the shape measurement error is a function of both size and magnitude and increases beyond $\sigma_{meas} = 0.1$ for MAG_AUTO > 24.5 and $d < 7$. Third, we compare this empirical determination of σ_{meas} to the theoretical one derived in § 5.5. We find that the theoretical model of § 5.5 does remarkably well in predicting σ_{meas} as a function of size and magnitude.

Fig. 18.—Intrinsic shape noise as a function of photometric redshift. Galaxies have first been selected to have measurement errors less than 0.1 ($d > 7$ and MAG_AUTO < 24.5), so that the scatter in the observed shear closely matches the intrinsic shape noise. A linear fit to the shape noise as a function of redshift reveals a flat distribution with a mean of $\sigma_{int} = 0.26$. [*See the electronic edition of the Supplement for a color version of this figure.*]

Fig. 17.—Observed scatter in the shear as a function of F814W magnitude and size. The scatter is a combination of intrinsic shape noise, σ_{int}, and a shape measurement error, σ_{meas}. The shape measurement error is determined by a theoretical model and tested using galaxies from overlapping regions. Shaded regions indicate the rms width of the measurement error distribution. The intrinsic shape noise appears to increase slightly as a function of magnitude but is independent of size. At faint magnitudes and small galaxy sizes, the shear scatter increases rapidly due to large measurement errors. The dashed vertical line indicates the size and magnitude cut that we apply in order to construct Fig. 18. [*See the electronic edition of the Supplement for a color version of this figure.*]

Thus confident in the validity of this model, we adopt it for subsequent derivations. Finally, using equation (22) and the shear measurement error σ_{meas}, we extract the intrinsic shape noise of our galaxy sample as a function of size, magnitude, and redshift. The results are shown in Figures 17 and 18.

As can be seen in Figure 17, large measurement errors lead to an increase of $\sigma_{\tilde{\gamma}}$ at small sizes and faint magnitudes. The intrinsic shape noise, however, appears to change little with either size of magnitude and remains constant at a value of $\sigma_{int} \sim 0.26$. The slight apparent increase of the intrinsic shape noise at fainter magnitudes is probably due to the simplified measurement error

estimator that we are using (indeed, the overlaps indicate slightly higher errors). From this analysis, we can draw the following conclusions:

1. The intrinsic shape noise varies little from $z = 0$ to $z = 3$. Deep space-based weak-lensing surveys will therefore confront equivalent intrinsic shape noise floors as their shallower counterparts.

2. Measurement errors lead to an increase in the shape noise as a function of size and magnitude. A joint improvement in both imaging quality and shape measurement methodology will lead to shape noises that are closer to the intrinsic floor of 0.26.

3. We have yet to explore whether the intrinsic shape noise varies as a function of galaxy type. If so, shear measurements could be improved by incorporating a galaxy-type discriminant into the weighting scheme. This will be the focus of a future paper.

8. CONCLUSION

We have carefully constructed a weak-lensing catalog from 575 ACS/*HST* tiles, the largest space-based survey to date. We have established the quality of this catalog and analyzed the COSMOS redshift distribution, showing broad agreement with other photometric redshift surveys out to $z \sim 1$. The photometric redshifts are currently limited by the lack of deep *K*-band data but will rapidly improve as this data soon becomes available. Shapes have been measured for over 3.9×10^5 galaxies and corrected for distortions induced by PSF and CTE effects. Simulations have been used in order to calibrate our shear measurement method and a STEP-like analysis has been performed, demonstrating our ability accurately measure shear with negligible additive and multiplicative bias. The effective number density of the COSMOS weak-lensing catalog is 66 galaxies per arcmin^2 (48 when we consider only those with accurate photometric redshifts). A large fraction of these galaxies are at $z > 1$, making COSMOS a powerful probe of

the dark matter distribution and its evolution from $z = 1$ to the present day. The COSMOS survey is also of foremost importance for the preparation and design of future wide field space-based lensing missions. Regarding the design of such missions, our main conclusions from working with the COSMOS data are the following.

1. Understanding and correcting for the time varying PSF and calibrating CTE effects were two of the most difficult challenges encountered with the COSMOS data. Reducing these two systematic effects should be a key specification in the design of next-generation telescopes and instruments.

2. Because (1) small galaxies are not as readily detected from the ground and (2) larger areas are masked out in the ground-based data than in the ACS imaging, we lose 27% of our source sample when we select only those with accurate photometric redshifts.

3. The effective number density of galaxies is a very sensitive function of survey depth and resolution. The capability to resolve and accurately measure the shapes of very *small*, *faint* galaxies will be key in obtaining number densities of over 66 galaxies arcmin^{-2}.

4. Finally, we have derived the intrinsic shape noise of the galaxy sample and demonstrated that it remains fairly constant as a function of size, magnitude, and redshift. Weak-lensing measurements with deep space-based imaging are therefore on par with more shallow imaging in terms of the intrinsic shape noise floors that they must overcome.

The COSMOS weak-lensing data described in this paper has already been used to measure cosmological parameters and to demonstrate the feasibility of the ''tomography'' technique (Massey et al. 2007c). Striking weak-lensing mass maps of the COSMOS field have been made that reveal tantalizing evidence of a complex interplay between the baryon and the dark matter distribution (Massey et al. 2007a). Yet more lensing analyses are underway, including a galaxy-galaxy lensing and a group-galaxy lensing study that will undoubtedly lead to a better understanding of the relationship between baryonic and dark matter structures and of its evolution over cosmic time.

REFERENCES

Bacon, D. J., Goldberg, D. M., Rowe, B. T. P., & Taylor, A. N. 2006, MNRAS, 365, 414
Bacon, D. J., Refregier, A., Clowe, D., & Ellis, R. S. 2001, MNRAS, 325, 1065
Bacon, D. J., et al. 2005, MNRAS, 363, 723
Bardeau, S., Kneib, J.-P., Czoske, O., Soucail, G., Smail, I., Ebeling, H., & Smith, G. P. 2005, A&A, 434, 433
Bartelmann, M., & Schneider, P. 2001, Phys. Rep., 340, 291
Baugh, C. M., & Efstathiou, G. 1993, MNRAS, 265, 145
Bernstein, G., & Jain, B. 2004, ApJ, 600, 17
Bernstein, G. M., & Jarvis, M. 2002, AJ, 123, 583
Bertin, E., & Arnouts, S. 1996, A&AS, 117, 393
Capak, P., et al. 2004, AJ, 127, 180
———. 2007, ApJS, 172, 99
Casertano, S., et al. 2000, AJ, 120, 2747
Coe, D., Benítez, N., Sánchez, S. F., Jee, M., Bouwens, R., & Ford, H. 2006, AJ, 132, 926
Erben, T., Van Waerbeke, L., Bertin, E., Mellier, Y., & Schneider, P. 2001, A&A, 366, 717
Ford, H. C., et al. 2003, in Proc. SPIE 4854, 81
Goldberg, D. M., & Bacon, D. J. 2005, ApJ, 619, 741
Goldberg, D. M., & Leonard, A. 2006, preprint (astro-ph/0607602)
Hetterscheidt, M., Simon, P., Schirmer, M., Hildebrandt, H., Schrabback, T., Erben, T., & Schneider, P. 2006, A&A, submitted (astro-ph/0606571)
Heymans, C., & Heavens, A. 2003, MNRAS, 339, 711
Heymans, C., et al. 2005, MNRAS, 361, 160
———. 2006, MNRAS, 368, 1323
Hoekstra, H., Yee, H. K. C., & Gladders, M. D. 2004, ApJ, 606, 67
Hoekstra, H., et al. 2006, ApJ, 647, 116
Hu, E. M., Cowie, L. L., Capak, P., McMahon, R. G., Hayashino, T., & Komiyama, Y. 2004, AJ, 127, 563
Ilbert, O., et al. 2006, A&A, 457, 8411
Impey, C. D., Sprayberry, D., Irwin, M. J., & Bothun, G. D. 1996, ApJS, 105, 209
Jain, B., & Taylor, A. 2003, Phys. Rev. Lett., 91, 1302
Jarvis, M., Jain, B., Bernstein, G., & Dolney, D. 2006, ApJ, 644, 71
Jee, M. J., White, R. L., Benítez, N., Ford, H. C., Blakeslee, J. P., Rosati, P., Demarco, R., & Illingworth, G. D. 2005, ApJ, 618, 46
Kaiser, N., Squires, G., & Broadhurst, T. 1995, ApJ, 449, 460 (KSB)
King, L. J. 2005, A&A, 441, 47

Kleinheinrich, M., et al. 2005, A&A, 439, 513
Kneib, J.-P., et al. 2003, ApJ, 598, 804
Koekemoer, A. M. 2007, ApJS, 172, 196
Koekemoer, A. M., Fruchter, A. S., Hook, R. N., & Hack, W. 2002, in Proc. STScI Symp., The 2002 *HST* Calibration Workshop, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), 2002, 337
Krist, J. 2005, Instrument Science Rep. ACS 2003-06 (Baltimore: STScI)
Krist, J., & Hook, R. 2004, The Tiny Tim User's Tuide version 6.3 (Baltimore: STScI), http://www.stsci.edu/software/tinytim/tinytim.pdf
Lilly, S. J., et al. 2007, ApJS, 172, 70
Lombardi, M., et al. 2005, ApJ, 623, 42
Lupton, R., Gunn, J. E., Ivezić, Z., Knapp, G. R., & Kent, S. 2001, in ASP Conf. Ser. 238, Astronomical Data Analysis Software and Systems X, ed. F. R. Harnden, Jr., F. A. Primini, & H. E. Payne (San Francisco: ASP), 269
Mandelbaum, R., Seljak, U., Kauffmann, G., Hirata, C. M., & Brinkmann, J. 2006, MNRAS, 368, 715
Massey, R., Bacon, D., Refregier, A., & Ellis, R. 2002, in ASP Conf. Ser. 283, A New Era in Cosmology, ed. N. Metcalfe & T. Shanks (San Francisco: ASP), 193
Massey, R., Refregier, A., Conselice, C. J., David, J., & Bacon, J. 2004a, MNRAS, 348, 214
———. 2004b, MNRAS, 348, 214
Massey, R., et al. 2007a, Nature, 445, 286
———. 2007b, MNRAS, 376, 13
———. 2007c, ApJS, 172, 239
Mobasher, B., et al. 2007, ApJS, 172, 117
Natarajan, P., Kneib, J.-P., Smail, I., & Ellis, R. S. 1998, ApJ, 499, 600
Okura, Y., Umetsu, K., & Futamase, T. 2006, preprint (astro-ph/0607288)
Peterson, B. A., Ellis, R. S., Kibblewhite, E. J., Bridgeland, M. T., Hooley, T., & Horne, D. 1979, ApJ, 233, L109
Refregier, A., Rhodes, J., & Groth, E. J. 2002, ApJ, 572, L131
Réfrégier, A., et al. 2006, in Proc. SPIE 6265, 58
Rhodes, J., Refregier, A., Collins, N. R., Gardner, J. P., Groth, E. J., & Hill, R. S. 2004, ApJ, 605, 29
Rhodes, J., Refregier, A., & Groth, E. J. 2000, ApJ, 536, 79 (RRG)
Rhodes, J. D., Refrigier, A., & Groth, E. J. 2001, ApJ, 552, L85
———. 2007, ApJS, 172, 203
Rix, H.-W., et al. 2004, ApJS, 152, 163

Robin, A. C., Reylé, C., Derrière, S., & Picaud, S. 2003, A&A, 409, 523
———. 2004, A&A, 416, 157
———. 2007, ApJS, 172, 545
Schrabback, T., et al. 2006, preprint (astro-ph/0606611)
Scoville, N. Z., et al. 2007, ApJS, 172, 1
Semboloni, E., et al. 2006, A&A, 452, 51

Sheldon, E. S., et al. 2004, AJ, 127, 2544
Taniguchi, Y., et al. 2007, ApJS, 172, 9
Tyson, J. A., Wenk, R. A., & Valdes, F. 1990, ApJ, 349, L1
Williams, R., et al. 1998, BAAS, 30, 1366
Williams, R. E., et al. 1996, AJ, 112, 1335
Wolf, C., et al. 2004, A&A, 421, 913