

SHAPES AND SHEARS, STARS AND SMEARS: OPTIMAL MEASUREMENTS FOR WEAK LENSING

G. M. BERNSTEIN AND M. JARVIS

Department of Astronomy, University of Michigan, Ann Arbor, MI 48109; garyb@astro.lsa.umich.edu, jarvis@astro.lsa.umich.edu

Received 2001 July 23; accepted 2001 October 11

ABSTRACT

We present the theoretical and analytical bases of optimal techniques to measure weak gravitational shear from images of galaxies. We first characterize the geometric space of shears and ellipticity and then use this geometric interpretation to analyze images. The steps of this analysis include measurement of object shapes on images, combining measurements of a given galaxy on different images, estimating the underlying shear from an ensemble of galaxy shapes, and compensating for the systematic effects of image distortion, bias from point-spread function (PSF) asymmetries, and “dilution” of the signal by the seeing. These methods minimize the ellipticity measurement noise, provide calculable shear uncertainty estimates, and allow removal of systematic contamination by PSF effects to arbitrary precision. Galaxy images and PSFs are expressed as “Laguerre expansions,” a two-dimensional generalization of the Edgeworth expansion, making the PSF correction and shape measurement relatively straightforward and computationally efficient. We also discuss sources of noise-induced bias in weak-lensing measurements—selection biases, and “centroid” biases arising from noise rectification—and provide a solution for these and previously identified biases.

Key words: gravitational lensing — methods: data analysis — techniques: image processing

On-line material: color figures

1. INTRODUCTION

Gravitational lensing is a powerful tool for studying the distribution of matter in the universe because photons are deflected by all forms of matter, regardless of luminosity, composition, or dynamical state. Dramatic manifestations of lensing—multiple images, Einstein rings, and giant arcs, so-called *strong* lensing—provide much information on the highest overdensities in the universe, namely, rich galaxy clusters, cores of individual galaxies, or collapsed objects. To characterize the more typical mass structures, or those without a fortuitously aligned bright background source, we may use *weak* gravitational lensing, in which we analyze the low-order distortions of the ubiquitous background galaxies in order to infer the mass distribution. Weak gravitational lensing signals are extraordinarily subtle, even by astronomical standards: one seeks a shear (or magnification) of the galaxy images amounting to a few percent at most, more typically 0.2%–1% in current studies. Because the undistorted image is not observable, the lensing distortions must be detected as a perturbation to the intrinsic distribution in galaxy shapes (or sizes), which have variation of 30% or more, giving a signal-to-noise ratio (S/N) of $\sim 1/30$ from observation of a single galaxy. Hence a very large number of galaxies must be observed before the weak lensing becomes detectable over this intrinsic *shape noise*. Weak-lensing analyses could not even be attempted until automated means of measuring very large numbers of galaxy shapes became available (Valdes, Tyson, & Jarvis 1983; Tyson et al. 1984). Furthermore, optical and atmospheric distortions in a typical sky image cause coherent shape (and size) distortions that can masquerade as a lensing signal. Such systematic errors are 1%–10% in a typical image, up to 50 times larger than the weak-lensing signals. A means to remove this contamination is crucial; the necessary analyses can only be conducted with well-calibrated, linear detectors.

Successful detection of a weak-lensing signal did not occur until CCD images of sufficient depth and field were available (Tyson, Valdes, & Wenk 1990), and early detections were of the $\approx 10\%$ shears that are found in the inner regions of rich clusters of galaxies (Fahlman et al. 1994; Bonnet, Mellier, & Fort 1994; Smail et al. 1995). In regions of strong shear, the S/N is sufficiently high that a map of the lensing mass can be created (Kaiser & Squires 1993). Mellier (1999) includes a review of results from cluster lensing studies.

With the increase in collecting area of CCD imagers, sufficient background galaxies can be measured to allow convincing detection of smaller shear signals around weaker overdensities: around individual weak clusters (Fischer et al. 1997) or collections of galaxy groups (Hoekstra et al. 2001); and around individual galaxies (Fischer et al. 2000; Smith et al. 2001; Wilson et al. 2001b). Most dramatically, lensing signals on random lines of sight, caused by the background matter fluctuation spectrum, have now been detected and are one tool for “precision cosmology” (Wittman et al. 2000; Van Waerbeke et al. 2000; Bacon et al. 2001; Wilson, Kaiser, & Luppino 2001a). As technology has advanced, weaker and weaker shears have become detectable under the shape noise, sometimes as small as a few tenths of a percent (e.g., Fischer et al. 2000; Jarvis et al. 2002). As a consequence, the demands for rejection of systematic errors have become more stringent. In many current weak-lensing publications, it is clear that the uncorrected systematic effects are only slightly smaller than the signals under study. It is therefore fair to say that, at present, it is the analysis techniques, rather than the ability to collect galaxy images, that bar the way to higher precision in many weak-lensing studies.

This paper describes the techniques for extraction of weak-lensing signals from imaging data, which we have developed over the past few years to meet these increasing demands. As described below, our efforts focus on the shear

rather than the magnification of the galaxy images by the lens, and hence we are measuring galaxy ellipticities. The desiderata for a weak-lensing methodology include the following:

1. Shapes of individual galaxy images are determined with the highest possible accuracy in the presence of measurement noise on the image.
2. Each shape measurement should have a known error distribution.
3. Individual galaxy shapes should be combined to yield an estimate of the underlying lens shear with maximal S/N.
4. The shear estimator should have an error level and a calibration that can be derived directly from the data without recourse to Monte Carlo simulations.
5. The galaxy shapes should be corrected for the systematic biases due to the point-spread function (PSF) to arbitrary precision.
6. The scheme must allow for a PSF that varies continuously across the image and is different in each exposure.

Given the intrinsic floor on weak-lensing accuracy because of shape noise, one might ask why we should expend much effort on goal 1, which is to minimize the effects of measurement noise—normally, we consider that once the ellipticity measurement noise σ_e is $\ll \sigma_{\text{SN}} \equiv \langle e^2/2 \rangle^{1/2} \approx 0.3$, further gains do not increase the shear estimation accuracy—the error σ_δ on the lensing distortion δ will just become $\sigma_{\text{SN}}/\sqrt{N}$, with N the number of measured galaxies. We note first that the sky density of galaxies scales with apparent magnitude m as $10^{\alpha m}$, with $0.3 \lesssim \alpha \lesssim 0.4$. If we can cut the shape-measurement error for a given image noise level, then we can either use fainter galaxies in our lensing measurement (increasing N) or cut the required exposure time. Second, note that convolution with a PSF suppresses the measured lensing signal *and* the intrinsic shape noise. Hence the level to which we aim to reduce σ_e must, for poorly resolved galaxies, be well below the canonical 0.3. Thirdly, we will see in § 4.2 that it may be possible to measure the shear to an accuracy much better than $\sigma_{\text{SN}}/\sqrt{N}$, in cases where the distribution of intrinsic galaxy shapes in the ellipticity plane has a cusp or pole at $e = 0$. In simple cases, such as a population of circular disk galaxies, the accuracy to which we can measure the applied shear can increase without limit as the measurement noise is decreased.

The need for traceable uncertainties is also critical, as weak lensing is used to measure the power spectrum of mass fluctuations in the universe. In this application, the measurement uncertainties (including shape noise) contribute to the power spectrum and must be accurately estimated and subtracted to reveal the true cosmic power spectrum. Of course, an accurate calibration is also necessary for most applications to precision cosmology; if one must rely on simulated data for the calibration, there is always the danger that the simulations do not properly incorporate some aspect of the real world.

Finally, the need for removal of the systematic PSF ellipticities to arbitrary precision is extremely strong. In the course of this paper, we will try to describe other approaches to the problem and compare with our own. The methods in most common current use (e.g., Kaiser, Squires, & Broadhurst 1995, hereafter KSB) are formally valid only in certain special cases of PSF. While heuristic adjustment and testing

has demonstrated that the method works to nominal accuracy in more general cases, the absence of a generally valid method is troubling. A formally exact PSF correction scheme has been put forward by Kaiser (2000, hereafter K00), which is based upon a Fourier-domain calculation of the effects of shear and of PSF convolution. Our approach will be to decompose the image and the PSF into a vector over orthogonal polynomials and treat the deconvolution as a matrix operation carried to desired order. A very similar approach has been independently put forth by Refregier (2001).

This is a long-winded paper, likely to be read in detail only by practitioners of weak lensing. A more casual reading will be beneficial to those who wish to understand the methods and limitations of past and future weak-lensing analyses. Some of the techniques we develop may be useful beyond the weak-lensing analysis, for example, our deconvolution method (§ 6.3.5) and the methods for rapid convolution with spatially varying kernels (§ 7). As discussed by Refregier (2001), our orthogonal function decomposition can be a useful means for compression of galaxy images.

The paper outline is as follows: The following section describes the mathematical space occupied by ellipticities and shears. Understanding the geometry of this space makes it easier to see how our (and other) measurement schemes work. In § 3, we describe a scheme that uses our geometric conceptualization of ellipticity to produce measurements with maximal S/N in images with infinitesimal PSF; a formula for the resultant uncertainty in each ellipticity is also derived. Next, § 4 discusses several schemes for combining shape measurements of a given galaxy from different exposures and/or filter bands, to obtain the shape estimate that again offers the best possible S/N and a closed-form error estimate. Section 5 describes the means to combine shape estimates from different galaxies to form an optimal estimate of the underlying lensing shear. In the absence of measurement noise, this takes a simple closed form; in the presence of measurement noise, some approximations must be made to obtain a closed form for the calibration and error of the shear, and hence we do not fully satisfy goal 4 above. Section 6 is a very extensive discussion of the effects of the PSF on the image, other approaches to the problem, and our method for optimal extraction of the intrinsic shape in this case. In this section we introduce the Laguerre decomposition technique. Section 7 uses the Laguerre formalism to construct convolution kernels that can add symmetries to the PSF of an image; this is one means of removing the ellipticity biases due to the PSF. Section 8 discusses two very important effects that can give rise to biased lensing measurements *even when a perfect deconvolution for PSF effects is possible*. It is likely that these biases are present in all previously published data.

Finally, § 9 puts together all the methods developed in the paper in a flowchart form, describing how raw image data are converted into optimized, calibrated lensing shear data. We reserve for a succeeding paper (Jarvis et al. 2002) the detailed discussion of the code that implements these methods, and a verification of its performance on real and simulated data. In appendices to this paper, we present the formulae for invoking various transformations on the Laguerre decomposition representation of an image and derive some approximate PSF correction formulae that were used for the analyses of Smith et al. (2001) but which are superseded by the full Laguerre methodology.

2. GEOMETRY OF SHAPE AND SHEAR

2.1. Linear Approximation to Lensing

The goal of weak gravitational lensing studies is to infer a distant gravitational potential via the distortions that the potential's deflection of light imparts upon the population of galaxies in the background. The lensing is fully characterized by the map $\mathbf{u}(\mathbf{x})$ from the observed angular position \mathbf{x} to the source angular position \mathbf{u} . The surface brightness observed at \mathbf{x} is equal to that which would have been observed at \mathbf{u} in the absence of the lens. For an individual background galaxy that is not near a lensing caustic, the map can be accurately approximated by a Taylor expansion

$$\mathbf{u} = \frac{d\mathbf{u}}{d\mathbf{x}}\mathbf{x} + \mathbf{u}_0. \quad (2.1)$$

The displacement \mathbf{u}_0 carries no information (unless the source is multiply imaged), because the source plane is unobservable. The 2×2 amplification matrix has a unique decomposition of the form

$$\frac{d\mathbf{x}}{d\mathbf{u}} = \mu \mathbf{S} \mathbf{R}, \quad (2.2)$$

where \mathbf{R} is an orthogonal matrix (rotation), \mathbf{S} is a symmetric matrix with unit determinant (shear), and μ is a scalar magnification.

The rotation \mathbf{R} is not useful for lensing studies because the unlensed orientation is not known, and the ensemble of background galaxies should be isotropic, and hence any collective statistic should be unchanged by rotation. Furthermore, the rotation is absent in the limits of single-screen or weak lensing.

The magnification μ increases the angular size by μ and the galaxy flux by μ^2 . While the unlensed quantities are not observable, the magnification is still detectable because the mean flux and size of the population will shift. The magnification also reduces the sky-plane density of sources by μ^2 . The magnification thus modulates the number versus magnitude relations for a given class of background galaxies, in a manner that depends upon the size/magnitude/redshift distribution of the original population.

The shear \mathbf{S} has 2 degrees of freedom, corresponding to the ellipticity and position angle imparted on a circular source galaxy. For weak lensing, this shear is undetectable on a single galaxy because the unlensed shape is not necessarily circular and is not observed. The collective distribution of galaxy shapes is assumed to be intrinsically isotropic, and the applied shear breaks this symmetry, rendering it detectable and measurable.

Both the shear and the magnification thus produce measurable effects on the ensemble of galaxies and can in theory be used to quantify the potential. Shear measurements have been used for numerous quantitative studies, but magnification methods still yield, at best, marginal detections (Dye et al. 2001). There are several factors that favor the shear method: First, the two effects of magnification (increased flux and reduced areal density) push the counts of background sources in opposite directions, weakening the signal. More importantly, the shear is manifested as a variation in the mean orientation of galaxy shapes, and this mean is zero in the absence of lensing; the magnification signal is a modulation of $N(m)$ or some other nonzero quantity. It is always far easier to measure a small change from zero than a small

change in a nonzero quantity. For example, exploitation of the magnification effect in the weak-lensing regime would require absolute photometry to much better than 0.01 mag accuracy. Magnification measurements, on the other hand, give a direct measure of the projected mass, whereas mass reconstructions from shear data are degenerate under the addition of a constant-density mass sheet. Hence, magnification data are very useful when there is no a priori means of determining the mean mass overdensity in the image.

Henceforth we will ignore the magnification effect and describe how to optimally measure the shear \mathbf{S} .

2.2. Parameterizations of Shear

2.2.1. Diagonal Shears

The simplest shear matrix is a small perturbation aligned with the coordinate axes:

$$\mathbf{S}_\eta = \begin{pmatrix} 1 + \frac{1}{2}\eta & 0 \\ 0 & 1 - \frac{1}{2}\eta \end{pmatrix}, \quad \eta \ll 1. \quad (2.3)$$

The effect of this transformation upon a circular source-plane object is to induce an elongation along the x -axis, creating an elliptical image with axis ratio $q \equiv b/a = 1 - \eta$. We can use this matrix as a generator for the full family of diagonal shear matrices with arbitrary η to obtain

$$\mathbf{S}_\eta = \begin{pmatrix} e^{\eta/2} & 0 \\ 0 & e^{-\eta/2} \end{pmatrix}, \quad -\infty < \eta < \infty. \quad (2.4)$$

The set of diagonal shear matrices forms a group under simple matrix multiplication. The operation is commutative and clearly corresponds to simple addition of the η -parameters:

$$\mathbf{S}_{\eta_3} = \mathbf{S}_{\eta_2} \times \mathbf{S}_{\eta_1} \quad \Leftrightarrow \quad \eta_3 = \eta_2 + \eta_1. \quad (2.5)$$

For this reason we will call η the *conformal shear* and will find it a useful parameterization of shear. Other common parameterizations of shear include the axis ratio q , the *distortion* δ (Miralda-Escudé 1991), and the *reduced shear* $g = \gamma/(1 - \kappa)$ (Schneider & Seitz 1995), which are related to η via

$$q \equiv b/a = e^{-\eta}, \quad (2.6)$$

$$\delta \equiv \frac{a^2 - b^2}{a^2 + b^2} = \tanh \eta, \quad (2.7)$$

$$g \equiv \frac{1 - q}{1 + q} = \tanh\left(\frac{\eta}{2}\right). \quad (2.8)$$

Bonnet & Mellier (1995) define a further set of shear parameterizations, also easily expressed in terms of η :

$$e \equiv 1 - q = 2e^{-\eta/2} \sinh\left(\frac{\eta}{2}\right),$$

$$\epsilon \equiv \frac{1 - q^2}{1 + q^2} = \tanh \eta,$$

$$\delta_B \equiv \frac{1 + q^2}{2q} = \cosh \eta,$$

$$\tau \equiv \epsilon \delta_B = \sinh \eta.$$

Note that for $\eta \ll 1$, the parameters η , e , ϵ , τ , and the distortion δ are all equal. Note also that most other author's

formulations of shear do not define the matrix \mathbf{S} to have unit determinant and so do not form a group.

2.2.2. General Shear

A general (nondiagonal) shear matrix can be decomposed into a diagonal shear and rotations as

$$\begin{aligned} \mathbf{S}_{\eta,\beta} &= \mathbf{R}_\beta \mathbf{S}_\eta \mathbf{R}_{-\beta} \\ &= \begin{pmatrix} \cosh \frac{1}{2}\eta + \cos \theta \sinh \frac{1}{2}\eta & \sin \theta \sinh \frac{1}{2}\eta \\ \sin \theta \sinh \frac{1}{2}\eta & \cosh \frac{1}{2}\eta - \cos \theta \sinh \frac{1}{2}\eta \end{pmatrix}. \end{aligned} \quad (2.9)$$

$\mathbf{S}_{\eta,\beta}$ transforms a circular source to an ellipse with axis ratio $q = e^{-\eta}$ at position angle $\beta = \theta/2$. The shear can be represented as a two-dimensional vector

$$\boldsymbol{\eta} \equiv (\eta_+, \eta_\times) \equiv (\eta \cos \theta, \eta \sin \theta). \quad (2.10)$$

Likewise, a shear may be represented as a two-dimensional distortion $(\delta_+, \delta_\times)$, etc. A shear $(\eta_+, 0)$ creates ellipses oriented to the x - or y -axis, while $(0, \eta_\times)$ aligns circular sources to axes at 45° to the coordinate axes. The shear $\boldsymbol{\eta}$ is not a vector in the image space but, rather, is a vector in a non-Euclidean shear space that we describe below.

The full set of shear matrices do not form a group under matrix multiplication, as $\mathbf{S}_{\eta_2} \mathbf{S}_{\eta_1}$ may be asymmetric (two-screen lenses can effect a rotation for this reason). But we can form a group with an addition operation for two-dimensional shears defined as

$$\boldsymbol{\eta}_3 = \boldsymbol{\eta}_2 \oplus \boldsymbol{\eta}_1 \Leftrightarrow \mathbf{S}_{\boldsymbol{\eta}_3} \mathbf{R} = \mathbf{S}_{\boldsymbol{\eta}_2} \mathbf{S}_{\boldsymbol{\eta}_1}, \quad (2.11)$$

where \mathbf{R} is the unique rotation matrix that allows $\mathbf{S}_{\boldsymbol{\eta}_3}$ to be symmetric. The geometric meaning of the shears is preserved, since \mathbf{R} will leave a circular source unchanged. The simplest expression of the composition operation in terms of components is

$$\cosh \eta_3 = \cosh \eta_2 \cosh \eta_1 + \sinh \eta_2 \sinh \eta_1 \cos(\theta_2 - \theta_1), \quad (2.12a)$$

$$\sinh \eta_3 \sin(\theta_3 - \theta_2) = \sinh \eta_1 \sin(\theta_1 - \theta_2). \quad (2.12b)$$

Note that the second equation is *not* symmetric in the two operands, and hence the shear matrix group is non-Abelian. The identity element is $\eta = 0$, and the inverse of $\boldsymbol{\eta} = (\eta_+, \eta_\times)$ is $-\boldsymbol{\eta} = (-\eta_+, -\eta_\times)$. The addition formula in terms of distortion components is derivable from equation (2.12) and is given by Miralda-Escudé (1991):

$$\begin{aligned} \delta_{3+} &= \frac{\delta_{1+} + \delta_{2+} + (\delta_{2\times}/\delta_2^2)[1 - \sqrt{1 - \delta_2^2}](\delta_{1\times}\delta_{2+} - \delta_{1+}\delta_{2\times})}{1 + \boldsymbol{\delta}_1 \cdot \boldsymbol{\delta}_2}, \\ \delta_{3\times} &= \frac{\delta_{1\times} + \delta_{2\times} + (\delta_{2+}/\delta_2^2)[1 - \sqrt{1 - \delta_2^2}](\delta_{1+}\delta_{2\times} - \delta_{1\times}\delta_{2+})}{1 + \boldsymbol{\delta}_1 \cdot \boldsymbol{\delta}_2}. \end{aligned} \quad (2.13)$$

We omit the derivations of these equations, which are straightforwardly but tediously executed by composing the transformation matrices. A more elegant derivation follows from noting that the transformation equation (2.9) trans-

forms the complex plane as¹

$$z \rightarrow z \cosh \frac{1}{2}\eta +; \bar{z} e^{i\theta} \sinh \frac{1}{2}\eta. \quad (2.14)$$

It will be useful to consider the limit in which $\delta_2 \ll 1$:

$$\left. \begin{aligned} (d\boldsymbol{\delta} \oplus \boldsymbol{\delta})_+ &\approx \delta + (1 - \delta^2)d\delta_+ \\ (d\boldsymbol{\delta} \oplus \boldsymbol{\delta})_\times &\approx d\delta_\times \end{aligned} \right\} d\delta \ll 1, \quad \delta_\times = 0. \quad (2.15)$$

If we instead make $\delta_1 \ll 1$, the asymmetry of shear addition is manifested as a change to the azimuthal component formula:

$$\left. \begin{aligned} (\boldsymbol{\delta} \oplus d\boldsymbol{\delta})_+ &\approx \delta + (1 - \delta^2)d\delta_+ \\ (\boldsymbol{\delta} \oplus d\boldsymbol{\delta})_\times &\approx \sqrt{1 - \delta^2}d\delta_\times \end{aligned} \right\} d\delta \ll 1, \quad \delta_\times = 0. \quad (2.16)$$

2.3. The Shear Manifold

We define a metric distance between two points $\boldsymbol{\eta}_3$ and $\boldsymbol{\eta}_1$ in shear space as

$$s \equiv |\boldsymbol{\eta}_3 - \boldsymbol{\eta}_1| = |\eta_2|, \quad \boldsymbol{\eta}_2 \equiv \boldsymbol{\eta}_3 \oplus (-\boldsymbol{\eta}_1). \quad (2.17)$$

The differential form of the metric can be derived by specializing equations (2.12a) and (2.12b) to the case $ds = \eta_2 \ll 1$, $\theta_1 = 0$, $\theta_2 = \theta$, yielding

$$\eta_3 = \eta_1 + ds \cos \theta, \quad (2.18)$$

$$\theta_3 \tanh \eta_1 = ds \sin \theta, \quad (2.19)$$

which means that the metric is

$$ds^2 = (\eta_3 - \eta_1)^2 + (\tanh^2 \eta_1)(\theta_3 - \theta_1)^2 \quad (2.20)$$

$$= d\eta^2 + \tanh^2 \eta d\theta^2 \quad (2.21)$$

$$= (1 - \delta^2)^2 d\delta^2 + \delta^2 d\theta^2. \quad (2.22)$$

Note that the η version of the metric has the normal Euclidean form for the radial component, and the δ parameterization has the Euclidean form for the tangential component of the metric, but neither representation gives a fully Euclidean metric—the shear space is curved. The two-dimensional shear manifold defined by this metric can be embedded in Euclidean 3-space, as illustrated in Figure 1. This geometric depiction of shear is helpful in understanding the transformations of shears. Near $\eta = 0$, the surface is tangent to the Euclidean plane, so small shears add with Euclidean componentwise addition. The shear-space surface then curves upward, and as the conformal radius η grows large, the surface approaches a cylinder of unit radius about the z -axis. If we project the shear surface onto the $z = 0$ plane, the radius vector in this plane is equal to the distortion δ . The $\boldsymbol{\delta}$ vector is confined within the unit circle.

2.4. Definition of Shape

A *shear* is a transformation of the image plane; we next need a quantity to describe the *shape* of an arbitrary galaxy image. Let I represent some object whose isophotes are a family of similar ellipses. We can simply parameterize the

¹ We thank the anonymous referee for this derivation.

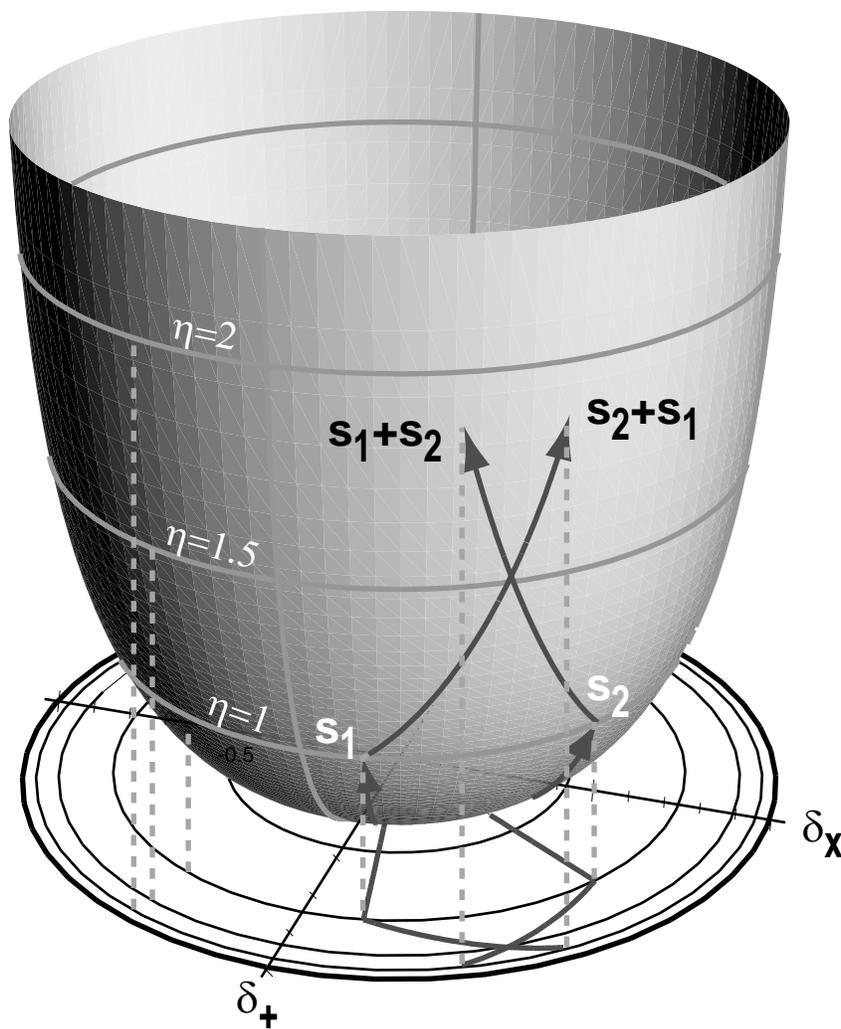


FIG. 1.—The shaded surface is an embedding of the shear manifold into Euclidean space. The radius vector along this surface is the conformal shear η ; the radius upon projection onto the x - y plane is the distortion δ (or ellipticity e). At small η the manifold is tangent to the δ -plane, and at large η approaches the unit cylinder. Two shear vectors s_1 and s_2 of length $\eta = 1$ are plotted from the origin, both on the true shear manifold and in the δ -plane. The result of adding the two vectors is also plotted; displacements do not commute in this non-Euclidean space.

shape of I by the shear η , which produces this object from some object I_0 having circular isophotes, i.e.,

$$I_\eta = S_\eta I_0 \Rightarrow S_{-\eta} I_\eta = I_0. \tag{2.23}$$

We could thus call η the *conformal shape* of the object and can think of a given ellipse as a location on the shear manifold. More commonly the distortion is used to define the shape; an object is said to have *ellipticity* e if a shear with distortion $\delta = -e$ makes it circular. We will use the symbol e since this quantity agrees with the traditional second-moment definition of ellipticities for truly elliptical objects. Equation (2.23) makes it obvious how an ellipse with shape η_1 will be transformed under the action of a shear η_2 ; we simply add the shear to the shape using the addition rules of shear space (eq. [2.12]):

$$S_{\eta_2} I_{\eta_1} = I_{\eta_3}, \quad \text{with} \quad \eta_3 = \eta_2 \oplus \eta_1. \tag{2.24}$$

Likewise we can also say that a distortion δ maps the ellip-

ticity $e \rightarrow \delta \oplus e$. In general, an applied shear may be viewed as a shift of all shapes along the shear manifold. We will use e to represent the *shape* of an image, whereas δ represents a *shear*, which is a transformation of the image plane. The shape and shear spaces, however, transform identically under an applied shear.

A real galaxy has some image intensity distribution $I(x)$, which may not have elliptical isophotes; we would like to define a shape for an arbitrary image. By analogy to equation (2.23), all that is needed is some definition of a “round” image. Let $\mathbf{M}(I)$ be any measurement applied to the image that has the simple property that if $\mathbf{M}(I) = 0$, then $\mathbf{M}(S_\eta I) \neq 0$ for any nonzero shear η . Then the condition $\mathbf{M}(I) = 0$ is our definition of roundness, and we can assign a unique shape η to an arbitrary object $I(x)$ by the condition

$$\mathbf{M}(S_{-\eta} I) = 0. \tag{2.25}$$

Any shape defined by such a rule clearly transforms under

an applied shear, just as an ellipse does, namely, via equation (2.24). With any roundness measure \mathbf{M} we therefore have a definition of shapes and their mapping under shears that follow the rules of addition in shear space. We do not attempt to prove that the solution to equation (2.25) exists or is unique.

2.5. Shear in Fourier Space

For consideration of the effects of convolution upon sheared images, it will be useful to ponder the action of shears in Fourier space. We first note that shearing an image $I(\mathbf{x})$ by $\boldsymbol{\eta}$ is equivalent to shearing its Fourier transform $\tilde{I}(\mathbf{k})$ by $-\boldsymbol{\eta}$. For a diagonal shear,

$$\widetilde{\mathbf{S}}_{\boldsymbol{\eta}} I(k_x, k_y) = \frac{1}{2\pi} \int d^2x I(e^{-\eta/2}x, e^{\eta/2}y) \exp[2\pi i(k_x x + k_y y)] \quad (2.26)$$

$$= \frac{1}{2\pi} \int d^2x I(x, y) \exp[2\pi i(e^{\eta/2}k_x x + e^{-\eta/2}k_y y)] \quad (2.27)$$

$$= (2\pi)^{-1} \tilde{I}(e^{\eta/2}k_x, e^{-\eta/2}k_y). \quad (2.28)$$

A nondiagonal shear must also be equivalent to an opposing shear in k -space, since rotation of the real-space function corresponds to rotation in k -space. We can therefore just as easily define a galaxy's shape by a roundness criterion in k -space as in real space. This is useful when considering finite resolution (§ 6) or when analyzing interferometric images with limited Fourier coverage (Chang & Refregier 2002).

3. OPTIMAL MEASUREMENTS (WITHOUT SEEING)

In this section, we derive an optimum method of measuring galaxy shapes in the case where the angular resolution and sampling of the instrument are assumed to be perfect. In § 6, we will treat the more complex case of seeing-convolved images.

3.1. The Ideal Test for Roundness

We have defined the shape of an image in equation (2.25) by asking what coordinate shear is needed to make the object appear round. We need to choose a measurement \mathbf{M} that detects with the highest possible S/N any small departure of the image from its round state $\mathbf{M} = 0$. It can be shown that, under some sensible simplifying assumptions about \mathbf{M} , the solution of equation (2.25) becomes equivalent to finding the best least-squares fit of an elliptical-isophote model to the galaxy image.

We assume first that the measurement will be a linear function of $I(\mathbf{x})$. Any nonlinear method will prove extremely difficult to apply to the case in which the image has been convolved with a PSF—it will be hard to use the measurements of bright stars to correct the shapes of faint galaxies for convolution. The most general form of \mathbf{M} is then

$$\mathbf{M}(I) = \int d^2x \mathbf{W}(\mathbf{x}) I(\mathbf{x}), \quad (3.1)$$

where \mathbf{W} is some weight function. The weight will be two-dimensional, as the measurement must test for

departure from roundness in both the η_+ and η_\times directions in shear space.

We consider first the weight component to detect a small change in η_+ . We can decompose the image $I(r, \theta)$ into multipole elements $I_m(r)$ via

$$I(\mathbf{x}) = I(r, \theta) = \sum_{m=-\infty}^{\infty} I_m(r) e^{im\theta}, \quad (3.2)$$

$$I_m(r) = \frac{1}{2\pi} \int_0^{2\pi} d\theta I(r, \theta) e^{-im\theta}. \quad (3.3)$$

We are interested in the change in our measurement upon mapping of the image $I(\mathbf{x}) \rightarrow \tilde{I}(\mathbf{x}) = I(\mathbf{S}\mathbf{x})$, where \mathbf{S} is a shear of amplitude $\eta \ll 1$ oriented on the x -axis, as in equation (2.3). The quantity for which we wish to optimize the S/N can thus be written as

$$\delta M = M(\tilde{I}) - M(I) = \sum_{m=-\infty}^{\infty} \int_0^{\infty} r dr w_m(r) [\tilde{I}_m(r) - I_m(r)], \quad (3.4)$$

with w_m an arbitrary radial function for each multipole. A little bit of algebra yields the transformation of multipoles

$$\tilde{I}_m(r) - I_m(r) = \frac{1}{4}\eta[(m-2)I_{m-2} - rI'_{m-2} - (m+2)I_{m+2} - rI'_{m+2}] + O(\eta^2), \quad (3.5)$$

where the primes denote derivatives with respect to r . For an object with truly circular isophotes, we have $I_m(r) = 0$ for $m \neq 0$, and the only effect of the shear is to induce a quadrupole term $\tilde{I}_2 = -\eta r I'_0/4$. For objects without perfect circular symmetry, there are terms beyond the monopole. But for an object to be thought of as “round,” the monopole term I_0 should dominate the higher multipoles. The monopole is also the only term guaranteed to be positive for all galaxies. Hence the largest effect of the shear will be to alter the $m = \pm 2$ quadrupole intensities (which are conjugates of each other, as I is real). The optimal sensitivity to small shear should therefore weight only the quadrupole term:

$$M(I) = \int_0^{\infty} r dr I_2(r) r^2 w(r) = \frac{1}{2\pi} \iint r dr d\theta I(r, \theta) w(r) r^2 e^{-2i\theta}. \quad (3.6)$$

It is clear that this quadrupole test is the optimal linear measurement for objects with circular symmetry; for more general shapes the shear has effects on other multipoles that can be measured and used to enhance S/N (related to the suggestions of Refregier 2001). But this would require a knowledge of $I_4(r)$ and the other multipoles to construct the ideal formula; we settle on the simple quadrupole as the best general solution, as we are always guaranteed that I_0 is present and positive for any real source. The measurement of some weighted quadrupole is also the normal definition of ellipticity for weak-lensing measurements (Miralda-Escudé 1991).

Combining equations (3.4), (3.5), and (3.6), we obtain

$$\delta M = \frac{-\eta}{4} \int r dr r^2 w(r) (rI'_0 + 4I_4 + rI'_4). \quad (3.7)$$

The noise in the measure of M can be derived in two limits: the most common case will be sky-dominated observations, for which the variance of the flux in area A is nA , where n is

the number of sky photons per unit area (it is assumed that I is in units of photons). In this case we have, from equation (3.6), the variance of each component of δM :

$$\text{Var}(M) = \frac{n}{4\pi^2} \iint r dr d\theta r^4 w^2(r) \cos^2 2\theta. \quad (3.8)$$

If we ignore the I_4 terms in equation (3.7) as being dominated by I_0 terms, then the choice of weight function that optimizes the detectability of the shear η is

$$w_{\text{opt}}(r) = \frac{-I'_0}{r} = -\frac{1}{r} \frac{dI_0}{dr}. \quad (3.9)$$

With this optimal weight, the variance of the measurement M would lead to an error in each component of η equal to

$$\sigma_\eta^2 = \frac{4n}{\pi} \left\{ \int r^3 dr [I'_0(r)]^2 \right\}^{-1}. \quad (3.10)$$

If the object is much brighter than the night sky, then the noise is no longer uniform and the optimization becomes

$$\text{Var}(M) = \frac{1}{4\pi^2} \iint r dr d\theta r^4 w^2(r) I(r) \cos^2 2\theta, \quad (3.11a)$$

$$w_{\text{opt}}(r) = \frac{-I'_0}{r I_0} = -\frac{1}{r} \frac{d \ln I_0}{dr}, \quad (3.11b)$$

$$\sigma_\eta^2 = \frac{4}{\pi} \left[\int r^3 dr (I'_0)^2 / I_0 \right]^{-1}. \quad (3.11c)$$

3.1.1. Gaussian Objects

An elliptical Gaussian object, when sheared to be circular, will obey

$$I(r, \theta) = I_0(r) = \frac{f}{2\pi\sigma^2} e^{-r^2/2\sigma^2}. \quad (3.12)$$

In the sky-limited case the optimal weight is the same Gaussian:

$$w_{\text{opt}}(r) = -I'_0/r \propto e^{-r^2/2\sigma^2}. \quad (3.13)$$

Note that the optimal weight for shape measurement is in this case equal to the optimal filter for detection, i.e., a matched filter. If we define the detection significance ν as the S/N for detection of the object with the matched filter, we find

$$\nu^2 = \frac{[\int dA w(r) I(r)]^2}{\int dA n w^2(r)} = \frac{f^2}{4\pi n \sigma^2}, \quad (3.14)$$

$$\sigma_\eta^2 = 16\pi n \sigma^2 / f^2 = (2/\nu)^2. \quad (3.15)$$

We therefore end up with the simple result that the error in each component of the shape is 2 over the detection significance.

The above derivations assume that the center and the size σ of the Gaussian are known in advance. If there were a sky filled with Gaussian galaxies, we likely would not know in advance the size and location of each. We can determine the centroid in the usual manner by forcing the weighted first moments to vanish:

$$\int dA w I r e^{i\theta} = 0. \quad (3.16)$$

The weight for centroiding does not necessarily have to match that used for the shape measurement, but it is convenient to do so. The proper size σ_w for the weight can be forced to match the size σ of the object by requiring the significance to be maximized:

$$0 = \frac{\partial \nu}{\partial \sigma_w} \propto \int dA w I \left(1 - \frac{r^2}{\sigma_w^2}\right). \quad (3.17)$$

In the limit of a Gaussian with low background noise, equations (3.11a)–(3.11c) apply, and the optimal weight is *uniform*. The detection significance in this case is just $\nu = \sqrt{f}$ (with the flux f in photons), and we find again that the standard error in η is equal to $2/\nu$. In practice this situation can never be realized because the weight extends to infinity, and at large radii the sky or read noise will dominate the shot noise from the galaxy, and neighboring objects will impinge upon the integrations. We will henceforth confine our discussion to background-limited observations.

3.1.2. Exponential or Other Profiles

To obtain an optimized measurement of a real galaxy, we would have to measure its radial profile and construct a custom optimized weight using equation (3.9). The majority of galaxies are spirals or dwarfs, which are typically described by exponential profiles:

$$I(r) \propto e^{-\alpha r} \Rightarrow w_{\text{opt}}(r) = e^{-\alpha r} / r. \quad (3.18)$$

This weight diverges at the origin, though all the necessary integrals of the weight are convergent. If the galaxy is truly cusped in the center, then the intensity near the center is very sensitive to small shears and is weighted heavily.

In practice it is simpler to adopt a weight function that is universal (up to a scale factor), especially at low S/N, where attempts to measure each individual profile would be pointless. There are a number of reasons to prefer a Gaussian weight:

1. The Gaussian drops very quickly at large radii, minimizing interference from neighboring objects. Integrals of all moments are convergent.

2. Weights with central divergences or cusps are difficult to use in data with finite sampling, and they also amplify the effects of seeing on the galaxy shapes. The Gaussian is flat at $r = 0$.

3. Gaussian weights are analytically convenient, allowing many useful formulae to be rendered in closed form.

4. Gaussian weights allow construction of the family of orthonormal basis functions that we will use in later sections to compensate measured shapes for finite resolution.

5. The Gaussian is not far from optimal for most galaxy shapes. For a well-resolved galaxy with an exponential profile, the Gaussian weight measures η with only 7% higher noise than the optimal weight in equation (3.18). In the presence of seeing, the difference between the Gaussian and optimal weight is even smaller. Recall that *any* weight we choose yields a valid definition of roundness and hence of shape; the Gaussian just incurs a small penalty in noise level.

The procedure for measuring galaxy shapes is therefore as follows:

1. Estimate a shape η for the image I and apply the shear $S_{-\eta}$ to obtain \tilde{I} .

2. Iterate the center and size of the Gaussian weight function until the centroid condition (eq. [3.16]) and the maximum significance condition (eq. [3.17]) are satisfied.

3. Compute the second moments with the Gaussian weight function

$$M(I) = \int dA w(r) I(r, \theta) r^2 e^{-2i\theta}, \quad w(r) = e^{-r^2/\sigma^2}. \quad (3.19)$$

4. If the real and imaginary parts of M are zero, then η is the shape of the object. If not, then we use the measured M to generate another guess for η and return to step 1.

The process is mathematically equivalent to measuring the second moments of $I(\mathbf{x})$ with an elliptical Gaussian weight and iterating the weight ellipticity, center, and size until they match the measured object shape. It is therefore an *adaptive* second-moment measurement. The method is also mathematically equivalent to finding the elliptical Gaussian that provides the best least-squares fit to the image.

3.2. Uncertainties in Shape Estimates

Once we have settled upon a weight of the form $w = \exp(-r^2/2\sigma^2)$, we can integrate δM from equation (3.7) by parts and use equation (3.8) to calculate the variance in η . We will again ignore the I_4 terms; this means we may have a small tendency to over- or underestimate our shape errors if galaxies tend to be boxy or disk. We first define the weighted flux f_w , significance ν , and weighted radial moments $\langle Ir^m \rangle_w$ as

$$f_w = \int dA w I, \quad (3.20)$$

$$\nu^2 = f_w^2 / \text{Var}(f_w) = f_w^2 / \pi n \sigma^2, \quad (3.21)$$

$$\langle Ir^m \rangle_w = f_w^{-1} \int dA w I r^m. \quad (3.22)$$

The condition for optimal significance (eq. [3.17]) is

$$\langle Ir^2 \rangle_w = \sigma^2, \quad (3.23)$$

and under this condition the variance in each component of the shape is

$$\sigma_\eta^2 = \frac{4\pi n \sigma^2}{f_w^2 (1 - a_4)^2} + O(\nu^{-4}) \quad (3.24)$$

$$= \frac{4}{\nu^2 (1 - a_4)^2} + O(\nu^{-4}), \quad (3.25)$$

$$a_4 \equiv \frac{\langle Ir^4 \rangle_w}{2\sigma^4} - 1. \quad (3.26)$$

The quantity a_4 is a form of kurtosis, which is zero for a Gaussian image. The terms of order ν^{-4} arise from errors in the centroid determination and are discussed further in § 8.

The procedure for measuring an object of shape η requires applying a shear $-\eta$ to the image coordinates \mathbf{x} to produce a coordinate system \mathbf{x}' in which the object appears round. The uncertainties in equation (3.24) apply in this sheared coordinate system. Because the object is round in this frame, there is no preferred direction in the shear space and the uncertainty region is circular, with an uncertainty of $\pm\sigma_\eta$ on each component (η'_+ , η'_x) measured in the sheared frame. We must reapply a shear $+\eta$ to restore our measurement to the original coordinate system. This process is illustrated in Figure 2. The η -coordinate transformations

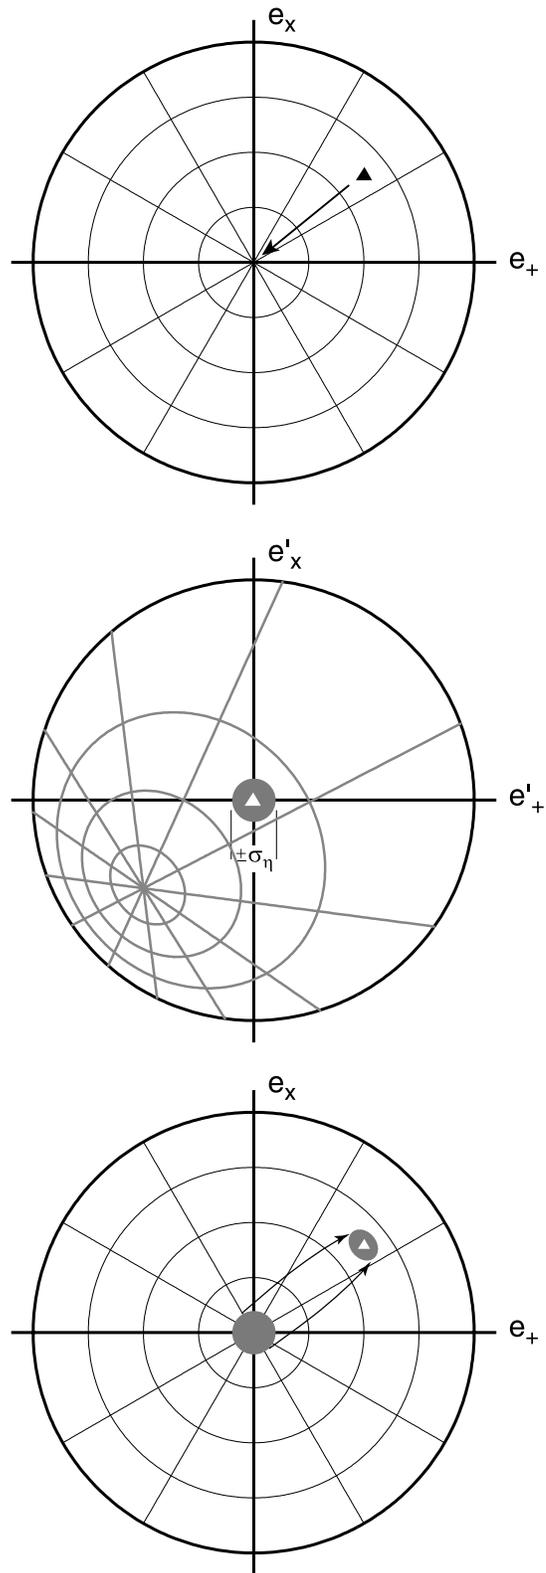


FIG. 2.—Scheme for ellipticity measurement and its errors illustrated schematically. *Top*: The triangle marks the true shape of a target galaxy, located in the e -plane. The shape is determined by shearing the image until the galaxy appears round. In the ellipticity plane, we are moving the object along the vector to the origin. *Middle*: Location of the target (and the original coordinate grid) in the e -plane after application of the shear that makes it round. The shaded region represents the uncertainty region for the shape in the sheared view—the error region must be circular because the image is round. *Bottom*: Finally we undo the applied shear, shifting the target and the error region back to the original shape. This mapping, however, shrinks the error region by a factor $1 - e^2 [(1 - e^2)^{1/2}]$ in the radial [azimuthal] axes.

defined in equations (2.12a)–(2.21b), in the limit of $\eta_1 = \sigma_\eta \ll 1$, indicate that the uncertainty region on η will be elliptical, with a shrunken principal axis in the circumferential direction of the shear manifold:

$$\sigma_\theta = \frac{\sigma_\eta}{\sinh \eta} \Rightarrow ds = \sigma_\theta \tanh \eta = \sigma_\eta \operatorname{sech} \eta. \quad (3.27)$$

If we instead use equation (2.16), we can find the uncertainty ellipse in the ellipticity plane to be

$$\sigma_e = (1 - e^2)\sigma_\eta, \quad e\sigma_\theta = \sqrt{1 - e^2}\sigma_\eta. \quad (3.28)$$

So on the e unit circle, the uncertainty ellipse shrinks radially by $1 - e^2$ and tangentially by $(1 - e^2)^{1/2}$ as we transport the error region from the origin back to the original ellipticity e .

Note that our derivation assumes that the noise characteristics of the image are unchanged when we apply a shear. This is true in the background-limited case because the noise spectrum is flat, and our shear matrices \mathbf{S} have unit determinant. For an image that has been smoothed or deconvolved, the power spectrum will have structure and the shear will alter the noise statistics. We will discuss this in § 6 in the context of finite image resolution.

3.3. Comparison with Other Methods

Galaxy ellipticities for weak lensing were first determined by computing unweighted second moments of the intensity (Tyson et al. 1990). If the moment integrals are taken to infinity, then the measured ellipticities transform under shear using the addition rules in equation (2.13), and furthermore, the correction for PSF effects is extremely simple. It is clearly impractical, however, to carry the integrals to infinity, since neighboring objects will interfere, the noise is divergent, and as noted in K00, many common PSFs have divergent second moments. So the initial methods generally used some sort of isophotal cutoff to the moments. This has the disadvantage of creating moments that are nonlinear in the object flux. An alternative would be to use unweighted moments within a fixed circular aperture, but as noted by KSB, the noise properties of unweighted moments are far from optimal.

In the KSB method, the measured ellipticity \hat{e}_{KSB} is computed from the second moments measured with a circular Gaussian weight with size selected to maximize the detection significance ν . (A different weight function was suggested by Bonnet & Mellier 1995.) The distinction between the KSB method and ours is that KSB always apply a weight that is circular in the original image plane; in our adaptive method, the weight is circular in the sheared image plane that makes the object round. Or, as viewed in the original image coordinates, the weight is an ellipse with shape iterated to match that of the object. This distinction has two consequences: First, our adaptive method yields lower uncertainties for noncircular objects because the weight is a better match to the image. This effect is minor, though, for objects with $e \lesssim 0.5$, and a minority of images are more elliptical than this.

The second, more important advantage over the KSB method is that our definition of shape via equation (2.25) guarantees that our measured ellipticities \hat{e} are transformed by an applied shear via equations (2.13). The circular-weight \hat{e}_{KSB} does not have this property—indeed, for an

object with elliptical isophotes, \hat{e}_{KSB} does not equal the true ellipticity. In the KSB method, it is therefore necessary to calculate a “shear polarizability” for each object, describing the response of \hat{e}_{KSB} to a small shear δ ; this polarizability depends upon the radial profile and $m = 4$ moments of the object. The “polarizability” of our measured ellipticities is just the $\delta \ll 1$ limit of equations (2.13):

$$e \rightarrow \frac{e + \delta}{1 + \delta \cdot e} \approx e + \delta - e(\delta \cdot e). \quad (3.29)$$

This transformation rule and the shapes of our uncertainty regions in ellipticity space arise solely from the geometry of the shear manifold and are independent of the details of the galaxy images. This will simplify the following discussions of methods to derive a shear from an ensemble of measured galaxy shapes.

We implement the adaptive weighted moments scheme in the program ELLIPTO, described further in Jarvis et al. (2002).

4. COMBINING EXPOSURES

In a typical observing program, a given background galaxy is imaged in a number of different exposures, in one or more bandpasses. This is done to increase exposure time, permit rejection of cosmic rays, and/or gather color information. Multiple exposures can also reduce systematic effects by placing data for a given galaxy on different parts of the detector and in different seeing conditions.

We hence encounter the question of how to combine data on a given galaxy in different images to an optimal single measure of the shape. There are two possible approaches:

1. Measure the shape on each exposure, then create a weighted average of the *measurements* as the final shape.
2. Register and average the *images*, then measure the object on the combined image.

We first consider which offers the lowest noise on the final shape. Consider the task of combining N exposures, with the object having significance ν_0 on each exposure. Following equation (3.24), the uncertainty in the shape of a nearly round object measured from a single image will be

$$\sigma_{\eta,0}^2 = 4\nu_0^{-2} + C'[\operatorname{Var}(x_0)]^2 = 4\nu_0^{-2} + C\nu_0^{-4}. \quad (4.1)$$

The second term is the uncertainty due to centroiding error, and C and C' are constants of order unity. If we average measurements (method 1), then we decrease σ_η by \sqrt{N} . If we average images (method 2), then we increase ν by \sqrt{N} . The net error on the shape in the two cases is then

$$\sigma_\eta = \frac{2}{\sqrt{N}\nu_0} \times \begin{cases} \sqrt{1 + C/2\nu_0^2}, & \text{method 1,} \\ \sqrt{1 + C/2N\nu_0^2}, & \text{method 2.} \end{cases} \quad (4.2)$$

The two methods are equivalent, except for the centroiding noise. If ν_0 is not $\gtrsim 5$, then averaging images will produce better accuracy on η . Keep in mind that (a) the galaxy will not even be *detected* on the individual exposures unless $\nu_0 \gtrsim 3$, and (b) the galaxy is useless for weak lensing unless $\sigma_\eta \lesssim 0.5$, which requires $N\nu_0^2 \gtrsim 16$. When $N \lesssim 3$, the centroiding penalty is small for any object that will be useful, so a combined image is extraneous. When $N \gtrsim 5$, there are many galaxies detectable on the summed image that are not

detectable on the individual images, and a summed image has detectability and centroid-noise advantages.

There is a compromise, “method 1.5,” which has the practical advantages (delineated below) of method 1, while retaining the small S/N edge of method 2: that is to create a summed image and use it for object detection and centroid determination, so that $\text{Var}(x_0) \approx 1/N\nu_0^2$. Then this centroid is used to measure shapes on individual exposures, and the shape errors are equivalent to method 2. In practice, we will combine deconvolved Laguerre coefficients (§ 6.3) rather than measured shapes.

There are several reasons why it may be preferable to average catalogs instead of images:

1. Correction of shapes for PSF effects is paramount and only possible if the PSF is constant or slowly varying across the image. If the different exposures in a summed image overlap only partially, then the PSF (and noise level) will jump discontinuously as one crosses the boundaries of component exposures. It is therefore preferable to correct for PSF effects on an exposure-by-exposure basis. If the PSF is very stable (e.g., a space telescope) or if the exposures all have nearly the same pointing (a single deep field), then a summed image will have well-behaved PSF variations.

2. For the smallest objects, the exposures with the best seeing will contain nearly all the useful shape information and should be weighted heavily (Fischer & Kochanski 1994). Large objects are, on the other hand, measured equally well in every exposure. Averaging catalogs allows one to adjust the weights of different exposures on an object-by-object basis, whereas this is not possible when combining images.

3. If there are exposures in different bands, then the optimal weighting of the exposures is dependent upon the color of the object. This is easily done when averaging cataloged shapes but not easily done by summing images.

4. Creation of a summed image requires registration and interpolation of pixels. The latter process smooths the noise field and causes subtle variations in the PSF, both of which complicate later analyses.

5. An especially pernicious hazard to creating a summed image is that slight misregistration of the component images will cause coherent elongations of the images, which if not corrected will mimic a lensing signal. This is discussed by KSB; in theory such effects are handled by a proper PSF correction scheme. This is a danger for method 1.5 as well.

Some practical advantages to method 2, combining images, are as follows:

1. The data storage and processing requirements can be lower for a single combined image if N is large.

2. In method 1, outliers (from cosmic rays) are rejected on an object-by-object basis, whereas in method 2 rejection is pixel-by-pixel. If the galaxies are very oversampled and the cosmic-ray rate is high, method 2 could salvage the uncontaminated parts of galaxy images that method 1 discards.

For the simplest circumstances (a single-filter stack of images with common pointing), image averaging is easier and has few drawbacks. For multifilter or mosaicked data, catalog averaging is needed. The hybrid method 1.5 is best for such cases, though more work. In the rest of this section we detail procedures for each method.

4.1. Combining Images

There are standard tools for combining exposures into a single image. We remark here upon a few special considerations when doing this for weak-lensing observations.

First, accurate registration is paramount. Our scheme for image registration is described in Jarvis et al. (2002).

Second, the use of median algorithms is commonplace but dangerous. Proper correction for PSF effects will require that the images of bright stars have precisely the same PSF as do the faint galaxies. But with a median algorithm, the bright, high-S/N stars will be constructed with a PSF that is a *median* of all the exposures. The images of faint objects, however, will tend toward a PSF that is the *mean* of all the exposures, because the noise fluctuations will dominate PSF variations. The final PSF will therefore vary with magnitude. A σ -clipping average is much preferred over the median for the necessary task of cosmic-ray rejection when combining.

Similarly, one must be careful about rejecting saturated pixels. There will be many stars that saturate only on the best-seeing exposures; if the saturated pixels are rejected, these stars will have final PSFs that are broader than the PSF for faint objects. One must take care to ignore stars that are saturated in *any one* of the exposures.

4.2. Combining Shape Measurements

Suppose that a given galaxy has been measured to have ellipticity e_i in images $i \in \{1, 2, \dots, N\}$. We desire the e that best estimates the true ellipticity of the object. Using the results of § 3.2, we see that in the absence of PSF distortions, the minimum variance estimate of e will be that which minimizes the χ^2 given by

$$\chi^2 = \sum_i (\mathbf{e} \ominus \mathbf{e}_i) \mathbf{S}_i^{-1} (\mathbf{e} \ominus \mathbf{e}_i). \quad (4.3)$$

Here $\mathbf{e} \ominus \mathbf{e}_i$ is equivalent to $\mathbf{e} \oplus (-\mathbf{e}_i)$, where \oplus corresponds to the addition operator introduced in equation (2.11) and \mathbf{S}_i is a covariance matrix, which is $\sigma_\eta^2 \mathbf{I}$ in simple cases.

Note that if the e_i are measured in different filters, then the galaxy may have no single well-defined ellipticity. By “best estimate,” then, we must mean that which offers the best sensitivity to a weak-lensing distortion, and the minimum variance combination of the e_i is still the desired quantity.

The \ominus is a nonlinear operator, so we could use a nonlinear minimization algorithm to find the value of e at which χ^2 is minimized. However, this is both impractical for time considerations and unnecessary, since the values of σ_η are usually small. Thus, we can linearize the subtraction operator

$$\mathbf{e} \ominus \mathbf{e}_i = \mathbf{T}_i (\mathbf{e} - \mathbf{e}_i) + O((\mathbf{e} - \mathbf{e}_i)^2). \quad (4.4)$$

\mathbf{T} can be derived from equations (2.13):

$$\mathbf{T} = \frac{1}{1 - e^2} \left[I - \frac{1 - \sqrt{1 - e^2}}{e^2} \begin{pmatrix} e_x^2 & -e_x e_+ \\ -e_x e_+ & e_+^2 \end{pmatrix} \right]. \quad (4.5)$$

The linearized χ^2 becomes

$$\chi^2 = \sum_i (\mathbf{e} - \mathbf{e}_i) \mathbf{\Sigma}_i (\mathbf{e} - \mathbf{e}_i), \quad (4.6)$$

$$\mathbf{\Sigma}_i \equiv \mathbf{T}_i^T \mathbf{S}_i^{-1} \mathbf{T}_i, \quad (4.7)$$

which has a minimum at

$$\begin{aligned} \mathbf{e} &= \left(\sum_i \boldsymbol{\Sigma}_i \right)^{-1} \left(\sum_i \boldsymbol{\Sigma}_i \mathbf{e}_i \right), \\ \text{Cov}(\mathbf{e}) &= \left(\sum_i \boldsymbol{\Sigma}_i \right)^{-1}. \end{aligned} \quad (4.8)$$

This is a standard least-squares solution for the mean of the \mathbf{e}_i given covariances $\boldsymbol{\Sigma}_i$ in a Euclidean \mathbf{e} -space. In the simple case of $\mathbf{S} = \sigma_\eta^2 \mathbf{I}$, the expression for $\boldsymbol{\Sigma}_i$ simplifies considerably, to

$$\boldsymbol{\Sigma} = \frac{1}{\sigma_\eta^2(1-e^2)^2} \begin{pmatrix} 1-e_x^2 & e_x e_+ \\ e_x e_+ & 1-e_+^2 \end{pmatrix}, \quad (4.9)$$

which is equivalent to equation (3.28). However, when we apply corrections for PSF dilution, we will find that the covariance matrix is more generally an ellipse with axes aligned in the radial and tangential directions. That is,

$$\boldsymbol{\Sigma} = \mathbf{R}_\theta \begin{pmatrix} \sigma_e^2 & 0 \\ 0 & e^2 \sigma_\theta^2 \end{pmatrix} \mathbf{R}_{-\theta}, \quad (4.10)$$

where \mathbf{R}_θ is a rotation matrix with $\theta = \arctan(e_x/e_+)$. In either case the minimization (eq. [4.8]) is numerically straightforward, and we are left with an uncertainty ellipse for the mean $\bar{\mathbf{e}}$. It is wise to implement some outlier rejection algorithm in this process as well.

5. ESTIMATING SHEAR FROM A POPULATION OF SHAPES

Now we presume to have measured ellipticities \mathbf{e}_i for a set of N distinct galaxies, with known measurement uncertainties for each. Our final task is to create a statistic $\hat{\boldsymbol{\delta}}$ from the \mathbf{e}_i that best estimates the lensing distortion $\boldsymbol{\delta}$ that has been applied to this ensemble. There are three main effects that must be considered in constructing the estimator: First, the \mathbf{e}_i respond differently to an applied distortion $\boldsymbol{\delta}$, as embodied by equation (3.29) for true ellipticities, or by shear polarizabilities for the KSB estimators, so we need to know the *responsivity* $\mathcal{R} \equiv \partial \hat{\boldsymbol{\delta}} / \partial \boldsymbol{\delta}$ of our statistic. Second, the variety of ellipticities in the parent (unlensed) galaxy population causes *shape noise* in the shear estimate. In most weak-lensing projects this is the dominant random error, and we wish to minimize its effects. Third, there is *measurement error* in each ellipticity, which we also wish to minimize in our shear estimator.

Most practitioners have adopted a simple arithmetic mean of e_+ and e_x as estimators for the applied distortion (e.g., Fischer & Tyson 1997, hereafter FT97). Using the weak-distortion equation (3.29), it is easy to see that in the absence of measurement error, this estimator has a responsivity $\mathcal{R} = 1 - \sigma_{\text{SN}}^2$ and a variance $\text{Var}(\hat{\boldsymbol{\delta}}_+) = \sigma_{\text{SN}}^2/N$, where we have defined the shape noise $\sigma_{\text{SN}}^2 \equiv \langle e_+^2 \rangle$ (the \times -component has the same properties and shape noise).

Others have realized, however, that rare, highly elliptical galaxies have too much influence on the arithmetic mean and should be deweighted. Cutoffs on $|e|$ (Bonnet & Mellier 1995) or other weighting functions $w(e)$ (Van Waerbeke et al. 2000) have been applied to the ellipticities and tested with simulations, but without any sort of analytic optimization

or justification. Lombardi & Bertin (1998) consider the optimization of a general weighted sum of second moments (rather than ellipticities); this unfortunately couples the ellipticity measurement to the distribution of *sizes* of the galaxies and leads them to consider only weights that are power-law functions of the moments.

Hoekstra, Franx, & Kuijken (2000) present a weighting scheme that incorporates both measurement error and the shape noise, and K00 gives a detailed discussion of optimal weighting for distortion measurements. Both are similar to our method in many respects, which we comment upon at the end of this section.

5.1. Without Measurement Error

We start with an unlensed background galaxy population with ellipticities distributed within the unit circle according to

$$dN/N = P(e') d^2 e' = P(e') e' de' d\theta'. \quad (5.1)$$

A fundamental assumption of weak lensing is that the background is isotropic so that the unlensed population can have P depend only upon the amplitude of e' , not its orientation. The effect of a distortion $\boldsymbol{\delta}$ is to map the background population to a new, anisotropic distribution $P_\delta(\mathbf{e})$, as illustrated in Figure 3. We are given a sample of N galaxies from the new distribution, and our task is to estimate the $\boldsymbol{\delta}$ that gave rise to the distribution from the original P .

One approach is to find the value of $\boldsymbol{\delta}$ that maximizes the likelihood of the observed \mathbf{e}_i . This is true when

$$0 = \frac{\partial}{\partial \delta_+} \sum \log P_\delta(\mathbf{e}_i) = \frac{\partial}{\partial \delta_+} \sum \log P(|-\boldsymbol{\delta} \oplus \mathbf{e}_i|). \quad (5.2)$$

A similar condition holds for δ_\times . These equations define the maximum likelihood $\boldsymbol{\delta}$ even for strong distortions—though there is not in general a closed-form solution for strong $\boldsymbol{\delta}$.

For weak lensing ($\delta \ll 1$), equation (3.29) and the conservation of number can be used to derive $P_\delta(\mathbf{e})$ to first order in $\boldsymbol{\delta}$:

$$P_\delta(\mathbf{e}) d^2 \mathbf{e} = P(e') d^2 e' \quad (5.3)$$

$$= P(|-\boldsymbol{\delta} \oplus \mathbf{e}|) \left| \frac{\partial \mathbf{e}'}{\partial \mathbf{e}} \right| d^2 \mathbf{e} \quad (5.4)$$

$$\approx P \left[e - \boldsymbol{\delta} \cdot \mathbf{e} \left(\frac{1-e^2}{e} \right) \right] (1 + 3\boldsymbol{\delta} \cdot \mathbf{e}) d^2 \mathbf{e}, \quad (5.5)$$

$$P_\delta(\mathbf{e}) = P(e) \left[1 + \boldsymbol{\delta} \cdot \mathbf{e} \left(3 - \frac{1-e^2}{e} \frac{d \log P}{de} \right) \right], \quad (5.6)$$

$$\log P_\delta(\mathbf{e}) = \log P(e) + \boldsymbol{\delta} \cdot \mathbf{e} \left(3 - \frac{1-e^2}{e} \frac{d \log P}{de} \right). \quad (5.7)$$

It can further be shown that the maximum likelihood estimator for $\boldsymbol{\delta}$ takes the form

$$\hat{\boldsymbol{\delta}} = \frac{1}{N} \sum \mathbf{e}_i \left(3 - \frac{1-e^2}{e} \frac{d \log P}{de} \right). \quad (5.8)$$

The parenthesized expression is thus a weight function for combining ellipticities into a distortion. We can also show that this weight function is optimal in terms of S/N for weak distortions, as follows. Let us create an estimator $\hat{\boldsymbol{\delta}}$ that is a

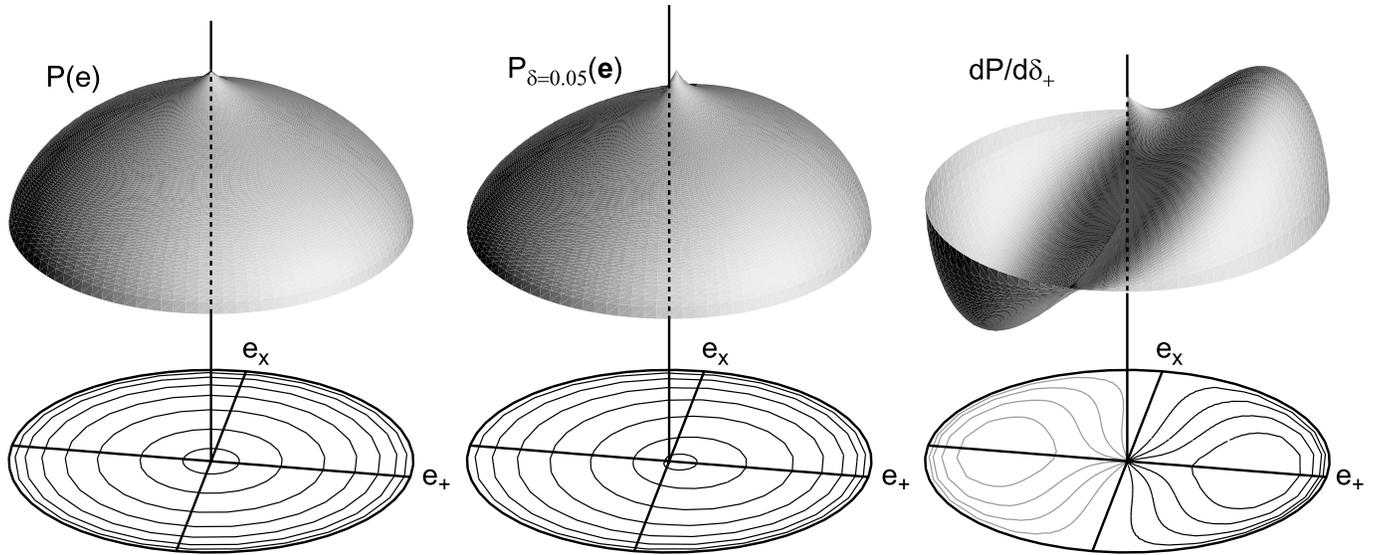


FIG. 3.—*Left*: Model for intrinsic distribution $P(e)$ of galaxy shapes over the e -plane—it must have circular symmetry. When each galaxy is sheared by $\delta_+ = 0.05$, the galaxy distribution shifts to the right as in the middle panel. *Right*: Change in population under the applied distortion; this is the signal that we wish to detect. Shape noise arises from the Poisson fluctuations in the population, which is proportional to the left panel’s $P(e)$. The optimal weight for δ_+ determination is the ratio of the right to the left panel.

general weighted sum of the ellipticities,

$$\hat{\delta} = \frac{\sum w(e_i)e_i}{\sum w(e_i)} = \frac{\int d^2e w(e)eP_\delta(e)}{\int d^2e w(e)P(e)}. \quad (5.9)$$

The response of this statistic to a small applied shear is

$$\mathcal{R} \equiv \frac{\partial \hat{\delta}_+}{\partial \delta_+} = \frac{\sum [w'(\partial e / \partial \delta_+)e_+ + w \partial e_+ / \partial \delta_+] \sum w}{(\sum w)^2} - \frac{\sum we_+ \sum w' \partial e / \partial \delta_+}{(\sum w)^2} \quad (5.10)$$

$$= \frac{\sum [w(1 - e_+^2) + w'(e_+^2/e)(1 - e^2)]}{\sum w}, \quad (5.11)$$

where in the last line, we have dropped terms linear in e_+ or e_x , which average to zero over an isotropic population. With an isotropic population, the derivative $\partial \hat{\delta}_x / \partial \delta_x = \mathcal{R}$ as well, and the off-diagonal elements of $\partial \hat{\delta} / \partial \delta$ are zero.

We may use equation (5.11) to calculate the response of any weighted estimator by summing over the *observed* e_i , because the small difference between observed and intrinsic distributions does not alter \mathcal{R} to first order. In the case where we have some analytic form for $P(e)$, we may replace the sums with integrals over the distribution to obtain

$$\langle \mathcal{R} \rangle = \frac{\int d^2e w(e)e_+(\partial P_\delta / \partial \delta_+)}{\int d^2e w(e)P(e)} \quad (5.12)$$

$$= \frac{\int d^2e w(e)e_+^2 P(e)[3 - (1 - e^2)e^{-1} d \log P / de]}{\int d^2e w(e)P(e)}. \quad (5.13)$$

In the absence of measurement noise, the variance in $\hat{\delta}$ is due to shot noise. Assuming that the background galaxies obey Poisson statistics and their shapes are randomly assigned, we can propagate the Poisson errors through

equation (5.9) to get the expected error:

$$\text{Var}(\hat{\delta}_+) = \frac{\int d^2e w^2(e)e_+^2 P(e)}{N[\int d^2e w(e)P(e)]^2} \quad (5.14)$$

$$= \sum w^2(e)e_+^2 / [\sum w(e)]^2. \quad (5.15)$$

In equation (5.15) it is assumed that the sum is over a sufficiently large ensemble of background galaxies to sample the distribution $P(e)$. Any weak-lensing measurement has thousands of background galaxies, so this gives a direct estimate of the error in the shear.

The optimal weight is that which minimizes $\text{Var}(\delta_+) / \mathcal{R}^2$, which is

$$w_{\text{opt}}(e) \propto 3 - \frac{1 - e^2}{e} \frac{d \log P}{de} \quad (5.16)$$

$$= 3 + 2 \frac{d \log P}{d \log(1 - e^2)} \quad (5.17)$$

$$\Rightarrow \sigma_\delta = \left\{ N \int d^2e P(e)e_+^2 \left[3 + 2 \frac{d \log P}{d \log(1 - e^2)} \right]^2 \right\}^{-1/2}, \quad (5.18)$$

where the last line gives the optimized error in $\hat{\delta}_+ / \mathcal{R}$, which is our calibrated estimate of the distortion. Equation (5.16) reproduces the maximum likelihood solution in equation (5.8). This may be compared with the distortion uncertainty for equal weighting $w = 1$,

$$\sigma_\delta = \frac{1}{\sqrt{N}} \left(\frac{\langle e_+^2 \rangle^{1/2}}{1 - \langle e_+^2 \rangle} \right) = \frac{1}{\sqrt{N}} \frac{\sigma_{\text{SN}}}{1 - \sigma_{\text{SN}}^2}. \quad (5.19)$$

We first see that a simple arithmetic mean of the ellipticities is the optimum estimator only if $P \propto (1 - e^2)^\alpha$ for some exponent α . For the real galaxy population, there can be a significant gain in accuracy through the use of w_{opt} over

equal weighting. An extreme case is a population of randomly oriented circular disks, for which

$$P(e) = \frac{1}{2\pi e} (1-e)^{-1/2} (1+e)^{-3/2}, \quad (5.20)$$

$$\Rightarrow w_{\text{opt}}(e) = (1+e)/e^2. \quad (5.21)$$

With $w = 1$, we would have $\sigma_\delta = 0.590/\sqrt{N}$. The optimal weight diverges at $e \rightarrow 0$ to take advantage of the extreme sensitivity of $P(e)$ to distortion near $e = 0$. The integral in equation (5.18) in fact diverges at $e = 0$, driving σ_δ to zero—which would be a significant improvement over the equal-weighting case! Unfortunately, any small measurement error or departures from circularity for the disks will smooth out the central spike in $P(e)$, creating a finite value for σ_δ .

Figure 4 shows the $P(e)$ measured for well-measured galaxies in the CTIO lensing survey (Jarvis et al. 2002). These shape histograms are derived from 230,000 galaxies that are well resolved ($R > 0.4$, under the definitions in Appendix C) and have errors on the intrinsic ellipticity of $\sigma_\eta \lesssim 0.03$ —primarily galaxies of magnitude $17 < m_R < 22$. The shape of $P(e)$ is observed to be highly dependent upon the surface brightness (SB) of the galaxies. The low-SB galaxies show the rise at $e > 0.8$ expected of a disk population, but there the distribution drops at $e > 0.95$ instead of diverging—this reflects the finite thickness of the disks. There is also no pole at $e = 0$ for the low-SB galaxies, showing that the disks are

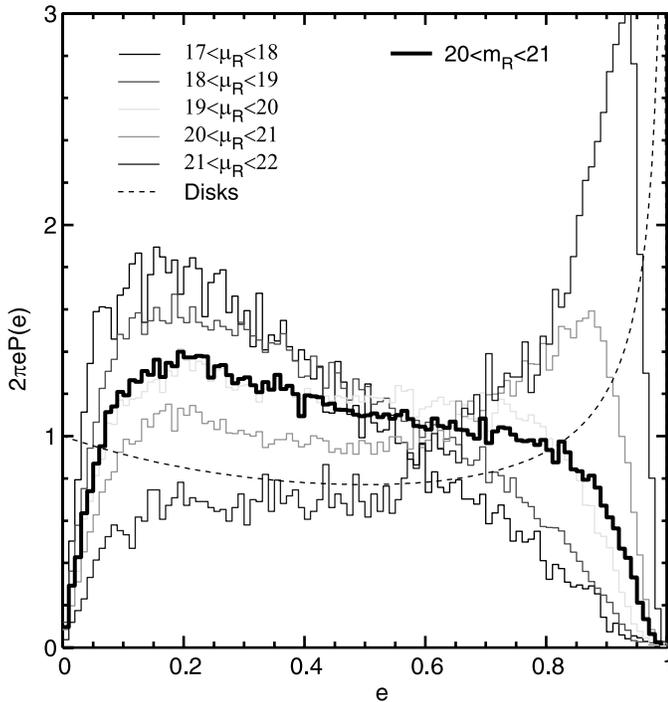


FIG. 4.—Distribution of intrinsic ellipticities for modestly bright galaxies ($m_R \approx 20$); we plot $2\pi e P(e)$, rather than $P(e)$, as the nature of the population is more apparent. The distribution is highly dependent upon surface brightness μ_R , presumably reflecting the difference between spheroid- and disk-dominated galaxies. The dashed line is the distribution for an isotropic population of circular disks. The high- μ galaxies are more useful for distortion measurements. The heavy histogram combines all surface brightnesses in the magnitude range $20 < m_R < 21$. Though it is difficult to tell from this plot, $P(e)$ is finite and increasing as $e \rightarrow 0$. Optimal weighting takes advantage of this structure to reduce the noise in the distortion measurement.

not perfectly circular. The high-SB galaxies are presumably early types, since there are very few with $e > 0.5$. While the value of $P(e = 0)$ increases with surface brightness, it always remains finite, but with $dP/de < 0$. The ideal-weight equation (eq. [5.16]) therefore grows as $1/e$ as $e \rightarrow 0$, but the contribution to the S/N does not diverge at zero, as for perfect disks. None of the $P(e)$ curves is well fitted by a single Gaussian or power law.

The intrinsic ellipticity variance σ_{SN} varies from 0.30 to 0.49 between the highest and lowest SB bins. The optimally weighted distortion S/N per galaxy for high-SB, early types is 2–3 times higher than for the lowest-SB galaxies, indicating the desirability of incorporating some galaxy type discriminant—surface brightness, color, or concentration—into the weighting scheme. The only requirement is, of course, that the discriminant be independent of ellipticity. It seems likely that $P(e)$ will vary substantially with magnitude.

The heavy histogram in Figure 4 combines all well-measured galaxies with $20 < m_R < 21$, which we henceforth use as a representative measure of the real galaxy population. The distribution has $\sigma_{\text{SN}} = 0.38$, which would lead to $\sigma_\delta = 0.44/\sqrt{N}$ for an unweighted average. The optimal weighting gives $\sigma_\delta = 0.33/\sqrt{N}$; the weighting therefore gives a gain equivalent to a 1.8-fold increase in N . The gain in telescope time is at least as large. This gain is reduced, however, in the presence of measurement noise, which will tend to wash out the sharp feature in $P(e)$, as discussed next.

We reiterate two favorable results of this section: First, the responsivity \mathcal{R} and the variance of $\hat{\delta}$ can be expressed exactly as direct sums over the observed population, for arbitrary choice of w ; there is no need for a calculation of polarizabilities or recourse to simulated images. Second, we note that the variance in $\hat{\delta}$ can be significantly below the canonical σ_{SN}^2/N if the ellipticity distribution $P(e)$ has structure that is not washed out by measurement noise.

5.2. With Measurement Error

A galaxy image with true ellipticity e will be measured at some ellipticity \tilde{e} , with a probability distribution of $p(\tilde{e}|e)$. We consider a population of galaxies all having the same significance ν and resolution parameter R (see § 6 and Appendix C), so that they all have a common $p(\tilde{e}|e)$. The measured distribution of ellipticities observed under distortion δ will then be

$$\tilde{P}_\delta(\tilde{e}) = \int d^2e P_\delta(e) p(\tilde{e}|e), \quad (5.22)$$

where P_δ is the distribution of true ellipticities as in the previous section. The symmetry of P and p in the ellipticity plane guarantees that in the absence of distortion, the measured distribution $\tilde{P}(\tilde{e})$ must again depend only upon the magnitude, not the direction, of the measured ellipticity.

Equation (5.22) is not strictly a convolution, because the measurement function $p(\tilde{e}|e)$ may depend upon e and not simply upon $(\tilde{e} - e)$ —for example, equation (3.28) describes how the error ellipses contract as e departs from zero, even if the significance of the detection is held fixed. In § 6, we will show that the behavior of the measurement error is different when the effects of PSF smearing upon the image are important, so at this point we will consider $p(\tilde{e}|e)$ to be, most generally, some kind of Gaussian whose 1σ ellipse depends

only upon the magnitude e . An important point is that the functional form of $p(\tilde{e}|e)$ is unchanged by an applied distortion, since p is determined by ν and R , which are unchanged by a pure shear.

Another fact to keep in mind is that, with finite resolution and noise, it is possible to measure $\tilde{e} > 1$, if the image noise makes the object appear smaller than the PSF in some dimension. Our formulae should therefore be tractable even for $\tilde{e} > 1$, and we cannot simply discard such measurements without contemplating the consequences.

We proceed as in the previous section, by assuming a distortion estimator of the form

$$\hat{\delta} = \frac{\sum w(\tilde{e}_i)\tilde{e}_i}{\sum w(\tilde{e}_i)} = \frac{\int d^2\tilde{e} w(\tilde{e})\tilde{P}_\delta(\tilde{e})\tilde{e}}{\int d^2\tilde{e} w(\tilde{e})\tilde{P}(\tilde{e})}, \quad (5.23)$$

$$\Rightarrow \langle \mathcal{R} \rangle = \frac{\partial \hat{\delta}_+}{\partial \delta_+} = \frac{\int d^2\tilde{e} [w(\tilde{e})\tilde{e}_+ \int d^2e p(\tilde{e}|e)(\partial P_\delta(e)/\partial \delta_+)]}{\int d^2\tilde{e} w(\tilde{e})\tilde{P}(\tilde{e})} \quad (5.24)$$

$$= \int d^2\tilde{e} \left[w(\tilde{e})\tilde{e}_+ \int d^2e P(e)p(\tilde{e}|e)e_+ \times \left(3 - \frac{1 - e^2}{e} \frac{d \log P}{de} \right) \right] / \int d^2\tilde{e} w(\tilde{e})\tilde{P}(\tilde{e}); \quad (5.25)$$

$$\text{Var}(\hat{\delta}_+) = \frac{\int d^2\tilde{e} w^2(\tilde{e})\tilde{e}_+^2 \tilde{P}(\tilde{e})}{N[\int d^2\tilde{e} w(\tilde{e})\tilde{P}(\tilde{e})]^2} \quad (5.26)$$

$$= \sum w^2(\tilde{e})\tilde{e}_+^2 / \left[\sum w(\tilde{e}) \right]^2, \quad (5.27)$$

$$\Rightarrow w_{\text{opt}}(\tilde{e}) = \frac{1}{\tilde{P}(\tilde{e})} \int d^2e P(e)p(\tilde{e}|e) \frac{e_+}{\tilde{e}_+} \left(3 - \frac{1 - e^2}{e} \frac{d \log P}{de} \right). \quad (5.28)$$

Given functional forms for the *intrinsic* distribution $P(e)$ and the uncertainty function p , we could use equation (5.28) to derive an optimal weight and use equations (5.25) and (5.26) to get the responsivity and noise for the estimator using this or any weight function. In most cases these integrals will not have analytic solutions.

The bracketed quantity under the integral in equation (5.13) is a galaxy's responsivity to shear, which depends upon the *intrinsic* shape. Equation (5.25) is the average of this responsivity for the galaxies with some *measured* shape. The measurement noise can cause these two quantities to differ; in other words, a naive determination of the responsivity is biased by the measurement noise. KSB-based methods will also suffer a calibration error due to this effect; binning the polarizabilities in parameter space can reduce the *noise* in the polarizability but will not remove biases. Precision cosmology will require that such calibration issues be addressed—there are unfortunately no cosmic calibration standards for shear.

The need for $P(e)$ in the above formulae is an unfortunate complication, since \tilde{P} is the directly observed quantity. Note that the variance of the estimator can be expressed as a closed sum over the observed shapes (eq. [5.27]), but the responsivity cannot. A precise calibration of the resultant shear-mass maps requires, therefore, that P be estimated either by deconvolving the observed \tilde{P} with the error distribution p , or by recourse to higher quality images that give P directly.

Derivation of the optimal weight also requires knowledge of the intrinsic P , but we can explore some generic cases and make some approximations that give workable methods.

5.2.1. Approximate Form for Responsivity with Errors

We wish to have a form for \mathcal{R} as a sum over the observed objects and applied weights, as in equation (5.11), for the case of finite measurement errors. Toward that end we can take the derivative of equation (5.23), which is greatly simplified if we assume that the measurement *error*, i.e., $\tilde{e}_i - e_i$, does not have any first-order dependence on δ . While not strictly valid, it is a good approximation. In this case

$$\mathcal{R} = \sum \left[w(1 - \langle e_+^2 \rangle_{\tilde{e}}) + \frac{\tilde{e}_+^2}{\tilde{e}} \frac{dw}{d\tilde{e}} \times (1 - \langle e_+^2 \rangle_{\tilde{e}} - \langle e_+e_\times \rangle_{\tilde{e}} \tilde{e}_\times / \tilde{e}_+) \right] / \sum w, \quad (5.29)$$

where the angle brackets indicate an average of the true quantity at a given measured value, e.g.,

$$\langle e_+^2 \rangle_{\tilde{e}} = \int d^2e p(e|\tilde{e})e_+^2 \quad (5.30)$$

$$= \frac{\int d^2e p(\tilde{e}|e)P(e)e_+^2}{\int d^2e p(\tilde{e}|e)P(e)}. \quad (5.31)$$

Note that the weight function w may depend upon \tilde{e} directly, but also indirectly through some dependence in its covariance matrix Σ . If the measurement error function $p(\tilde{e}|e)$ and the intrinsic distribution $P(e)$ have circular symmetry, then we must be able to write

$$\langle e_+^2 \rangle_{\tilde{e}} = k_0(\tilde{e}) + k_1(\tilde{e})\tilde{e}_+^2, \quad (5.32)$$

where k_0 and k_1 are functions only of the magnitude, not the direction, of \tilde{e} . We must also have $\langle e_+e_\times \rangle_{\tilde{e}} = k_1(\tilde{e})\tilde{e}_+\tilde{e}_\times$. Further manipulation, taking advantage of the isotropy of the parent population, yields

$$\mathcal{R} = \sum \left[w \left(1 - k_0 - \frac{k_1\tilde{e}^2}{2} \right) + \frac{\tilde{e}}{2} \frac{dw}{d\tilde{e}} (1 - k_0 - k_1\tilde{e}^2) \right] / \sum w. \quad (5.33)$$

This form for \mathcal{R} depends only upon the observed quantities and the chosen weight scheme, except through the two functions k_0 and k_1 , which we will approximate below. The resemblance to equation (5.11) is clear. With this equation and some integration by parts, we may also derive a form for the optimal weighting function:

$$w_{\text{opt}}(\tilde{e}, \Sigma(\tilde{e})) = 3k_1 + \frac{1}{\tilde{e}} \frac{dk_0}{d\tilde{e}} + \tilde{e} \frac{dk_1}{d\tilde{e}} - \frac{1 - k_0 - k_1\tilde{e}^2}{\tilde{e}} \frac{d \log \tilde{P}}{d\tilde{e}}. \quad (5.34)$$

5.2.2. Special Case: Gaussians

In general, the functions $k_0(\tilde{e})$ and $k_1(\tilde{e})$ must be calculated numerically using a presumed underlying $P(e)$ for the background population, but analytic solutions are possible in the case of a Gaussian $P(e)$ with variance σ_{SN}^2 in each component (the shape noise) and a constant measurement error σ_e^2 on each component. We find that both k_0 and k_1 are

independent of \tilde{e} :

$$\begin{aligned} k_0 &= (1-f)\sigma_{\text{SN}}^2, & k_1 &= f^2, \\ f &= \frac{\sigma_{\text{SN}}^2}{\sigma_{\text{SN}}^2 + \sigma_e^2}. \end{aligned} \quad (5.35)$$

The quantity f is the fraction of the total ellipticity variance that is attributable to shape noise. When the measurement noise is small, $f \approx 1$, the ideal weight is close to $(1 - \tilde{e}^2)/(\sigma_{\text{SN}}^2 + \sigma_e^2)$. This is quite similar to the weight adopted by Hoekstra et al. (2000).

For the Gaussian case, $d \log \tilde{P}/d\tilde{e}$ is also quite simple, so equation (5.33) can be used. In our surveys to date (Smith et al. 2001), we have adopted a weight that results from optimizing the Gaussian case.

5.2.3. Practical, Nearly Ideal Approximation

We obtain an approximation to the correct responsivity \mathcal{R} and resultant ideal weight if we adopt the constant k_0 and k_1 functions in equations (5.35) even for non-Gaussian $P(e)$ distributions. The shape noise σ_{SN}^2 may be defined as the assumed variance of the underlying e_+ and e_x and may be found by subtracting the measurement noise from the observed $\langle \tilde{e}^2 \rangle$. The measurement noise σ_e^2 is known for each galaxy using the methods of this paper; since the covariance matrix for \tilde{e} is generally anisotropic, some representative scalar must be selected.

With these guesses for k_0 and k_1 in hand, \mathcal{R} may be estimated with a sum over the *observed* galaxies using equation (5.33) for any chosen weight function.

For the real-universe shape distributions measured in the CTIO survey (Jarvis et al. 2002), we find that the following “easy” weight function offers very close to optimal distortion measurements:

$$w = [e^2 + (1.5\sigma_\eta)^2]^{-1/2}, \quad (5.36)$$

where σ_η is the shape uncertainty that the object would have were it circular (cf. eq. [3.28]).

We can check the accuracy of our approximations numerically for chosen $P(e)$ and $p(\tilde{e}|e)$ functions. We examine the case in which $P(e)$ is that shown in Figure 4 for galaxies with $20 < m_R < 21$, and the measurement error follows equation (3.28). We find that the weight function given by equation (5.34) is in fact very close to optimal for all noise levels, even when the simple approximations (eq. [5.35]) are used for the k -functions. The “easy” weight equation (eq. [5.36]) also performs nearly optimally, so most applications could use this weight and need not attempt to determine $P(e)$.

A more critical question is whether the approximations (eq. [5.35]) yield a proper estimate of the calibration factor \mathcal{R} when used with equation (5.33). Figure 5 shows how the simple \mathcal{R} estimator compares with the correct value in equation (5.25) for our choice of underlying distribution and the “easy” weights (eq. [5.36]). The approximate form yields a responsivity correct to better than 5% for $\sigma_\eta \lesssim 0.4$. It is clear from the figure that some detailed knowledge of the underlying $P(e)$ distribution will be needed in order to calibrate lensing measurements to the 1% level.

The bottom panel of Figure 5 shows the potential advantage of optimal weighting. When the measurement error is $\gtrsim 0.2$, there is little difference between various weighting schemes. For an unweighted distortion estimator, the accu-

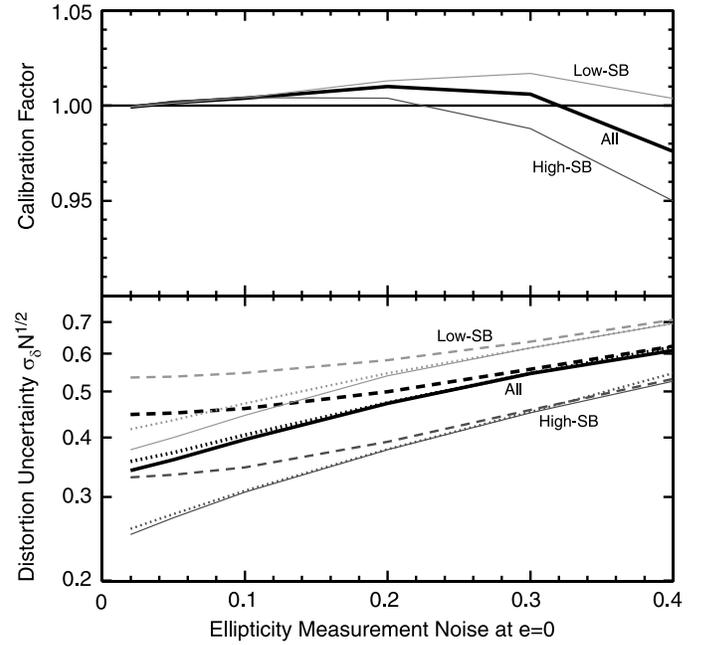


FIG. 5.—*Top*: Approximate, simplified closed-sum estimate of the calibration factor \mathcal{R} (eq. [5.33]) relative to the exact form (eq. [5.25]) as a function of the ellipticity measurement noise at $e = 0$. The weight function is the “easy” eq. (5.36) and the factors k_0 and k_1 adopt the simple heuristic approximation (eq. [5.35]). The heavy curve is for the population of $20 < m_R < 21$ galaxies, the top curve is for a low-SB sample ($20 < \mu_R < 21$), and the bottom curve for a high-SB sample ($17 < \mu_R < 19$), for which the intrinsic $P(e)$ distributions are plotted in Fig. 4. The simple formulation yields a calibration accurate to 5% or better in all cases, but 1% accuracy is difficult to achieve. *Bottom*: Uncertainty in the distortion determination when the galaxy shapes are combined with optimum weights (solid lines), the “easy” weights (eq. [5.36]; dotted lines), and equal weighting (dashed lines). The line weights code the galaxy sample, as above. Note that with optimal or easy weighting, the distortion errors continue to shrink even when measurement error is well below the canonical shape noise of 0.3.

racy levels out as the measurement noise drops below $\sigma_e \approx 0.2$. When optimal weights are used, however, the distortion errors continue to drop as the measurement error is pushed toward zero—the optimal weights take advantage of the $e = 0$ cusp in the shape distribution. Our “easy” weight scheme recovers nearly all this potential gain.

To summarize, a practical method of weighting and calibrating the response in the presence of measurement noise is as follows:

1. Determine the underlying $\sigma_{\text{SN}}^2 = \langle e_+^2 \rangle$ of the intrinsic distribution. The measurement error σ_η (or the full covariance matrix) of each galaxy is known from the formulae in previous sections.
2. Approximate the quantities k_0 and k_1 with equation (5.35).
3. Choose a weight function, for example, the “easy” form equation (5.36) or, preferably, the optimal form equation (5.34), which can be derived from the observed distribution $\tilde{P}(\tilde{e})$.
4. The distortion estimator is the sum form in equation (5.23). The variance in the estimator is the sum form in equation (5.27).
5. The estimator (and variance) must be scaled by the responsivity, which is calculated with equation (5.33).

If more is known about the intrinsic shape distribution, then more accurate functions for k_0 and k_1 may be derived and used in the sums.

5.3. Additional Weighting Considerations

The optimal weight may depend upon parameters other than the observed ellipticity \bar{e} . It must, for example, depend upon the measurement errors as described above. If intrinsic shape distribution $P(e)$ depends upon galaxy type, for example, then it may be advantageous to have different weight functions for each type—as long as galaxy type can be determined independent of e . Elliptical galaxies may be distinguished from spirals, for example, by a concentration ratio such as the b_{22} coefficient described below.

The expected shear will depend upon the redshift z of the background galaxy, and hence there may be a z -dependent weight determined from photometric redshift estimates of the background population. More crudely, the apparent magnitude may be taken as an indicator of z and used in the weight formulation. The use of such additional weights depends upon the problem being addressed; see Smith et al. (2001) as an example.

5.4. Relation to Previous Methods

The optimal weighting scheme of K00 differs from ours in several respects. It assumes a different measure of e , which does not follow our geometric transformation relations, so the mean polarizability of galaxies must be calculated in some parameter space, e.g., of flux, size, and e . Binning or smoothing in some such space is common to most of the KSB-based weighting schemes as well (Erben et al. 2001; Hoekstra et al. 2000). The variance of the estimators $\hat{\delta}_i$ are calculated within each bin, and then weights are chosen inversely to the variance of each bin to create a minimum variance estimator. There are three contributors to the variance from each bin:

1. The intrinsic ellipticities of galaxies within the bin are drawn at random from the parent distribution. If $|e|$ is one parameter of the space, then only the direction θ is allowed to vary within the bin. This is of course the shape noise.
2. The measured shape of a given galaxy is drawn from the measurement error distribution.
3. The number of galaxies within the bin is drawn from a Poisson distribution. If $|e|$ is a parameter, then this Poisson noise includes some elements of the shape noise (1) and measurement noise (2).

All three of these effects are important to optimization; all are included in our formulation and (implicitly) in that of K00, so we expect them to be essentially equivalent in the long run—this is not the case, though, for some of the heuristic or parameter-space weight formulations in the literature. The virtue of our scheme is that the nature of the weight function is apparent given the intrinsic shape distribution $P(e)$ and the measurement errors, and there is no need to choose a parameter space for weight selection and polarizability smoothing. Our formulation tells us when further parameters might be desirable, namely, when $P(e)$ changes significantly.

6. MEASUREMENTS WITH FINITE RESOLUTION

The preceding sections outline a method for optimal recovery of weak distortion from galaxy images and rigorous estimation of the uncertainties on these shears, for the case when the detector views the galaxies with perfect resolution. Unfortunately, the finite resolution of real observa-

tions has a strong effect upon shape measurements in every weak-lensing observation to date, even those using the *Hubble Space Telescope* (*HST*). Finite resolution produces two deleterious effects:

1. A PSF that is not circularly symmetric can induce ellipticities on the images, thus breaking the intrinsic isotropy of the background galaxy population and mimicking a lensing distortion. This is a *bias* induced by asymmetric PSFs. Since present-day weak-lensing surveys are seeking distortion signals well below 1%, measured shapes must be compensated for even the slightest asymmetry in the PSF with some *smear* correction.
2. Convolution by a circularly symmetric PSF will make most galaxies appear rounder, driving $e \rightarrow 0$. This is therefore a *dilution* of the true lensing signal. While this mechanism cannot create a lensing signal where there is none, it can misleadingly modulate the lensing signal or cause a calibration error in the inferred mass distribution.

Most (but not all) approaches to PSF corrections treat the bias and dilution effects in separate steps. To our knowledge, *all published weak-lensing observations have incomplete PSF correction, leaving systematic distortion errors of a fraction of a percent or higher*. While in most cases these residual systematic effects have not altered the validity of the authors' conclusions, the proper correction of PSF effects is presently the largest barrier to the use of weak lensing for precision cosmology.

In § 6.1, we review some existing approaches to these problems, in § 6.2 we contemplate how one would ideally wish to approach the problem, and in further sections we develop two means of implementing a nearly ideal approach: one that treats bias and dilution separately and another that corrects both problems simultaneously with a limited form of deconvolution.

6.1. Existing Approaches to PSF Corrections

6.1.1. The Unweighted Ideal

Unweighted second moments of galaxies are ideal measures of ellipticity—not only do the derived ellipticities transform according to the rules of shear space, but correction for PSF effects is in principle quite simple because the unweighted second moment of the image is just the sum of the original moment and the PSF moment. Thus, by subtracting the PSF moments from the measured moments we simultaneously correct for bias and dilution and obtain the image shape. In the case where the PSF is round, the dilution of a true (preseeing) image-plane ellipticity e^i to an observed (postseeing) ellipticity e^o is described by the exact equation

$$e^i = e^o/R, \quad (6.1)$$

$$R \equiv \frac{\langle r^2 \rangle_i}{\langle r^2 \rangle_i + \langle r^2 \rangle_*} = 1 - \frac{\langle r^2 \rangle_*}{\langle r^2 \rangle_o}. \quad (6.2)$$

The *resolution parameter* R is determined by the unweighted second radial moment of the measured image $\langle r^2 \rangle_o$ relative to that of the PSF $\langle r^2 \rangle_*$. Two things to note: first, the error ellipse on the dilution-corrected, preseeing ellipticity e^i is magnified by $1/R$ from the original measurement error equation (3.28) and is further stretched in the radial direction by the uncertainty in R itself. Thus, the error ellipse is no longer simply described by a single σ_η . Second, equation (6.1) can give rise to $|e^i| > 1$, if the noise makes the galaxy

look smaller than the PSF about some axis. We cannot arbitrarily discard such measurements without creating a bias in our mean shear. These two phenomena are common to all modes of PSF dilution correction.

This blissfully simple dilution correction is spoiled by two major problems: First, as discussed above, unweighted second moments have divergent noise properties and for this and other reasons are not practical shape estimators. An equally serious problem noted by K00 is that the second moments themselves are divergent for many realistic PSFs. Further, many galaxies follow de Vaucouleurs profiles, for which the second moment converges very slowly.

The simple formulae (6.1)–(6.2) are still valid under the special circumstances that the object and PSF are both Gaussians. The post-PSF object is again a Gaussian, and deconvolution of Gaussians is a simple subtraction of second moments. Hence any shape-measuring algorithm that extracts the proper ellipticity for a Gaussian ellipsoid would allow PSF dilution correction via equation (6.1) in this limited (and unrealistic) case.

Some early weak-lensing measurements (Valdes et al. 1983) adopt second-moment subtraction as a means of PSF correction, despite the fact that this method is not exact when isophotally bounded or weighted moments are used, and the images are not Gaussian. This would not suffice, however, for the more sensitive measurements being done today.

6.1.2. Heuristic Methods

In the case of unweighted moments or Gaussians, equation (6.1) would indicate that a regression of the lensing signal against $\langle r^2 \rangle_o^{-1}$ would yield a distortion free of PSF effects as $\langle r^2 \rangle_o^{-1} \rightarrow 0$. Mould et al. (1994) have attempted to measure very weak shears in the presence of PSF effects using such a regression (though against $\langle r \rangle^{-1}$, in which case a linear relation is not expected). Even with weighted moments, we expect the PSF dilution and bias to decrease as the object becomes well resolved, so there is some basis to this method, even if it is not exact. Other problems, however, are that the distortion is not likely to be the same for all sizes of galaxy, as they likely lie at different distances. Also, the regression will lead to substantially higher noise than a more direct dilution correction.

Another approach to the dilution correction is exemplified by FT97, who attempt no analytic correction, instead calibrating the dilution effect by measuring simulated background galaxies that have been subjected to the same distortion, seeing, sampling, and noise as the real images. Such a simulation is an essential test of any weak-lensing methodology. The difficulty with sole reliance upon simulated data is that the result is extremely sensitive to one’s ability to exactly match the size-magnitude distribution of the true galaxy population, because the dilution correction is a strong function of size (as in eq. [6.1]) in the typical regime of slightly resolved galaxies. Further, as we show below, the dilution correction depends upon detailed higher order moments of the galaxy images, which would be very difficult to simulate faithfully. One alternative is to use high-resolution, high-S/N images from *HST* instead of simulated galaxies—but the total sky area imaged to sufficient S/N by *HST* is a tiny fraction of a square degree, too small for rigorous calibration tests. It would be preferable to have an ana-

lytic correction for dilution and use the simulated data to spot-check the accuracy of the analytic method.

6.1.3. Perturbative Methods

A step beyond the unweighted-mean approximation to the bias correction is taken by KSB and by FT97. Both make the assumption that the anisotropy of the PSF can be described as a small anisotropic convolution applied to a larger, circularly symmetric PSF. In this case, the effect of the tiny asymmetric deconvolution upon the weighted second moments of a given image can be expressed as a fourth-order weighted moment of the image, which KSB christen the *smear polarizability*. Given the smear polarizability of an image and a measure of the anisotropy of the PSF, the measured second moments are corrected analytically for the PSF bias.

The FT97 method differs in that the correction for PSF anisotropy is applied to the image rather than to the measured moments: a minimal 3×3 convolution kernel is created, which will “circularize” the PSF. The galaxy shapes are measured after this kernel is applied to the image.

The primary drawback to these methods (K00) is that the approximation upon which they are based often fails: a typical diffraction-limited PSF in no way resembles a small convolution to a round PSF, and even a simple aberration such as coma creates PSFs that violate this condition.

These methods have features, however, that we wish to retain in any improved formulation. They are easily adapted to a PSF that varies across the image—the requisite moments of the PSF are measured wherever a star falls upon the image and interpolated to the location of each galaxy. The FT97 method, by fixing the image, frees us from having to measure the higher order moments that make up the smear polarizability, though at the expense of a slight reduction in image resolution, an increase in the image noise, or both.

The perturbative methods correct only the PSF bias, not the dilution, because the dilution cannot be considered a small perturbation in any extant data set. FT97 calibrate the dilution with simulations, as mentioned above. The original KSB work made use of empirical calibration as well, as the shear polarizability (cf. § 3.3) measures only the susceptibility of the image to a distortion that might be applied *after* the PSF convolution. Wilson, Cole, & Frenk (1996) suggest an empirical calibration by deconvolving the real images, applying a shear, reconvolving, and remeasuring to determine the response. Luppino & Kaiser (1997) introduce the *preseeing shear polarizability*, which approximates the susceptibility of the KSB weighted moments to a shear applied *before* the PSF. The preseeing shear polarizability is, to its lowest order, similar to the resolution factor R introduced above for unweighted moments. There are, however, fourth-order moments of galaxy and stellar shapes involved, and all are measured with Gaussian weights so the noise does not diverge.

The KSB method, updated to use the preseeing shear polarizability, is exact for Gaussians and in the limit of a compact anisotropy kernel, but it is not exact in the general case (K00). There are additional ambiguities regarding the appropriate size of the Gaussian weight to be applied when measuring the PSF moments (Hoekstra et al. 1998). The updated KSB method is in wide use, and several papers have investigated how accurately it performs for simulated gal-

axes and PSFs that are not Gaussian (Erben et al. 2001; Bacon et al. 2001). While it is clear that in many circumstances KSB is good enough, we would prefer to understand and overcome its limitations.

We will demonstrate below that the KSB and FT97 bias corrections are the first terms in a series expansion of the deconvolution of the galaxy image.

Rhodes, Refregier, & Groth (2000) investigate the KSB method in some detail, carrying forth the transformation equations to higher order than do KSB. The PSF corrections, however, still require substantial approximations. Below we construct a hierarchy of *all* the weighted moments of the image and give general formulae for their transformations under shear, convolution, and other operations.

6.1.4. Deconvolution Methods

Kuijken (1999) suggests that one determine shear by summing the images of all the galaxies in some cell and then comparing this high-S/N summed galaxy image with the PSF. The summed galaxy image should approach circular symmetry in the source plane; the shape of the summed image can thus be modeled relatively simply as a circular source with arbitrary radial profile, sheared by the lens to have elliptical isophotes, then convolved with the PSF. He adopts a flexible parametric representation of this mean radial profile. A candidate profile with candidate lens distortion can be convolved with the known PSF and compared with the measured mean image. The radial profile and distortion are then adjusted to give a best fit. This method can be viewed as a limited form of deconvolution: the only characteristic of the deconvolved image one cares about is the ellipticity; the multipoles beyond quadrupole are irrelevant to the measurement and are discarded.

In Kuijken's method, the higher order multipoles are discarded by averaging over many galaxies. This may be difficult in practical situations, where the PSF and/or the distortion signal are varying across the field too rapidly to gather sufficient background galaxies to sum. Another drawback is that summing galaxies' images may not be the optimal way to combine their information on the shear. Kuijken suggests applying the method to individual galaxies given these potential problems, which amounts to an assumption that the galaxy ellipticity is constant with radius. The accuracy of this method is not discussed in detail. But the method does have the strong advantage of being able to cope with arbitrary PSF behavior and simultaneously removing both the bias and dilution effects from the measurement.

6.1.5. Commutator Method

K00 introduces a new approach to PSF correction, deriving from the PSF an operator that can be applied to the *observed* image to effect the transformation of applying a given *preseeing* shear. This operator is derived by considering the commutation of the shear and convolution operators in Fourier space. With the operator in hand, the response to a preseeing shear can be determined directly from the postseeing image, given sufficiently detailed knowledge of the PSF.

The K00 formulation is the first, to our knowledge, to offer an exact correction for PSF bias and dilution. We take a different approach below and then offer some comparisons of the two approaches.

6.2. Optimal Methods in the Presence of a Convolution

In § 2.5, we noted that a real-space shear is equivalent to an opposing shear of the Fourier-space image, so that we can conduct the roundness test in Fourier space. One virtue of using a Gaussian weight function for our roundness test, as in equation (3.19), is that this test takes the exact same form in k -space, namely,

$$M(I) = 0 \quad \Leftrightarrow \quad 0 = \tilde{M}(\tilde{I}) = \int d^2k e^{-k^2\sigma^2/2} I(\mathbf{k}) k^2 e^{-2i\phi}. \quad (6.3)$$

The only difference is that the Gaussian weight has width $1/\sigma$ in k -space; i.e., a broader weight in real space is narrower in Fourier space. Furthermore, our assumed uniform white noise in real space transforms to uniform white noise in Fourier space. Therefore the entire derivation of the optimal weight in § 3.1 could have been executed in k -space without any change in the result.

The effect of the PSF convolution is to suppress the image by some transfer function $\tilde{T}(\mathbf{k})$. With perfect knowledge of the PSF, we can deconvolve the observed image \tilde{I}^o to retrieve the image-plane transform $\tilde{I}^i = \tilde{I}^o/\tilde{T}$ (here we mean the image plane of the gravitational lens, before the application of seeing). A deconvolution would remove the PSF bias entirely, but the noise is no longer homogeneous in k -space, having been amplified (perhaps infinitely) by $1/\tilde{T}$. If the PSF is close to circularly symmetric, then \tilde{T} is nearly independent of direction. If we now imagine making our roundness test on the deconvolved k -space image, we adapt equations (3.7)–(3.9) to give

$$\delta\tilde{M} = \frac{\eta}{4} \int k dk k^2 w(k) k \frac{d\tilde{I}_0^i}{dk}, \quad (6.4)$$

$$\text{Var}(M) = \frac{n}{4\pi^2} \int k dk d\phi k^4 w^2(k) \frac{\cos^2 2\phi}{\tilde{T}^2(k)}, \quad (6.5)$$

$$\Rightarrow w_{\text{opt}}(k) \propto \frac{-\tilde{T}^2(k) d\tilde{I}_0^i}{k dk}. \quad (6.6)$$

The optimal filter is therefore narrower in k -space than both $d\tilde{I}/d(k^2)$ and $\tilde{T}(k)$. Hence, in real space, on the deconvolved image, the optimal filter is broader than both the object and the PSF. This means that we should, sensibly, restrict our roundness test to the region of k -space that (1) has signal in the true galaxy image and (2) is not suppressed below the noise by the convolution.

Our strategy might then be to create some kind of deconvolved image and apply a Gaussian weight to test for roundness in the deconvolved k -space (which is equivalent to using a Gaussian weight in real space). For a Gaussian object with preseeing size σ_i and Gaussian seeing with size σ_* , the optimal weight in *deconvolved* k -space is a Gaussian with size $(2\sigma_*^2 + \sigma_i^2)^{-1/2}$. Such a roundness test is equivalent to a roundness test on the *observed, real-space* image with a Gaussian weight of size $(\sigma_i^2 + \sigma_*^2)^{1/2} = \sigma_o^2$. So the optimal size of the weight is again matched to the size of the *observed* image.

Recall that our algorithm for measuring δ requires that we shear the coordinates until the object appears round ($\tilde{M} \rightarrow 0$). We want to apply this shear to the deconvolved image. If the object is not round to begin with, then in this sheared coordinate system the transfer function $\tilde{T}(\mathbf{k})$ will no longer have azimuthal symmetry, which will invalidate the

above derivation of the optimal weight. We will still have a valid measurement of the shape of the deconvolved object, but possibly with suboptimal noise level. The increase in noise is a second-order effect, however, so we will not bother to reoptimize the roundness test for this asymmetric noise spectrum.

When the transfer function is anisotropic, the noise spectrum in the deconvolved image is also anisotropic. There are subtle second-order effects, described in § 8, that bias the orientation of the measured shear in the presence of an anisotropy in the noise. Such a noise anisotropy is present after any deconvolution of an anisotropic PSF. A noise anisotropy is also present in the *observed* image, if the PSF is anisotropic and we are not strictly sky limited. Furthermore, K00 points out a selection bias that can creep into the shear measurements even with a perfectly unbiased shape-measurement algorithm. We will discuss means to defeat these biases in § 8.

Applying a Gaussian weight in deconvolved k -space leads, formally, to divergent noise if $\tilde{T}(k) = 0$ for some k . This is a real problem, as any finite-sized telescope must produce a transfer function that is identically zero beyond some critical k_c . The deconvolved image hence has infinite noise for $k > k_c$, while the Gaussian weight remains finite. As we shear the k -space to make our source appear round, the “wall” of infinite noise moves inward to $e^{-\eta/2}k_c$. In order for our method to remain feasible, our deconvolution algorithm must not attempt to fully deconvolve those portions of k -space at or near k_c , so that the noise remains bounded. There is hence a balance to be struck in executing the deconvolution: we want to carry out the deconvolution to sufficiently high order that the effects of the PSF upon the source ellipticity are removed, but we do not want to deconvolve high-order details that will increase the noise. It would seem, intuitively, that this is possible, since the ellipticity we seek, the Gaussian-weighted quadrupole moment, is a low-order characteristic of the deconvolved image.

The method we describe below is based upon an expansion of the image and PSF into hierarchies of Gaussian-weighted moments—essentially an eigenfunction decomposition. Convolution corresponds to a matrix operation on the vector of moments. The moment vector, and hence the convolution matrix, are formally infinite, but we can choose to truncate the description at some order that we believe to contain all the useful information on the image ellipticity. The convolution matrix is then finite and can be inverted, and the deconvolution executed as a matrix operation. The high-order moments are not deconvolved, so the noise in the deconvolved moments remains finite and, in fact, close to optimally small. This moment-based method reduces the deconvolution (as well as other transformations) to a matrix multiplication, which can be executed on an exposure-by-exposure basis, even for objects with very low S/N on a single exposure, and is therefore very practical for the purpose of weak-lensing measurements.

6.3. The Laguerre Expansion

6.3.1. Definitions

The simplicity of the formulae for deconvolution in the special case of Gaussian objects, plus the utility of the Gaussian weight for shape measurements, led us to seek a description of the image in some Gaussian-based expansion. To maintain the simplest form for convolutions, we look for

a decomposition of our images into components that are eigenfunction of the Fourier operator. Such functions will also be eigenfunctions of $(-\nabla^2 + r^2)$, and hence we are led to the eigenfunctions of the two-dimensional quantum harmonic oscillator (QHO).² In one dimension, the QHO eigenfunctions are each a Gaussian times a Hermite polynomial. The Edgeworth expansion, familiar to many astronomers, is a decomposition into one-dimensional QHO eigenfunctions. The two-dimensional version we call the “Laguerre expansion,” using the QHO eigenfunctions

$$I(r, \theta) = \sum_{p, q \geq 0} b_{pq} \psi_{pq}^\sigma(r, \theta), \quad (6.7)$$

$$\psi_{pq}^\sigma(r, \theta) = \frac{(-1)^q}{\sqrt{\pi}\sigma^2} \sqrt{\frac{q!}{p!}} \left(\frac{r}{\sigma}\right)^m e^{im\theta} e^{-r^2/2\sigma^2} L_q^{(m)}\left(\frac{r^2}{\sigma^2}\right) \quad (p \geq q), \quad (6.8)$$

$$\psi_{qp}^\sigma = \bar{\psi}_{pq}^\sigma, \quad (6.9)$$

$$m \equiv p - q. \quad (6.10)$$

$L_q^{(m)}(x)$ are the Laguerre polynomials, which are defined by the generating function

$$(1 - z)^{-q-1} \exp\left(\frac{xz}{z-1}\right) = \sum_{m=0}^{\infty} L_m^{(q)}(x) z^m \quad (6.11)$$

(Abramowitz & Stegun 1965). The Laguerre polynomials satisfy many recurrence relations; the following provide a way to calculate them rapidly:

$$L_0^{(m)}(x) = 1, \quad (6.12)$$

$$L_1^{(m)}(x) = (m+1) - x, \quad (6.13)$$

$$(q+1)L_{q+1}^{(m)}(x) = [(2q+m+1) - x]L_q^{(m)}(x) - (q+m)L_{q-1}^{(m)}(x). \quad (6.14)$$

A QHO with wave function ψ_{pq} has angular momentum $m = p - q$. We will also make use of the quantum number $N = p + q$, which is the excitation energy of the state. Any two of $\{N, m, p, q\}$ suffice to specify the state. The intensity multipole functions $I_m(r)$ are composed from the ψ_{Nm} for $N = |m|, |m| + 2, |m| + 4, \dots$. The polynomial in ψ_{pq} has terms up to order N in r , and ψ_{pq} can also be expressed as the Gaussian times a (complex) polynomial of order N in x and y . A few low-order ψ_{pq} are plotted in Figure 6.

Refregier (2001) has independently introduced the application of QHO eigenfunctions to galaxy shape analysis. Some of the results and ideas presented in this section are presented therein, and they are applied in Refregier & Bacon (2001) with different notation. Refregier (2001) also presents useful relations for a Cartesian-based family of two-dimensional QHO eigenfunctions. We make use only of the polar family, which are eigenstates of the angular momentum and hence have strong rotational symmetry.

The Laguerre functions have many properties we will find useful. As eigenfunctions of the harmonic oscillator, they are orthonormal, up to the factor σ , which we introduce in

² Thanks again to the anonymous referee for providing this logic.

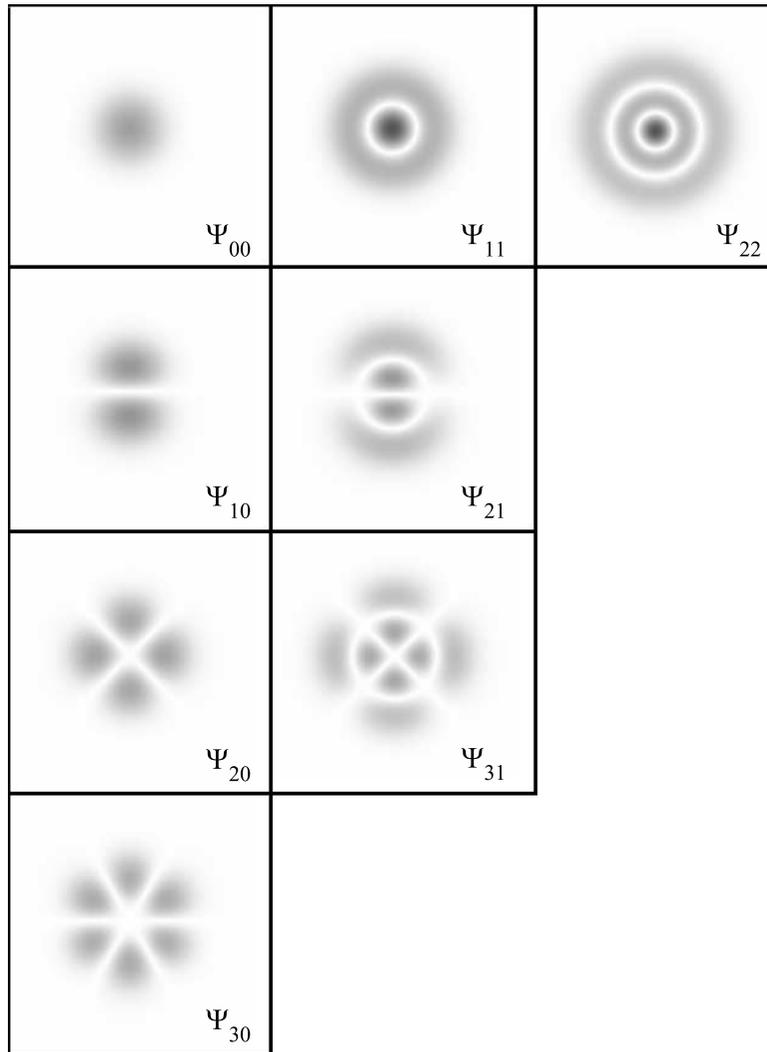


FIG. 6.—First few of the orthogonal functions ψ_{pq} . Only the real parts are plotted. The characteristics are familiar from their use as eigenfunctions of the two-dimensional quantum harmonic oscillator: each has $2m = 2|p - q|$ azimuthal nodes and $N - 2m$ radial nodes, where $N = p + q$ is the number of quanta in the state. The ψ_{20} component is most important, as it responds in first order to shear, and hence its absence is our test for “roundness.”

order to give the b_{pq} units of flux:

$$\int d^2x \psi_{pq}^\sigma(\mathbf{x}) \bar{\psi}_{p'q'}^\sigma(\mathbf{x}) = \frac{1}{\sigma^2} \delta_{pp'} \delta_{qq'} , \quad (6.15)$$

$$\Rightarrow b_{pq} = \sigma^2 \int d^2x I(\mathbf{x}) \bar{\psi}_{pq}^\sigma(\mathbf{x}) . \quad (6.16)$$

Thus, b_{pq} is a Gaussian-weighted moment of the intensity. Since I is real, we must have $b_{pq} = \bar{b}_{qp}$. When measuring the b_{pq} from an image with uniform white noise, equation (6.16) yields a covariance matrix for the b_{pq} that is diagonal:

$$\text{Cov} (b_{pq} \bar{b}_{p'q'}) = n\sigma^2 \delta_{pp'} \delta_{qq'} , \quad (6.17)$$

where n is the number of counts per unit area. The variance in b_{pq} is shared between the real and imaginary parts, except for b_{pp} , which must be real. Nonuniform noise, e.g., shot noise from the galaxy itself, produces a more complicated covariance matrix, as described by Refregier (2001). The significance ν of detection with the Gaussian filter is

given by

$$\nu^2 = b_{00}^2 / n\sigma^2 . \quad (6.18)$$

Our algorithm for measuring object shapes requires finding the coordinate system in which the centroid is zero (eq. [3.16]) and the roundness test (eq. [3.6]) yields zero. When disregarding seeing, we also set the weight size σ by maximizing the significance (eq. [3.17]). These conditions are succinctly stated by the Laguerre coefficients:

$$\text{Centroid} \quad b_{10} = 0 ; \quad (6.19)$$

$$\text{Roundness} \quad b_{20} = 0 ; \quad (6.20)$$

$$\text{Significance} \quad b_{11} = 0 . \quad (6.21)$$

The first two of these equations involve both real and complex components; the third is real. We must satisfy these equations by translation, shear, and dilation of the object (or of the underlying coordinate system). These operations

can be expressed as transformation matrices acting upon the vector $\mathbf{b} = \{b_{pq}\}$. The determination of shape is thus efficiently executed by measuring \mathbf{b} in the original coordinate frame, converting equation (6.16) to a sum over pixels. Then we can iterate to a solution of our three conditions by manipulating \mathbf{b} , and there is no need to return to the pixel data.

In the case of finite resolution, we wish to satisfy equations (6.19) and (6.20) for the deconvolved image. We can express the deconvolution as a matrix operation on \mathbf{b} as well. So we need to find the matrix equivalents of the translation, dilation, shear, and convolution transformations.

6.3.2. Raising and Lowering Operators, Point Transformations

The raising and lowering operators for the two-dimensional harmonic oscillator wave functions have the properties

$$\begin{aligned} a_p^\sigma &= \frac{1}{2} \left[\frac{x - iy}{\sigma} + \sigma \left(\frac{\partial}{\partial x} - i \frac{\partial}{\partial y} \right) \right], & a_p^\sigma \psi_{pq}^\sigma &= \sqrt{p} \psi_{(p-1)q}^\sigma, \\ a_p^{\sigma\dagger} &= \frac{1}{2} \left[\frac{x + iy}{\sigma} - \sigma \left(\frac{\partial}{\partial x} + i \frac{\partial}{\partial y} \right) \right], & a_p^{\sigma\dagger} \psi_{pq}^\sigma &= \sqrt{p+1} \psi_{(p+1)q}^\sigma, \\ a_q^\sigma &= \frac{1}{2} \left[\frac{x + iy}{\sigma} + \sigma \left(\frac{\partial}{\partial x} + i \frac{\partial}{\partial y} \right) \right], & a_q^\sigma \psi_{pq}^\sigma &= \sqrt{q} \psi_{p(q-1)}^\sigma, \\ a_q^{\sigma\dagger} &= \frac{1}{2} \left[\frac{x - iy}{\sigma} - \sigma \left(\frac{\partial}{\partial x} - i \frac{\partial}{\partial y} \right) \right], & a_q^{\sigma\dagger} \psi_{pq}^\sigma &= \sqrt{q+1} \psi_{p(q+1)}^\sigma. \end{aligned} \quad (6.22)$$

Note that the commutators of the operators are all zero except for $[a_p, a_p^\dagger] = [a_q, a_q^\dagger] = 1$. These operators can be used to formulate the transformation matrices we need. Consider first a dilation of the image I by factor of $1 + \mu$ with $\mu \ll 1$ to a new image I' :

$$I'(x, y) = I((1 - \mu)x, (1 - \mu)y) \quad (6.23)$$

$$\approx \left[1 - \mu \left(x \frac{\partial}{\partial x} + y \frac{\partial}{\partial y} \right) \right] I \quad (6.24)$$

$$= [1 - \mu(a_p^\sigma a_q^\sigma - a_p^{\sigma\dagger} a_q^{\sigma\dagger} - 1)] I \quad (6.25)$$

$$= \sum b_{pq} [(1 + \mu) \psi_{pq} - \mu \sqrt{pq} \psi_{(p-1)(q-1)} + \mu \sqrt{(p+1)(q+1)} \psi_{(p+1)(q+1)}], \quad (6.26)$$

$$\Rightarrow b'_{pq} = (1 + \mu) b_{pq} - \mu \sqrt{(p+1)(q+1)} b_{(p+1)(q+1)} + \mu \sqrt{pq} b_{(p-1)(q-1)}. \quad (6.27)$$

An infinitesimal dilation μ thus mixes coefficients $(N \pm 2, m)$ into the (N, m) coefficient—as expected, multipole order m is conserved by the transformation. The dilation transformations are thus generated by the matrix

$$\mathbf{d} = 1 - a_p^\sigma a_q^\sigma + a_p^{\sigma\dagger} a_q^{\sigma\dagger}, \quad (6.28)$$

and the transformation matrix \mathbf{D} for a finite dilation e^μ can be expressed as

$$\mathbf{D}_\mu = \exp(\mu \mathbf{d}). \quad (6.29)$$

The finite dilation will preserve m as the generator does but will mix terms of $N + 2j$ together, with j any integer. The matrix elements for the finite dilation are derivable in closed form and in a rapid recursion form; the latter are derived in Appendix A.

The generator \mathbf{s} for the shear transformation is also easily expressed in terms of the raising and lowering operators. For a shear $\eta \ll 1$ along the x -axis,

$$I'(x, y) = I((1 - \eta/2)x, (1 + \eta/2)y) \quad (6.30)$$

$$\approx \left[1 - \frac{\eta}{2} \left(x \frac{\partial}{\partial x} - y \frac{\partial}{\partial y} \right) \right] I \quad (6.31)$$

$$= (1 + \eta \mathbf{s}) I, \quad (6.32)$$

$$\mathbf{s} = \frac{1}{4} [(a_p^\dagger)^2 + (a_q^\dagger)^2 - a_p^2 - a_q^2], \quad (6.33)$$

$$\begin{aligned} \Rightarrow b'_{pq} &= b_{pq} + \frac{1}{4} \eta [\sqrt{p(p-1)} e^{-2i\beta} b_{(p-2)q} \\ &\quad + \sqrt{q(q-1)} e^{2i\beta} b_{p(q-2)} \\ &\quad - \sqrt{(p+1)(p+2)} e^{2i\beta} b_{(p+2)q} \\ &\quad - \sqrt{(q+1)(q+2)} e^{-2i\beta} b_{p(q+2)}], \end{aligned} \quad (6.34)$$

$$\mathbf{S}_\eta = \exp(\eta \mathbf{s}). \quad (6.35)$$

In equation (6.34) we have inserted the phase factors that result from a shear at arbitrary position angle β . To leading order, the shear mixes in states with p or q changed by ± 2 ; in the finite case, the shear transformation mixes together b_{pq} for which p and q each change by any even number. Note that this equation generalizes the weighted-moment transformations in Rhodes et al. (2000) to all orders. Note also that the KSB method identifies the quantity b_{20}/b_{00} as the ellipticity, so the “shear polarizability” is compactly expressed using equation (6.34) with $p = 2$ and $q = 0$.

An infinitesimal translation of the image by (x_0, y_0) yields

$$I'(x, y) = I(x - x_0, y - y_0) \quad (6.36)$$

$$\approx \left(1 - x_0 \frac{\partial}{\partial x} - y_0 \frac{\partial}{\partial y} \right) I \quad (6.37)$$

$$= (1 + \mathbf{z} \mathbf{t} - \bar{\mathbf{z}} \mathbf{t}^\dagger) I, \quad (6.38)$$

$$\mathbf{z} \equiv (x_0 + iy_0)/\sigma, \quad (6.39)$$

$$\mathbf{t} = \frac{1}{2} (a_q^\dagger - a_p), \quad (6.40)$$

$$\begin{aligned} \Rightarrow b'_{pq} &= b_{pq} + \frac{1}{2} \mathbf{z} (\sqrt{q} b_{p(q-1)} - \sqrt{p+1} b_{(p+1)q}) \\ &\quad + \frac{1}{2} \bar{\mathbf{z}} (\sqrt{p} b_{(p-1)q} - \sqrt{q+1} b_{p(q+1)}), \end{aligned} \quad (6.41)$$

$$\mathbf{T}_z = \exp(\mathbf{z} \mathbf{t} - \bar{\mathbf{z}} \mathbf{t}^\dagger). \quad (6.42)$$

To leading order the translation changes p or q by ± 1 , and the finite translation \mathbf{T}_z mixes all the states together. Appendix A gives the matrix elements for all the finite transformations.

The “smear polarizability” coefficients of KSB could be quickly derived at this point by noting that an infinitesimal convolution along the x -direction can be expressed as an average of two images translated by $\pm x_0$, which will lead to a smear generator that is second order in the raising and lowering operators. Appendix C gives a closely related derivation.

6.3.3. Shape Measurement from Laguerre Decomposition

Our algorithm for shape measurements (with perfect resolution) is to find the translation \mathbf{z} , dilation μ , and shear η that must be applied to the image to satisfy the conditions in equations (6.19)–(6.21). If \mathbf{b} is the vector of Laguerre coefficients for the image I , we can write

$$\mathbf{b}' = \mathbf{M} \mathbf{b} = (\mathbf{S}_\eta \mathbf{D}_\mu \mathbf{T}_z) \mathbf{b}, \quad (6.43)$$

and the elements of \mathbf{b}' with $(p, q) = (1, 0)$, $(1, 1)$, and $(2, 0)$ must vanish. This is in general a complex nonlinear equation, but for $\{\mu, \eta, z\} \ll 1$, the linearized equations (6.27), (6.34), and (6.41) yield

$$b'_{10} = b_{10} + (b_{10} - \sqrt{2}b_{21})\mu - \frac{1}{4}\sqrt{2}b_{12}\bar{\eta} - \frac{1}{4}\sqrt{6}b_{30}\eta - \frac{1}{2}\sqrt{2}b_{20}z + \frac{1}{2}(b_{00} - b_{11})\bar{z}, \quad (6.44)$$

$$b'_{11} = b_{11} + (b_{00} + b_{11} - 2b_{22})\mu - \frac{1}{4}\sqrt{6}b_{13}\bar{\eta} - \frac{1}{4}\sqrt{6}b_{31}\eta + \frac{1}{2}(b_{10} - \sqrt{2}b_{21})z + \frac{1}{2}(b_{01} - \sqrt{2}b_{12})\bar{z}, \quad (6.45)$$

$$b'_{20} = b_{20} + (b_{20} - \sqrt{3}b_{31})\mu + \frac{1}{4}\sqrt{2}(b_{00} - b_{22})\bar{\eta} - \frac{1}{2}\sqrt{3}b_{40}\eta - \frac{1}{2}\sqrt{3}b_{30}z + \frac{1}{2}(\sqrt{2}b_{10} - b_{21})\bar{z}, \quad (6.46)$$

writing $\eta = \eta_+ + i\eta_x$. These linear equations may be solved explicitly for the desired transformation coefficients η , μ , and z . The solution appears much simpler if the object is nearly round, such that $b_{pq} \ll b_{00}$ for $p \neq q$. In this case the transformation parameters are, to leading order,

$$\eta = \frac{-2\sqrt{2}b_{02}}{b_{00} - b_{22}}, \quad (6.47)$$

$$z = \frac{-2b_{01}}{b_{00} - b_{11}}, \quad (6.48)$$

$$\mu = \frac{-b_{11}}{b_{00} - 2b_{22}}. \quad (6.49)$$

From these equations it is clear that b_{20} , b_{10} , and b_{11} are the primary carriers of information on shape, centroid, and size, respectively. Since the shape component η_+ is $(\eta + \bar{\eta})/2$, its uncertainty is, to leading order,

$$\sigma_{\eta}^2 \equiv \text{Var}(\eta_+) = \frac{2[\text{Var}(b_{20}) + \text{Cov}(b_{20}b_{02}) + \text{Var}(b_{02})]}{(b_{00} - b_{22})^2} \quad (6.50)$$

$$= \frac{4n\sigma^2}{b_{00}^2} \left(1 - \frac{b_{22}}{b_{00}}\right)^{-2}, \quad (6.51)$$

$$\Rightarrow \sigma_{\eta} = \frac{2}{\nu} \left(1 - \frac{b_{22}}{b_{00}}\right)^{-1}. \quad (6.52)$$

Here we have made use of equation (6.17) for the covariance matrix in the case of white noise, and equation (6.18) for the definition of the significance ν . We see that the Laguerre expansion easily reproduces the earlier result in equation (3.24), with the a_4 parameter in that equation being simply the strength of the b_{22} coefficient.

With a more tedious procedure we may derive the solution for η to second order in b_{pq}/b_{00} . We take $b_{pp} \ll b_{00}$ for $p > 0$ this time in order to simplify the expression

$$\begin{aligned} \eta \approx & -2\sqrt{2}\frac{b_{02}}{b_{00} - b_{22}} - 2\sqrt{3}\frac{b_{20}b_{04}}{b_{00}^2} \\ & + \frac{2b_{01}^2 - 2\sqrt{6}b_{10}b_{03} - 2\sqrt{2}b_{01}b_{12}}{b_{00}^2} \\ & - 2\sqrt{6}\frac{b_{11}b_{31}}{b_{00}^2}. \end{aligned} \quad (6.53)$$

The terms in the first line arise from solving for $\eta = 0$; the terms in the second and third lines arise when simultaneously solving the centroid and size constraints, respectively.

From these terms, one can derive the $O(\nu^{-4})$ terms in the shape uncertainty if desired. But more importantly, this expression uncovers a very significant source of bias in shape measurements: the presence of the b_{pq} in second order means that it is possible for noise to be rectified in the determination of \mathbf{e} . If the noise is anisotropic—as is the case after correction for an anisotropic PSF—then \mathbf{e} will be biased. We will explore this in more detail in § 8.

6.3.4. Fourier Transforms and Convolution

The observed galaxy intensity $I^o(\mathbf{x})$ is the convolution of the image-plane intensity $I^i(\mathbf{x})$ and the stellar PSF $I^*(\mathbf{x})$:

$$I^o(\mathbf{x}) = I^i(\mathbf{x}) \circ I^*(\mathbf{x}) = \int d^2x' I^i(\mathbf{x}') I^*(\mathbf{x} - \mathbf{x}'). \quad (6.54)$$

The postseeing, preseeing, and stellar images can be expressed as vectors \mathbf{b}^o , \mathbf{b}^i , and \mathbf{b}^* , respectively, of coefficients over our eigenfunction sets $\psi_{p_0q_0}^{\sigma_o}$, $\psi_{p_1q_1}^{\sigma_i}$, and $\psi_{p_2q_2}^{\sigma_*}$. The convolution can then be expressed as a matrix relation

$$\mathbf{b}^o = \mathbf{C}(\mathbf{b}^*)\mathbf{b}^i, \quad (6.55)$$

$$b_{p_0q_0}^o = \sum C_{p_0q_0}^{p_1q_1p_2q_2} b_{p_1q_1}^i b_{p_2q_2}^*. \quad (6.56)$$

We will effect the deconvolution by inverting a truncated version of the matrix \mathbf{C} . Its coefficients are determined by the relation

$$\psi_{p_1q_1}^{\sigma_i} \circ \psi_{p_2q_2}^{\sigma_*} = \sum C_{p_0q_0}^{p_1q_1p_2q_2} \psi_{p_0q_0}^{\sigma_o}. \quad (6.57)$$

The convolution is more easily expressed in k -space. With the usual definitions (denoted as “system 3” by Bracewell 1978),

$$\tilde{I}(\mathbf{k}) = \frac{1}{2\pi} \int d^2x I(\mathbf{x}) e^{-i\mathbf{k} \cdot \mathbf{x}}, \quad (6.58)$$

$$I(\mathbf{x}) = \frac{1}{2\pi} \int d^2k \tilde{I}(\mathbf{k}) e^{i\mathbf{k} \cdot \mathbf{x}},$$

the convolution becomes a multiplication, and the matrix coefficients in equation (6.57) can also be expressed as

$$2\pi \tilde{\psi}_{p_1q_1}^{\sigma_i} \tilde{\psi}_{p_2q_2}^{\sigma_*} = \sum C_{p_0q_0}^{p_1q_1p_2q_2} \tilde{\psi}_{p_0q_0}^{\sigma_o}. \quad (6.59)$$

We now make use of another remarkable property of the Laguerre-exponential eigenfunctions, which is that *they are their own Fourier transforms*. First we note that ψ_{00}^{σ} is a two-dimensional Gaussian, and the transform (eq. [6.58]) is easily executed to yield a new Gaussian,

$$\tilde{\psi}_{00}^{\sigma} = \frac{1}{\sqrt{\pi}} e^{-k^2\sigma^2/2}. \quad (6.60)$$

The functional form of $\tilde{\psi}_{00}(k)$ matches that of $\psi_{00}(x)$, except that we must send $\sigma \rightarrow 1/\sigma$. Next we can use the definition of the raising operators (eq. [6.22]) to note that

$$\widetilde{a_p^{\sigma\dagger}} \tilde{\psi} = \frac{-i}{2} \left[\sigma(k_x + ik_y) - \frac{1}{\sigma} \left(\frac{\partial}{\partial k_x} + i \frac{\partial}{\partial k_y} \right) \right] \tilde{\psi} \equiv \tilde{a}_p^{\sigma\dagger} \tilde{\psi}. \quad (6.61)$$

Thus, the k -space raising operator has the same form as the x -space raising operator, save the $\sigma \rightarrow 1/\sigma$ transformation and an additional factor of $-i$. The same is true of $\tilde{a}_q^{\sigma\dagger}$. Since ψ_{00} and the raising operators are each unchanged by the

Fourier transform, it must be true that all the ψ_{pq} are their own Fourier transforms as well, with appropriate scaling factors of i and σ . More precisely, we have

$$\tilde{\psi}_{pq}^\sigma(k, \phi) = \frac{(-i)^m}{\sqrt{\pi}} \sqrt{\frac{q!}{p!}} (k\sigma)^m e^{im\phi} e^{-k^2\sigma^2/2} L_q^{(m)}(k^2\sigma^2) \quad (p \geq q) \quad (6.62)$$

(cf. eq. [6.8]). So the problem of convolving two of our eigenfunctions is reduced to the simpler problem of multiplying the same two eigenfunctions. We may reach several conclusions immediately:

1. The convolution of $\psi_{N_i m_i}^{\sigma_i}$ with $\psi_{N_* m_*}^{\sigma_*}$ will produce an observed image with azimuthal order $m_o = m_i + m_*$. Furthermore, the multipole phase of the observed image must be the sum of the phases for the original and PSF images' coefficients.

2. If we choose the scale size σ_o for the observed image eigenfunctions such that

$$\sigma_o^2 = \sigma_i^2 + \sigma_*^2, \quad (6.63)$$

then the convolution contains components $\psi_{N_o m_o}^{\sigma_o}$ only for $N_o \leq N_i + N_*$. Recall also that we must have $N_o \geq m_o = m_i + m_*$, so it must be true that $\psi_{p_i q_i}^{\sigma_i} \circ \psi_{p_* q_*}^{\sigma_*} \propto \psi_{(p_i+p_*) q}^{\sigma_o}$ for $q = 0$.

In Appendix B we give a recursion relation to calculate any of the elements $C_{p_o q_o}^{p_i q_i p_* q_*}$ that we need to calculate the convolution and, thus, the deconvolution.

6.3.5. Deconvolution, Noise Amplification, and Truncation

With the formulae for the convolution matrix \mathbf{C} in hand, we can investigate the nature of the trade-off between the fidelity of the deconvolution—i.e., the extent to which effects of the PSF upon the shape are removed—and the noise level of the deconvolution. We must truncate the \mathbf{b} -vectors at some finite order N to implement the deconvolution; higher N will remove more of the PSF effects but also increase the noise level. We can illustrate this phenomenon by considering the simplest case of convolution by a unit Gaussian PSF, $b_{pq}^* = 2\sqrt{\pi}\delta_{p0}\delta_{q0}$. Since the PSF has only $m_* = 0$ terms, we will have $m_o = m_i$, so the convolution matrix is block diagonal, and we may deconvolve each m independently. If we choose $\sigma_o^2 = \sigma_i^2 + \sigma_*^2$, then the convolution matrix coefficients are zero for $N_o > N_i$, and hence each block matrix is also upper triangular. This makes the inversion easy. Using the results of Appendix B, defining a deconvolution parameter $D = \sigma_i^2/\sigma_o^2 = 1 - \sigma_*^2/\sigma_o^2$, and using equation (6.17) for the covariance of the measured moments, the deconvolved value of b_{00}^i and its variance are

$$b_{00}^i = \sum_{p=0}^{\infty} \left(\frac{1-D}{D}\right)^p (-1)^p b_{pp}^o, \quad (6.64)$$

$$\text{Var}(b_{00}^i) = n\sigma_o^2 \sum_{p=0}^{\infty} \left(\frac{1-D}{D}\right)^{2p} = n\sigma_i^2 \frac{D^2}{2D-1} \quad \left(D > \frac{1}{2}\right). \quad (6.65)$$

In practice, we could not measure the b_{pp}^o to infinite p and we must truncate the sum (eq. [6.64]) at some finite p . The more terms we include, the more accurately we describe the deconvolved b_{00}^i , but each additional term in the deconvolu-

tion adds more noise. For $D > \frac{1}{2}$, the added noise in successive terms drops as p increases, and we could in principle do a complete deconvolution of b_{00}^i with finite variance. For $D \leq \frac{1}{2}$, however, the noise in the deconvolved moment diverges for $p \rightarrow \infty$, and we are forced to truncate the deconvolution matrix.

For ellipticity measurements, we are primarily interested in the b_{20}^i moment. For this case of a Gaussian PSF, the general b_{m0}^i moment deconvolution is

$$b_{m0}^i = \sum_{q=0}^{\infty} \frac{1}{D^{m/2}} \left(\frac{1-D}{D}\right)^q (-1)^q \sqrt{\binom{q+m}{m}} b_{(q+m)q}^o \quad (6.66)$$

$$\Rightarrow \text{Var}(b_{m0}^i) = \frac{n\sigma_o^2}{D^m} \sum_{q=0}^{\infty} \left(\frac{1-D}{D}\right)^{2q} \binom{q+m}{m} \quad (6.67)$$

$$= n\sigma_i^2 D \left(\frac{D}{2D-1}\right)^{m+1} \quad \left(D > \frac{1}{2}\right). \quad (6.68)$$

We note that these moments also demonstrate the properties that they are noisier than the observed moments, and infinitely so if $D < \frac{1}{2}$ and the deconvolution is not truncated. The noise increases as m increases, as we would expect, since higher spatial frequencies must be recovered.

For a purely Gaussian PSF with $D > \frac{1}{2}$, it is possible to complete the deconvolution with finite noise. A real PSF must have components beyond b_{00}^* , since the transfer function vanishes above a critical k -value. When we invert the convolution matrix for such a PSF we will find that the variance series akin to equation (6.65) would diverge for any $D < 1$. The best truncation value for the matrix will depend upon the form of the PSF and the accuracy to which we demand correction of PSF effects. We will examine some specific cases in Jarvis et al. (2002).

6.3.6. Analytic or Kernel Correction for PSF Bias?

There are two general means of eliminating the shape *bias* induced by the PSF. One alternative is to measure the anisotropy of the PSF carefully and apply analytic corrections to the measured objects, as occurs naturally within the deconvolution framework described in the preceding paragraphs. The KSB formalism contains an approximate analytic correction. The other method is to convolve the image with a spatially varying kernel that removes the anisotropy from the PSF, as first demonstrated by FT97 and further advocated in K00. Removal of the PSF bias is the most critical task for a weak-lensing method: the PSF *dilution* is only a calibration error or signal modulation, but the PSF *bias* introduces a first-order false signal to the lensing analysis.

Ideally, analytic correction is preferred. The convolution kernel will always degrade the image resolution to some extent. Many PSFs cannot be “rounded out” unless the kernel is of size comparable to the PSF itself (K00). Furthermore, the convolution will perturb the noise spectrum of the image, complicating error estimation. The kernel method, however, can be faster for simple kernels than an involved analytic correction, especially if eliminating the bias is more important than calibrating the dilution for the task at hand.

The kernel method may offer peace of mind. The demands for rejection of the PSF bias are very stringent: PSF ellipticities of 10% are not uncommon, yet state-of-the-art lensing surveys require systematic errors below 0.1%, so the method we choose must reduce the PSF bias by over 2

orders of magnitude. If the kernel method can make the PSF truly round, then symmetry principles preclude any artificial coherent ellipticity. We flesh out this statement in § 7. We may have more trust in the symmetry principle than we do in an analytic correction, especially when finite sampling is taken into account. Note, however, that the symmetry principle requires that the noise power spectrum also be isotropic; § 8 below demonstrates how anisotropic noise can bias shape measurements through *centroid bias*.

We therefore will describe an alternative procedure: eliminate the bias with a kernel correction; measure the shapes, either isotropizing the noise or making a correction for centroid bias; then correct them for dilution with a heuristic formula such as equation (6.1). In § 7, we describe a rounding-kernel methodology, and in Appendix C we derive a higher order version of the dilution correction. These methods were used by Fischer et al. (2000), Smith et al. (2001), and the forthcoming CTIO lens survey (Jarvis et al. 2002).

6.4. Pixelated Data

Real astronomical data are integrated into pixels and sampled at finite intervals, and the continuum limit that we have assumed in our analysis is not strictly valid. A formal description of the process is that the PSF is altered by convolution with the pixel response function (PRF) to produce the *effective PSF* (ePSF), which is then sampled at the pixel pitch a —see Lauer (1999) or Bernstein (2002) for further elaboration. If the pixel pitch is coarser than the Nyquist interval of the PSF ($\lambda/2D$ for a diffraction-limited image), then there is aliasing and we cannot unambiguously reconstruct the original image or the true b_{pq} coefficients.

In the case where the pixel size a is small compared with the PSF size, the b_{pq} may be estimated by converting the integral equation (6.16) into a sum over pixels. But the formally correct procedure is to fit the vector of sampled pixel fluxes I_i with a model galaxy with free b_{pq}^σ . This is a standard minimization problem for the χ^2 parameter,

$$\chi^2 = \sum_i \frac{[I_i - \sum_{pq} b_{pq}^\sigma \psi_{pq}^\sigma(\mathbf{x}_i)]^2}{\text{Var}(I_i)}. \quad (6.69)$$

Since the model is linear, the solution for b_{pq}^σ and its covariance matrix has a closed form, but of course the number of b_{pq}^σ coefficients allowed in the model must be less than the number of pixels being fitted. With the best-fit coefficients and their covariance matrix in hand, we may proceed with the methods outlined above. It is the ePSF rather than the PSF that has convolved the source image, but since we in fact measure the ePSF for stars, the deconvolution procedures are unchanged. Dithered exposures can be handled by extending the pixel sum over multiple exposures, as long as the PSF is unchanging. Poorly sampled images will manifest the aliasing as strong degeneracies in the solution for the b_{pq} .

A potential time-saving procedure would be to fit the pixel data to a Laguerre expansion that is already convolved with the local PSF, minimizing

$$\chi^2 = \sum_i \frac{[I_i - \sum_{p_0 q_0} \psi_{p_0 q_0}^\sigma(\mathbf{x}_i) \sum_{p_i q_i} C_{p_0 q_0}^{p_i q_i} b_{p_i q_i}^i]^2}{\text{Var}(I_i)}. \quad (6.70)$$

Here C is the Laguerre-coefficient convolution matrix for the local PSF. The intrinsic coefficients b_{pq}^i are the

unknowns in this linear fit. This is very reminiscent of the method of Kuijken (1999), except that our model for the intrinsic galaxy shape may depart from circular symmetry, and our Laguerre formalism will allow the fit to proceed as a matrix solution, with matrices rapidly built by recursion. Aside from its suitability to poorly sampled data, this direct fitting approach has two further advantages over the deconvolution matrix method of the previous section. First, when the noise is not flat and equation (6.17) does not apply, the direct fitting method more directly produces the full covariance matrix for the \mathbf{b}^i vector. Second, the direct fit may make use of images that are partially contaminated by invalid pixels, e.g., cosmic rays, since they may be excluded from the χ^2 sum. It is this method that we consider most promising for real data.

7. ROUNDING-KERNEL METHODS

In this section, we develop a method for producing a convolution kernel that symmetrizes the PSF to some desired degree. As discussed above, this is a potentially efficient means of reducing or eliminating PSF bias from the shear determination. By taking advantage of the Laguerre decomposition, the derivation and application of the spatially dependent kernel can be efficiently implemented.

7.1. The Target Transfer Function

As discussed in § 6.2, the effect of seeing is to convolve the initial image, I^i , with a PSF transfer function T to produce an observed image, I^o . One could theoretically remove all the effects of seeing by convolving the observed image with a kernel K whose Fourier transform $\tilde{K} = 1/\tilde{T}$ to produce a final image, I^f .

However, as discussed above, this is impossible to carry out in practice, because $\tilde{T}(\mathbf{k}) = 0$ for all k above some critical value. We can avoid this problem by ignoring the high-order moments of I^o and I^f . This is straightforward when one does an eigenfunction expansion of the images using the eigenfunctions introduced in § 6.3, since the expansion can be truncated at some value of $N = p + q$. This captures most of the real information about the PSF without introducing the noise from the higher order moments. After the convolution, we then have a target transfer function, T' , which is not exactly a delta function, but which can be made to match a delta function up to some order N . Most generally we have

$$I^o = T \circ I^i, \quad (7.1)$$

$$T = \sum_{pq} b_{pq}^* \psi_{pq}^\sigma(r, \theta), \quad (7.2)$$

$$I^f = T' \circ I^i = K \circ T \circ I^o, \quad (7.3)$$

$$T' = \sum_{pq} b_{pq}^{*'} \psi_{pq}^\sigma(r, \theta). \quad (7.4)$$

For the ideal case of $T' = \delta(\mathbf{x})$,

$$b_{pq}^{*'} = [(-1)^p / \sqrt{\pi}] \delta_{pq}. \quad (7.5)$$

Thus, if T' satisfies equation (7.5) up to some cutoff order N , I^f will be approximately identical to I^i up to the same order. If we can find a kernel K that produces this target transfer function, T' , we will be able to remove all the effects of see-

ing as well as possible given our ignorance of the high-order terms.

How stringently do we need to satisfy equation (7.5)? Less strict requirements on T' make it easier to find an appropriate, compact kernel. Our present goal is to create a transfer function T' that does not produce any shear bias. If the original scene I' is unlensed, and we represent the shear measurement process as some operator $\delta(I)$, then the isotropy of the universe guarantees that $\delta(I) = 0$ (we consider δ as a complex number $\delta_+ + i\delta_\times$). Our demand on the transfer function is that

$$\delta(I') = \delta(T' \circ I') = 0. \quad (7.6)$$

Most generally, the results of the shear measurement process can be expanded as a power series in the coefficients of the T' PSF:

$$\delta(T' \circ I') = \sum_{p_1 q_1} a_{p_1 q_1} b_{p_1 q_1}^{*'} + \sum_{\substack{p_1 q_1, \\ p_2 q_2}} a_{p_1 q_1, p_2 q_2} b_{p_1 q_1}^{*'} b_{p_2 q_2}^{*'} + \dots, \quad (7.7)$$

where the a_i are some coefficients that depend upon the measurement process, the image characteristics, and the size of the PSF. We now examine the consequence of rotating the image and the PSF by some angle β . The measured shear and the T' coefficients behave as

$$\delta(I') \rightarrow e^{2i\beta} \delta(I'), \quad b_{pq}^{*'} \rightarrow e^{(p-q)i\beta} b_{pq}^{*'} \quad (7.8)$$

The individual galaxies in the original scene I' are not invariant under rotation, but any statistical measure of their *collective* properties must be invariant under rotation, i.e., I' is invariant under rotation in the same sense that the universe is isotropic. Therefore the a_i coefficients of equation (7.7) are unaffected by the rotation, and to satisfy the conditions of equation (7.8), we must have

$$\begin{aligned} a_{p_1 q_1} &= 0 & \text{for } p_1 - q_1 \neq 2, \\ a_{p_1 q_1, p_2 q_2} &= 0 & \text{for } p_1 - q_1 + p_2 - q_2 \neq 2, \\ & \vdots & \end{aligned} \quad (7.9)$$

Thus, only PSF terms with $m = p - q = 2$ can cause a shear bias, to first order. The primary goal of our kernel, therefore, will be to set

$$b_{pq}^{*'} = 0 \quad (m = p - q = 2). \quad (7.10)$$

With this condition satisfied, shear bias can only be of order $(b_{pq}^{*'})^2$, where $p \neq q$. Higher order a 's are nonzero only if $m_1 + m_2 + \dots = 2$. Many of these elements are zero as well; for example, by considering the invariance of $\delta(I)$ under infinitesimal translation, one can demonstrate that $a_{10,21}$ vanishes, but $a_{10,10}$ must exist. Satisfying equation (7.10) does not, therefore, guarantee the elimination of shear bias to all orders.

For absolute assurance that a shear bias is absent, we can set

$$b_{pq}^{*'} = 0 \quad [p - q \neq 0 \pmod{4}], \quad (7.11)$$

i.e., ensure that the PSF has fourfold symmetry. In this case the condition $\sum m_i = 2$ can never be satisfied for nonvanishing coefficients of T' . Note that it would also suffice to enforce *any* m -fold rotational symmetry on the PSF beyond $m = 2$, e.g., the diffraction pattern of a triangular secondary support structure cannot by itself cause shear bias.

In practice we produce a kernel K that enforces the condition in equation (7.10) up to some order N . We can set $b_{10}^{*'} = 0$ by appropriate choice of PSF center. The remaining shear biases must be of order $(b_{21}/b_{00})^2$, $b_{41}b_{01}/b_{00}^2$, etc., which are generally quite small.

7.2. Components of the Kernel

Since we are going to convolve the kernel with a pixelated image, we must construct the kernel as a two-dimensional array instead of a continuous function. The simplest kernel is the identity kernel, composed of an array of 0's with a single 1 in the middle. This kernel conserves flux; we want a flux-conserving kernel, so we only consider kernels that are the identity kernel plus a kernel for which the elements sum to zero.

Next note that taking the derivative of an image is usually approximated by a discrete difference. For example,

$$\frac{\partial I}{\partial x} = \frac{I(x+dx) - I(x-dx)}{2dx}. \quad (7.12)$$

Another way to write this equation is as a convolution:

$$\frac{\partial I}{\partial x} = \begin{pmatrix} 0 & 0 & 0 \\ -\frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & 0 \end{pmatrix} \circ I. \quad (7.13)$$

Similarly,

$$\frac{\partial I}{\partial y} = \begin{pmatrix} 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 \\ 0 & -\frac{1}{2} & 0 \end{pmatrix} \circ I. \quad (7.14)$$

In fact, all partial derivatives of any order in x and y can be represented as a convolution by a zero-sum kernel. First- and second-order derivatives can be effected using 3×3 kernels; third- and fourth-order derivatives require 5×5 kernels, and so on. One can therefore think of the kernel as being made up of a sum of these derivatives $(\partial/\partial x)^i (\partial/\partial y)^j$ including, of course, the identity kernel, $(\partial/\partial x)^0 (\partial/\partial y)^0$.

The eigenfunction expansion of our images actually suggests a slightly different set of components for the kernel. Namely,

$$K = \sum_{ij} k_{ij} D_{ij}, \quad (7.15)$$

$$D_{ij} = \left(\frac{\partial}{\partial x} + i \frac{\partial}{\partial y} \right)^i \left(\frac{\partial}{\partial x} - i \frac{\partial}{\partial y} \right)^j \quad (7.16)$$

$$= \sigma^{-(i+j)} (a_q^\sigma - a_p^{\sigma\dagger})^i (a_p^\sigma - a_q^{\sigma\dagger})^j. \quad (7.17)$$

These components make it easy to use the raising and lowering operators to determine how a given kernel will act on an image. Note, however, that the D_{ij} are complex, which means that to end up with a real image after the convolution, we require that $k_{ji} = \bar{k}_{ij}$.

The 3×3 kernel components are

$$D_{10} = \begin{pmatrix} 0 & \frac{1}{2}i & 0 \\ -\frac{1}{2} & 0 & \frac{1}{2} \\ 0 & -\frac{1}{2}i & 0 \end{pmatrix}, \quad (7.18)$$

$$D_{01} = \overline{D_{10}}, \quad (7.19)$$

$$D_{20} = \begin{pmatrix} -\frac{1}{2}i & -1 & \frac{1}{2}i \\ 1 & 0 & 1 \\ \frac{1}{2}i & -1 & -\frac{1}{2}i \end{pmatrix}, \quad (7.20)$$

$$D_{02} = \overline{D_{20}}, \quad (7.21)$$

$$D_{11} = \begin{pmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{pmatrix}. \quad (7.22)$$

We note at this point that the FT97 method is equivalent to use of just the D_{20} and D_{02} kernel elements. The 5×5 and larger kernel components can be found by convolving the 3×3 components. For instance, $D_{30} = D_{10} \circ D_{20}$, and $D_{41} = D_{11} \circ D_{30}$.

7.3. Calculating the Kernel

7.3.1. The Kernel for Infinitesimal Pixels

The kernels given by equations (7.18)–(7.22) are only equal to the continuous derivatives (eq. [7.16]) to lowest order in $1/\sigma_*$. Typically, the PSF is only a few pixels in size, so this is not that good of an approximation. However, it is a good place to start, as most of the technique will apply to the case of finite pixels.

Combining equation (7.15) with the definition of T' (eq. [7.4]),

$$\mathbf{b}^{*'} = \sum_{ij} k_{ij} \mathbf{D}_{ij} \mathbf{b}^*. \quad (7.23)$$

The operator matrix \mathbf{D}_{ij} , in the limit of infinitesimal pixels, is easily calculated using equations (7.17) and (6.22), and there is a fast recursion:

$$\begin{aligned} \mathbf{D}_{00} \mathbf{b}^* &= \mathbf{b}^*, \\ \mathbf{D}_{(i+1)j} \mathbf{b}^* &= \frac{1}{\sigma_*} (a_q^{\sigma_*} - a_p^{\sigma_* \dagger}) (\mathbf{D}_{ij} \mathbf{b}^*), \\ \mathbf{D}_{i(j+1)} \mathbf{b}^* &= \frac{1}{\sigma_*} (a_p^{\sigma_*} - a_q^{\sigma_* \dagger}) (\mathbf{D}_{ij} \mathbf{b}^*). \end{aligned} \quad (7.24)$$

Since \mathbf{b}^* is measured and the \mathbf{D}_{ij} are fixed matrices, we have a matrix equation for the unknown kernel coefficients k_{ij} , given the chosen constraints on $\mathbf{b}^{*'} (e.g., eq. [7.10])$:

$$\mathbf{M} \mathbf{k} = \mathbf{b}^{*'}, \quad (7.25)$$

where $\mathbf{k} = \{k_{ij}\}$ and the ij row of \mathbf{M} is given by $(\mathbf{D}_{ij} \mathbf{b}^*)^T$. Thus, given \mathbf{b}^* and $\mathbf{b}^{*'}$, one calculates \mathbf{M} using equation (7.24) and then simply solves equation (7.25) for \mathbf{k} .

The problem with this method is that $\mathbf{b}^{*'}$ will not usually be completely specified. Neither equation (7.10) nor equation (7.11) fully constrains $\mathbf{b}^{*'}$. The easiest way to deal with this is to set $k_{pq} = 0$ for each unspecified $b_{pq}^{*'}$. Then the kernel will be as simple as possible while still satisfying all the requirements for $\mathbf{b}^{*'}$. A more sophisticated technique for dealing with this issue is described below in § 7.4.

7.3.2. The Kernel with Pixelization

The same general technique applies to the case of finite pixels in that we want to solve equation (7.25) for \mathbf{k} . The only difference is the calculation of $\mathbf{D}_{ij} \mathbf{b}^*$. In this case, we must use equations (7.18)–(7.22) rather than equation (7.17).

Consider the version of D_{10} given in equation (7.18):

$$\begin{aligned} \mathbf{D}_{10} \mathbf{b}^*(x, y) &= \frac{1}{2} [\mathbf{b}^*(x+1, y) - \mathbf{b}^*(x-1, y)] \\ &\quad + \frac{1}{2} i [\mathbf{b}^*(x, y+1) - \mathbf{b}^*(x, y-1)] \end{aligned} \quad (7.26)$$

$$= \frac{1}{2} (\mathbf{T}_{z_1} \mathbf{b}^* - \mathbf{T}_{-z_1} \mathbf{b}^*) + \frac{1}{2} i (\mathbf{T}_{z_2} \mathbf{b}^* - \mathbf{T}_{-z_2} \mathbf{b}^*), \quad (7.27)$$

where $z_1 = 1/\sigma$, $z_2 = i/\sigma$, and \mathbf{T}_z is defined in § A1.

Calculations for the other \mathbf{D}_{ij} are similar.

7.3.3. The Kernel in a Distorted Frame of Reference

Most telescopes, especially those with large fields of view, have fairly significant distortion. To deal with this correctly, the shape measurements should be made in the undistorted world coordinates rather than the chip coordinates. The kernel's pixel grid, however, is still in the distorted coordinate system. The calculation of $\mathbf{D}_{ij} \mathbf{b}^*$ must therefore take into account the two different coordinate systems.

If the world coordinates are (u, v) and the pixel coordinates are (x, y) then the correct values for z_1 and z_2 are

$$z_1 = \frac{1}{\sigma} \left(\frac{\partial u}{\partial x} + i \frac{\partial v}{\partial x} \right), \quad z_2 = \frac{1}{\sigma} \left(\frac{\partial u}{\partial y} + i \frac{\partial v}{\partial y} \right). \quad (7.28)$$

7.4. Minimizing the Noise from the Kernel Convolution

The motivation for having a compact kernel is to minimize the noise added by the convolution. A large kernel will use data with significant noise but little signal, adding to the noise in the convolved image.

Therefore, let us consider the noise in the convolved image. Define $K(m, n)$ to be the total convolution mask:

$$K(m, n) = \sum_{ij} k_{ij} D_{ij}(m, n), \quad (7.29)$$

$$I'(x, y) = \sum_{mn} K(m, n) I^o(x+m, y+n). \quad (7.30)$$

If the noise in the image is dominated by sky noise, then we can define n^2 to be the variance in each pixel of the observed image. The noise in each pixel of the convolved image is

$$n_f^2 = \sum_{mn} K(m, n)^2 n^2 = n^2 \sum_{mn} \left[\sum_{ij} k_{ij} D_{ij}(m, n) \right]^2. \quad (7.31)$$

In § 7.3.1, we set some kernel coefficients to zero to account for the unspecified components of $\mathbf{b}^{*'}$. Really, each of these coefficients is arbitrary, so the solution to equation (7.25) will be

$$\mathbf{k} = \mathbf{k}_0 + \mathbf{A} \mathbf{Y}, \quad (7.32)$$

where \mathbf{k}_0 is a specific solution, \mathbf{A} is a matrix, and \mathbf{Y} is an arbitrary vector.

trary vector. The dimension of \mathbf{Y} is equal to the number of unspecified components of \mathbf{b}^{**} .

We can then find the particular \mathbf{Y} that minimizes the noise added to the image by the convolution. For this derivation $\mathbf{D}(m, n)$ is the vector of each $D_{ij}(m, n)$. So, $\mathbf{K}(m, n) = \mathbf{D}(m, n)\mathbf{k}$:

$$\begin{aligned} 0 &= \frac{\partial n_r^2}{\partial \mathbf{Y}} \\ &= \frac{\partial}{\partial \mathbf{Y}} \left\{ n^2 \sum_{mn} [\mathbf{D}(m, n)\mathbf{k}]^2 \right\} \\ &= 2n^2 \sum_{mn} [\mathbf{D}(m, n)\mathbf{k}][\mathbf{D}(m, n)\mathbf{A}], \end{aligned} \quad (7.33)$$

$$0 = \sum_{mn} \{ \mathbf{D}(m, n)(\mathbf{k}_0 + \mathbf{A}\mathbf{Y})[\mathbf{D}(m, n)\mathbf{A}] \}, \quad (7.34)$$

$$\begin{aligned} \sum_{mn} [\mathbf{D}(m, n)\mathbf{A}] [\mathbf{D}(m, n)\mathbf{A}\mathbf{Y}] \\ = - \sum_{mn} [\mathbf{D}(m, n)\mathbf{A}][\mathbf{D}(m, n)\mathbf{k}_0]. \end{aligned} \quad (7.35)$$

This is now a matrix equation that can be solved for \mathbf{Y} , which then gives the solution for \mathbf{k} that minimizes the noise.

It turns out that if this method is implemented exactly as described, one ends up with a fairly large kernel, which is essentially a smoothing filter. While this solution will work, it is not exactly what we want from the kernel. We would rather have a smaller kernel that minimally changes the image. So, rather than minimize the noise in I^f , we minimize the noise in $I^f - I^o$. In other words, we leave out the D_{00} term in the vector products of equation (7.35). This results in fairly compact kernels that vary smoothly across the image.

7.4.1. Additional Kernel Components

There are nine free parameters in a general real 3×3 kernel. Equations (7.18)–(7.22) define five kernels. The identity kernel is another. Thus, there are three more independent 3×3 kernels we could construct. Two of these can be made to approximate discrete versions of the derivatives D_{10} and D_{01} . The other can be made to approximate a discrete version of D_{11} . These alternate versions are

$$\text{Alt } D_{10} = \begin{pmatrix} \frac{1}{4}(-1+i) & 0 & \frac{1}{4}(1+i) \\ 0 & 0 & 0 \\ \frac{1}{4}(-1-i) & 0 & \frac{1}{4}(1-i) \end{pmatrix}, \quad (7.36)$$

$$\text{Alt } D_{01} = \overline{\text{Alt } D_{10}}, \quad (7.37)$$

$$\text{Alt } D_{11} = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & -2 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} \end{pmatrix}. \quad (7.38)$$

We can construct alternate higher order kernels in the same way as we constructed the regular higher order kernels, namely, $\text{Alt } D_{30} = \text{Alt } D_{10} \circ D_{20}$, $\text{Alt } D_{41} = \text{Alt } D_{11} \circ D_{30}$, etc.

These extra kernel components are useful if one is minimizing the noise as described above, since they add extra degrees of freedom for the minimization and therefore can result in a smaller, less noisy kernel.

7.5. Interpolating Across an Image

The above discussion explains how to find the appropriate kernel given a particular PSF. However, in a real image, the PSF varies across the chip. Thus, the kernel will also vary across the chip.

There are two potential ways to deal with this. One can find the appropriate kernel for each star in the image. Then fit the kernel components k_{ij} as functions of (x, y) . Alternatively, one can fit the coefficients b_{pq}^* of the measured PSF decomposition as functions of (x, y) and then solve for the appropriate kernel at each point.

We choose the first method in our analysis for two reasons. First, the derived kernel can be directly applied to each star to make sure that it really does make the star round. Occasionally, a star will have significant high-order components due to crowding or an uncorrected cosmic ray. When this happens, the derivation above fails, and we reject this kernel from the fit. It is a little cleaner to recognize these outliers with kernel interpolation than with PSF interpolation.

The second reason to prefer fitting the kernel rather than the PSF is computational efficiency. The appropriate kernel must be calculated at each pixel. It is significantly faster to evaluate a function than to solve a matrix equation. On a $2K \times 4K$ chip, there are 8 million kernel evaluations. Both methods gain significantly by using a locally linear approximation to the spatial variation, but there is still a significant difference in computation for the two methods.

The kernel scheme described here allows a substantial speed gain over a typical convolution method. Fourier methods are fastest for large convolutions, but they are not practical for spatially varying kernels. In our scheme, the output image can be written as

$$I^f = \sum_{ij} k_{ij}(x, y) * (D_{ij}I^o). \quad (7.39)$$

Instead of calculating the kernel at each output pixel, we can instead produce the images $D_{ij}I^o$ (which are constant-kernel convolutions) and then accumulate sums of these images with spatially varying weights $k_{ij}(x, y)$. More importantly, the $D_{ij}I^o$ can be produced by recursive application of the 3×3 kernels. Hence, convolution by a 7×7 kernel can, for example, be reduced to three successive applications of 3×3 kernels.

Note that interpolation of the PSF elements across the image is required for the analytic deconvolution described in § 6.3.5.

7.6. Dilution Correction

Once a kernel has been applied to the image to symmetrize the PSF, the measured galaxy ellipticities need to be scaled by some resolution factor R to account for the PSF dilution. Appendix C describes the scheme we use for estimating R , which closely resembles the KSB “smear polarization” derivation.

8. CENTROID AND SELECTION BIASES

Even if the asymmetries of the PSF have been perfectly removed by a deconvolution or other method, there are two effects that can cause the estimate of the mean shape to be biased in the direction of the original PSF orientation. The first is a *selection bias*, first noted by K00. The second is a

measurement bias that arises in the presence of anisotropic noise or anisotropic PSF, which we term the *centroid bias*. K00 discusses an effect whereby the errors in centroid appear in second order as biases in the ellipticity and concludes that such errors are probably negligible. We show below that this bias is significant and is, in fact, just one of a class of noise rectification biases that can occur.

8.1. The Selection Bias

Kaiser’s selection bias operates as follows in the presence of an anisotropic PSF: if the PSF is elongated in the x -direction ($e_+^* > 0$), then objects with intrinsic shape $e_+ < 0$ cover a larger area after PSF convolution than do objects with intrinsic $e_+ > 0$. On the observed image, therefore, such objects have both lower surface brightness and lower significance ν . As most detection algorithms involve some cut in surface brightness or ν , the detected population will be biased toward $e_+ > 0$. The mean e_+ of the population will be biased, therefore, even if all the detected objects are perfectly corrected for the PSF anisotropy. K00 demonstrates that the bias will scale roughly as $e^* \sigma_x^2 / \sigma_y^2 \nu^2$, where ν is the detected significance.

The selection bias may be defeated by careful definition of the sample of target galaxies. The key is to produce some significance statistic $\tilde{\nu}$ that is independent of the shape of the object. For an image of flux f covering area A in an image with white-noise density n , the significance is normally $\nu \propto f/(nA)^{1/2}$. But the observed area A is shape-dependent in an anisotropic fashion if the PSF is elliptical. We may instead define $\tilde{\nu} \propto f/(n\tilde{A})^{1/2}$, where \tilde{A} is the object’s area on a version of the image that has been deconvolved, or at least had the PSF rounded by a convolution kernel, as described in the previous section. We then define the target sample by some cutoff limit $\tilde{\nu}_{\min}$ to this “isotropic” significance.

Detection algorithms generally have some statistic that is used as a cutoff. For KSB’s IMCAT, this is ν . For SExtractor and FOCAS, this is the number of pixels N above some isophotal threshold in a filtered version of the image. To eliminate the selection bias, we make a scatter plot of the detection statistic, e.g., N , versus the “isotropic” significance $\tilde{\nu}$. The selection bias is eliminated if we choose our cutoff $\tilde{\nu}_0$ sufficiently high that no members of the selected population approach the cutoff N_{\min} of the detection statistic. This is illustrated in Figure 7. In this way we ensure that our selected population is free of any anisotropic selection criterion that may have been created by the original detection algorithm.

8.2. Centroid Bias

The centroid bias, discussed in K00, can be qualitatively explained as follows: An error δx in centroid determination along the x -axis inflates the measured second moment I_{xx} by some amount $\propto \delta x^2$, whereas a y -axis error inflates I_{yy} . If the centroid errors are isotropic, the mean effect upon e_+ is zero. But if the x centroid errors $\langle \delta x^2 \rangle$ exceed those in y , there is a net tendency to measure a positive e_+ . If the PSF has $e_+^* > 0$, then the PSF—and consequently the mean galaxy image—is more extended in the x -direction, and hence centroids are less accurate in x than in y . The noise in centroid estimation therefore pushes the mean measured e in the direction of e^* . The centroid errors scale as ν^{-1} , and therefore the centroid bias in e scales roughly as $e^* \nu^{-2}$. Faint galaxies at low ν have a much larger centroid bias than the

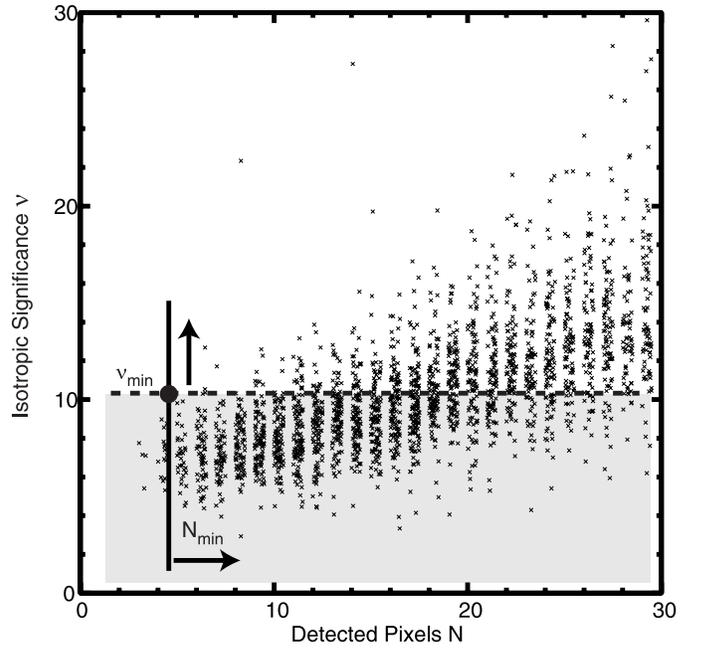


Fig. 7.—The scatter plot of isotropic significance $\tilde{\nu}$ vs. pixels above the SExtractor thresholds, N_{pix} , indicates the criteria necessary to avoid the selection bias. To avoid selection bias, we want the galaxy selection to be entirely on the basis of the shape-independent statistic $\tilde{\nu}$, and not influenced by the potentially shape-biased SExtractor threshold $N_{\min} = 5$ (some galaxies with $N < 5$ are present as a result of object splitting). We must therefore set our $\tilde{\nu}_{\min}$ threshold at ≈ 10 , which is the lowest value for which the galaxy population does not extend to $N < N_{\min}$. Galaxies in the shaded area are therefore excluded.

bright, high- ν stars used as PSF templates, so the effect is not properly removed by the KSB “smear polarizability” corrections.

Deconvolving the image eliminates the PSF anisotropy but does not eliminate the centroid bias, because the deconvolved image will have anisotropic noise. There will be more noise power along k_x than k_y after deconvolution, and hence there will still be anisotropic centroid errors and a bias toward the original e^* .

The same situation arises if we apply a convolution kernel to symmetrize the PSF. This kernel will smooth the image slightly in the direction perpendicular to the PSF, reducing the noise level in that direction somewhat. The convolved image will therefore again have a higher centroid error along the original PSF axis, as well as a consequent bias in e .

A quantitative understanding of the centroid bias can be gained from equation (6.53), the second-order expression for η in the simultaneous solution for shape and centroid (and possibly size). Let us presume that each b_{pq} has a true value plus some measurement noise δb_{pq} . The measurement noise has mean value of zero, since b_{pq} is a linear function of the intensity. If the object is intrinsically round, centered, and the weight is properly sized, then we have $b_{00} = b_{10} = b_{20} = 0$, and we would indeed measure $\eta = 0$. In the presence of noise, however, the second-order terms in equation (6.53)—for instance, the term $\propto b_{10}^2/b_{00}^2$ in the second line—may have nonzero mean values. If $\langle b_{10}^2 \rangle = \text{Var}(b_{10}) \neq 0$, there will be a nonzero $\langle \eta \rangle$.

Note that there are a number of second-order terms in the η solution, arising not just from the solution for object center, but also from the solutions for object size and ellipticity. We apply the term “centroid bias” to all these effects.

According to equation (6.17), $\text{Var}(b_{10}) = 0$ in the presence of a white-noise spectrum $P(k) = n$ (recall that b_{10} is a complex number). If we have a PSF of size σ_* and ellipticity e^* along the x -axis, we may attempt to round out the kernel by smoothing along the y -axis a little bit. This will produce a noise power spectrum $P(k) \approx n(1 - 2e^*\sigma_*^2 k_y^2)$. The k_y^2 term in the noise power produces a nonzero value for $\text{Var}(b_{10})$, and also nonzero covariances between all the other second-order elements of equation (6.53). The leading term in the expression for centroid bias takes the form

$$\langle e \rangle = K \frac{e^* \sigma_*^2}{\nu^2 \sigma_o^2} \approx \frac{K e^* (1 - R)}{\nu^2}, \quad (8.1)$$

where K is a constant, e_o is the ellipticity measured on the convolved image, σ_o is the size observed on this image, and R is the resolution factor of § 6. Note that the bias is in the measured shape, before correction for PSF dilution, which would add a factor R in the denominator. The value of K depends upon some higher order moments of the typical galaxy, and upon the details of the PSF correction and measurement procedure—in some cases K can be negative. Note also that the functional dependence of the centroid bias is essentially the same as that for the selection bias, and hence they are difficult to distinguish.

Figure 8 demonstrates the existence of the centroid bias in a very simple numerical simulation, in which we convolve a circular Gaussian source with an elliptical Gaussian weight, then measure the ellipticity and centroid with a fixed-size circular Gaussian weight. It is clear that the mean measured ellipticity depends upon the significance (i.e., S/N level) as described by equation (8.1) with $K = -2$. If the weight size is iterated to maximize significance, the bias increases to $K \approx -6$. For other measurement algorithms or galaxy shapes, K will differ, so we must determine K empirically.

The bias in equation (8.1) will be equal to the shape-noise uncertainty in the mean of N galaxies when

$$\frac{K e^* (1 - R)}{\nu^2} \approx \frac{0.3R}{\sqrt{N}} \Rightarrow N \approx \left(\frac{0.3}{K e^*} \frac{R}{1 - R} \right)^2 \nu^4. \quad (8.2)$$

For typical PSF ellipticities and K -values, with galaxies somewhat resolved ($R \approx 0.5$) and detected at $\nu = 10$, we find the bias equal to the shape noise when $N \sim 10^4$. The effect clearly cannot be ignored in present-day surveys. If the PSF ellipticity varies, then the bias will induce false power into the distortion power spectrum, with a total fluctuation power of $\approx \langle (e^*)^2 \rangle K^2 / \nu^4$, for $R \approx 0.5$. If we take the rms value of e^* to be about 3%, take $|K| \approx 5$, and $\nu = 8$, we see that the bias power is about 5% of the typical cosmic signal, for which the rms distortion is $\approx 1\%$. In unfavorable parts of the power spectrum, the ratio of PSF bias power to cosmic signal will be worse; hence, the bias cannot be ignored for cosmic shear studies that aim to move beyond mere detection to true precision measurements.

There are several possible strategies for defeating the centroid bias. The simplest is to find K empirically and then apply an appropriate bias correction to each object using equation (8.1). Another approach, useful if one has applied a PSF-rounding kernel to the image, is to add noise back to the image to recreate the original isotropic noise spectrum. If both the PSF and the noise spectrum are isotropic, there is no way the mean shape can be biased. Unfortunately, creating the appropriate noise field to add to the image is com-

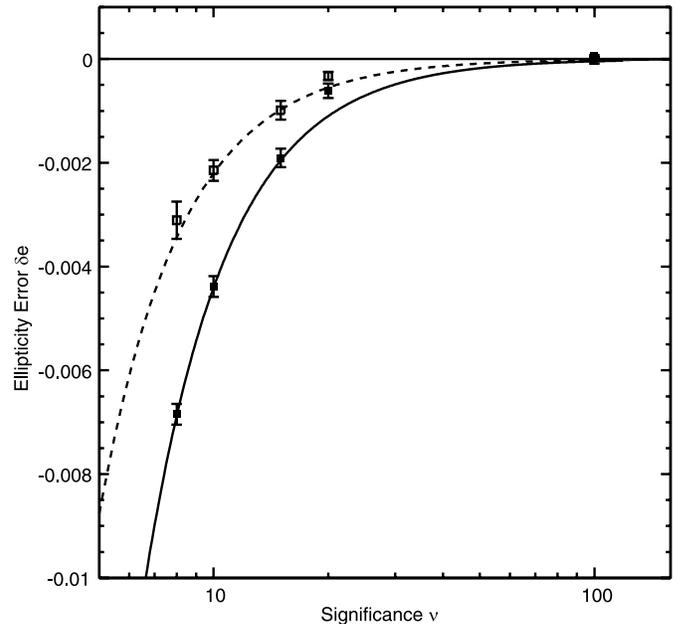


Fig. 8.—Effect of centroid bias, demonstrated by a simple numerical test in which a circular Gaussian is measured in the presence of an elliptical PSF and white noise. High-significance detections are observed at the correct ellipticity and would yield $e = 0$ when corrected for the PSF ellipticity. But the measured ellipticity drops as the significance ν decreases. The test results are well described by eq. (8.1) with $K = -2$, which is plotted as the solid [dashed] line for the tests with $e^*(1 - R) = 0.2$ [0.1].

putationally expensive for all but the simplest convolution kernels.

Symmetrizing the noise is easier if we have used the Laguerre decomposition method to deconvolve the images. If we have properly propagated the original covariance matrix for the b_{pq} , then we know $\text{Var}(b_{10})$ —and all the other relevant covariances—in the deconvolved image. We can add noise to the deconvolved b_{pq} elements in order to zero out the asymmetric elements of the covariance matrix.

9. PROCEDURES FOR SHAPE MEASUREMENT WITH PSF CORRECTION

We now have all the necessary tools for a procedure for measuring shapes and shear in the presence of a PSF convolution. Figure 9 is a flowchart for this procedure, and we elaborate on each step below. First note that there are two branch points. The first is the decision whether to measure the shapes on a summed image or to measure PSF-corrected moment information from each exposure and combine the moments. The former is easier if all exposures cover the same sky area, but the latter is recommended whenever the coverage is interlaced on the sky. The second branch point is deciding whether to use an analytic or a kernel correction for PSF bias. The remainder of this section delineates the overall data reduction procedure for weak-lensing measurements.

9.1. Exposure Processing

The steps listed in this section must be performed on each exposure. For mosaic detectors, these steps should be done separately for each channel of the detector if the PSF may change abruptly at the boundaries between detectors.

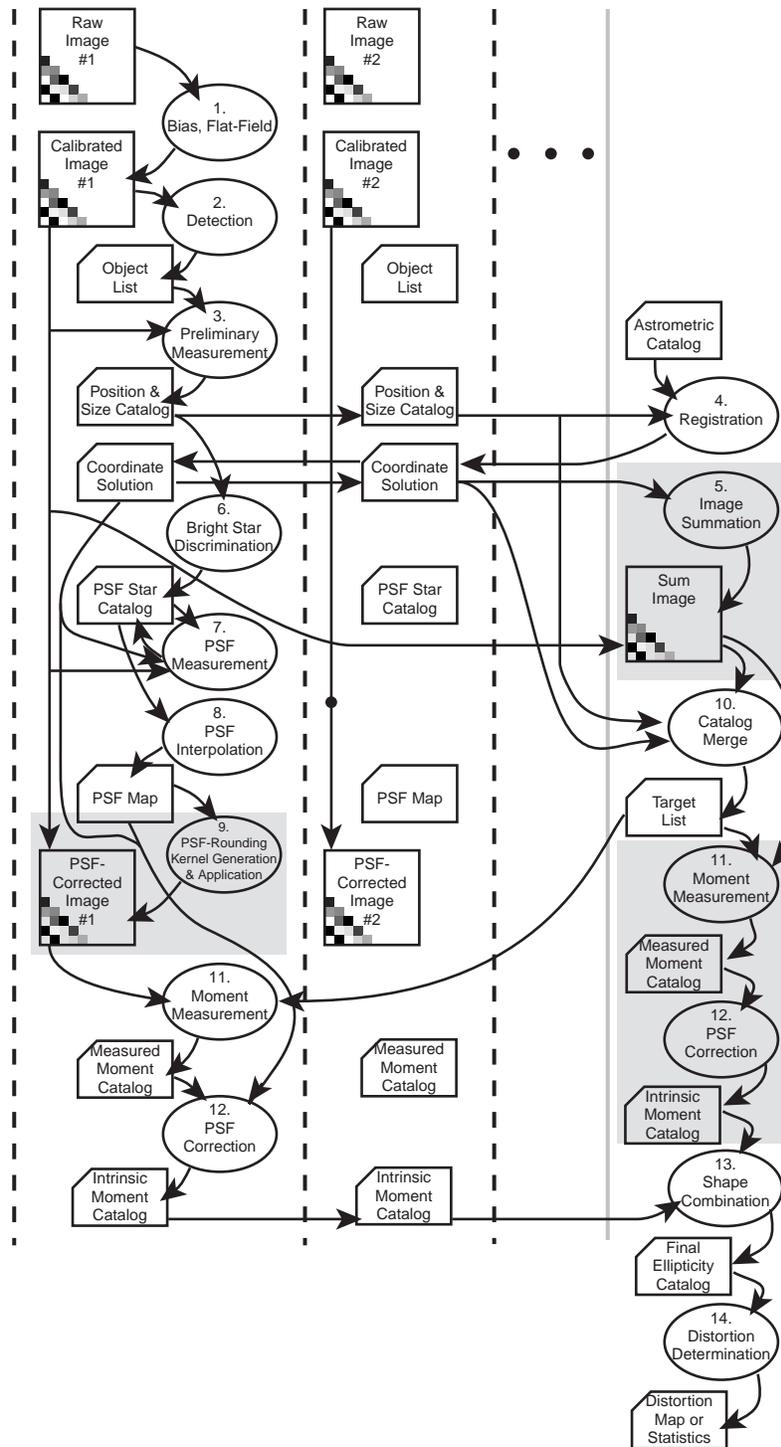


FIG. 9.—Illustration of the data flow from raw images through a final lensing distortion map or statistic, with each step described in § 9. Shaded regions represent steps that may not be used in all circumstances. Implementation details will be presented in Jarvis et al. (2002).

1. *Bias subtraction and flat-fielding.*—These can be done in the usual fashion. Note that field distortion leads to significant variation in pixel area in many CCD cameras. If the flat-field image is, as usual, obtained by exposing a source of uniform brightness, then the flattened pixel data are a properly calibrated map of the *surface brightness* of the sky rather than a map of the *flux* from the sky. This is actually

what we want, though it means that simple aperture photometry will lead to incorrect flux estimates.

2. *Object detection.*—Available detection packages such as FOCAS (Valdes 1982), SExtractor (Bertin & Arnouts 1996), and ProFit (P. Fischer 2001, private communication) can be used to identify all objects on the exposure with significance $\nu \gtrsim 3$ and produce a catalog with preliminary

position, size, and ellipticity estimation. Each of these packages also produces a useful estimate of total magnitude, which we will preserve. The choice of object detection package is not critical, because the objects that will have useful shape information have sufficiently high significance that any decent detection algorithm will work.

3. *Preliminary measurement.*—Our shape-measurement algorithm is applied to each detected object in order to obtain an accurate Gaussian-weighted estimate of centroid, size (i.e., the σ that maximizes significance), and shape. Objects with saturated or bad pixels are flagged at this point.

4. *Registration.*—The map from pixel coordinates \mathbf{x}_p to the global celestial coordinate grid \mathbf{x} is established by fitting the measured centroid of the bright objects to a collective catalog (e.g., the average of all individual exposures' catalogs) and/or an astrometric catalog. This step is critical, as even slight misregistrations will produce a systematic coherent ellipticity in the summed images. The photometric offsets between images are determined at this point as well.

5. *Image summation (optional).*—The exposures can now be remapped to a common grid and photometric scale and summed to give a deeper image. If the flattened images are surface brightness maps as described in step 1, then a simple interpolation can be done when remapping the images—no Jacobian flux corrections are needed. Recall the caveats about image combination in § 4.1. If shapes are to be measured from the summed image, then we must at this point reconstruct the preliminary catalog by executing steps 2 and 3 on the summed image.

6. *Identification of stars.*—As many unsaturated stellar images as possible should be identified on the image for use in PSF determination. Stars are typically identified on the size-magnitude plane. This is a difficult task to execute without excessive human intervention; we want our star-finding algorithm to be flexible enough to identify stars successfully in the presence of a spatially varying PSF, but strict enough to ensure that galaxies are not counted as stars. Failure on either count will lead to an erroneous estimate of the PSF size and shape on some part of the exposure, which will lead to a false feature in the shear field. Our algorithms for star identification are described in Jarvis (2002). The algorithms are particularly suited to identifying stars in the presence of a spatially varying PSF size.

7. *PSF measurement.*—The shape-measurement algorithms are now applied to the identified stars. The integrals in equation (6.16) give \mathbf{b}^* at the location of each star. Note that the integral can readily be executed in the global coordinate system \mathbf{x} , because we know the map from pixel coordinates and because pixel data are already in terms of surface brightness I . Thus, the effects of image distortion are removed at this step.

8. *PSF interpolation.*—With the PSF \mathbf{b}^* measured at the locations of the stars, we need to fit a function $\mathbf{b}^*(\mathbf{x})$ that describes the PSF at any location on the image. We use a polynomial to describe the variation of each b_{pq}^* across the image. Note that for complex or undersampled PSFs (such as those on *HST*), interpolation of the b_{pq}^* components is much easier and more accurate than interpolating a pixel-wise representation. In any case, this step is again critical: the interpolation scheme must be flexible enough to follow PSF variation, but it must remain well behaved at the image edges and in regions where PSF stars are sparse. There must

also be some form of outlier rejection for PSF stars that have cosmic rays or near neighbors contaminating the measurement, but a true excursion of the PSF in some part of the image must not be rejected.

9. *PSF-rounding kernel (optional).*—At this point we derive and apply the kernel to remove the anisotropy from the PSF, if desired. Our method for doing so is described in § 7. After the convolution it is necessary to repeat steps 7 and 8 so as to have an updated map of the PSF.

9.2. Object Measurement

We next outline the procedures that yield a shape estimate for every object in the field.

10. *Creation of target catalog.*—A list of all targets for shape measurement must be compiled. If there are fewer than about five exposures of each object, then a master target list can be produced by taking the union of the individual exposure catalogs produced in step 2. One can use a very low detection threshold on the individual exposures and then guard against noise detections by demanding coincidence on two or more exposures. If there are more than about five exposures per object, then it will be necessary to run a detection algorithm on a summed image in order to find all the potentially useful target galaxies. The target catalog should include an estimate of the centroid (in global coordinates) and some observed size s_o of each object. To avoid Kaiser's selection bias (§ 8), the criterion for acceptance into the catalog must be some shape-independent statistic.

The following steps are performed for each object in the target catalog:

11. *Measurement of observed moments.*—Given the global coordinates of the object, one can determine all the exposures on which it should appear. Note that for low-significance objects, this may include exposures for which the object was missed in the preliminary detection of step 2. We then use equation (6.16) to measure the Laguerre expansion \mathbf{b}^o of the image as observed on each exposure. Once again the integrals are performed in the global coordinate system to remove the effects of optical distortions. The integrals require choice of a centroid and a size parameter σ_o for the basis functions ψ_{pq}^o . We may use the approximate centroid from the target catalog and set the weight size σ_o equal to the typical observed object size s_o from the target catalog. The covariance matrix for the \mathbf{b}^o vector has the diagonal form given in equation (6.17) for sky-dominated galaxies. Measurements contaminated by saturated pixels or other defects may be rejected at this step.

12. *Correction for PSF effects.*—There are two methods available here. The first sequence may be used if the rounding kernel has been applied to remove the PSF bias. In this case, if the significance on the individual exposure is $\gtrsim 3$, we may proceed as follows:

a) *Shape measurement.*—The image can be shifted, dilated, and sheared until it passes the centroid, significance, and roundness criteria embodied by equations (6.19)–(6.21). This yields an optimal measure of the observed ellipticity e^o . If the object is of significance $\nu_0 \gtrsim 3$ on the individual exposure, then we can solve for centroid and size on each exposure independently and still ignore the higher order terms in the uncertainty of equation (4.1).

b) *Centroid bias correction.*—Each measured ellipticity must be corrected for the centroid bias (§ 8) using an empirical value of K in equation (8.1).

c) *Dilution correction.*—The PSF is interpolated to the position of this object to determine the resolution parameter R (eq. [6.1]) or the higher order version derived in Appendix C. The correction for dilution to give the image-plane (preseeing) ellipticity is then simply $e^i = e^o/R$. The uncertainty ellipse for e (eq. [3.28]) is therefore also scaled by $1/R$. Note that it is possible to obtain $e^i > 1$ if the noise makes the target appear smaller than the PSF.

d) *Averaging of exposures.*—The e^i from the collection of exposures are averaged using the weighting procedures of § 4.2. Some form of outlier rejection is necessary to remove objects contaminated by cosmic rays or other defects.

If we are to perform the PSF bias correction analytically, or if the significance per exposure is low, then it is best to average the deconvolved moments rather than deriving e^i for each exposure:

a) *Deconvolution of Moments.*—The PSF decomposition \mathbf{b}^* is interpolated to the position of the object, and the form of the convolution matrix $\mathbf{C}(\mathbf{b}^*)$ is calculated using the coefficients in Appendix B. This matrix is truncated at some order $p + q \leq N$ and inverted; the deconvolved (preseeing) Laguerre decomposition is $\mathbf{b}^i = \mathbf{C}^{-1}\mathbf{b}^o$. Since this is a linear operation on \mathbf{b} , the covariance matrix for \mathbf{b}^i can be propagated from the simple diagonal covariance matrix for \mathbf{b}^o (eq. [6.17]). There is a subtlety involved in the choice of weight scale σ_o ; In § 6.2, we determined that the ideal weight scale for the *deconvolved* image is $\sigma_i^2 = s_i^2 + \sigma_*^2 = s_o^2$, where s_i and s_o are the sizes of the pre- and postseeing objects and σ_* is the size of the PSF. The typical value of s_o was placed in the target catalog in step 1. According to equation (6.63), our deconvolution formulae are simplest if we choose the weight scale by

$$\sigma_o^2 = \sigma_i^2 + \sigma_*^2 = s_o^2 + \sigma_*^2. \quad (9.1)$$

Thus in step 1 we in fact want to use a weight scale somewhat larger than the s_o that maximizes the significance of detection.

b) *Combination of moments.*—From each exposure we have estimated a \mathbf{b}^i deconvolved moment vector, with known covariance matrix. We average these vectors to obtain a single best estimate of \mathbf{b}^i . We have a choice of weights to apply in producing the average; an obvious choice is to weight each exposure inversely with $\text{Cov}(b_{20}^i, \bar{b}_{20}^i)$, since b_{20} carries most of the shape information. Again, some form of outlier rejection is necessary at this step.

c) *Symmetrization of noise.*—Certain elements of the propagated covariance matrix for the \mathbf{b} -vector must be zero in order to avoid noise anisotropies that produce centroid bias (§ 8). This can be achieved by selectively adding Gaussian noise to various elements of the combined moment vector.

d) *Determination of shape.*—Given the average of \mathbf{b}^i from all exposures, we find the translation and shear that must be applied to satisfy the centroid and roundness conditions $b_{10} = b_{20} = 0$. The formulae of Appendix A are used for this. The covariance matrix for \mathbf{b}^i can be propagated through the transformations as well. The uncertainty in the shape $\boldsymbol{\eta}$ is then the square root of $\text{Var}(b_{20}\bar{b}_{20})/(b_{00}^2)$. Note

that we do not want to maximize the significance by dilating to set $b_{11} = 0$. Our optimization criterion is that the shape have minimal error after our transformations. One could dilate the image to satisfy this desire, but our choice of σ_i should already have us close to the optimum, according to the arguments of § 6.2.

After either of these procedures, we have a measure of the deconvolved shape along with its uncertainty.

13. *Combination of different wavelengths.*—If we have imaged the field in a variety of filters, then we will have obtained a shape measurement in each filter. The galaxy's shape and moments may depend upon wavelength, so we do not want to average together images or moments measured in different filters. We can, however, use the methods of § 4.2 to produce a wavelength-averaged e^i that is maximally sensitive to shear. Weighting each filter by the error in its measured shape ensures that we obtain the most sensitivity from each galaxy regardless of its color.

After completion of all these steps, we have a catalog of all objects in the field, specifying their location, magnitude, optimal shape measurement, and shape uncertainty.

9.3. Determination of Shear

With the shape catalog in hand, we are close to the scientific goals. The remaining step is the following:

14. *Generation of shear data.*—The target galaxies are binned by position, etc., into subsets for which we wish to determine a shear. The shapes (and their uncertainties) may have to be rotated, e.g., into tangential coordinates about some mass center, depending upon the shear statistic under study. The formulae of § 5 take the collection of shapes and uncertainties and allow us to create an optimal shear estimate, as well as to propagate the uncertainties in shape to the shear measurement.

10. CONCLUSIONS

We have attempted to produce, as rigorously as possible, an end-to-end methodology for measuring gravitational distortion that has optimally low noise, with calibration and noise levels derivable entirely from the observations themselves. We have succeeded in many, but not all, aspects of the problem:

1. The measurement of individual galaxy shapes appears to be optimal and has traceable noise characteristics; indeed, there even appears to be a straightforward way to handle undersampled data and retain proper covariance information for the intrinsic Laguerre coefficients b_{pq}^i (§ 6.4). The Gaussian weights underlying the Laguerre decomposition are nearly optimal for sky-dominated exponential-profile galaxies. We may wish to reexamine this scheme for the case in which the galaxy is brighter than the sky background.

2. The correction of measured moments for the distorting and diluting effects of the PSF can be effected to arbitrary accuracy using the Laguerre decomposition methods. There will be a trade-off between elimination of systematic errors—which pushes toward inclusion of higher order terms in the deconvolution—and the minimization of measurement noise from high-frequency terms. It is not clear

whether the Laguerre method is optimal with, for example, Airy PSFs with sharp cutoffs in k -space, but the method should be better than those yet applied.

3. We have identified methods to work around two sources of bias that arise from PSF ellipticities even in the presence of perfect deconvolution.

4. Measurements of galaxy shapes from different exposures or filter bands can be optimally combined with standard least-squares techniques, since we know the e uncertainty from each individual exposure.

5. The determination of lensing distortion from the ensemble of galaxy shapes has a straightforward, optimal, and calibratable solution in the case of no measurement noise. In the presence of measurement noise, however, it is necessary to have some knowledge of the noiseless shape distribution to get the calibration factor exactly correct. Our approximations, however, seem to suffice to obtain accuracies of 5% or better.

A detailed performance comparison of our methods with other authors' is beyond the scope of this paper. We can guess, however, that the reduction of measurement noise relative to a carefully weighted implementation of KSB will be slight, perhaps a factor of 1.5. But our methods, like those of K00 and Refregier (2001), are formally valid for any reasonable PSF, and hence we expect to have much-reduced systematic errors. Indeed, it is likely that the biases of § 8 have not been extensively tested because they were lost under the larger errors in PSF correction (as also noted by K00).

Our methods share elements with many previous proposals. An aesthetic difference is that we retain the geometric meaning of ellipticity by, in effect, using adaptively shaped weights instead of fixed circular weights. This makes the “shear polarizability” a purely geometric effect. As a consequence we can examine the $P(e)$ distributions and find ways to exploit surface brightness, b_{22} , or color information to separate the spheroid galaxies and weight them more heavily to reduce the shape noise.

It will be of interest to see how our method compares with the commutator method of K00. One would hope that the two independent methods could be applied to the same data set and yield the same results, bolstering our confidence in these very difficult measurements.

In a succeeding paper (Jarvis et al. 2002), we will present some of the implementation details for the analytical methods here and test the methods on real and simulated data. Fischer et al. (2000), Smith et al. (2001), and Wittman et al. (2000) make use of portions of this methodology, so the systematic-error tests and calibration tests in those papers already serve as demonstrations. Upcoming precision measurements of cosmic shear will make even greater demands upon the systematic-error reduction and accurate calibration that our methods offer.

This work was supported by grant AST 96-24592 from the National Science Foundation. We would like to thank Phil Fischer, Deano Smith, Tony Tyson, Jordi Miralda-Escudé, and Dave Wittman for many discussions on these methodologies, and help in implementing various image processing algorithms. Thanks also to all our collaborators who have waited several years for a coherent explanation of the methods we are using.

APPENDIX A

POINT TRANSFORMATION FOR LAGUERRE EXPANSIONS

In § 6.3.2, we derived the mapping matrices for the vector \mathbf{b} of Laguerre expansion coefficients when the underlying image I is transformed by an infinitesimal translation, dilation, or shear. In this appendix, we derive the coefficients of the mapping matrices for finite transformations. We have implemented these transformations as methods for C++ classes that represent the Laguerre expansions.

A1. TRANSLATION

We defined the transformation \mathbf{T}_z on the image by

$$\mathbf{T}_z I(x, y) = I(x - x_0, y - y_0), \quad (\text{A1})$$

$$z \equiv (x_0 + iy_0)/\sigma. \quad (\text{A2})$$

We first define a translated raising operator $\hat{a}_p^{\sigma\dagger}$ by

$$\hat{a}_p^{\sigma\dagger}(\mathbf{T}_z I) \equiv \mathbf{T}_z(a_p^{\sigma\dagger} I) \quad (\text{A3})$$

$$\Rightarrow \hat{a}_p^{\sigma\dagger} = a_p^{\sigma\dagger} - \frac{1}{2}z. \quad (\text{A4})$$

The second line is apparent from examining the form of the raising operators in equation (6.22). We decompose the images into our eigenfunctions as $I = \sum b_{pq}\psi_{pq}^\sigma$ and $\mathbf{T}_z I = \sum b'_{pq}\psi_{pq}^\sigma$; then we can express \mathbf{T}_z as a matrix operation on \mathbf{b} :

$$\mathbf{b}' = \mathbf{T}_z \mathbf{b}, \quad (\text{A5})$$

$$b'_{p'q'} = \sum T_{p'q'}^{pq} b_{pq}, \quad (\text{A6})$$

$$\mathbf{T}_z \psi_{pq}^\sigma = \sum T_{p'q'}^{pq} \psi_{p'q'}^\sigma, \quad (\text{A7})$$

$$\Rightarrow T_{p'q'}^{pq} = \sigma^2 \int d^2x (\mathbf{T}_z \psi_{pq}^\sigma) \bar{\psi}_{p'q'}^\sigma. \quad (\text{A8})$$

A direct integration of equation (A8) yields the first coefficient

$$T_{00}^{00} = e^{-|z|^2/4}. \quad (\text{A9})$$

Applying the translated raising operator (eq. [A4]) to the definition of $T_{p'q'}^{pq}$ (eq. [A7]) yields the recursion relation

$$\sqrt{p'} T_{(p'-1)q'}^{pq} - \frac{1}{2}z T_{p'q'}^{pq} = \sqrt{p+1} T_{p'q'}^{(p+1)q}. \quad (\text{A10})$$

The same procedure using the q raising operator and the lowering operators gives the recursions

$$\sqrt{q'} T_{p'(q'-1)}^{pq} - \frac{1}{2}\bar{z} T_{p'q'}^{pq} = \sqrt{q+1} T_{p'q'}^{p(q+1)}, \quad (\text{A11})$$

$$\sqrt{p'+1} T_{(p'+1)q'}^{pq} - \frac{1}{2}\bar{z} T_{p'q'}^{pq} = \sqrt{p} T_{p'q'}^{(p-1)q}, \quad (\text{A12})$$

$$\sqrt{q'+1} T_{p'(q'+1)}^{pq} - \frac{1}{2}z T_{p'q'}^{pq} = \sqrt{q} T_{p'q'}^{p(q-1)}. \quad (\text{A13})$$

These two relations allow us to generate any $T_{p'q'}^{pq}$ recursively from T_{00}^{00} . In particular, using the first two we have

$$T_{00}^{pq} = \frac{(-1/2)^{p+q}}{\sqrt{p!q!}} z^p \bar{z}^q e^{-|z|^2/4}. \quad (\text{A14})$$

One can derive a closed-form expression for the general $T_{p'q'}^{pq}$ using further recursion, but the expression involves a double sum and is not particularly illuminating. The most

efficient algorithm for computing all the coefficients is to note that the generator for translation (eq. [6.38]) separates into p - and q -dependent components. The general matrix element can therefore be expressed as

$$T_{p'q'}^{pq} = f(p, p') \bar{f}(q, q'), \quad (\text{A15})$$

with the above recursion relations for $T_{p'q'}^{pq}$ leading to

$$f(p, 0) = \frac{(-z/2)^p}{\sqrt{p!}} e^{-|z|^2/8}, \quad (\text{A16})$$

$$f(p, p' + 1) = [\sqrt{p}f(p-1, p') + \frac{1}{2}\bar{z}f(p, p')]/\sqrt{p'+1}. \quad (\text{A17})$$

A2. DILATION

The dilation operation is defined by

$$\mathbf{D}_\mu I(x, y) = I(e^{-\mu}x, e^{-\mu}y). \quad (\text{A18})$$

As for the translation, we can define a dilated raising operator and use it to derive a recursion relation for the coefficients of \mathbf{D} :

$$\begin{aligned} \hat{a}_p^{\sigma\dagger}(\mathbf{D}_\mu I) &= \mathbf{D}_\mu(a_p^{\sigma\dagger} I), \\ \Rightarrow \hat{a}_p^{\sigma\dagger} &= \cosh \mu a_p^{\sigma\dagger} - \sinh \mu a_p^\sigma, \end{aligned} \quad (\text{A19})$$

$$\Rightarrow \sqrt{p+1} \mathbf{D}_{p'q'}^{(p+1)q} = \cosh \mu \sqrt{p'} \mathbf{D}_{(p'-1)q'}^{pq} \quad (\text{A20})$$

$$- \sinh \mu \sqrt{q'+1} \mathbf{D}_{p'(q'+1)}^{pq}. \quad (\text{A21})$$

Using the recursion operator and its q equivalent, we can generate any desired coefficient from \mathbf{D}_{pq}^{00} . With a direct integration analogous to equation (A8), we derive

$$\mathbf{D}_{pq}^{00} = e^\mu \operatorname{sech} \mu (\tanh \mu)^p \delta_{pq}. \quad (\text{A22})$$

In deriving this, we make use of the identity

$$L_q^{(m)}(\alpha x) = \sum_{k=0}^q \binom{q+m}{q-k} \alpha^k (1-\alpha)^{q-k} L_k^{(m)}(x) \quad (\text{A23})$$

(Abramowitz & Stegun 1965). In fact, with this identity one can derive any $\mathbf{D}_{p'q'}^{pq}$ by direct integration, but the closed form is a sum over k that is again not particularly useful, as the recursion relation (eq. [A21]) is a faster way to generate the coefficients.

A3. SHEAR

A shear η oriented on the x -axis gives the transformation

$$\mathbf{S}_\eta I(x, y) = I(e^{-\eta/2}x, e^{\eta/2}y). \quad (\text{A24})$$

When we define the transformed raising operator $\hat{a}_p^{\sigma\dagger}$ as for the translation and dilation, we find the operator and consequent recursion relation to be

$$\hat{a}_p^{\sigma\dagger} = \cosh(\frac{1}{2}\eta) a_p^{\sigma\dagger} - \sinh(\frac{1}{2}\eta) a_p^\sigma$$

$$\Rightarrow \sqrt{p+1} \mathbf{S}_{p'q'}^{(p+1)q} = \cosh(\frac{1}{2}\eta) \sqrt{p'} \mathbf{S}_{(p'-1)q'}^{pq} \quad (\text{A25})$$

$$- \sinh(\frac{1}{2}\eta) \sqrt{p'+1} \mathbf{S}_{(p'+1)q'}^{pq}. \quad (\text{A26})$$

The recursion again simplifies by noting that the shear generator in equation (6.33) separates into p -operands and q operands. Hence the matrix elements must be expressible as

$S_{p'q'}^{pq} = f(p, p')f(q, q')$. We can determine the function $f(p, 0)$ by direct integration of the S_{00}^{p0} matrix element, which yields (for p even)

$$S_{00}^{p0} = \frac{\sqrt{p!}}{(p/2)!} \operatorname{sech}\left(\frac{\eta}{2}\right) \left[\frac{-\tanh(\eta/2)}{2}\right]^{p/2}, \quad (\text{A27})$$

$$\Rightarrow S_{00}^{pq} = \frac{\sqrt{p!q!}}{(p/2)!(q/2)!} \operatorname{sech}\left(\frac{\eta}{2}\right) \left[\frac{-\tanh(\eta/2)}{2}\right]^{(p+q)/2}. \quad (\text{A28})$$

The coefficients vanish if p or q is odd. The recursion relation then generates any desired coefficient. For a shear oriented at position angle β , the coefficients acquire an additional phase factor $\exp[i(p' - q' - p + q)\beta]$.

APPENDIX B

CONVOLUTION OF LAGUERRE EXPANSIONS

We wish to derive the coefficients that express the convolution of two eigenfunctions as a new sum over eigenfunctions, as defined in equation (6.57). This will be easier if we work in k -space, with the convolution turned into a multiplication as in equation (6.59), with the k -space eigenfunctions given in equation (6.62).

A rapidly computable recursive formulation of the coefficients is again derivable from the raising operator. From the form of the k -space raising operator (eq. [6.61]), we see that if $\sigma_o^2 = \sigma_i^2 + \sigma_*^2$, then

$$\begin{aligned} \sigma_o \tilde{a}_p^{\sigma_o\dagger} \tilde{I}_o &= (\sigma_i \tilde{a}_p^{\sigma_i\dagger} \tilde{I}_i) \tilde{I}_* + \tilde{I}_i (\sigma_* \tilde{a}_p^{\sigma_*\dagger} \tilde{I}_*) , \\ \Rightarrow \sigma_* \sqrt{p_*+1} C_{p_o q_o}^{p_i q_i (p_*+1) q_*} &= \sigma_o \sqrt{p_o} C_{(p_o-1) q_o}^{p_i q_i p_* q_*} \quad (\text{B1}) \\ &\quad - \sigma_i \sqrt{p_i+1} C_{p_o q_o}^{(p_i+1) q_i p_* q_*}. \end{aligned}$$

$$(\text{B2})$$

An equivalent manipulation with the lowering operator yields the recursion relation

$$\sigma_o \sqrt{p_o} C_{(p_o+1) q_o}^{p_i q_i p_* q_*} = \sigma_i \sqrt{p_i} C_{p_o q_o}^{(p_i-1) q_i p_* q_*} + \sigma_* \sqrt{p_*} C_{p_o q_o}^{p_i q_i (p_*-1) q_*}. \quad (\text{B3})$$

These recursion relations and their q equivalent will allow us to derive any coefficient if we know $C_{p_o q_o}^{p_i q_i 00}$, i.e., if we can calculate the effect of multiplication by the Gaussian $\tilde{\psi}_{00}^{\sigma_o}$. This is straightforward if we recall that $m_o = m_i$ for $m_* = 0$ and use equation (A23):

$$\begin{aligned} 2\pi \tilde{\psi}_{p_i q_i}^{\sigma_i} \tilde{\psi}_{00}^{\sigma_*} &= 2\sqrt{\pi} \sum_{p_o=0}^{p_i} \sum_{q_o=0}^{q_i} \sqrt{\frac{p_i! q_i!}{p_o! q_o!}} \\ &\quad \times \frac{D^{(p_o+q_o)/2} (1-D)^{q_i-q_o}}{(q_i-q_o)!} \tilde{\psi}_{p_o q_o}^{\sigma_o}, \end{aligned} \quad (\text{B4})$$

$$\begin{aligned} \Rightarrow C_{p_o q_o}^{p_i q_i 00} &= 2\sqrt{\pi} \sqrt{\binom{p_i}{p_o} \binom{q_i}{q_o}} D^{(p_o+q_o)/2} \\ &\quad \times (1-D)^{(p_i-p_o+q_i-q_o)/2}, \end{aligned} \quad (\text{B5})$$

$$D \equiv \frac{\sigma_i^2}{\sigma_o^2} = \frac{\sigma_i^2}{\sigma_i^2 + \sigma_*^2} = 1 - \frac{\sigma_*^2}{\sigma_o^2}. \quad (\text{B6})$$

The parameter D is the ‘‘deconvolution ratio,’’ with $D = 1$ being the limit of perfect resolution and $D = 0$ meaning a

PSF much broader than the image. This expression and the recursion relations are both separable into p - and q -dependent terms, so we can simplify the computation of the matrix elements by using the expression

$$C_{p_o q_o}^{p_i q_i p_* q_*} = 2\sqrt{\pi} \left[\sqrt{\frac{p_i! p_*!}{p_o! \Delta!}} G(p_o, p_i, p_*) \right] \left[\sqrt{\frac{q_i! q_*!}{q_o! \Delta!}} G(q_o, q_i, q_*) \right], \quad (\text{B7})$$

$$\Delta \equiv p_i + p_* - p_o = q_i + q_* - q_o \geq 0. \quad (\text{B8})$$

Terms for which the conditions in equation (B8) are not met are zero. The recursion relations and specific values given above for C can be recast for the function G as follows:

$$G(p_o + 1, p_i, p_*) = \frac{\sigma_i}{\sigma_o} G(p_o, p_i - 1, p_*) + \frac{\sigma_*}{\sigma_o} G(p_o, p_i, p_* - 1), \quad (\text{B9})$$

$$G(0, p_i, p_*) = (-1)^{p_*} \binom{p_i + p_*}{p_i} \left(\frac{\sigma_i}{\sigma_o} \right)^{p_*} \left(\frac{\sigma_*}{\sigma_o} \right)^{p_i}, \quad (\text{B10})$$

The symmetry between initial and PSF images is clear in these equations. There is a consequent closed form for $G(p_o, p_i, p_*)$, but it is again not particularly illuminating, and the recursive form is stable and faster for computation.

In the case where the PSF is a unit-flux Gaussian, equation (B5) can be used to give the observed \mathbf{b}^o in terms of the intrinsic decomposition:

$$b_{p_o q_o}^o = \sum_{j=0}^{\infty} D^{(p_o + q_o)/2} \sqrt{\binom{p_o + j}{p_o} \binom{q_o + j}{q_o}} \times (1 - D)^j b_{(p_o + j)(q_o + j)}^i \quad (p_o \geq q_o). \quad (\text{B11})$$

Note that the convolution matrix \mathbf{C} is in this case block-diagonal, as the m states do not mix, and also upper triangular, as $p_o, q_o \leq p_i, q_i$ for nonzero elements. The inverse (deconvolution) matrix can in this case be expressed in closed form:

$$b_{p_i q_i}^i = \sum_{j=0}^{\infty} D^{-(p_i + q_i)/2} \sqrt{\binom{p_i + j}{p_i} \binom{q_i + j}{q_i}} \times \left(\frac{1 - D}{D} \right)^j (-1)^j b_{(p_i + j)(q_i + j)}^o. \quad (\text{B12})$$

APPENDIX C

SECOND-ORDER FORMULAE FOR PSF DILUTION

Here we describe a refinement to the resolution parameter R defined in equation (6.1). The observed ellipticity is R times the true ellipticity in the special cases of unweighted moments or when the PSF is a circular Gaussian and the galaxy an elliptical Gaussian. In these cases the factor R can be expressed as

$$R = 1 - s_*^2/s_o^2, \quad (\text{C1})$$

$$s^2 \equiv \langle (x^2 + y^2)/2 \rangle = \sigma^2 \cosh \eta. \quad (\text{C2})$$

In the second line we have assumed that the methods of § 3 have been used to shear the object by η to produce something that is “round” under a Gaussian weight of optimal size σ .

Here we derive a form for R that is applicable to the case in which the galaxy has homologous elliptical isophotes with $\eta \ll 1$, but with an arbitrary radial profile. In the language of our Laguerre coefficients, we have $b_{11} = 0$ by proper choice of σ , and b_{pp} are arbitrary for $p \geq 2$. By equation (6.34), this slightly out-of-round object has $b_{20} = \sqrt{2}\eta(b_{00} - b_{22})/4$. All odd-indexed coefficients are zero.

We also take the PSF size σ_* to be small compared with the intrinsic object size σ_i . The action of this small isotropic convolution \mathbf{C}_{σ_*} on the moments is the same as a transformation in which the image I is replaced by the average of two versions displaced by $\pm\sigma_*$ in the x -direction, followed by a similar infinitesimal spread in the y -direction. Defining $z = \sigma_*/\sigma$ and using a second-order version of the generator equation (6.38), we can show that

$$\mathbf{C}_{\sigma_*} = \frac{1}{4}(\mathbf{T}_z + \mathbf{T}_{-z})(\mathbf{T}_{iz} + \mathbf{T}_{-iz}) \quad (\text{C3})$$

$$\approx 1 + \frac{1}{2}z^2(a_p^\dagger a_q^\dagger + a_p a_q - p - q - 1). \quad (\text{C4})$$

When this convolution acts upon our original slightly elliptical object, the resulting object has Laguerre coefficients \mathbf{b}' with

$$b'_{00} = b_{00}(1 - z^2/2), \quad (\text{C5})$$

$$b'_{20} = b_{20}(1 - 3z^2/2), \quad (\text{C6})$$

$$b'_{22} = b_{22}(1 - 5z^2/2). \quad (\text{C7})$$

We then need the value η' of the shear that will make this new object appear round. According to equation (6.47), this will give

$$\eta' = \frac{\eta(b_{00} - b_{22})(1 - 3z^2/2)}{1 - b_{00}(z^2/2) - b_{22}(1 - 5z^2/2)} \quad (\text{C8})$$

$$\Rightarrow R \equiv \frac{\eta'}{\eta} = 1 - z^2 \left(1 + \frac{2b_{22}}{b_{00} - b_{22}} \right) + O(z^4) \quad (\text{C9})$$

$$\approx 1 - \frac{\sigma_*^2}{s^2}, \quad (\text{C10})$$

$$s^2 \equiv \sigma^2 \frac{b_{00} - b_{22}}{b_{00} + b_{22}}. \quad (\text{C11})$$

Note that the kurtosis measure a_4 defined in equation (3.24) is the same as the b_{22}/b_{00} that appears here. We make the *Ansatz* that the correct form for R in the case of finite dilution or finite e is the Gaussian form equation (C1) with the kurtosis term added:

$$e_i = e_o/R, \quad (\text{C12})$$

$$R = 1 - s_*^2/s_o^2, \quad (\text{C13})$$

$$s^2 \equiv \frac{1 - a_4}{1 + a_4} \sigma^2 \cosh \eta. \quad (\text{C14})$$

Note that we apply the kurtosis correction to the PSF size measure s_*^2 in the same way as for the object to give a well-behaved correction for poorly resolved objects.

REFERENCES

- Abramowitz, M., & Stegun, I. A., ed. 1965, *Handbook of Mathematical Functions* (New York: Dover)
- Bacon, D. J., Refregier, A., Clowe, D., & Ellis, R. S. 2001, *MNRAS*, 325, 1065
- Bernstein, G. 2002, *PASP*, 114, 98
- Bertin, E., & Arnouts, S. 1996, *A&AS*, 117, 393
- Bonnet, H., & Mellier, Y. 1995, *A&A*, 303, 331
- Bonnet, H., Mellier, Y., & Fort, B. 1994, *ApJ*, 427, L83
- Bracewell, R. N. 1978, *The Fourier Transform and Its Applications* (2d ed.; New York: McGraw-Hill)
- Chang, T.-C., & Refregier, A. 2002, *ApJ*, in press
- Dye, S., Taylor, A. N., Thommes, E. M., Meisenheimer, K., Wolf, C., & Peacock, J. A. 2001, *MNRAS*, 321, 685
- Erben, T., Van Waerbeke, L., Bertin, E., Mellier, Y., & Schneider, P. 2001, *A&A*, 366, 717
- Fahlman, G., Kaiser, N., Squires, G., & Woods, D. 1994, *ApJ*, 437, 56
- Fischer, P., Bernstein, G., Rhee, G., & Tyson, J. A. 1997, *AJ*, 113, 521
- Fischer, P., & Kochanski, G. P. 1994, *AJ*, 107, 802
- Fischer, P., et al. 2000, *AJ*, 120, 1198
- Fischer, P., & Tyson, J. A. 1997, *AJ*, 114, 14 (FT97)
- Hoekstra, H., Franx, M., & Kuijken, K. 2000, *ApJ*, 532, 88
- Hoekstra, H., et al. 2001, *ApJ*, 548, L5
- Hoekstra, H., Franx, M., Kuijken, K., & Squires, G. 1998, *ApJ*, 504, 636
- Jarvis, M. 2002, in preparation
- Jarvis, M., Bernstein, G., Fischer, P., & Smith, D. 2002, in preparation
- Kaiser, N. 2000, *ApJ*, 537, 555 (K00)
- Kaiser, N., & Squires, G. 1993, *ApJ*, 404, 441
- Kaiser, N., Squires, G., & Broadhurst, T. 1995, *ApJ*, 449, 460 (KSB)
- Kuijken, K. 1999, *A&A*, 352, 355
- Lauer, T. R. 1999, *PASP*, 111, 227
- Lombardi, M., & Bertin, G. 1998, *A&A*, 330, 791
- Luppino, G. A., & Kaiser, N. 1997, *ApJ*, 475, 20
- Mellier, Y. 1999, *ARA&A*, 37, 127
- Miralda-Escudé, J. 1991, *ApJ*, 380, 1
- Mould, J., Blandford, R., Villumsen, J., Brainerd, T., Smail, I., Small, T., & Kells, W. 1994, *MNRAS*, 271, 31
- Refregier, A. 2001, *MNRAS*, submitted (astro-ph/0105178)
- Refregier, A., & Bacon, D. 2001, *MNRAS*, submitted (astro-ph/0105179)
- Rhodes, J., Refregier, A., & Groth, E. J. 2000, *ApJ*, 536, 79
- Schneider, P., & Seitz, C. 1995, *A&A*, 294, 411
- Smail, I., Ellis, R. S., Fitchett, M. J., & Edge, A. C. 1995, *MNRAS*, 273, 277
- Smith, D. R., Bernstein, G. M., Fischer, P., & Jarvis, M. 2001, *ApJ*, 551, 643
- Tyson, J. A., Valdes, F., Jarvis, J. F., & Mills, A. P., Jr. 1984, *ApJ*, 281, L59
- Tyson, J. A., Valdes, F., & Wenk, R. A. 1990, *ApJ*, 349, L1
- Valdes, F. 1982, *Proc. SPIE*, 331, 465
- Valdes, F., Tyson, J. A., & Jarvis, J. F. 1983, *ApJ*, 271, 431
- Van Waerbeke, L., et al. 2000, *A&A*, 358, 30
- Wilson, G., Cole, S., & Frenk, C. S. 1996, *MNRAS*, 280, 199
- Wilson, G., Kaiser, N., & Luppino, G. A. 2001a, *ApJ*, 556, 601
- Wilson, G., Kaiser, N., Luppino, G. A., & Cowie, L. L. 2001b, *ApJ*, 555, 572
- Wittman, D. M., Tyson, J. A., Kirkman, D., Dell'Antonio, I., & Bernstein, G. 2000, *Nature*, 405, 143